

IBM Spectrum Control
Version 5.2.14

User's Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 501.

This edition applies to version 5, release 2, modification 14 of IBM Spectrum Control (product numbers 5725-F93, 5725-G33, 5725-Y23, and 5725-Y24) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC27-6588-05.

© **Copyright IBM Corporation 2001, 2017.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this guide vii

Who should read this guide.	vii
Publications	vii
Accessing publications online	vii
IBM Redbooks	ix
Translation.	ix
Providing feedback about publications.	x
IBM Spectrum Control Service Management Connect community	x
Contacting IBM Software Support	x
Reporting a problem	xi
Conventions used in this guide	xi

New for IBM Spectrum Control Version

5.2.14 xiii

Summary of changes for IBM Spectrum Control	xx
---	----

Chapter 1. Getting started 1

Managing resources	4
Configuring and administering	7
Navigation	8

Chapter 2. Videos and blogs for IBM Spectrum Control 15

Chapter 3. Product licenses. 21

Functions that are available based on role and license	23
License compliance.	29
Configuring storage systems to determine license compliance	33

Chapter 4. Product overview 37

Architecture	37
Interfaces for IBM Spectrum Control	38
Starting IBM Spectrum Control	39
Navigation	40
Replication products	45
Key concepts	47
Data collection	47
Performance monitoring and troubleshooting	47
Applications, departments, and general groups	49
Storage optimization	51
Cloud configuration	52
Provisioning storage	57
Storage reclamation.	60
IBM Spectrum Control VASA provider	62
vSphere Web Client extension for IBM Spectrum Control	63
Alerts	64
Reporting	65
Rollup servers	66
Units of measurement for storage data	69
Role-based authorization	70

Fabrics and zones	72
Agents	75
IBM Spectrum Control REST API	76

Chapter 5. Scenarios. 85

Optimizing the performance of storage virtualizers	85
Redistributing volume workloads to balance pools	86
Analyzing and re-tiering volumes in pools on tier 1	87
Collocating volumes	90
Monitoring capacity usage at different levels of a business hierarchy	91
Comparing storage usage in each department	93
Using applications and subcomponents to monitor capacity and space usage.	95
Viewing storage capacity and usage trends	96
Limiting a set of storage to provision in the vSphere Web Client extension	98
Viewing NPIV connections between server ports and switch ports in a fabric	99
Exporting performance data for a SAN Volume Controller system	100
Comparing the performance of storage systems	102
Troubleshooting performance	103
Monitoring IBM Spectrum Scale performance.	105

Chapter 6. Managing resources 107

Resources that you can monitor	107
Adding resources	112
Required user roles for monitoring resources	112
Adding storage systems	115
Adding fabrics and switches	120
Adding hypervisors	122
Adding servers.	125
Adding rollup servers	132
Removing resources	132
Collecting data	133
Collecting asset and status data	134
Collecting performance data	140
Collecting information about shares on storage systems	147
Collecting information about the sizes of GPFS snapshots.	147
Alerting and event processing.	148
Alert severities	150
Defining alerts	151
Viewing and modifying alert definitions	185
Viewing and modifying alert notification settings	187
Viewing alerts	190
Acknowledging alerts	193
Removing alerts	193
Viewing information about resources	194
How information is organized.	194

Dashboard view	197
Viewing information about top-level resources	200
Viewing information about internal, object, and related resources	201
Monitoring the status and condition of resources	203
How the condition of a resource is determined	204
Viewing the overall condition of resources.	208
Viewing the condition of specific types of resources	208
Viewing the status of resources	209
Acknowledging the condition and status of resources	210
Monitoring vaults in IBM Cloud Object Storage	212
Monitoring the performance of resources	217
Viewing performance information	219
Best practice performance guidelines	234
Viewing performance alerts.	235
Exporting performance data for storage systems and fabrics	237
Monitoring the capacity of resources	238
Viewing capacity information	240
Viewing the capacity of external storage	251
Viewing capacity alerts and violations	254
Monitoring and administering applications, departments, and general groups.	254
Applications.	261
Departments	270
Application and department hierarchies	277
Exporting information about applications and departments.	278
General groups.	279
Viewing the data paths of resources	287
Resources that you can view in a data path	288
Navigating in a data path	292
Viewing information about top-level resources	292
Viewing information about internal resources	292
Zooming in	292
Zooming out	293
Changing the layout	293
Saving the display settings	293
Printing a view.	293
Exporting a data path to a file.	294
Provisioning storage	294
Configuring IBM Spectrum Control for provisioning.	294
Provisioning storage with the IBM Spectrum Control GUI.	311
Provisioning storage with the vSphere Web Client extension	316
Optimizing storage tiering	319
Investigating the capacity of tiered storage	321
Setting the tier level of storage pools	322
Analyzing tiering by servers	322
Analyzing tiering by hypervisors.	323
Analyzing tiering by storage virtualizers	323
Analyzing tiering by storage pools	323
Analyzing tiering by volumes	323
Tiering volumes by I/O density and I/O rate	323
Modifying the criteria for analyzing tiering	324
Optimizing storage pools	325
Balancing pools	325

Modifying the criteria for balancing pools	326
Criteria for identifying the pools that require balancing.	327
Reclaiming storage	328
Viewing storage reclamation	329
Viewing storage reclamation by storage system	330
Excluding volumes from reclamation analysis	331
Transforming and migrating volumes	331
Modifying the period for analyzing performance data	332
Modifying the properties of resources	333
Opening the management GUI for a resource.	335
Opening the DS Storage Manager GUI for a DS8000	335
Exporting information to a file	336
Customizing lists of resources, tasks, and alerts	337
Filtering lists	338
Sorting lists	338
Showing, hiding, and reordering columns in lists	339
Managing tasks.	339
Viewing tasks	340
Managing tasks for provisioning	344
Managing tasks for tiering storage, balancing pools, and transforming storage	346
Renaming tasks	351
Deleting tasks	352
Planning copy data resources	352
Monitoring copy data resources	353

Chapter 7. Reporting 355

Creating chargeback reports	357
Chargeback reports	358
Creating consumer reports	359
Consumer reports	360
Creating summary reports of the storage capacity	361
Configuring the email server	362
Capacity metrics for reports	362
Capacity metrics for chargeback and consumer reports	362
Creating resource reports	365
Creating resource reports about capacity shortages in pools	366
Creating resource reports about the input/output performance of volumes	367

Appendix A. Reference 369

Alerts	369
Triggering conditions for alerts	370
Triggering actions and notifications for alerts	401
Performance data	403
Performance metrics	403
Capacity data	469
Capacity metrics for block storage systems	469
Capacity metrics for file storage systems	482
Capacity metrics for object storage systems	484
Capacity metrics for tiers	485
Cinder driver for IBM Spectrum Control	485
Planning to deploy the Cinder driver	486
Setting up the Cinder driver environment	488

Requirements for installing the Cinder driver	489
Cinder driver package contents	490
Installing the Cinder driver.	491
Uninstalling the Cinder driver.	493
Configuring the Cinder driver.	493
Learn more about the Cinder driver.	494
Limitations and known issues for the Cinder driver	494
Opening IBM Spectrum Control on Windows operating systems	495
Opening IBM Spectrum Control GUIs and CLIs	495
Accessing administration tools	497
Appendix B. Accessibility features for IBM Spectrum Control	499
Notices	501
Privacy policy considerations	503
Trademarks	503
Glossary	505
A	505

B	505
C	505
D	506
E	506
F	506
G	506
H	507
J.	507
L	507
M	507
N	507
P	507
R	508
S	508
T	509
V	509
Z	509
Index	511

About this guide

IBM Spectrum Control™ manages storage infrastructure by centralizing, automating, and simplifying the management of complex and heterogeneous storage environments. This guide provides task-oriented user information that helps you to obtain optimal product performance.

Who should read this guide

This publication is intended for administrators and users who want to learn how to use IBM Spectrum Control to monitor and manage storage resources and infrastructure.

Administrators and users must be familiar, or become familiar, with the following concepts:

- Storage Area Network (SAN)
- Simple Network Management Protocol (SNMP)
- Storage provisioning, optimization, performance, tiers, and cloud configuration
- IBM® DB2®
- Jazz™ for Service Management
- IBM Tivoli® Common Reporting
- vSphere Web Client extension and VASA

Publications

A number of publications are provided with IBM Spectrum Control.

IBM Tivoli Storage Productivity Center for Replication is no longer installed with IBM Spectrum Control. For current Tivoli Storage Productivity Center for Replication users who upgrade to IBM Spectrum Control, Tivoli Storage Productivity Center for Replication will remain available. See IBM Knowledge Center for any earlier versions of Tivoli Storage Productivity Center for Replication publications.

The following section describes how to access these publications online.

Accessing publications online

Information about installing, configuring, using, upgrading, and uninstalling IBM Spectrum Control and related products is available online.

Use Table 1 to view and download these publications. Translated documents are available for some products.

Table 1. Locations of publications for IBM Spectrum Control and related products

Product	Online location
IBM Cognos® Business Intelligence	http://www.ibm.com/support/knowledgecenter/SSEP7J
IBM Copy Services Manager	http://www.ibm.com/support/knowledgecenter/SSESK4

Table 1. Locations of publications for IBM Spectrum Control and related products (continued)

Product	Online location
IBM DB2 Database for Linux UNIX and Windows	http://www.ibm.com/support/knowledgecenter/SSEPGG/welcome
IBM FlashSystem™ 840	http://www.ibm.com/support/knowledgecenter/ST2NVR/com.ibm.storage.flashsystem.doc/Flashsystem_840_welcome.htm
IBM FlashSystem V840 Enterprise Performance Solution	http://www.ibm.com/support/knowledgecenter/ST2HTZ/
IBM FlashSystem 900	http://www.ibm.com/support/knowledgecenter/STKMQB/
IBM FlashSystem A9000	http://www.ibm.com/support/knowledgecenter/STJKMM
IBM FlashSystem A9000R	http://www.ibm.com/support/knowledgecenter/STJKN5
IBM FlashSystem V9000	http://www.ibm.com/systems/storage/flash/v9000/overview.html
IBM Scale Out Network Attached Storage (IBM SONAS)	http://www.ibm.com/support/knowledgecenter/STAV45/landing/sonas_151_kc_welcome.html
IBM Spectrum Accelerate™	http://www.ibm.com/support/knowledgecenter/STZSWD/welcome
IBM Spectrum Control	http://www.ibm.com/support/knowledgecenter/SS5R93/ In the navigation pane, click IBM Spectrum Control for the appropriate release. To obtain PDF documents, in the navigation pane, click IBM Spectrum Control > Printable documentation .
IBM Spectrum Scale™	http://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html
IBM Storwize® V3500	http://www.ibm.com/support/knowledgecenter/STLM6B
IBM Storwize V3700	http://www.ibm.com/support/knowledgecenter/STLM5A
IBM Storwize V5000	http://www.ibm.com/support/knowledgecenter/STHGJ
IBM Storwize V7000	http://www.ibm.com/support/knowledgecenter/ST3FR7
IBM Storwize V7000 Unified	http://www.ibm.com/support/knowledgecenter/ST5Q4U
IBM System Storage® DS3500	http://www.ibm.com/support/docview.wss?uid=ssg1S7004887
IBM System Storage DS3700	http://www.ibm.com/support/docview.wss?uid=ssg1S7005174
IBM System Storage DS5020	http://www.ibm.com/support/docview.wss?uid=ssg1S7004951

Table 1. Locations of publications for IBM Spectrum Control and related products (continued)

Product	Online location
IBM System Storage DS6000™	http://www.redbooks.ibm.com/abstracts/sg246781.html The DS6000 series is no longer supported. For information, see the preceding publication.
IBM System Storage DS8000®	http://www.ibm.com/support/knowledgecenter/ Go to the preceding website and search for the DS8000 series system that you want.
IBM System Storage DS® Open Application Programming Interface publications	http://www.ibm.com/support/docview.wss?uid=ssg1S7001160 Use these publications for information about how to install, configure, and use the CIM agent.
IBM System Storage N series	http://www.ibm.com/support/knowledgecenter/nseries/kc_welcome_nseries.html For more information about IBM System Storage N series, see http://www.ibm.com/systems/storage/network/software/ .
IBM System Storage SAN Volume Controller	http://www.ibm.com/support/knowledgecenter/STPVGU
IBM XIV® Storage System	http://www.ibm.com/support/knowledgecenter/STJTAG
IBM WebSphere® Application Server	http://www.ibm.com/support/knowledgecenter/SSEQTP/mapfiles/product_welcome_was.html
IBM WebSphere Application Server Liberty	https://www.ibm.com/support/knowledgecenter/SSEQTP_liberty/as_ditamaps/was900_welcome_liberty.html
Jazz for Service Management	http://www.ibm.com/support/knowledgecenter/SSEKCU/welcome
VMware vSphere	http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html

IBM Redbooks

The IBM Redbooks® are publications about specialized topics.

You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page at <http://www.redbooks.ibm.com>.

Translation

Translated publications are available from IBM Knowledge Center, which is available in certain translated languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a translated version, the information is displayed in English, which is the default language. Translations of the PDFs are available when the information is translated.

Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information.

To provide comments or suggestions about the product documentation, go to the IBM Knowledge Center for IBM Spectrum Control: <http://www.ibm.com/support/knowledgecenter/SS5R93/>. Click **Feedback** at the bottom of any page and complete the form.

IBM Spectrum Control Service Management Connect community

Connect, learn, and share with Service Management professionals: product support technical experts who provide their perspectives and expertise.

Access Service Management Connect at <https://www.ibm.com/developerworks/servicemanagement/>. Use Service Management Connect in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and developers of IBM products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about IBM and the Storage Management community.
- Read blogs to benefit from the expertise and experience of others.
- Use wikis and forums to collaborate with the broader user community.

Contacting IBM Software Support

You can contact IBM Software Support by phone, and you can register for support notifications at the technical support website.

- Go to the IBM Spectrum Control technical support website at http://www.ibm.com/support/entry/portal/product/tivoli/ibm_spectrum_control/ibm_spectrum_control_standard_edition.

To receive future support notifications, sign in under **Subscribe to support notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for IBM Spectrum Control technical support website updates.

- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to <https://www.ibm.com/planetwide/>.

You can also review the *IBM Software Support Handbook*, which is available at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

The support website offers extensive information, including a guide to support services; frequently asked questions (FAQs); and documentation for all IBM Software products, including Redbooks and white papers. Translated documents are also available for some products.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information might also be needed to access various online services available on the website. See “Reporting a problem” on page xi.

Reporting a problem

Provide the IBM Support Center with information about the problems that you report.

Have the following information ready when you report a problem:

- The IBM Spectrum Control version, release, modification, and service level number.
- The communication protocol (for example, TCP/IP), version, and release number that you are using.
- The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.
- The exact text of any error messages.

Conventions used in this guide

Information is given about the conventions that are used in this publication.

This publication uses several conventions for special terms and actions, and for operating system-dependent commands and paths.

The following typeface conventions are used in this publication:

Bold

- Flags that display with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

Italic

- Variables
- Values that you must provide
- New terms
- Words and phrases that are emphasized
- Titles of documents

monospace

- Commands and command options
- Flags that display on a separate line
- Code examples and output
- Message text
- Names of files and directories
- Text strings that you must type, when they display within text
- Names of Oracle Java™ methods and classes
- HTML and XML tags that display like this, in monospace type

For syntax notations, remember the following details.

- In AIX®, the prompt for the root user is #.
- In AIX and Linux, the commands are case-sensitive, so you must type commands exactly as they are shown.

New for IBM Spectrum Control Version 5.2.14

New features and enhancements are available in IBM Spectrum Control Version 5.2.14.

The following features and enhancements are available for all product licenses:

- “Enhancements for IBM Cloud Object Storage”
- “Configuring LDAP made easier” on page xv
- “Viewing port delay performance metrics on IBM Spectrum Virtualize” on page xvii
- “Viewing host connections that are associated with a host or cluster” on page xvii
- “Viewing capacity metrics for read-intensive flash drives on IBM Storwize” on page xix
- “Uploading logs automatically” on page xix


The “Securing the connection to email servers for sending alert notifications and reports by email” on page xvi feature and the Choosing the capacity metric that you want to show in reports feature is included in the IBM Spectrum Control Advanced Edition license and in the IBM Spectrum Control Advanced Select Edition license.

Enhancements for IBM Cloud Object Storage

Monitoring cloud data consumption by applications

You can now add COS Vaults to applications and general groups so you can track how much space is being consumed by logically related vaults and how that consumption trends over time.

For example, logically related vaults might be vaults in IBM Cloud Object Storage that are used by a Sales database within your organization. You can create an application model in IBM Spectrum Control to monitor how much capacity on the vaults is being consumed by that database.



The screenshot shows the 'Applications' page in IBM Spectrum Control. It features a table with columns for Name, Block Capacity (GiB), Block Allocated Space (GiB), File Capacity (GiB), and Object Capacity (GiB). The 'Object Capacity (GiB)' column is highlighted with a red box. The table contains one row for the 'NA Sales app'.

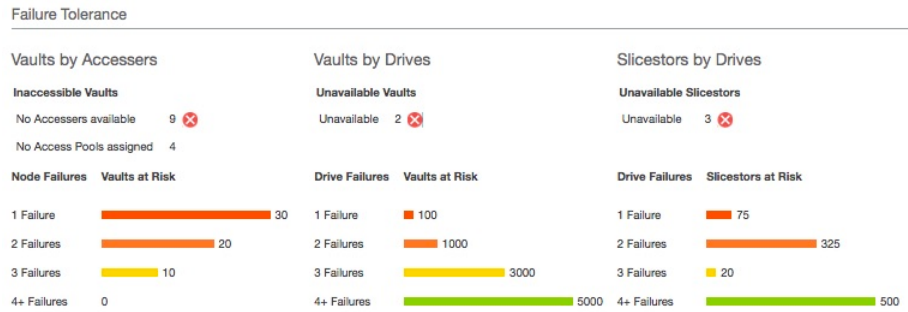
Name	Block Capacity (GiB)	Block Allocated Space (GiB)	File Capacity (GiB)	Object Capacity (GiB)
NA Sales app	0.00	0.00	14,850	40,519.68

Monitoring access risk and storage risk for COS Vaults and COS Slicestor[®] nodes

The Vaults by Accessers chart is enhanced and new charts are added to enable you to monitor the availability of all your COS Vaults and COS Slicestor nodes.

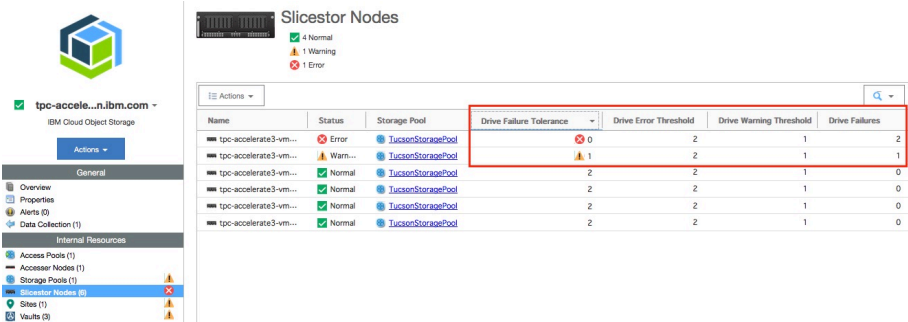
The Vaults by Drives and Slicestors by Drives charts are shown only for IBM Cloud Object Storage with firmware version 3.10.0.117 or later.

To view the charts, go to **Storage > Object Storage Systems** and double-click an IBM Cloud Object Storage system.



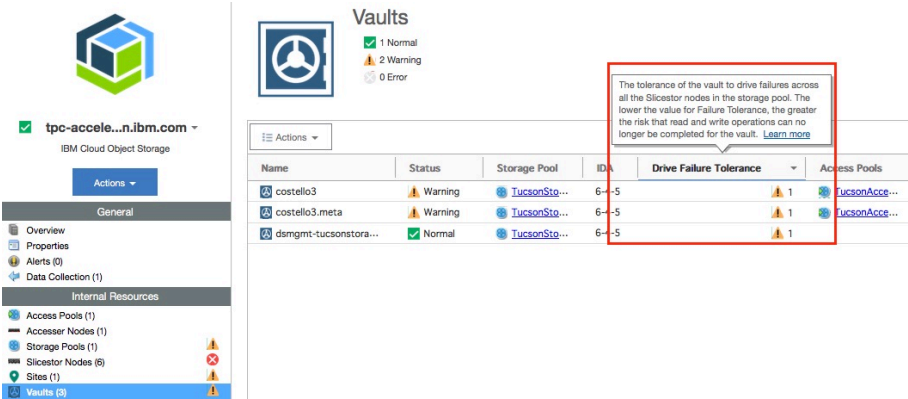
Viewing information about the tolerance of COS Slicestor nodes to drive failures

You can now view the number of failed drives in the COS Slicestor nodes and the tolerance of the nodes to more drive failures. The lower the value for drive failure tolerance, the greater the storage risk for the node. A value of 0 means that the node can no longer store vault data.



Viewing information about the tolerance of COS Vaults to drive failures

You can now view the tolerance of COS Vaults to drive failures across all the COS Slicestor nodes in the storage pool. A value of 0 for the Drive Failure Tolerance column means that write operations cannot be completed for the COS Vault.

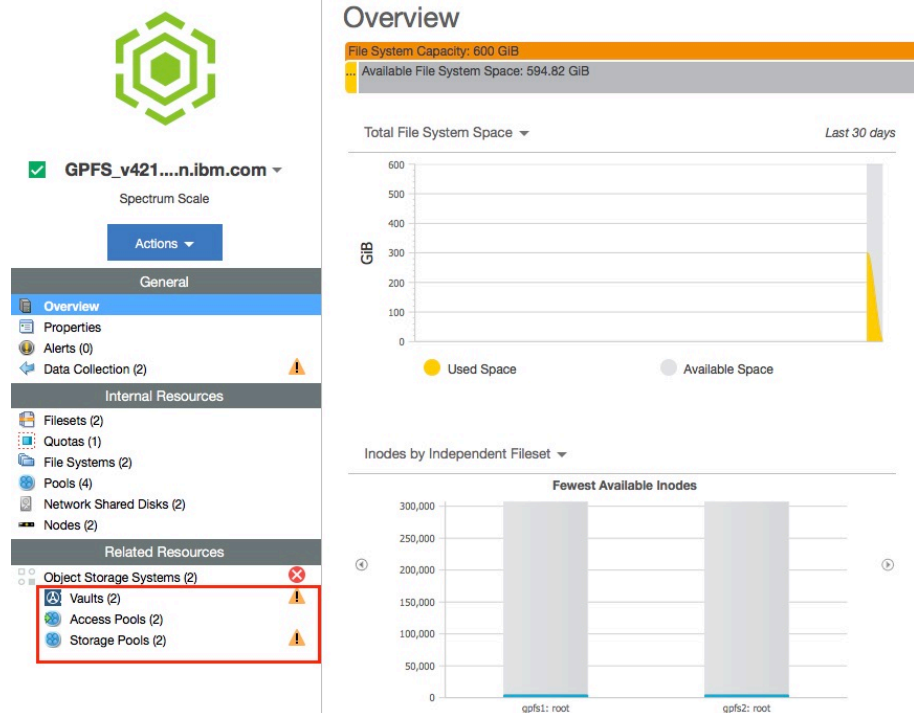


See Calculating the failure tolerance for vaults.

Viewing information about the internal resources of IBM Cloud Object Storage that are used by IBM Spectrum Scale clusters

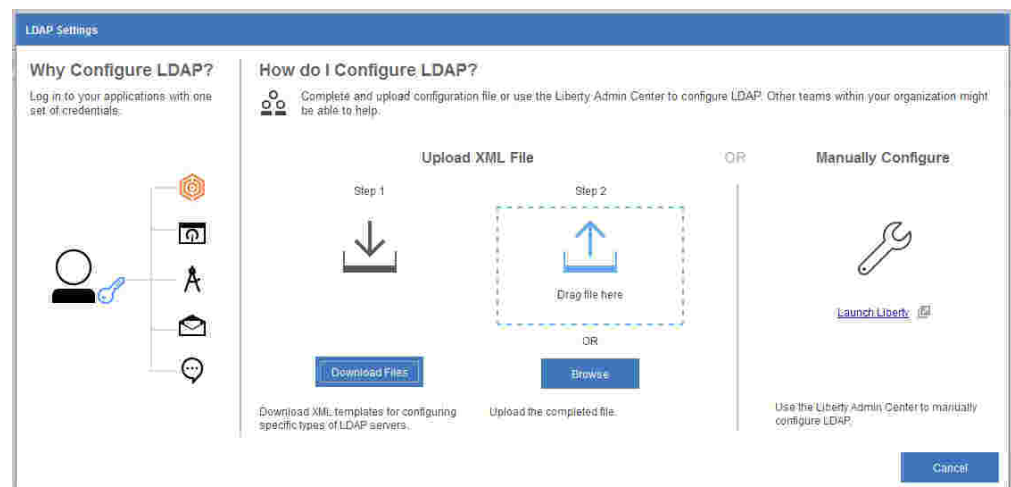
You can now view the COS Vaults, access pools, and storage pools that are

used by IBM Spectrum Scale storage systems.



Configuring LDAP made easier

The user experience is enhanced within the GUI for integration with an external LDAP repository. IBM Spectrum Control is now available to a larger set of users and groups and you are able to log in to IBM Spectrum Control with one set of credentials. Your existing LDAP configurations will be maintained.



See Adding an LDAP repository to the federated repositories for IBM Spectrum Control.

Securing the connection to email servers for sending alert notifications and reports by email

Email Server Setup

To send reports, you must set up the email server. The same email server is used to send reports and alert notifications. You can also set up and change the settings for the email server on the [Settings > Alert Notifications](#) page.

Mail server:	<input type="text"/>
Port:	<input type="text" value="25"/>
User name:	<input type="text"/>
Password:	<input type="password"/>
Reply to address:	<input type="text"/>

IBM Spectrum Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server. For example, if your email server requires authentication and TLS to send emails, then IBM Spectrum Control uses the authentication credentials that you provide when you set up the email server, and uses TLS to establish the connection to your email server. Because the email server that you set up is automatically trusted, you don't have any additional management tasks such as importing the server certificate into the keystores that are used by IBM Spectrum Control.

See [Configuring the email server](#)

Choosing the capacity metric that you want to show in reports

The screenshot shows the IBM Spectrum Control interface. The top navigation bar includes Home, Storage, Servers, Network, Groups, Advanced Analytics, Reports, and Settings. The Reports tab is selected. Below the navigation bar, there is a section for Email Reports. On the left, there is a form to configure the report. On the right, there is a summary of the report data.

Email Reports Configuration:

- Title: Showing consumption in GB
- To: anyone@anyone.com
- Subject: Application Consumption Report in GB
- Frequency: Now
- Configure Report: Calculate Capacity by Allocated Space (selected) or Assigned Space
- Capacity Unit: GB
- Storage Consumer: Type Applications, Consumer SRA

Summary:

Metric	Value	Unit
Total Cost	2544.52	
Used Capacity	2996.60	GB
Primary	2597.12	GB
FlashCopy	107.39	GB

You can now choose the capacity metric when you create or edit chargeback and consumer reports. So, you can generate reports that show storage consumption in the capacity metric of your choice whether it is PB or PiB, TB or TiB, or GB or GiB.

See [Creating chargeback reports](#) and [Creating consumer reports](#).

Viewing port delay performance metrics on IBM Spectrum Virtualize™

You can now view the following port delay metrics for storage systems that run IBM Spectrum Virtualize:

- Port Delay Time
- Port Delay I/O Percentage

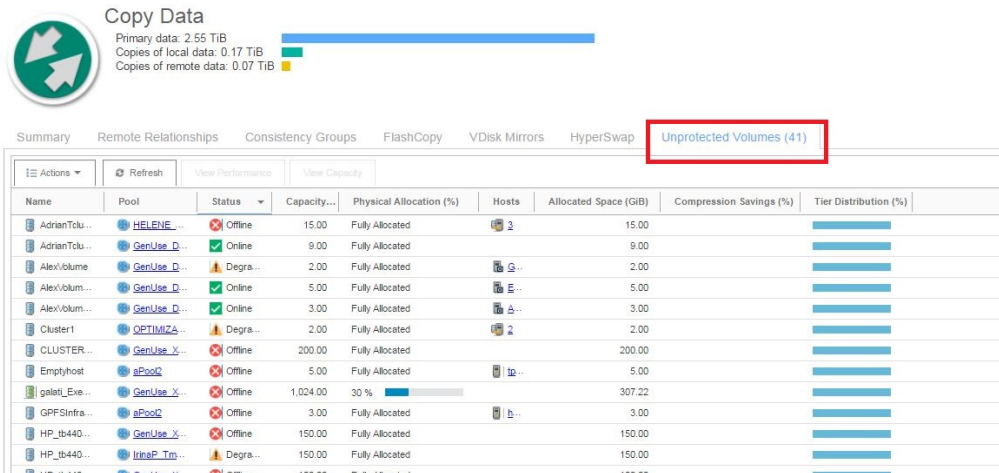
Use these metrics to investigate whether lack of buffer credits is causing performance problems if the Zero Buffer Credit Timer metric is not available.

Viewing host connections that are associated with a host or cluster

For DS8000 storage systems, you can now see the hosts and host clusters that host connections belong to.

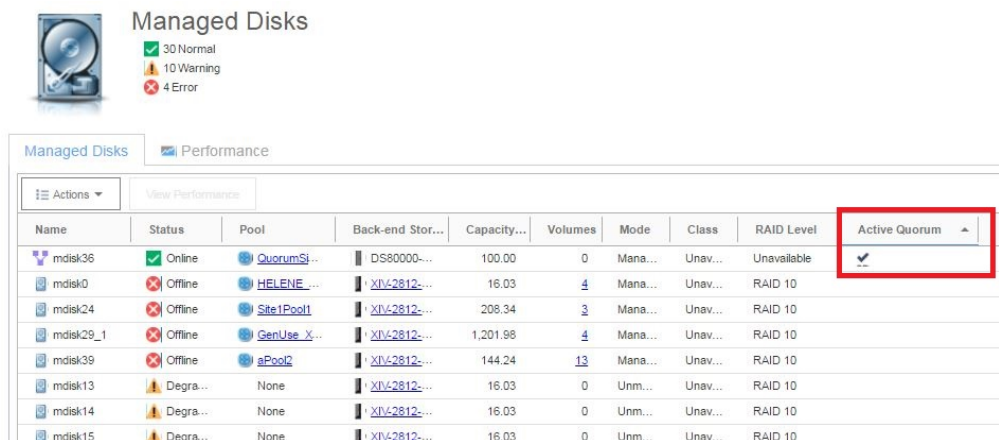
View Unprotected Volumes

You can now see a list of all the primary volumes that are not protected by a copy data relationship. Use this information to identify volumes that might need copy data protection.



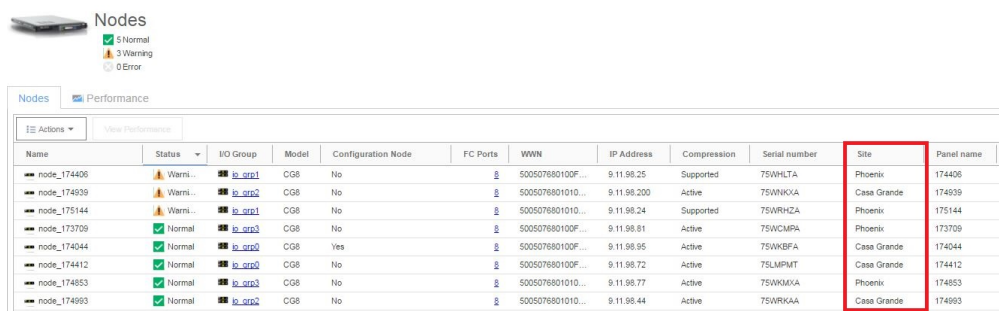
View Quorum Disks

You can now see which Managed Disks are active Quorum disks.



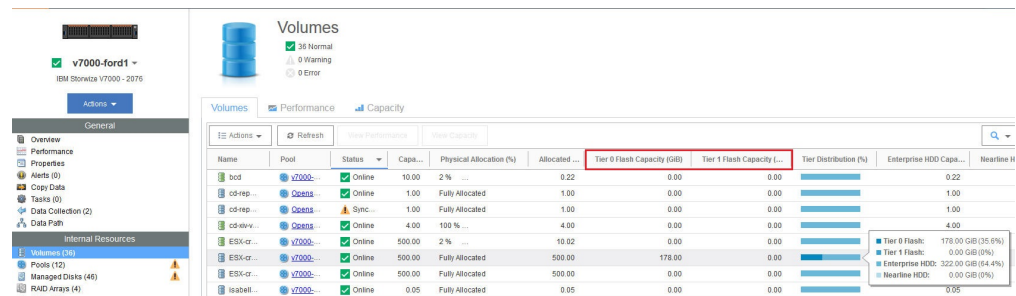

View the site name of each node in a stretched cluster

You can now see the site name of each node in a stretched cluster. The site name is populated by probes. If the site name is not configured, the site column is left blank.



Viewing capacity metrics for read-intensive flash drives on IBM Storwize

You can now view information about Storwize flash drives by tier in storage pools and volumes. For storage pools, you can view the capacity and available capacity for solid-state drives (Tier 0) and for read-intensive flash storage (Tier 1). For volumes, you can view the capacity for solid-state drives (Tier 0) and for read-intensive flash storage (Tier 1).



Uploading logs automatically

Getting support for technical issues has never been easier. Go to **Help > Get Support** to open PMRs and collect system logs. Uploading those logs to IBM Software Support is now just a click away. All you have to do is enter a PMR number, an email address, and click **Collect and Upload Log** to automatically start the transfer.

113-ui.ibm.com

5.2.14.0

20170527-2001

Actions

General

Overview

Alerts (0)

Get Support

Components

Component Servers

Database

Option 1: Automatic Upload

Enter your PMR number to collect and upload the current product logs to IBM support by using encrypted FTP transfer. The email address that you specify receives a confirmation email when the logs are received.

*Blue Diamond users must use the manual upload option for security purposes.

PMR Number: *

Don't have a PMR?

Email: *


Collect and Upload Log

Option 2: Manual Upload

Collect log and upload later. Choose this if you don't wish to use automatic FTP transfer.

Collect and Download Log

Cancel


 Opening a PMR, collecting logs, and uploading to IBM Software Support

Enhanced support for products and platforms

The following products and platforms are supported:


- Minimum supported version of DB2 is Version 10.5.
- DB2 Version 11.1 is shipped with IBM Spectrum Control Version 5.2.14.

Limitations and known issues

For information about limitations and known issues that might occur when you use IBM Spectrum Control Version 5.2.14, see  <http://www.ibm.com/support/docview.wss?uid=swg21999891>.

Summary of changes for IBM Spectrum Control

View a summary of the new features and enhancements that were added in previous versions of IBM Spectrum Control.

- “What's new in 5.2.13”
- “What's new in 5.2.12” on page xxi
- “What's new in 5.2.11” on page xxi
- “What's new in 5.2.10” on page xxii
- “What's new in 5.2.9” on page xxiii
- “What's new in 5.2.8” on page xxiii
-  What's new in 5.2.7

What's new in 5.2.13

The following features and enhancements were available in IBM Spectrum Control 5.2.13:

- View more information about the capacity, space usage, and performance of EMC. Support for versions of EMC SMI-S Provider or EMC Solutions Enabler that are compliant with SMI-S 1.6 provides a more representative view of EMC VMAX, EMC VNX, and EMC VNXe storage systems.
- Gain insights into the external storage that is used by IBM Spectrum Scale:
 - View capacity values for the external storage that clusters and file systems use. This view provides insight into the total amount of data that is migrated from all file systems in a storage system, or the amount of data that is migrated from specific file systems.
 - Check whether you have enough internal space to recall data that was migrated to external storage.
- For storage systems that run IBM Spectrum Virtualize, you can now see the cluster that host connections belong to.
- Enhancements for IBM Cloud Object Storage:
 - View information about the hard disk drives, solid-state drives, and flash modules in COS Slicestor nodes.

- In the Vaults by Accessers chart, view the vaults that cannot be accessed and the vaults that are at risk of access failure.
- On the new VDisk Mirrors and HyperSwap[®] pages in the Copy Data panel, you can analyze storage device relationships across your block storage environment in IBM Spectrum Control to identify how redundancy is affecting available capacity.
- You exclude the cost of storage in chargeback and consumer reports when you want to highlight storage consumption and not storage cost.

View all the new features in 5.2.13

What's new in 5.2.12

The following features and enhancements were available in IBM Spectrum Control 5.2.12:

- Monitor the capacity, space usage, and performance of FlashSystem A9000 and FlashSystem A9000R storage systems.
- Create consumer reports that show the capacity and the cost of the block storage that an application, department, hypervisor, or physical server uses.
- Report on allocated versus assigned capacity in chargeback and consumer reports.
- Analyze replication and FlashCopy relationships across your block storage environment
- View the amount of external storage from cloud services and other storage providers that is being used by the file systems in a IBM Spectrum Scale cluster. For external storage that is provided by IBM Cloud Object Storage, you can also view capacity information, including the percentage of space that is being used.
- Monitor sites, mirrors, and vault quotas for IBM Cloud Object Storage. For example, you can:
 - Monitor the capacity and space usage of sites. See which COS Slicestor and COS Accesser[®] nodes are located at each site and which sites cannot tolerate COS Accesser node failures.
 - View the space quotas that are configured for vaults and see whether any vaults are violating the soft or hard quota limits.
 - View information about mirrors, such as the pair of vaults in each of the mirrors, and the number of COS Accesser nodes that are configured to access the mirrors. You can also see the storage pools and sites for the vaults in the mirror.
- Define alerts that notify you of status changes on the IBM Spectrum Scale nodes that are configured as cloud gateways.
- Customize the analysis period for analyzing tiering and balancing pools to the days of the week and hours of the day of your choice.
- Access information about your resources by using the Representational State Transfer (REST) API for IBM Spectrum Control. You can use this information to generate custom configuration and performance reports.

View all the new features in 5.2.12

What's new in 5.2.11

The following features and enhancements were available in IBM Spectrum Control 5.2.11:

- Changes to alerts:
 - Create alerts for a set of storage resources that are assigned to an application.
 - Define alerts for a set of resources in custom groups that you create.
 - Combine multiple alert conditions from multiple resources in a single, custom alert.
- Create, schedule, and send chargeback reports by email to make the owners and managers aware of the cost of the capacity and the amount of capacity that is used by their applications, departments, hypervisors, and servers.
- View key performance metrics for IBM® FlashSystem™ 840 and IBM FlashSystem 900 storage systems:
 - Response times
 - I/O rates and data rates
 - The condition of the storage system
- Review new capacity charts to detect capacity shortages and investigate trends in storage growth for your tiers.
- Analyze replication relationships across your block storage environment.
- Monitor the status, capacity, and space usage of IBM Cloud Object Storage.
- Identify which file systems in a Spectrum Scale cluster use external storage that is provided by cloud services or other storage providers.
- Enable automated probe scheduling to simplify the process of scheduling probes and help avoid excessive load on the product server.
- Determine whether the total capacity of the storage systems that IBM Spectrum Control manages exceeds your product license entitlement.
- When you complete a fresh installation of the product, the number of random listening ports is reduced. Information about the ports is documented and you can change the Device server ports after you complete the installation or perform an upgrade.
- Manually uninstall IBM Tivoli® Storage Productivity Center for Replication from IBM Spectrum Control servers on your computer.
- IBM DB2® Version 10.5 Fix Pack 7 is now included with IBM Spectrum Control.

View all the new features in 5.2.11

What's new in 5.2.10

The following features and enhancements were available in IBM Spectrum Control 5.2.10:

- Identify key performance metrics that are outside of a standard range on IBM System Storage® SAN Volume Controller and IBM Storwize® storage systems.
- Export performance data for a resource to a compressed file.
- Review new capacity charts to detect capacity shortages and investigate trends in storage growth for your block, file, and object storage resources.
- Review the capacity and space usage that the tiered storage in your environment consumes. You can see when your tiered storage will run out of space and the weekly growth rate in storage usage for each tier.
- View the values for the capacity of virtualized storage and non-virtualized storage to understand how capacity is used when storage virtualization is implemented in your storage environment.
- Use rollup reporting to combine capacity and status information from multiple instances of IBM Spectrum Control™. You can use this information to gain a

network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.

- Group resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location.
- IBM Spectrum Control V5.2.10 now supports only Storage Management Initiative (SMI) data sources, also called CIM agents, for managing switches from Brocade Communications Systems, Inc.
- View the resource types that were explicitly added to an application or subcomponent.

View all the new features in 5.2.10

What's new in 5.2.9

The following features and enhancements were available in IBM Spectrum Control 5.2.9:

- Reclaim unused storage capacity and more effectively allocate volume space on your block storage systems.
- View the amount of storage that can be reclaimed in each of the storage systems that are monitored.
- Use the volume reclamation recommendations to reclaim unused storage space and use your storage more efficiently. By reclaiming the volumes, you can reduce costs by recycling existing storage space instead of purchasing new storage media.
- IBM Spectrum Control™ now displays the switch blade number, port number, and blade names just as they were defined in other tools, such as Brocade Network Advisor.
- IBM DB2® Version 10.5 Fix Pack 6 is now included with IBM Spectrum Control.

View all the new features in 5.2.9

What's new in 5.2.8

The following features and enhancements were available in IBM Spectrum Control 5.2.8:

- The name of the product was changed from IBM Tivoli Storage Productivity Center to IBM Spectrum Control.
- Comprehensive improvements to alerting capabilities:
 - You can now define alerts for all the key attributes of a resource, including attributes for status, configuration, capacity, and performance.
 - To avoid receiving too many emails and disruptive notifications for some alerts, you can now choose a new suppression option where only one notification is sent.
 - If you want to define more than one alert for the same attribute but with different criteria and settings, you can duplicate that alert.
- View the performance of clusters, nodes, and file systems in IBM Spectrum Scale.
- Monitor object storage systems on IBM Spectrum Scale.
- Monitor the capacity and space usage of IBM FlashSystem 840 and IBM FlashSystem 900 storage systems.

- View charts that show the current and historical capacity of your block storage pools.
- IBM Spectrum Control is now a native, 64-bit application for all server operating systems.
- Simple Network Management Protocol (SNMP) Version 3 is now the preferred protocol for the management of switches and fabrics from Cisco Systems, Inc., and is an optional protocol for devices from Brocade Communications Systems, Inc. and other vendors.
- View information about the performance of inter-switch connections.
- View information about Brocade switches that are in Access Gateway mode. You can also view information about the routes from ports that are connected to a hypervisor, server, or storage system to ports that are connected to fabrics.
- View information about Brocade F_port trunks and about NPV links.
- You can now more easily identify the internal resources that are encountering problems in your environment.
- Compare how much capacity is managed by IBM Spectrum Control with the total amount of capacity that you are entitled to manage.
- View information about the capacity for each of the class drives (SSD, Nearline HDD, and Enterprise HDD), that Easy Tier can use to tier or distribute volume extents.
- To consolidate the functions of IBM Spectrum Control into a single GUI, the stand-alone GUI was removed.

View all the new features in 5.2.8

Chapter 1. Getting started

Learn about the key tasks for setting up IBM Spectrum Control for managing the resources and infrastructure in a storage environment. The getting started guide provides a list of tasks in sequential order for helping you to learn how to install, configure, and use IBM Spectrum Control.

Complete the following tasks to install and set up IBM Spectrum Control for monitoring and managing storage resources. Click the included links to learn more information about each task.

Table 2. Tasks for getting started with IBM Spectrum Control








Task	Description	 Learn more
Planning for installation and upgrades.	<p>The time that you invest in planning directly affects your ease of implementation. Planning includes the following tasks:</p> <ul style="list-style-type: none">• Ensuring that your systems meet hardware and software requirements• Becoming familiar with the operating systems, storage resources, agents, and communication protocols that can be used with IBM Spectrum Control• Ensuring that ports are open for use• Understanding authentication and role-based authorization• Understanding the requirements for specific types of resources	<ul style="list-style-type: none">•  Hardware requirements•  Software requirements.•  TCP/IP ports used by IBM Spectrum Control•  Planning for IBM Spectrum Control authentication and authorization•  Planning for resources <p>For more information about the software and resources that are supported by IBM Spectrum Control, see  IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446.</p>

Table 2. Tasks for getting started with IBM Spectrum Control (continued)








Task	Description	 Learn more
Installing or upgrading IBM Spectrum Control.	<p>If you have a previous version of IBM Spectrum Control installed, you can upgrade it to the current version. If you do not have a previous version, you can use the following methods to install it:</p> <p>Single-server environment Install IBM Spectrum Control and its components on a single server.</p> <p>Multiple-server environment Install IBM Spectrum Control on multiple servers if you are monitoring large storage environments, where one server is not sufficient to manage all the IBM Spectrum Control components.</p> <p>When you install IBM Spectrum Control, the following components are also installed:</p> <p>DB2 (required) IBM Spectrum Control stores the information that it collects about storage resources in a DB2 database.</p> <p>Jazz for Service Management and Tivoli Common Reporting (optional) Install these components if you want to use IBM Spectrum Control reports in the web-based GUI.</p>	<p>You can run the installation process as a GUI wizard or through a command line in silent mode.</p> <ul style="list-style-type: none">  Upgrading and migrating  Installing DB2  Installing IBM Spectrum Control in a single-server environment  Installing IBM Spectrum Control in a multiple-server environment  Install Jazz for Service Management and Tivoli Common Reporting
Starting the IBM Spectrum Control GUI.	The IBM Spectrum Control GUI runs in a web browser and can be accessed from anywhere that you have a web browser and connectivity to a network.	“Starting IBM Spectrum Control” on page 39
Adding resources for monitoring.	<p>Before you can monitor the resources in your environment, you must add them to IBM Spectrum Control. You can add the following resources for monitoring:</p> <ul style="list-style-type: none"> Storage systems Switches Fabrics Servers Hypervisors 	“Adding resources” on page 112
Deploying Storage Resource agents on servers (optional).	<p>You can deploy Storage Resource agents on servers that you want to monitor. Deploy agents if you want to complete the following tasks:</p> <ul style="list-style-type: none"> Collect storage usage and trending information about servers (scans) Collect database application information Discover the fabrics that are visible to a server Enforce storage policies Collect zone information 	<p>For more information about Storage Resource agents, see “Adding servers” on page 125.</p> <p>For a list of the servers where you can deploy Storage Resource agents, see  IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446 and go to the <i>Agents, Servers, and Browsers</i> section.</p>

Table 2. Tasks for getting started with IBM Spectrum Control (continued)






Task	Description	 Learn more
Collecting data about resources.	<p>Determining the information that you want to gather about storage resources is critical to helping you implement a storage management strategy. When you determine the types of information you want to collect about your resources, you can select the data collection jobs that you want to run.</p> <p>Probes Probes collect status, asset, and capacity information about resources.</p> <p>Performance Monitors Performance monitors collect metrics that measure the performance of storage systems and switches.</p> <p>Tip: Probes and performance monitors are automatically scheduled when you add resources for monitoring.</p>	<ul style="list-style-type: none"> • “Adding resources” on page 112 • “Collecting data” on page 133
Configuring how alerts are received.	<p>In the web-based GUI, enable alert notifications. Alerts can define notification actions that send email, generate Simple Network Management Protocol (SNMP) traps, or generate IBM Tivoli Netcool®/OMNIBus events. To receive these types of alert notifications, you must configure IBM Spectrum Control to direct the notifications to a specific email address, SNMP destination, or Tivoli Netcool/OMNIBus server.</p>	 Configuring alert notifications
Specifying the alerts that are generated when a problem is detected on monitored storage.	<p>Determining when and how you are alerted to conditions or violations within your storage environment is important to helping you maintain and administer storage resources. Alerts are triggered by conditions that are detected during data collection and event processing.</p> <p>To set up and manage alerts, complete the following tasks:</p> <ul style="list-style-type: none"> • For each resource that you want to monitor for issues, specify the conditions and performance thresholds that can trigger alerts. • Determine how you want to be notified when alerts are triggered. • View, acknowledge, and remove alerts that are detected. For alerts that are triggered by performance threshold violations, you can view a chart of the historical performance of the resource that violated a threshold. 	<ul style="list-style-type: none"> • “Defining alerts” on page 151 • “Viewing and modifying alert notification settings” on page 187 • “Viewing all alerts” on page 191 • “Triggering conditions for alerts” on page 370

Table 2. Tasks for getting started with IBM Spectrum Control (continued)

Task	Description	 Learn more
Configuring data retention (optional).	Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control. By specifying the number of weeks for history retention, you can control the amount of data that is retained and available for historical analysis and charting. The longer that you retain data, the more informative your analysis, but the more storage space that is required to store that data. Default retention settings are provided.	 Configuring history and data retention
Configuring authentication and user access.	Determine the authentication mechanism and authorization level for users. The authentication mechanism determines the user groups to which you can assign roles. Roles determine the IBM Spectrum Control functions that are available to the users in a group. In the web-based GUI, you can assign the following roles to groups of users: <ul style="list-style-type: none"> • Administrator • Monitor • External Application 	 Authorizing users

Managing resources

Learn about the key tasks for getting started with monitoring and managing storage resources.

After installing IBM Spectrum Control and adding resources for monitoring, complete the following tasks to manage your storage environment.

Table 3. Tasks for monitoring and managing resources



Task	Description	 Learn more
Monitoring the performance of resources.	Collect and view key performance metrics about resources and specify alerts to be notified of threshold violations. Use metrics to measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.	<ul style="list-style-type: none"> • “Monitoring the performance of resources” on page 217 •  Reporting in IBM Spectrum Control

Table 3. Tasks for monitoring and managing resources (continued)





Task	Description	 Learn more
Monitoring the capacity of resources.	<p>View detailed information about the capacity of monitored resources in the web-based GUI. You can view capacity information for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual machines. For example, you can view the total amount of storage space that is committed to a volume, the amount of allocated space that is used by a volume, and the amount of space that is allocated to a volume and is not yet used.</p> <p>You can also use the optional Cognos Business Intelligence reporting tool to view and create capacity reports about multiple resources.</p> <p>Use the information that is collected about storage capacity to help with the capacity planning of the storage within your organization.</p>	<ul style="list-style-type: none"> • “Viewing information about resources” on page 194 •  Reporting in IBM Spectrum Control •  Click the Help icon in the web-based GUI to view descriptions of the capacity information that is displayed for resources.
Monitoring the status of resources.	<p>Monitor the status and condition of resources to identify potential problem areas in a storage environment. On the dashboard in the IBM Spectrum Control GUI, you can view a high-level summary of the statuses for each type of resource. From the dashboard, you can drill down to view the individual status for each monitored resource.</p>	<ul style="list-style-type: none"> • “How the condition of a resource is determined” on page 204 • Viewing the overall status of a storage environment • Viewing the status of individual resources • Viewing the status of internal resources
Viewing the relationships between resources.	<p>In the web-based GUI, you can generate a graphical or tabular representation of the relationships between host connections, servers, hypervisors, virtual machines, storage systems, and the fabrics through which they are connected. View data paths to monitor status and pinpoint problem areas in the path of monitored storage throughout your environment.</p> <p>You can also use the optional Cognos BI reporting tool to view reports about the relationships between resources. For example, you can view a report about the end-to-end relationship of all the servers in your environment to back-end storage.</p>	<ul style="list-style-type: none"> • Viewing data paths •  Viewing relationship reports in the Cognos BI reporting tool

Table 3. Tasks for monitoring and managing resources (continued)





Task	Description	 Learn more
Monitoring and managing applications and departments.	<p>Storage resources can be grouped together to represent data modeling concepts such as applications, subcomponents (member applications), departments and subdepartments to create a business hierarchy. The monitoring and management of applications and departments enables you to perform capacity trending, health monitoring, and performance troubleshooting for the storage resources in your environment.</p> <p>In IBM Spectrum Control, you can create applications and departments, add applications and departments, use a filtering criteria to add resources to the applications, remove related resources from within the context of applications, remove applications and departments, view key properties and view detailed storage information for applications and departments.</p>	<ul style="list-style-type: none"> • “Creating applications” on page 263 • “Adding subcomponents” on page 267 • “Adding applications to departments” on page 274 • “Using filters to assign resources to applications” on page 264 • “Removing resources from applications” on page 269 • “Viewing information about applications” on page 268 • “Creating departments” on page 272 • “Adding and creating subdepartments” on page 273 • “Viewing information about departments” on page 275
Identifying hot spots of high utilization.	<p>IBM Spectrum Control collects detailed information about storage system pools, including utilization values. The utilization value is an estimate of the average daily utilization of the physical resources such as controllers, disks, ports, and nodes that are associated with the pool.</p> <p>By scanning the utilization values for pools on a storage virtualizer, such as SAN Volume Controller, you can identify the hot spots or the pools that have high utilization values on the storage virtualizer. You can also determine the course of action that you take.</p>	<ul style="list-style-type: none"> • “Redistributing volume workloads to balance pools” on page 86 • “Analyzing and re-tiering volumes in pools on tier 1” on page 87 • “Collocating volumes” on page 90
Setting up self-service provisioning.	<p>Storage provisioning is the process of assigning storage to optimize the performance of a storage environment. To provision storage with IBM Spectrum Control, create service classes and capacity pools to classify and separate storage resources. After configuration, use the web-based GUI, vSphere Web Client extension, or other tool with those service classes and capacity pools to provision storage.</p> <p>In the web-based GUI, you are guided through the steps for provisioning volumes or shares. IBM Spectrum Control then recommends storage placement that is based on storage system free space and performance data. If you agree with the recommendation, IBM Spectrum Control can implement the provisioning.</p>	<p>“Provisioning storage” on page 57</p>

Table 3. Tasks for monitoring and managing resources (continued)

Task	Description	 Learn more
Provisioning storage in a virtual storage environment.	If you maintain virtual storage in a VMware vSphere environment, you can use the vSphere Web Client extension for IBM Spectrum Control to complete provisioning tasks and to view information about storage resources in the vSphere Web Client. The client includes details about performance and asset, copy services, and fabric information.	<ul style="list-style-type: none"> • “Provisioning storage with the vSphere Web Client extension” on page 316 • “Limiting a set of storage to provision in the vSphere Web Client extension” on page 98
Running reports and setting up report scheduling and delivery.	<p>Use reports to view detailed information about the condition, capacity, and performance of resources. IBM Spectrum Control provides multiple user interfaces for running reports.</p> <p>Cognos BI reporting tool (optional) You can use the optional Cognos BI reporting tool to run, schedule, and deliver predefined and custom reports. These reports include information about storage capacity, assets, and performance of storage systems, switches, fabrics, servers, hypervisors, and their internal resources.</p> <p>vSphere Web Client Use the vSphere Web Client extension to view reports about virtual environments and storage resources. You can view reports that are customized to use information from IBM Spectrum Control. The reports include fabric connections, storage mapping information, and performance metrics for storage systems.</p>	<ul style="list-style-type: none"> •  Reporting in IBM Spectrum Control •  Reporting in a vSphere environment

Related concepts:

Chapter 6, “Managing resources,” on page 107


Use IBM Spectrum Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

Configuring and administering

Learn about the key tasks for getting started with configuring and administering IBM Spectrum Control and its related components.


Complete any of the following tasks to customize IBM Spectrum Control for your environment and administer its components.

Table 4. Tasks for configuring and administering IBM Spectrum Control

Task	Description	 Learn more
Configuring IBM Spectrum Control.	<p>Configure IBM Spectrum Control according to the standards within your organization, including the following tasks:</p> <ul style="list-style-type: none"> • Customizing lists of resources, tasks, and alerts • Enabling reports in the optional Cognos Business Intelligence reporting tool • Configuring how alerts are received • Configuring data retention • Configuring authentication and user access • In the vSphere client, you can configure the vSphere client and VASA provider 	<ul style="list-style-type: none"> • “Customizing lists of resources, tasks, and alerts” on page 337 • If the Cognos BI reporting tool is not enabled, you must install Jazz for Service Management and IBM Tivoli Common Reporting. For more information, see Install Jazz for Service Management and Tivoli Common Reporting. • Configuring the vSphere Web Client extension for IBM Spectrum Control • Configuring the VASA provider
Administering DB2 and IBM Spectrum Control.	<p>After IBM Spectrum Control is installed and configured, complete the following administering tasks as needed:</p> <ul style="list-style-type: none"> • Administering resources and data sources • Starting and stopping IBM Spectrum Control servers • Checking the status of IBM Spectrum Control components • Packaging log files for IBM Software Support • Changing passwords for IBM Spectrum Control components • Administering DB2 and backing up the database repository 	<ul style="list-style-type: none"> • Administering resources and data sources • Starting and stopping the IBM Spectrum Control servers • Checking IBM Spectrum Control status • Opening a PMR, collecting logs, and uploading to IBM Software Support • Changing passwords • Administering DB2

Navigation


IBM Spectrum Control provides many of the functions for managing a storage environment. To access the functions in its GUI, use the menu bar at the top of the main window.

GUI element	Description
Page banner	<p>Use the banner at the top of every page to complete the following tasks:</p> <ul style="list-style-type: none"> • View the name of the product license that is active. • Access resource pages and main product functions from the menu bar. • View the ID and role of the user who is logged in. The role of a user determines the product functions that are available to that user. For more information about roles, see “Role-based authorization” on page 70. • Access the online help to view information about the currently displayed page and the overall product. • Enter and exit rollup mode by using the rollup mode icon . You must add at least one subordinate rollup server for the rollup mode icon to display in the banner.

GUI element	Description
Home	<p>Dashboard</p> <p>Use the dashboard to view the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about:</p> <ul style="list-style-type: none"> • Condition and usage of resources • Entities that consume storage on those resources • Number and status of unacknowledged alert conditions that are detected on the monitored resources • Most active storage systems in your environment. <p>Performance Monitors</p> <p>Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations.</p> <p>Alerts Use alerts to be notified when certain configuration, status, and performance conditions are detected on monitored resources.</p> <p>Tasks Manage the tasks that are created when you complete the following actions:</p> <ul style="list-style-type: none"> • Assign storage to servers, hypervisors, or a cluster. • Place volumes on the tiers that match the workload requirements of the volume. • Move or convert volumes. • Move volumes from overutilized pools to pools that are less used. <p>System Management</p> <p>View information about the overall condition of IBM Spectrum Control. You can view information about the servers on which the product is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.</p>

GUI element	Description
Storage	<p>Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems:</p> <ul style="list-style-type: none"> • Modify the schedules for data collection. • View the data path of connected resources. • Assign storage to servers, hypervisors, or a cluster. • Place volumes on the tiers that match the workload requirements of the volume. • Move or convert volumes. • Move volumes from overutilized pools to pools that are less used. • Add resources to a capacity pool. • Open the management GUIs for storage systems. • Update login credentials. • Acknowledge conditions and statuses.
Servers	<p>Add servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> • Deploy Storage Resource agents to servers for full monitoring. • Modify the schedules for data collection. • View the data path of connected resources. • Assign storage to servers, hypervisors, or a cluster. • Modify and upgrade Storage Resource agents. • Acknowledge conditions and statuses.
Network	<p>Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> • Modify the schedules for data collection. • View the data path of connected resources. • Setting or modifying the zoning policy. • Open the management GUI for a switch. • Acknowledge conditions and statuses.

GUI element	Description
Groups	<p>Add applications and departments to combine resources that are grouped by applications and departments. Applications and departments are used to monitor and view status, capacity, and performance information, and view details about their related resources. You can also complete the following actions to manage and administer applications and departments:</p> <ul style="list-style-type: none"> • Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting. • View the status of resources that make up the applications and departments. • Create resource filters to automatically add resources to applications. • Add resources directly to applications and add applications to existing applications and departments. • Add departments to existing departments. • View subcomponent (member application) information to support hierarchical levels of storage grouping. • Remove applications and departments. <p>Review the distribution of capacity in your storage environment to determine whether you have sufficient capacity to meet the demands of your tiered storage.</p>
Advanced Analytics	<p>Configure IBM Spectrum Control for provisioning, optimizing, and reclaiming storage. It guides you through the steps for managing your storage such as creating service classes and capacity pools to provision storage and creating tiers to optimize storage placement. You can complete the following tasks:</p> <ul style="list-style-type: none"> • Provision servers, hypervisors, or a cluster • Optimize storage by re-tiering volumes and balancing pools • Reclaim storage that is not being used

GUI element	Description
Reports	<p>Use the Email Reports option to create, configure, schedule, and send chargeback reports for applications, departments, hypervisors, and servers.</p> <p>Use the web-based GUI when you want to individually view details about resources, for example when you want to troubleshoot a resource.</p> <p>You can also use the optional Cognos Business Intelligence reporting tool to analyze multiple storage systems, switches, servers, and hypervisors. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor.</p> <p>Use the Cognos BI reporting tool to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports.</p> <p>When you click a Cognos BI reporting option, a separate browser tab is opened.</p> <p>For information about how to get started with the Cognos BI reporting tool, see  http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.14/com.ibm.spectrum.sc.doc/fqz0_c_getting_started_reports.html.</p>
Settings	<p>Complete the following actions to customize IBM Spectrum Control for your environment:</p> <ul style="list-style-type: none"> • Modify the authentication repository and assign IBM Spectrum Control roles to user groups. Roles determine the product functions that are available to users. • Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control. • Add and configure rollup servers so you can view capacity and status information about resources that are managed by your IBM Spectrum Control servers.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.

- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Chapter 2. Videos and blogs for IBM Spectrum Control

Watch videos, read blogs, and explore social media to learn more about how to use IBM Spectrum Control to manage your storage environment.

Videos

Important: The following videos were recorded by using multiple versions of IBM Spectrum Control (formerly known as Tivoli Storage Productivity Center), but are at least partially applicable for the current version:

Table 5. Videos about IBM Spectrum Control (English only)









Video	Summary
(1:08)  IBM Storage / Spectrum Control	Watch a brief overview of IBM Spectrum Control and learn how it can help you manage a storage environment.
(4:50)  IBM Spectrum Control: Technical Point of View	Watch a brief overview of IBM Spectrum Control and learn about some of the newest features.
(8:18)  Partnering with MicroStrategies	See how an IBM business partner uses IBM Spectrum Control to get a better view of storage environments.
(5:07)  IBM Spectrum Storage Software Training	Go to the YouTube channel for IBM Spectrum Storage Training to watch a series of videos about how to install and configure IBM Spectrum Control.
(5:08)  Simplify storage, deploy new applications, and control costs with IBM Spectrum Storage	See how IBM Spectrum Storage Suite can help you cope with data explosion, manage all types of storage, and quickly test and implement new cognitive applications, all while you maintain predictable, controlled costs.
(5:33)  How to monitor IBM Spectrum Scale with IBM Spectrum Control, V5.2.12	Learn how to use IBM Spectrum Control to monitor multiple IBM Spectrum Scale clusters, view cross cluster relationships, monitor Transparent Cloud Tiering between IBM Spectrum Scale and IBM Cloud Object Storage, and view relationships between Network Shared Disks and SAN attached storage.
(6:49)  Learn how to use IBM Spectrum Control to troubleshoot performance problems in your storage environment.	Learn how to use IBM Spectrum Control to troubleshoot an application and how to use the controls on the performance charts. Also, learn how to compare metrics across time ranges and identify the best or worst values for a performance metric.
(7:20)  How to use IBM Spectrum Control for chargeback reporting and reporting on storage consumption	Learn how to use IBM Spectrum Control for chargeback reporting and reporting on storage consumption; it will show how to set up items related to reports, for example, tiers and applications.

Table 5. Videos about IBM Spectrum Control (English only) (continued)














Video	Summary
(5:38)  Monitoring IBM Spectrum Scale performance with IBM Spectrum Control	Learn how to use IBM Spectrum Control to compare the performance of multiple IBM Spectrum Scale clusters and then analyze the clusters that have the heaviest I/O loads or the longest response times.
(1:57)  Gathering performance data in IBM Spectrum Control	Learn how to quickly gather performance data in IBM Spectrum Control by using the Create Performance Support Package option.
(6:08)  Monitoring block storage performance with IBM Spectrum Control	Learn how IBM Spectrum Control can be used to quickly monitor the performance of your entire storage environment.
(6:34)  Alerting capabilities in IBM Spectrum Control	Learn about the alerting capabilities in IBM Spectrum Control.
(6:55)  Using IBM Spectrum Control to monitor GPFS clusters	Learn how to monitor IBM Spectrum Scale clusters.
(10:30)  Troubleshooting Storage Infrastructure Performance	Learn how to monitor, compare, and troubleshoot performance on storage systems.
(10:48)  How to extract performance data	Learn how to extract performance data from IBM Spectrum Control using the GUI.
(10:12)  Using Tivoli Storage Productivity Center in a service provider environment	Learn how IBM Global Technology Services uses Tivoli Storage Productivity Center to manage storage in a service provider environment.

Table 6. Videos about IBM Spectrum Control and Cognos BI (English only)

Video	Summary
(8:29)  Cognos BI: Creating a tier capacity by location report	Learn how to create a capacity report of different tiers of storage across multiple data centers.
(8:00)  Cognos BI: Pre-defined reports	Learn how to use IBM Spectrum Control and Cognos BI to access pre-defined reports. These reports can help meet your SAN reporting needs in the areas of capacity, performance, and storage relationships.
(15:54)  Cognos BI: Setting up basic security for your reporting environment	Learn how to use the security features of Jazz for Service Management, Tivoli Common Reporting, and Cognos BI to provide some basic security for your reporting environment.
(4:39)  Cognos BI: Restricting user access to specific reports	Learn how to limit the reports and reporting capabilities that users can access within Cognos BI.
(8:03)  Cognos BI: Restricting which users can author and modify the directory structure that contains reports	Learn how to restrict which users can author and modify the directory structure that contains IBM Spectrum Control reports.

Blogs

IBM Spectrum Control has a strong presence on social media. Explore blogs and forums to read the perspectives of storage management experts and gain insights about how to use IBM Spectrum Control.

General blogs:

- IBM Spectrum Control blog series
- Storage Management blogs

Table 7. Blogs about IBM Spectrum Control V5.2.13

Blog	Summary
IBM Spectrum Control V5.2.13	Learn about the features that were introduced in IBM Spectrum Control V5.2.13.
Chargeback and Consumer Reports: Don't show me the money!	Learn how to create a custom report that sends you a daily summary of the capacity used by applications without including storage costs.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to see what it offers and learn about topics that you might not be aware of but could find useful.
IBM Spectrum Control virtual mailbox	Learn how much capacity is consumed by replication. As a storage administrator you are asked to provide a capacity report across your entire storage environment to your CTO. A critical part of capacity reporting is understanding how much capacity is consumed by replication.

Table 8. Blogs about IBM Spectrum Control V5.2.12

Blog	Summary
IBM Spectrum Control V5.2.12	Learn about the features that were introduced in IBM Spectrum Control V5.2.12.
Consumer Reports: For your eyes only!	Learn how to configure a report that shows the block capacity and cost of the block capacity for a single storage consumer. Each consumer gets to star in their own report.
Chargeback and Consumer Reports: Same price for the same storage	In chargeback and consumer reports, you want to ensure that you charge the same price for the same type of storage whether it is block storage, copy data, or file storage. And, if your block storage is tiered, you want to charge the same price for each tier to all of your storage consumers.
Chargeback and Consumer Reports: How much storage was really eaten?	When you create chargeback and consumer reports, you can charge for the block storage that is allocated to the storage consumer or the block storage that is assigned to the storage consumer.

Table 8. Blogs about IBM Spectrum Control V5.2.12 (continued)

Blog	Summary
Look at that IBM FlashSystem data reduction!	FlashSystem A9000 and FlashSystem A9000R can reduce the size of your data significantly by using technologies such as compression and data deduplication. Learn how to see exactly what sort of savings you are getting, and how that compares to savings on other storage systems.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to see what it offers and learn about topics that you might not be aware of but could find useful.
From the IBM Spectrum Control virtual mailbox: Can I get enhanced optimization results for storage tiering and pool balancing?	You can reduce storage costs by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers. Storage performance can be improved and used more efficiently by moving volumes with heavy workloads to the tiers that meet their workload requirements. If you implement the recommendations, you can ensure that volumes are placed on the tiers that best match the workload requirements of the volumes.

Table 9. Blogs about IBM Spectrum Control V5.2.11

Blog	Summary
IBM Spectrum Control V5.2.11	Learn about the features that were introduced in IBM Spectrum Control V5.2.11.
Chargeback Reports: Who's eating my storage?	Whether you manage or own applications, departments, hypervisors, or physical servers, you want to know how much storage is being consumed and how much it costs to maintain the storage that is consumed. And, you want that information in your inbox now, every week or month so that you can keep a close eye on storage costs and capacity consumption.
Capacity and compliance: Know your software license in IBM Spectrum Control	Your business organization just purchased a software license(s) for IBM Spectrum Control or any IBM product that includes IBM Spectrum Control such as IBM® Virtual Storage Center or IBM Spectrum Storage Suite and you want to assess your storage license situation.
Show your boss how your IBM FlashSystem compares	Learn how to compare the performance of your IBM FlashSystem to other storage systems. You can compare the response times and the I/O rates of IBM FlashSystem storage systems to other storage systems.
See how Cognos BI can help you create capacity and performance reports	Learn how to use the Cognos® Business Intelligence (BI) reporting tool to report on the condition, capacity, and performance of your storage resources.

Table 9. Blogs about IBM Spectrum Control V5.2.11 (continued)

Blog	Summary
From the IBM Spectrum Control hints and tips virtual mailbox: Is there a way to simplify the process of scheduling probes?	The automated probe option enables the creation of a probe and collects status and asset information more efficiently about your storage resources for monitoring.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to see what it has to offer and learn about topics that you might not be aware of but could find useful.
Navigating the new knowledge center	Learn how to quickly find information in the new IBM Knowledge Center and read about tips for getting the most out of its new interface.

Table 10. Blogs about IBM Spectrum Control V5.2.10

Blog	Summary
IBM Spectrum Control V5.2.10	Learn about the features that were introduced in IBM Spectrum Control V5.2.10.
Extracting bulk performance data for analysis was never easier!	Learn about how to export bulk performance data for a storage system or fabric to a compressed file. If you contact IBM Support to help you analyze a performance problem with a resource, you might be asked to send this file.
IBM Spectrum Control: Rollup reporting provides a dynamic view into a large enterprise environment	Learn about how to use rollup reporting to gain a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.

Table 11. Blogs about IBM Spectrum Control V5.2.9

Blog	Summary
IBM Spectrum Control V5.2.9	Learn about the features that were introduced in IBM Spectrum Control V5.2.9.

Table 12. Blogs about IBM Spectrum Control V5.2.8

Blog	Summary
IBM Spectrum Control V5.2.8	Learn about the features that were introduced in IBM Spectrum Control V5.2.8.
Monitoring the entire data path of object storage systems	Learn about how to monitor object storage implementation on IBM Spectrum Scale and troubleshoot along the entire data path.

Chapter 3. Product licenses

Multiple product licenses are available for IBM Spectrum Control. Each license determines the features that you can access from the IBM Spectrum Control GUI, command-line interface, and the API. You can review the features that are included with each license and determine which license meets your storage management needs.

IBM Spectrum Control Standard Edition and IBM Spectrum Control Standard Select Edition

In addition to many other features, these licenses provide the following features:

- Single dashboard view of the storage environment that you can use to manage storage systems, hypervisors, servers, and Fibre Channel fabrics.
- Performance monitoring for storage systems and Fibre Channel networks.
- Capacity and usage monitoring.
- File provisioning for IBM Storwize V7000 Unified and IBM Scale Out Network Attached Storage.
- Health and alerting for hypervisors, networks, servers, and storage systems.
- Capacity and performance of the storage that applications and departments use.
- Administrative tools, such as:
 - VMware vCenter plug-in, which you can use to manage and provision virtual storage in a vSphere environment, and to view reports about storage systems that are monitored by IBM Spectrum Control.
 - IBM Cognos(R) Business Intelligence reports, which provide reporting and analytics information.

These licenses include a license for IBM Copy Services Manager, which manages 2-site replication, 3-site replication, and advanced copy services. You can learn more about how to download and install Copy Services Manager at <http://www.ibm.com/support/knowledgecenter/SSESK4>.

IBM Spectrum Control Advanced Edition and IBM Spectrum Control Advanced Select Edition

These licenses contain all of the features of the IBM Spectrum Control Standard Edition license, and the following features:

- Analytics-driven tiering that automatically moves volumes to the most cost-effective tier.
- Storage pool balancing, block storage provisioning, file provisioning, and storage reclamation.
- Snapshot-based protection that is provided by IBM Spectrum Protect[™] Snapshot for the applications and databases that IBM Spectrum Control uses.
- Roll-up reporting, in which capacity data is combined from multiple instances of IBM Spectrum Control for reporting purposes.

- Chargeback reporting for showing the capacity and the cost of the storage that is used by applications, departments, hypervisors, and physical servers.
- Consumer reporting for showing the capacity and the cost of the block storage that is used by an application, department, hypervisor, and server.

The IBM Spectrum Control Standard Edition and the IBM Spectrum Control Advanced Edition licenses are priced by the number of storage capacity units (SCUs) that are needed to cover the total usable capacity that they manage. Total usable capacity is measured in tebibytes.

The IBM Spectrum Control Standard Select Edition and the IBM Spectrum Control Advanced Select Edition licenses are priced by the number of slots in the storage enclosures that they manage, where every 25 slots require one license.

To learn about the definitions of licensing metrics and examples of pricing for the IBM Spectrum Control licenses, go to <http://www.ibm.com/common/ssi/index.wss> and follow these steps:

1. Click **Search results**.
2. Under Information Type, click **Announcement letters**.
3. Under By Geo, click the geography that you want.
4. Under By Language, click the language that you want.
5. In the Search in box, click the down arrow and click **Letter Number**.
6. In the Search for box, enter the letter number of the announcement letter for the geography that you want (for example, 217-125 for the United States). Other announcement letter numbers are:
 - AP17-0121 Asia, Pacific
 - A17-0199 Canada
 - ZP17-0159 Europe, Middle East, Africa
 - JP17-0155 Japan
 - LP17-0158 Latin America

Note: The Announcement letter might not be available in the language of your choice.

7. Click **Search**. The search returns a link to the announcement letter. Click the link.
8. Open the PDF and search Charge metric, Licensing Metric Definitions, or Pricing Examples.

To view a complete list of devices that can be used with IBM Spectrum Control, go to <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

Table 13 shows the IBM Spectrum Control licenses and their names from previous releases.

Table 13. IBM Spectrum Control licenses and their previous names

Name of IBM Spectrum Control license	Name of Tivoli Storage Productivity Center Version 5.1 license	Name of Tivoli Storage Productivity Center Version 5.2.1 through 5.2.7 license
IBM Spectrum Control Standard Edition	Tivoli Storage Productivity Center	Tivoli Storage Productivity Center

Table 13. IBM Spectrum Control licenses and their previous names (continued)

Name of IBM Spectrum Control license	Name of Tivoli Storage Productivity Center Version 5.1 license	Name of Tivoli Storage Productivity Center Version 5.2.1 through 5.2.7 license
IBM Spectrum Control Standard Select Edition	Tivoli Storage Productivity Center Select	Tivoli Storage Productivity Center Select
IBM Spectrum Control Advanced Edition	Tivoli Storage Productivity Center Advanced or IBM SmartCloud® Virtual Storage Center	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
IBM Spectrum Control Advanced Select Edition	Tivoli Storage Productivity Center Advanced or IBM SmartCloud Virtual Storage Center	IBM SmartCloud Virtual Storage Center Storage Analytics Engine

Related reference:

“Functions that are available based on role and license”

Your IBM Spectrum Control role and product license determine the functions that are available in the product.

“License compliance” on page 29

Determine the capacity of the storage systems in the environment that IBM Spectrum Control manages. Check the records of the licenses that you purchased to discover how much additional storage can be added before you no longer comply with your license.

Functions that are available based on role and license

Your IBM Spectrum Control role and product license determine the functions that are available in the product.

Users who are assigned the Administrator role or the Monitor role can use product functions. The functions that are available depend on the role that is assigned to the user:

Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs. Users who are assigned the Monitor role do not have access to administrative functions, with the following exceptions:

- They can assign storage pools to tiers.
- They can be granted permission in a service class to provision storage by using the service class. If so, they can provision storage and create a provisioning task. Users can delete provisioning tasks that they create. If the service class specifies that administrator approval is not required, the users can run or schedule the provisioning tasks that they create.

To use some functions, you must have the IBM Spectrum Control Advanced Edition license. If you have the IBM Spectrum Control Standard Edition license or the IBM Spectrum Control Standard Select Edition license, the following functions are not available:

- Storage tier optimization.
- Pool balancing.
- Block storage provisioning. You are not able to provision volumes by provisioning storage or using the vSphere Web Client extension. However, there are no license restrictions for file storage provisioning, therefore, Network Attached Storage (NAS) file shares can be provisioned.
- Chargeback reporting

The following table outlines the functions available by role. Some functions are restricted to the Administrator role. Functions that are restricted to the IBM Spectrum Control Advanced Edition license are also shown. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 14. Role and license restrictions

Resource or object	Functions available by role	Functions that require the IBM Spectrum Control Advanced Edition
Performance monitors	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about the switch or storage system that is being monitored • Open performance monitor logs • Export performance monitor information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Start or stop a performance monitor • Schedule a performance monitor 	All functions are available, depending on role.
Alerts	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about alerts • Remove alerts • Acknowledge alerts • Enable and disable alerts • Export alert information to a file <p>Users who are assigned the Administrator role can also define alert definitions and edit alert notification settings.</p>	All functions are available, depending on role.

Table 14. Role and license restrictions (continued)

Resource or object	Functions available by role	Functions that require the IBM Spectrum Control Advanced Edition
Tasks	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about tasks • Open task logs • Export task information to a file <p>If a user assigned the Monitor role has permission to provision storage by using a service class, the user can delete provisioning tasks that the user creates. If the service class does not require administrator approval, the user can also schedule and run the provision tasks that the user creates.</p> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Run tasks • Schedule tasks • Delete tasks 	All functions are available, depending on role.
Storage systems	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about storage systems • Acknowledge storage system status • Open the management GUI or element manager for a storage system • Export storage system information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Add and remove storage systems • Schedule data collection jobs • Optimize storage tiering • Add storage systems to capacity pools • Administer connections • Determine license compliance 	Optimize storage tiering
Volumes	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about volumes • Acknowledge volume status • Export volume information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Transform storage volumes • Optimize storage tiering • Reclaim storage 	<ul style="list-style-type: none"> • Transform storage volumes • Optimize storage tiering • Reclaim storage

Table 14. Role and license restrictions (continued)

Resource or object	Functions available by role	Functions that require the IBM Spectrum Control Advanced Edition
Pools	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about pools • Acknowledge pool status • Export pool information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Add pools to capacity pools • Balance pools • Optimize storage tiering 	<ul style="list-style-type: none"> • Balance pools • Optimize storage tiering
Shares	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about shares • Export pool information to a file 	All functions are available, depending on role.
Servers	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about servers • Acknowledge server status • View and collect logs • Export server information to a file <p>If a user assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to servers.</p> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Add and remove servers • Schedule data collection jobs • Provision storage to servers • Modify Storage Resource agents 	<ul style="list-style-type: none"> • Provision block storage. <p>There are no license restrictions for provisioning file storage.</p> <ul style="list-style-type: none"> • Optimize storage tiering.
Rollup servers	<p>Users who are assigned the Administrator role can complete the following actions:</p> <ul style="list-style-type: none"> • Add and remove subordinate servers • Start a probe for a subordinate server • Open probe logs (also available to the Monitor user role) • Modify the connection information for a subordinate server 	The master rollup server must be running IBM Spectrum Control Advanced Edition.

Table 14. Role and license restrictions (continued)

Resource or object	Functions available by role	Functions that require the IBM Spectrum Control Advanced Edition
Hypervisors	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about hypervisors • Acknowledge hypervisor status • Export hypervisor information to a file <p>If a user who is assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to hypervisors.</p> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Add and remove hypervisors • Schedule data collection jobs • Provision storage to hypervisors • Administer connections 	<p>Provision block storage.</p> <p>There are no license restrictions for provisioning file storage.</p>
Switches	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about switches • Acknowledge switch status • Open the management GUI or element manager for a switch • Export switch information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Add and remove switches • Schedule data collection jobs • Administer connections • Enable automatic zoning 	<p>Enable automatic zoning</p>

Table 14. Role and license restrictions (continued)

Resource or object	Functions available by role	Functions that require the IBM Spectrum Control Advanced Edition
Fabrics	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about fabrics • Acknowledge fabric status • Export fabric information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Add and remove fabrics <p>Note: For Cisco fabrics, you must remove individual switches using the Switch panels. When the last switch of a fabric is removed, the fabric is removed.</p> <ul style="list-style-type: none"> • Schedule data collection jobs • Administer connections • Enable automatic zoning 	Enable automatic zoning
Service classes	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about service classes • Export service class information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Create service classes • Modify service classes • Delete service classes 	<p>Create block-storage service classes.</p> <p>There are no license restrictions for working with file-storage service classes.</p>
Capacity pools	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about capacity pools • Export capacity pool information to a file <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Create capacity pools • Modify capacity pools • Delete capacity pools 	All functions are available, depending on role.

Table 14. Role and license restrictions (continued)

Resource or object	Functions available by role	Functions that require the IBM Spectrum Control Advanced Edition
Applications and departments	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> • View information about applications and subcomponents • View information about departments and subdepartments • View related resources assigned to applications and departments <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> • Create applications and departments • Create filters to add resources • Modify filters to change resources • Remove filters • Add resources to applications directly • Remove resources from applications • Add applications as members of other applications • Create departments • Add departments to other departments • Add applications to departments 	All functions are available, depending on role.
Reports	<p>Users who are assigned the Administrator role can complete the following actions:</p> <ul style="list-style-type: none"> • Configure the email server • Create reports • Delete reports • Edit reports • View reports <p>Users, who are assigned the Monitor role, can view the reports that are created.</p>	All functions are available, depending on role.

Related concepts:

Chapter 3, “Product licenses,” on page 21

Multiple product licenses are available for IBM Spectrum Control. Each license determines the features that you can access from the IBM Spectrum Control GUI, command-line interface, and the API. You can review the features that are included with each license and determine which license meets your storage management needs.

License compliance

Determine the capacity of the storage systems in the environment that IBM Spectrum Control manages. Check the records of the licenses that you purchased to discover how much additional storage can be added before you no longer comply with your license.

Storage capacity that is managed by IBM Spectrum Control is calculated from the license settings that are automatically assigned for each storage system following a probe job. If a license setting is not automatically assigned, you can assign it by specifying the *license model* of the storage system. The license model (capacity or enclosure) is the type of IBM license that you purchased for the storage system:

- If the storage system is a capacity license model, you must also specify the category of storage that is used by its pools.
- If the storage system is an enclosure license model, you must also specify the number of enclosures that the storage system manages.

When you finish entering the settings, the managed capacity is displayed for the storage system. For details, see “Configuring storage systems to determine license compliance” on page 33.

Tip:

- To enter settings, you must be assigned the Administrator role. Users who are assigned a Monitor role can view the information, but cannot change it.
- To enter settings, click **Settings > License compliance**.

Restriction: If you accept the automatic assignment that follows a probe job, you cannot override that assignment for the following storage systems:

- IBM System Storage DS8000
- IBM FlashSystem family: FlashSystem 840, FlashSystem 900, FlashSystem V840, FlashSystem V9000, IBM FlashSystem A9000, and IBM FlashSystem A9000R
- IBM Spectrum Accelerate and IBM XIV Storage System
- IBM Storwize V3500, IBM Storwize V3700, IBM Storwize V5000, IBM Storwize V5010, IBM Storwize V5020, and IBM Storwize V5030
- IBM System Storage SAN Volume Controller
- IBM Storwize V7000 and IBM Storwize V7000 Unified
- IBM Cloud Object Storage systems

How managed capacity is calculated for storage systems

Managed capacity is the physical capacity of the storage systems that are being monitored and that is made available for storage consumption. The physical capacity that is included in managed capacity varies, depending on the type of storage system:

- For block storage systems, managed capacity is the physical capacity that is made available to pools.
- For file storage systems, managed capacity is the physical capacity that is made available to file systems.

Raw disk capacity and virtual capacity are not included in the calculation of managed capacity.

The method by which managed capacity is calculated varies among storage systems. Refer to the information that follows.

Storage virtualizers

The following storage systems are storage virtualizers:

- IBM System Storage SAN Volume Controller
- IBM Storwize V7000 and IBM Storwize V7000 Unified

- IBM FlashSystem V840 Enterprise Performance Solution and IBM FlashSystem V9000
- Hitachi USP V and Hitachi VSP

Important: In a virtualized environment, always probe the back-end subsystems after adding their virtualizer. This ensures that managed capacity is reported correctly on the License Compliance page.

The managed capacity is the sum of the capacity available in the pools of the storage system.

Restrictions: To avoid double-counting, capacity in child pools on SAN Volume Controller and Storwize systems, and in thin-provisioned pools on Hitachi USP V and Hitachi VSP systems is not included in the managed capacity.

You can configure Storwize V7000, Storwize V7000 Unified, Hitachi USP V, and Hitachi VSP storage systems as follows:

- As storage virtualizers that virtualize the storage capacity provided by other back-end storage systems.
- As back-end storage systems that provide their capacity to other virtualizers or GPFS™ file systems.
- As a storage system that is configured as back-end storage *and* is configured to virtualize back-end storage.

The capacity that is configured as back-end storage that is provided to other virtualizers or GPFS file systems, and is used in a pool by those virtualizers or file systems, is excluded from the managed capacity. Instead, that capacity is counted as virtualized capacity.

IBM Spectrum Scale, GPFS clusters, and IBM Scale Out Network Attached Storage (SONAS) storage systems

The managed capacity is the sum of the capacity available in the pools that the file systems consume.

Restriction: Network Shared Disks (NSDs) that are not assigned to a pool are not calculated as managed capacity.

For object storage systems on IBM Spectrum Scale systems that are configured for object storage, the managed capacity is not counted separately. The managed capacity is included in the count of managed capacity for the IBM Spectrum Scale storage system to avoid double counting.

Block storage systems

Block storage systems include non-IBM storage systems and the following IBM storage systems:

- IBM System Storage DS4000® and IBM System Storage DS5000™ (extends to NetApp E series)
- IBM System Storage DS6000 and IBM System Storage DS8000
- IBM FlashSystem 840 and IBM FlashSystem 900
- IBM FlashSystem A9000
- IBM FlashSystem A9000R
- IBM Spectrum Accelerate and IBM XIV Storage System
- IBM Storwize V3500, IBM Storwize V3700, IBM Storwize V5000, IBM Storwize V5010, IBM Storwize V5020, and IBM Storwize V5030

The managed capacity is the sum of the physical capacity of all of the storage pools.

Restrictions:

- The capacity of the following types of volumes is not included in the calculation of managed capacity:
 - Volumes that are assigned to storage virtualizers and are used in pools by those virtualizers. This capacity is counted in the managed capacity of the virtualizers.
 - Volumes that are assigned as NSDs to GPFS clusters, to IBM SONAS, or to IBM Spectrum Scale storage systems, and are used in pools by those systems. This capacity is counted in the managed capacity of those systems.
- For FlashSystem A9000 and FlashSystem A9000R, the managed capacity is the overall physical capacity of the storage system.

NetApp FAS series, V series, N series, and NearStore R200

The managed capacity is the sum of the physical capacity of all of the pools on the storage systems. Some space that is used for system management is not included in the calculation of managed capacity.

The capacity of the following types of volumes is not included in the calculation of managed capacity:

- Volumes that are assigned to storage virtualizers, and are used in pools by those virtualizers. This capacity is counted in the managed capacity of the virtualizers instead.
- Volumes that are assigned as NSDs to GPFS clusters, to IBM SONAS, or to IBM Spectrum Scale storage systems, and are used in pools by those systems. This capacity is counted in the managed capacity of those systems instead.

Similar to the Storwize V7000 Unified, NetApp storage systems can be configured to provide both block storage capacity and file storage capacity. If the NetApp storage system is added as both a block storage system and a file storage system, the block capacity is ignored to avoid double-counting.

IBM Cloud Object Storage systems

The managed capacity is the sum of all storage pools that are present in all storage systems of this type and that were added to IBM Spectrum Control. It is counted towards the Near-Line SAS category of managed capacity.

Other notes about the calculation of managed capacity, total capacity, and virtualized capacity

- For thin-provisioned volumes, only the allocated space of the volume is included in the virtualized capacity. Any available space in the pool to which the volume belongs is included in the managed capacity.
- Only pool space is considered for total capacity and managed capacity. If the storage system has excess disk space that is not in a pool, this space is ignored. The following examples are situations where a storage system might have disk space that is not in a pool:
 - A rank that is configured into a RAID array on a DS8000 where the RAID array is not yet allocated to an extent pool.
 - A managed disk on a SAN Volume Controller that is not yet added to a pool.

- Available but unconfigured space in a storage system that provides Capacity on Demand (CoD).
- An NSD on a GPFS cluster that is not yet added to a pool.
- Some storage systems use pool capacity for purposes other than providing capacity to volumes. For example, when you create a volume on a DS8000, some system management space is required. Similarly, on SAN Volume Controller systems, volume migrations can use pool capacity. In these cases, even if all volumes in a pool are virtualized, there might still be some remaining capacity in the pool that is counted as managed capacity against the storage system.

Related concepts:

Chapter 3, “Product licenses,” on page 21

Multiple product licenses are available for IBM Spectrum Control. Each license determines the features that you can access from the IBM Spectrum Control GUI, command-line interface, and the API. You can review the features that are included with each license and determine which license meets your storage management needs.

Related information:



Product licenses

Learn about the different licenses that are available for the product. Each license determines the functions that are accessible in the GUI, command-line interface, and the API.

Configuring storage systems to determine license compliance

Discover the capacity that IBM Spectrum Control manages so that you can determine whether you are in compliance with your license.

To determine the storage capacity that IBM Spectrum Control manages, you must confirm the type of IBM license that you purchased for each storage system (known as the *license model*).

If you do not know the type of IBM Spectrum Control license that you purchased for each storage system, you can skip the process and return to complete the configuration later.

Whether you log on to IBM Spectrum Control for the first time or after an upgrade, the procedure for configuring license compliance is similar.

1. Click **Settings** > **License Compliance**.

- If you are logging on to IBM Spectrum Control for the first time, the Select license model page is displayed, with no selections made in the License Model column. Go to step 2 on page 34.
- If you are logging on to IBM Spectrum Control after an upgrade, the Select license model page is displayed, with the selections from the previous release in the License Model column. If some of your storage systems were not classified by license model and pool category in the last release, you must classify them by clicking the **Assign licenses** link.

Important: For some storage systems that support the IBM Easy Tier[®] IBM feature, the capacity is not calculated when you upgrade from IBM Spectrum Control V5.2.10 or IBM Tivoli Storage Productivity Center. In that case, it is necessary to probe the storage systems again. Following an upgrade and before you open the License Compliance page for the first time, be sure to run a probe of all of your managed storage systems.

2. Choose the type of license model. For storage systems that are licensed by enclosure, enter the number of enclosures that the storage system uses to create and store volumes or logical information units (LUNs).

Tip: A storage enclosure that is managed by IBM Spectrum Control is an independently powered, channel-attached device that stores data on magnetic disks or solid-state drives. An example is a disk controller and its expansion units. Each expansion unit is a separate enclosure. It can be the main controller that houses disk or solid-state drives, or the expansion chassis that houses extra disk or solid-state drives to expand the total capacity of the storage system.

3. The IBM Spectrum Control software automatically assigns the category of disk for the pools:
 - If all automatic assignments are correct, click **Finish**.
 - If one or more assignments are incorrect, click the correct category for each pool and click **Finish**.

When you finish entering information, the License Compliance page is displayed. The page presents data about the storage in your environment. If the page shows uncategorized storage systems, you can click the **Assign licenses** link to enter its license model and drive category. The page indicates the total tebibytes (TiB) of capacity that IBM Spectrum Control manages. Values are shown as tebibytes (TiB). One TiB is equal to 2^{40} (1,099,511,627,776) bytes. To determine license compliance, compare the total TiB with your proof of entitlement that you received when you purchased your licenses.

The License Compliance page is divided into a summary area, followed by a table. The summary area includes the following information:

Capacity License Model

The total number of storage systems that are licensed by storage capacity.

Total TiB

The sum of the managed capacity of all storage systems that are licensed by the capacity license model. The sum is converted to tebibytes (TiB).

SCU The total number of storage capacity units that are needed to cover the tebibytes of capacity that IBM Spectrum Control manages.

Category 1, 2, or 3

For capacity license models, the breakdown of the storage environment by the category of drives that are used by the storage system. The breakdown is shown both as a bar chart and as tebibytes (TiB).

Category Info

For storage systems that are licensed by capacity rather than enclosure, a link that describes the categories of drives that are used by pools.

Enclosure License Model

The total number of storage systems that are licensed by the number of enclosures that IBM Spectrum Control manages.

Enclosures

The total number of enclosures that are managed by storage systems that are licensed by enclosure. This value applies if you have IBM Spectrum Control Standard Select Edition or IBM Spectrum Control Advanced Select Edition.

The table contains the following information:

Storage System

The name of the storage system which IBM Spectrum Control is managing.

License Model

The type of IBM Spectrum Control license that was purchased for the storage system. The license model is capacity or enclosure. A capacity license model is based on three categories of drives that are used by the storage system or is based directly on the capacity. An enclosure license model is based on the number of enclosures that IBM Spectrum Control manages.

SCU (Storage capacity unit) A unit of measure by which a storage system is licensed for IBM Spectrum Control. For Category 1 drives (Flash and Solid State Drives), 1 SCU equals 1.00 TB. For Category 2 drives (Serial Attached SCSI (SAS) Drives, Fibre Channel Drives, and Category 3 with advanced architectures for high-end performance), 1 SCU equals 1.18 TB. For Category 3 drives (Near-Line SAS and Serial ATA drives), 1 SCU equals 4.00 TB.

Category

For capacity license models, the breakdown of the storage environment by the three categories of drives that are used by the storage system. The breakdown is shown as tebibytes (TiB). For enclosure license models, the number of enclosures that each storage system manages.

Managed capacity

The physical capacity of the storage system that is being monitored and that is made available for storage consumption. The physical capacity that is included in managed capacity varies, depending on the type of storage system. For block storage systems, managed capacity is the physical capacity that is made available to pools. For file storage systems, managed capacity is the physical capacity that is made available to file systems.

Category 1 Capacity

The total Category 1 capacity that is licensed for the storage system.

Category 2 Capacity

The total Category 2 capacity that is licensed for the storage system.

Category 3 Capacity

The total Category 3 capacity that is licensed for the storage system.

Uncategorized Capacity

The total capacity that is not categorized for the storage system.

Total Capacity

The total amount of block and file storage space that is assigned to the pools on the storage system.

Virtualized Capacity

The total amount of block and file storage space that is assigned to, and is used in a pool by, a storage virtualizer or a GPFS file system.

Non-Virtualized Capacity

The total amount of block and file storage space that is managed by the storage system, excluding the virtualized capacity. This value is the same as the managed capacity.

Chapter 4. Product overview

IBM Spectrum Control provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance, and resources. It can reduce the complexity of managing a storage environment by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, performance troubleshooting, and capacity management.

Because it is policy-based, IBM Spectrum Control can detect potential problems and automatically adjust storage and resources that are based on the policies and actions that you define. For example, it can notify you when a server or storage system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues that are related to your stored data, you can use IBM Spectrum Control to prevent unnecessary system and application downtime.

IBM Spectrum Control provides the following benefits and functions:

- Simplifies the management of storage infrastructures
- Manages, administers, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached devices
- Monitors, manages, and controls (through zones) SAN fabric components
- Manages the capacity utilization and availability of file systems and databases

Architecture

The IBM Spectrum Control consists of several components that form the infrastructure of its storage-management functions.

Data server

This component is the control point for product scheduling functions, configuration, and event information. It also includes functions that schedule data collection and discovery for the Device server.

Device server

This component discovers, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

Web server

IBM Spectrum Control uses WebSphere Application Server Liberty as the web application server to host its GUI, storage management API for cloud, and VMware vSphere web client extension.

Alert server

This component manages the complex event processing that is related to alerting on the condition of resources and their attributes. The ability to detect and be notified about configuration, capacity, and performance changes within a storage environment is important to helping you maintain and administer storage resources.

Jazz for Service Management and Tivoli Common Reporting (optional)

If you plan to use the optional Cognos Business Intelligence reporting tool, you must install Jazz for Service Management and Tivoli Common Reporting.

Tivoli Common Reporting includes Cognos BI reporting tool. You can use this tool to create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors.

For large storage environments where one server is not sufficient to manage all the IBM Spectrum Control components, you can install the product on multiple servers. In this type of environment, Jazz for Service Management and Tivoli Common Reporting can run on a separate server.

Database

A single database instance serves as the repository for all IBM Spectrum Control components.

Agents

Storage Resource agent, CIM agents, and SNMP agents gather host, application, storage system, and SAN fabric information and send that information to the Data server or Device server.

GUI Use the IBM Spectrum Control GUI to manage the storage infrastructure in an enterprise environment.

CLI Use the command-line interface (CLI) to issue commands for key IBM Spectrum Control functions.

Interfaces for IBM Spectrum Control

IBM Spectrum Control provides multiple user interfaces for managing the storage infrastructure in an enterprise environment.

IBM Spectrum Control GUI

This interface runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can also report on the condition, capacity, and performance resources in the optional Cognos Business Intelligence reporting tool. You can access this interface from anywhere that you have a web browser and connectivity to a network.

Cognos BI reporting tool (optional)

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about IBM Spectrum Control. You can access the reports from the IBM Spectrum Control GUI, and work with the reports in the Cognos BI reporting tool. The Cognos BI reporting tool is included in Tivoli Common Reporting. For more information about how to use the Cognos BI reporting tool to view and create reports about IBM Spectrum Control, see [Reporting in IBM Spectrum Control](#).

Command-line interface

Use this interface to run IBM Spectrum Control commands from a command prompt. For more information about CLI commands, go to the product documentation at [Command-line reference](#).


Related concepts:

Chapter 6, “Managing resources,” on page 107

Use IBM Spectrum Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

Starting IBM Spectrum Control

You can start IBM Spectrum Control by opening a web browser and entering a web address for the IBM Spectrum Control logon page. For example, you might enter `https://storage.example.com:9569/srm`.

Before you start IBM Spectrum Control, ensure that you are using a supported web browser. For a list of web browsers that you can use with IBM Spectrum Control, see  IBM Spectrum Control interoperability matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>. In the *Agents, Servers and Browser* column, click the version of IBM Spectrum Control that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

Start the IBM Spectrum Control GUI to administer and monitor the condition, capacity, and relationships of the resources within your storage environment.

1. On a server running the Windows operating system, start **IBM Spectrum Control GUI**. If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

`https://host_name:port/srm`

In the preceding address, specify the following values:

host_name

The IBM Spectrum Control server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

port


The port number for IBM Spectrum Control. The default port number for connecting to IBM Spectrum Control by using the HTTPS protocol is 9569. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your IBM Spectrum Control administrator for the correct port number.

Tip: If you have a non-default port, check the value of the `WC_defaulthost_secure` property in `installation_dir/web/conf/portdef.props` file.

2. From the IBM Spectrum Control logon page, type your user name and password and click **Log in**. The IBM Spectrum Control GUI opens in the browser.

Navigation


IBM Spectrum Control provides many of the functions for managing a storage environment. To access the functions in its GUI, use the menu bar at the top of the main window.

GUI element	Description
Page banner	<p>Use the banner at the top of every page to complete the following tasks:</p> <ul style="list-style-type: none">• View the name of the product license that is active.• Access resource pages and main product functions from the menu bar.• View the ID and role of the user who is logged in. The role of a user determines the product functions that are available to that user. For more information about roles, see “Role-based authorization” on page 70.• Access the online help to view information about the currently displayed page and the overall product.• Enter and exit rollup mode by using the rollup mode icon . You must add at least one subordinate rollup server for the rollup mode icon to display in the banner.

GUI element	Description
Home	<p>Dashboard</p> <p>Use the dashboard to view the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about:</p> <ul style="list-style-type: none"> • Condition and usage of resources • Entities that consume storage on those resources • Number and status of unacknowledged alert conditions that are detected on the monitored resources • Most active storage systems in your environment. <p>Performance Monitors</p> <p>Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations.</p> <p>Alerts Use alerts to be notified when certain configuration, status, and performance conditions are detected on monitored resources.</p> <p>Tasks Manage the tasks that are created when you complete the following actions:</p> <ul style="list-style-type: none"> • Assign storage to servers, hypervisors, or a cluster. • Place volumes on the tiers that match the workload requirements of the volume. • Move or convert volumes. • Move volumes from overutilized pools to pools that are less used. <p>System Management</p> <p>View information about the overall condition of IBM Spectrum Control. You can view information about the servers on which the product is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.</p>

GUI element	Description
Storage	<p>Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems:</p> <ul style="list-style-type: none"> • Modify the schedules for data collection. • View the data path of connected resources. • Assign storage to servers, hypervisors, or a cluster. • Place volumes on the tiers that match the workload requirements of the volume. • Move or convert volumes. • Move volumes from overutilized pools to pools that are less used. • Add resources to a capacity pool. • Open the management GUIs for storage systems. • Update login credentials. • Acknowledge conditions and statuses.
Servers	<p>Add servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> • Deploy Storage Resource agents to servers for full monitoring. • Modify the schedules for data collection. • View the data path of connected resources. • Assign storage to servers, hypervisors, or a cluster. • Modify and upgrade Storage Resource agents. • Acknowledge conditions and statuses.
Network	<p>Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> • Modify the schedules for data collection. • View the data path of connected resources. • Setting or modifying the zoning policy. • Open the management GUI for a switch. • Acknowledge conditions and statuses.

GUI element	Description
Groups	<p>Add applications and departments to combine resources that are grouped by applications and departments. Applications and departments are used to monitor and view status, capacity, and performance information, and view details about their related resources. You can also complete the following actions to manage and administer applications and departments:</p> <ul style="list-style-type: none"> • Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting. • View the status of resources that make up the applications and departments. • Create resource filters to automatically add resources to applications. • Add resources directly to applications and add applications to existing applications and departments. • Add departments to existing departments. • View subcomponent (member application) information to support hierarchical levels of storage grouping. • Remove applications and departments. <p>Review the distribution of capacity in your storage environment to determine whether you have sufficient capacity to meet the demands of your tiered storage.</p>
Advanced Analytics	<p>Configure IBM Spectrum Control for provisioning, optimizing, and reclaiming storage. It guides you through the steps for managing your storage such as creating service classes and capacity pools to provision storage and creating tiers to optimize storage placement. You can complete the following tasks:</p> <ul style="list-style-type: none"> • Provision servers, hypervisors, or a cluster • Optimize storage by re-tiering volumes and balancing pools • Reclaim storage that is not being used

GUI element	Description
Reports	<p>Use the Email Reports option to create, configure, schedule, and send chargeback reports for applications, departments, hypervisors, and servers.</p> <p>Use the web-based GUI when you want to individually view details about resources, for example when you want to troubleshoot a resource.</p> <p>You can also use the optional Cognos Business Intelligence reporting tool to analyze multiple storage systems, switches, servers, and hypervisors. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor.</p> <p>Use the Cognos BI reporting tool to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports.</p> <p>When you click a Cognos BI reporting option, a separate browser tab is opened.</p> <p>For information about how to get started with the Cognos BI reporting tool, see  http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.14/com.ibm.spectrum.sc.doc/fqz0_c_getting_started_reports.html.</p>
Settings	<p>Complete the following actions to customize IBM Spectrum Control for your environment:</p> <ul style="list-style-type: none"> • Modify the authentication repository and assign IBM Spectrum Control roles to user groups. Roles determine the product functions that are available to users. • Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control. • Add and configure rollup servers so you can view capacity and status information about resources that are managed by your IBM Spectrum Control servers.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.

- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Replication products

Tivoli Storage Productivity Center for Replication and IBM Copy Services Manager, which is built on existing replication technology, are the replication tools available for use with IBM Spectrum Control.

Availability

Tivoli Storage Productivity Center for Replication is available under the following conditions:










- In Tivoli Storage Productivity Center V5.2.7 and earlier, Tivoli Storage Productivity Center for Replication is automatically available.
- In IBM Spectrum Control V5.2.8 and later, Tivoli Storage Productivity Center for Replication is only available if you upgrade from Tivoli Storage Productivity Center V5.2.7 or earlier.

IBM Copy Services Manager is available with every IBM Spectrum Control V5.2.8 or later license.

Information about Tivoli Storage Productivity Center for Replication

Information about how to use Tivoli Storage Productivity Center for Replication is available in the IBM Knowledge Center for Tivoli Storage Productivity Center V5.2.7. Use the following links to view this information:

Table 15. Information about Tivoli Storage Productivity Center for Replication

Task	Link to related information
Learn about Tivoli Storage Productivity Center for Replication and its key concepts	 Product overview
Plan for using the product.	 Planning
Starting Tivoli Storage Productivity Center for Replication	 Starting the GUI
View scenarios that illustrate how replication can be used.	 Scenarios
Manage replication components, including setting up the management servers, adding storage connections and paths, creating sessions and copy sets, setting up data replication, performing planned maintenance, and recovering from a disaster.	 Managing replication
Use CLI commands to access functions within the product.	 CLI commands
View information about configuration files	 Configuration files
View explanations of messages that might be generated when you use the product	 Messages
View product documentation	 ftp://public.dhe.ibm.com/software/tivoli/tpc/v527/

Information about IBM Copy Services Manager

IBM Copy Services Manager provides automation for 2-site replication, 3-site replication, and advanced copy services from a single point of control across multiple platforms.

Important:

- The license enablement files for Copy Services Manager are included with IBM Spectrum Control.
- The installation files for Copy Services Manager are available on IBM Fix Central.

To learn more about how to download and install it, see [🔗](#) Copy Services Manager.

Key concepts

This section contains a technical overview that will help you understand how IBM Spectrum Control works. An understanding of the concepts in this section will help you use IBM Spectrum Control effectively.

Data collection

To help you implement a storage management strategy, it is critical that you determine the information that you want to gather about the resources within your environment. You can schedule different data collection jobs depending on the types of information that you want to gather and the resources that you want to monitor.

Use probes and performance monitors to collect detailed information about resources. You must schedule and run these data collection jobs before you can complete other tasks within IBM Spectrum Control, such as viewing reports, setting alerts, viewing data paths, and managing storage systems.

You schedule probe and performance monitor jobs for a resource when you add the resource for monitoring. After you add the resource, you can modify the probe and performance monitor job schedules on the list and details pages for the resource. For example, to modify the probe job schedule for a storage system, go to the Storage Systems page or the Storage System details page in the web-based GUI.

Schedule data collection jobs to gather the following types of information:

Asset and status information

Use probes to collect asset, storage statistics, and status information about resources. You can run probes on all the resources that are monitored by IBM Spectrum Control.

Performance monitoring

Use performance monitors to collect metrics that measure the performance of switches and storage systems.

Performance monitoring and troubleshooting

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

To monitor the performance of resources and check for threshold violations, complete the following tasks:

- Add resources for monitoring and schedule data collection
- Define performance thresholds
- View and troubleshoot performance issues

Collect performance data

Before you can troubleshoot and view reports about performance, you must collect data about monitored resources. Performance monitors are data collection jobs that gather performance information about resources. This information includes metrics that measure the performance of the components within a resource. Metrics measure the performance characteristics of volumes, ports, and disks on storage systems and switches. IBM Spectrum Control provides many different metrics for measuring performance. For example, some key metrics for storage systems are I/O rate in I/O operations per second, data rate in MiB per second, and response time in milliseconds.

You can use metrics in IBM Spectrum Control to track growth or change in I/O rates, data rates, and response times. In many environments, I/O and data rates grow over time, and response times increase as those rates increase. This relationship can help with "capacity planning" for your storage. As rates and response times increase, you can use these trends to project when more storage performance and capacity is required.

Define performance thresholds

Performance thresholds alert you when the performance of a monitored resource falls outside of a specified range and might represent a potential problem. When you define a performance threshold for an internal component of a resource, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, a threshold violation occurs.

For example, you can define a performance threshold that alerts you when the overall back-end response time for a managed disk on a SAN Volume Controller exceeds a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

View and troubleshoot performance issues

After data collection and performance thresholds are configured, you can use the web-based GUI to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance of a resource fell outside of a specific range.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.

- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller storage system is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server for monitoring, IBM Spectrum Control automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, you can view the performance of the internal resources that are directly associated with the SAN storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.
- In the optional Cognos Business Intelligence reporting tool, you can also view and create performance reports about multiple resources.

Related concepts:

“Monitoring the performance of resources” on page 217

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

Related tasks:

“Adding resources” on page 112

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Related reference:

“Performance metrics” on page 403

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can view performance information or specify alerts for when the performance of a resource falls outside a specified range.

Applications, departments, and general groups

Use applications and departments to organize the resources in your storage environment into a hierarchy that matches the structure of your business organization. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. Use general groups to group your resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location in the GUI.

To manage and administer applications and departments, complete the following tasks:

- Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.
- View the status of resources that make up the applications and departments.
- Create resource filters to automatically add resources to applications.

- Add resources directly to applications
- Add applications to existing applications and departments.
- Add departments to existing departments.
- View subcomponent (member application) information to support hierarchal levels of storage grouping.

Applications

An application is a program or a project that consumes storage resources within an organization and interfaces with other enterprise groups that are important to the running of a business. Use IBM Spectrum Control to model the storage usage that is consumed in your environment by assigning the usage to applications to see the overall health status.

An application can be part of a department and have its own subcomponents (member applications) that are used to create a five level deep hierarchy. Applications that are grouped together can range from large line of business systems to specialized software, in a department, that runs on either client computers or servers. For example, an application might be an automated billing system within the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

The Applications page shows capacity information about the applications that are monitored by organization and interfaces with other enterprise groups that are important to the running of a business. Use IBM Spectrum Control to view and manage the resources that are defined to the applications. If the application is associated with a department, information about that department is also shown.

Departments

A department is a division within a business. Use IBM Spectrum Control to model the storage capacity that is consumed in your department for your business environment, in accordance with other department members.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or an application even if they do not belong in the same hierarchy. For example, in a collaboration scenario, a single IBM SAN Volume Controller might be shared by multiple departments.

The Departments page shows storage capacity information about the top level departments, the subdepartments and any applications that belong to the department that are monitored by IBM Spectrum Control.

General groups

Create general groups to quickly view information about storage resources that have common characteristics. For example, you might group the subset of ports on a SAN Volume Controller that are used for inter-node communication or the storage systems that are used by a critical business application.

You can organize your storage resources into a general group hierarchy. Organizing resources into general groups and their subgroups can be helpful when you want

to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group. For example, you can group the resources that are used by your production application so that you can monitor all the application resources and separately monitor the specific subgroups of storage systems and ports.

Storage optimization

To optimize the resources in your storage environment, you can create tasks to balance pools, re-tier volumes, or transform volumes that are in storage virtualizer pools.

Balancing the workload of volumes across storage pools

You can balance the workload of volumes across pools on the same tier. The pools are analyzed and recommendations are generated to move volumes from pools with high-activity values to pools with low-activity values.

Re-tiering volumes

You can re-tier volumes to balance pools or to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes that are based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

Transforming volumes

You can complete the following tasks for one or more volumes in storage virtualizer pools:

- Move volumes from one storage virtualizer pool to another pool in the same storage virtualizer pool.
- Move volumes in a storage virtualizer pool to a pool that is enabled for Easy Tier.
- Convert fully allocated volumes to thin-provisioned volumes and convert thin-provisioned volumes to fully allocated volumes.
- Convert fully allocated volumes to compressed volumes and convert compressed volumes to fully allocated volumes.

The pools are analyzed and recommendations are generated. For example, analysis occurs to ensure that there is sufficient space in the pool to convert volumes or add volumes.

Related tasks:

“Optimizing storage pools” on page 325

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Optimizing storage tiering” on page 319

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

“Managing tasks for tiering storage, balancing pools, and transforming storage” on page 346

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Related reference:

“Criteria for identifying the pools that require balancing” on page 327

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

Cloud configuration

To take advantage of the simplified provisioning and optimization capabilities in IBM Spectrum Control, some configuration is required.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, IBM Spectrum Control can analyze and optimize volume performance. To take advantage of these capabilities, create service classes and, optionally, capacity pools.

A tutorial of cloud configuration concepts is available in the GUI. In the menu bar, go to **Advanced Analytics > Cloud Configuration**, and click **Learn the Concepts**.

Service classes

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

Service classes simplify provisioning requests by representing a level or type of storage quality. When you are requesting storage, only the required capacity and service class must be specified. Before you can provision storage, you must create service classes that describe the capabilities and characteristics of the storage you want to be able to provision. Service classes can later be modified or deleted as the needs of your installation change.

Block-storage service classes

A block-storage service class describes attributes and capabilities of block storage resources. When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a storage pool for the volume.

File-storage service classes

A file-storage service class describes attributes of file storage resources. The file storage resources include Network Attached Storage (NAS) filers, Network Shared Disks (NSDs), and file systems. When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a file system or IBM SONAS NSD for the share. Space on an IBM SONAS NSD is allocated if a new file system is required or an existing file system requires more space.

Service class properties

A service class represents storage resources with common traits. These common traits are the service class properties, and collectively describe a particular quality of service. Before you create service classes, consider the levels of storage quality that you want to have available for provisioning requests. For block storage, for example, you can organize service classes by storage tier or RAID level. Some service class properties are used during provisioning to identify the best location for storage placement. Other service class properties determine how the storage is configured, and reflect the performance, reliability, or security considerations for the level of service.

In addition to the standard properties of a service class, you can create your own properties by using custom tags. Custom tags can represent any common trait that storage systems must have. For more information about creating your own properties by using custom tags, see “Custom tags” on page 55.

Block-storage service classes can specify the following properties:

- Storage tier or range of storage tiers
- RAID level
- Whether volume virtualization is required or not allowed. This property is available only when at least one IBM storage virtualizer is managed by IBM Spectrum Control. The IBM storage virtualizers include SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified.

If volume virtualization is required, the VDisk mirroring property specifies whether volumes are mirrored to a second storage pool, which uses different backend storage, on the same IBM storage virtualizer.

- Whether volumes must be thin provisioned or must not be thin provisioned. When a thin-provisioned volume is created in a DS8000, an extent space efficient (ESE) volume is created.

If thin provisioning is required, you can specify thin-provisioning configuration properties for the IBM storage virtualizers and IBM XIV Storage System. The following properties might be available:

- If the volume is created on an XIV, the locking behavior configuration property affects the volume that is created.
- If the volume is created on an IBM storage virtualizer, the following configuration properties affect the volume that is created:
 - Allocated space
 - Auto-expand
 - Warning level
 - Grain size

- Whether volumes are compressed. This property is applied only when an IBM storage virtualizer has at least one I/O group that contains a compressed volume. If there are no existing compressed volumes in an I/O group, IBM Spectrum Control will not create compressed volumes.

This property is available only when the thin provisioning property is enabled.

- A multipathing policy. If a server has a Storage Resource agent (SRA) and uses an IBM System Storage Multipath Subsystem Device Driver (SDD), the server multipathing policy is used to configure the driver. If provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control, the VMware vSphere multipathing policy is used to configure the multipath driver on the hypervisor.

If a multipathing policy is set, and fabrics are managed by IBM Spectrum Control, the following properties affect how the storage system is connected to the host:

- Whether the storage system and the client host are connected through fully redundant paths
- Number of paths between the storage system and the client host
- If at least one DS8000 Encryption Group is discovered by IBM Spectrum Control, the following properties are available and affect the storage placement and configuration of volumes:
 - Whether volumes must be encrypted
 - If volumes must be encrypted, the DS8000 encryption group to use

File-storage service classes have the following properties:

- Whether dedicated storage is required, or whether shared storage is allowed. In other words, whether the file system from which the NAS share is provisioned can contain other NAS shares.
- Fileset type. The fileset type is either an independent fileset with its own allocated nodes, or a dependent fileset allocated on the file system.
- Access path host name replacement. For more information, see “Special consideration for provisioning by using the VMware vSphere Web Client extension” on page 56.

Predefined service classes

To illustrate service classes, a set of predefined service classes are provided. From the Service Classes page in the web-based GUI, you can view any of these service classes for more information. Select a service class and then select **Actions > View / Modify**. To display help information about a service class property, complete the following steps:

1. Move the mouse pointer over the field or control.
A question mark icon is displayed next to the field.
2. Move the mouse pointer over the question mark icon to display an explanation of the property.

The following sample block-storage service classes are provided for illustration.

Gold The Gold service class represents the highest-performing storage resources for mission-critical applications.

Silver The Silver service class represents high-performing storage resources for applications in production.

Bronze
The Bronze service class represents standard storage resources for non-mission-critical applications.

The following sample file-storage service classes are provided for illustration.

NormalIsolation

The NormalIsolation service class describes normal isolation file storage by specifying that shared storage is allowed. The file system from which the NAS share is provisioned can contain other NAS shares.

EnhancedIsolation

The EnhancedIsolation service class describes enhanced isolation file

storage by specifying that dedicated storage is required. The file system from which the NAS share is provisioned cannot contain other NAS shares.

Custom tags

During provisioning, IBM Spectrum Control determines candidates for storage placement by comparing the requirements of a service class with the known capabilities, configuration, and performance of the available storage systems. You can think of the storage placement determination as a filtering process. IBM Spectrum Control filters the set of available storage resources against properties of the service class until only the storage resources that can satisfy all of the requirements remain.

In addition to the standard properties of a service class, you can create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

When you specify tags on a block-storage service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you specify tags on a file-storage service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

During provisioning, IBM Spectrum Control filters candidates for provisioning against the custom tags.

Before you create service classes, consider whether there are any special requirements for provisioning that are not addressed by the standard properties of the service class. If so, tag the appropriate storage resources to satisfy custom requirements of a service class. When you are defining the service class, specify the custom tags that are required.

Adding users to a service class

By default, only administrators can provision storage. However, an administrator can add users to a service class to grant them permission to provision by using the service class. In order for an administrator to grant permission to provision to a user, the user must be assigned to the Monitor or External Application role.

When you are granting users permission to provision by using a particular service class, you can also specify whether execution of provisioning plans created by these users requires administrator approval. The requirement for administrator approval applies to all non-administrative users who are added to the service class, and also applies to external applications such as the vSphere Web Client extension. Approval cannot be required for some non-administrative users and not for others.

Before you create service classes, consider whether you want to allow users to request their own storage. Consider whether you want user provisioning requests to require administrator approval.

Associating capacity pools with a service class

Before you create service classes, consider whether certain storage requests must always be satisfied from a particular set of resources. If so, you can associate one or more capacity pools with a service class. For example, suppose that a set of storage resources are allocated to a particular department of your business. All storage requests for that department must be satisfied by those storage resources. In this case, you can add all the storage resources to a capacity pool and associate that capacity pool with a service class. You can also add the members of the department to the service class. The department members can then request their own storage as needed. However, the requests can be satisfied only by the storage resources that are allocated to the department.

When you are configuring capacity pools and service classes, keep in mind that a storage resource can be a member of only one capacity pool. Adding storage resources to a capacity pool is not a viable approach if the resources are already members of another capacity pool. Consider specifying custom tags to achieve the same result. In the preceding example, suppose that some or all of the storage resources that are allocated to the department are already members of capacity pools. In this case, you can instead tag the storage systems and the service class with matching tags to achieve the same result. With either approach, only the storage systems that are allocated to the department are candidates for provisioning when storage is requested by using the service class.

Special consideration for provisioning by using the VMware vSphere Web Client extension

When a file share is provisioned by the VMware vSphere Web Client extension, an access path to the share is provided by IBM Spectrum Control. By default, the path includes the cluster name of the file storage system in which the share is created. A cluster name must be a valid host name that is registered with a Domain Name System (DNS) server. If the cluster name is not valid, the share is not accessible under the access path that is returned by IBM Spectrum Control.

To address this situation, you can set one of its custom tags for a file storage system to specify an access path host name replacement. A file-storage service class can specify which of the file storage system custom tags contains the access path host name replacement.

If file shares are provisioned by using the VMware vSphere Web Client extension, determine whether you must set access path host name replacements.

Capacity pools

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

The following types of storage resources can be grouped into capacity pools:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM SONAS Network Shared Disks (NSDs)

You can organize storage resources into capacity pools in any way that serves your business needs. For example, you can separate the storage resources that are allocated for separate divisions of your business into separate capacity pools. You can then track the storage use for each division separately, and restrict provisioning requests to the appropriate set of storage resources.

Tracking used and available space in a capacity pool

You can track the used and available space for any set of storage resources by adding the storage resources to a capacity pool. From the Capacity Pools pane of the GUI, you can view the total capacity of the resources, and monitor the following capacity measurements:

- Used file space
- Available file space
- Used block space
- Available block space

Restricting a provisioning request to a capacity pool

You can use capacity pools to define a set of storage resources from which provisioning requests must be satisfied. When you are provisioning storage or using vSphere Web Client extension, you can specify a capacity pool. If you do, the provisioning request is restricted to resources in the capacity pool. Only those resources are candidates for provisioning.

You can also associate a service class with capacity pools. If a service class is associated with capacity pools, provisioning requests for the service class must specify, and be constrained to, one of the associated capacity pools. Because you can also grant non-administrative users permission to provision storage by using the service class, associating capacity pools with the service class restricts the users to a specific set of resources.

Provisioning storage

IBM Spectrum Control guides you through the steps for provisioning storage volumes or network-attached storage (NAS) shares. You can provision storage volumes or NAS file shares to one or more servers, one or more hypervisors, or one cluster.

Storage requirements

When you set up provisioning, you must specify your storage requirements by defining a service class. Certain properties of the service class describe capabilities that storage resources must have in order to provide the service class. For example, when you provision volumes, a block-storage service class specifies properties such as a storage tier, a RAID level, and whether the pool must be able to encrypt or thin provision volumes.

A service class typically represents a quality or service. For example, IBM Spectrum Control provides predefined block-storage service classes named Gold, Silver, and Bronze. Gold represents the highest-performing storage resources, and Silver and Bronze represent lesser levels of storage quality.

The different levels of storage quality that are provided by each service class are defined by the service class properties. For example, each service class specifies a

different required storage tier that depends on the level of service that is required. Because the Gold service class is for mission critical applications, it sets thin provisioning off.

When IBM Spectrum Control guides you through the steps for provisioning, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require. If you are unfamiliar with the service classes, however, you can open a separate window to view service class information. If you have Administrator privileges, you can modify or create service classes.

Storage constraints

IBM Spectrum Control identifies the storage resources that can provide the capacity and the service class from a set of storage resources. This set might be all the storage resources that are known to IBM Spectrum Control, or it might be constrained to a subset of those resources. In particular, your site can create capacity pools. *Capacity pools* are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business.

A service class might restrict storage placement to one or more capacity pools. If the service class you are provisioning from restricts placement, you must select one of the allowed capacity pools. If the service class does not restrict storage placement, you can specify that candidates for provisioning can be selected from all storage resources that are known to IBM Spectrum Control, all storage resources in any capacity pool, or storage resources in a particular capacity pool. If you are unfamiliar with the capacity pools that you can select when you set up provisioning, you can open a separate window to view capacity pool information. If you have Administrator privileges, you can modify or create capacity pools.

If the fabrics are managed by IBM Spectrum Control, the set of storage resources might be further constrained by your zoning policy. If the zoning policy is not configured to automatic zoning, only storage systems with connectivity to the server or hypervisor are candidates for storage placement.

Provisioning volumes

Restriction: To provision volumes, you must have the IBM Spectrum Control Advanced Edition license.

To provision volumes, you must select one or more servers or hypervisors. You can request one or more volumes. For each volume, you specify a volume name, the capacity that is required, a service class, and, optionally, a capacity pool. For each volume, IBM Spectrum Control identifies the storage pools that can provide the capacity and the service class. From the set of pools that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the unallocated volume space in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is then given to systems that have available performance data.

On storage systems, volumes are allocated in increments of a set unit size that depends on the storage system type. Depending on the storage system that was

identified as the best location for the storage, the volumes that are allocated might be larger than the capacity you requested for the following reasons:

- On SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified, volumes are allocated in increments of an extent size. The extent size is set when the MDisk group is created. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On a DS8000, volumes are allocated in increments of a 1 GiB fixed extent size. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On an XIV, volumes are allocated in fixed increments of 16 GiB. The capacity that you requested is rounded up, if necessary, to complete the 16 GiB increment.

Provisioning shares

To provision a share, you must first select one or more servers or hypervisors. You specify the capacity that is required for the share, a service class, and information about how to export the share. IBM Spectrum Control identifies the file systems and Network Shared Disks (NSDs) that can provide the capacity and the service class. From the set of file systems and NSDs that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available space on the file system or NSD.

The provisioning task

When you complete the steps for setting up provisioning, IBM Spectrum Control creates a provisioning task for the share or volume. If you request multiple volumes, a separate task is created for each unique service class and capacity pool combination you specify in your volumes request. If a requested volume is not constrained to a capacity pool, the set of all available storage resources is considered the capacity pool. If you have Administrator privileges, you can save, run, or schedule provisioning tasks.

In the service class, an administrator can grant users permission to provision by using the service class. In the service class, an administrator can also specify whether scheduling or running provisioning tasks that are created by using the service class requires administrator approval. If you are a user with permission to provision by using the service class, and administrator approval is not required, you can save, run, or schedule the task. If administrator approval is required, you can only save the task. If you save the task, an administrator can later schedule or run it.

Although IBM Spectrum Control identifies the best location for storage when it creates the provisioning task, the implementation of the task might fail. For example, the implementation of the task might fail because of changes in the environment that occur after the task is created and before it is run. The implementation of the task might also fail because of storage system restrictions that are not considered when the task is created.

Viewing task details and logs

When a provisioning task is running, you can view its status in a details page. If you run the provisioning task immediately after you complete the steps for setting up provisioning, the details page is already displayed. You can also display details

of a provisioning task from the Tasks page. From the details page, you can open logs for the provisioning task. The logs show the steps that are taken by the task during processing and include detailed information about any warnings or errors.

Tracking

When you set up provisioning, you can specify a ticket identifier for tracking purposes. The ticket identifier is associated with the provisioning task, and with any volume or share that is created by the provisioning task.

The ticket identifier can be viewed by showing the Ticket column in the Volumes page or the Shares page. The ticket identifier can also be viewed in the properties notebook for a volume or share.

Tip: If you request multiple volumes, a separate provisioning task is created for each unique service class and capacity pool combination you specify in your volumes request. By specifying a ticket identifier when you request the volumes, you can easily track the set of provisioning tasks and volumes that are created by the single request.

Storage configuration

Whether you are provisioning volumes or shares, IBM Spectrum Control configures the storage and resources as specified by certain properties of the service class, and, in the case of volumes, according to your zoning policy.

When you provision volumes, IBM Spectrum Control configures volumes and resources according to properties of the service class:

- If the service class requires thin provisioning, the volume is created in a pool that can thin provision volumes. During provisioning, other properties of the service class are used to configure the thin provisioned volume for the particular type of storage system that contains the pool. Because you configure thin provisioned volumes differently for the different storage system types, the service class has a unique set of thin provisioning properties for each type.
- If the service class specifies a multipathing policy for servers, the IBM System Storage Multipath Subsystem Device Driver (SDD) on the server is configured to use that policy. For the SDD to be configured, a Storage Resource agent (SRA) must be running on the server. If a VMware vSphere multipathing policy is specified for hypervisors, the policy is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.

When you are provisioning volumes, if automatic zoning is enabled, new zones might be created to connect a server to the storage system. Existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more new zones are created between a host initiator port and a controller, node, or module port.


When you provision shares, IBM Spectrum Control configures at least one file access protocol for the file share.

Storage reclamation

You can use volume reclamation recommendations to reclaim unused storage capacity in your environment and use your storage more efficiently.

The block storage systems in your data center that you add for monitoring are regularly analyzed. You can view a list of the volumes that are allocated space but are not used for storing data. By reclaiming the volumes, you can recycle existing storage space instead of purchasing new storage media.

The following figure shows an overview of storage reclamation:



Reclamation recommendations which are based on capacity data are available after a device probe, while recommendations which are based on performance data are available after the specified analysis period. By default, capacity data for one day and performance metrics about the I/O activity for 14 days are collected. Volumes are identified as reclaimable if, for example, they have no I/O activity over the previous 14 days or if they are not assigned to servers.

For a more thorough analysis, you can change the number of days for which performance metrics are collected. If you shorten the analysis period, less data is used to determine I/O activity. When you extend or shorten the analysis period, it is also changed for other optimization features such as tiering and balancing.

To identify volumes that can be reclaimed, you enable performance monitoring for the block storage systems. Only block storage systems that support performance monitoring can be analyzed for I/O activity.

Restriction: Volumes are identifiable as reclaimable in IBM Spectrum Accelerate storage systems based on I/O activity. Information about the assignment of volumes to servers is not available.

Supported storage systems: Go to IBM Spectrum Control interoperability matrix, and then click the release number in the **Storage** column.

IBM Spectrum Control VASA provider

IBM Spectrum Control VASA provider enables the IBM Spectrum Control administrator to bring visibility of IBM Spectrum Control managed storage resources to a vCenter server administrator.

About vSphere API for Storage Awareness APIs Providers (VASA)

VASA enables coordination between VMware vCenter and storage arrays. It is designed to provide built-in storage insight to support VM storage provisioning, assist with storage troubleshooting and enable new DRS-related use cases for storage.

VASA benefits


For VMware vSphere users, the IBM Spectrum Control VASA provider improves the ability to monitor and automate storage-related operations in VMware environments.

A IBM Spectrum Control storage administrator can share certain IBM Spectrum Control alerts and storage capabilities for resources that are mapped to a vCenter server.

The IBM Spectrum Control VASA provider provides the underlying connectivity for alerts and information to be shared.

IBM Spectrum Control VASA provider architecture

This figure shows the IBM Spectrum Control VASA provider running on a IBM Spectrum Control connected to both the VMware vCenter server and the managed storage resources.



A IBM Spectrum Control storage administrator can share certain IBM Spectrum Control alerts for resources that are mapped to a vCenter server.

The IBM Spectrum Control VASA provider provides the underlying connectivity, and is automatically deployed and running after a IBM Spectrum Control installation. You must register a IBM Spectrum Control server to view its information in vCenter storage reports and views.

vSphere Web Client extension for IBM Spectrum Control

With the vSphere Web Client extension for IBM Spectrum Control, you can manage virtual storage in a vSphere environment. You can use the vSphere Web Client extension to complete provisioning tasks and to view information about storage resources in the vSphere Web Client.

Provisioning storage in the vSphere Web Client extension

You can use the vSphere Web Client extension to provision storage capacity on storage systems that are monitored by IBM Spectrum Control. You can provision file and block storage and use the service classes and capacity pools that you create in IBM Spectrum Control. Additionally, you can choose to create data stores on your storage volumes.

Viewing storage reports in the vSphere Web Client extension

You can use the vSphere Web Client extension to view information about the storage systems that are monitored by IBM Spectrum Control. You can view the mapping of your virtual storage resources to storage systems that are monitored by IBM Spectrum Control. You can view metrics for your storage systems. You can also view information about the fabric that is connected to a storage adapter. From the storage reports, you can click links to open the resources in the IBM Spectrum Control GUI.

vSphere supported locales

The locales that are supported by the vSphere Web Client extension for IBM Spectrum Control depend on the vSphere support. In vSphere 5.x, the following locales are supported:

- English: en_US
- German: de_DE
- Japanese: ja_JP
- Simplified Chinese: zh_CN
- French: fr_FR
- Korean: ko_KR

When a locale does not have a translated version, the vSphere Web Client extension is displayed in English, which is the default language.

Related concepts:

“Provisioning storage with the vSphere Web Client extension” on page 316
You can use the vSphere Web Client extension for IBM Spectrum Control to manage and monitor your virtual environment and storage devices.

Alerts

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to configuration and performance conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Alerts are generated when particular conditions are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource

For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource. Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem V840, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. In general, the following types of conditions can trigger alerts:

- An attribute or configuration of a resource changed
- The capacity of a resource fell outside a specified range
- The performance of a resource fell outside a specified range
- The storage infrastructure was changed, such as a new or removed resource
- Data is not being collected for a resource

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over used or under used.

IBM Spectrum Control provides many metrics for measuring performance and determining violations of the thresholds that you specify.

Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, IBM Tivoli Netcool/OMNIbus events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

Acknowledging alerts: Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, IBM Tivoli Netcool/OMNIbus events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable IBM Spectrum Control to listen to SNMP traps. The default port number is 162, and the default community is public.

Related concepts:

“Alerting and event processing” on page 148

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Related tasks:

“Acknowledging alerts” on page 193

Some alerts in IBM Spectrum Control are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

Reporting

Use reporting functions to view overview and detailed information about your storage.

You can view detailed capacity information for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual

machines. For example, you can view the total amount of storage space that is committed to a volume, the amount of allocated space that is used by a volume, and the amount of space that is allocated to a volume and is not yet used.

You can configure rollup servers to collect information from the resources they manage and roll that data up to master servers. This gives you a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.

You can also use the optional Cognos Business Intelligence reporting tool to view predefined reports and create custom reports about the resources managed by IBM Spectrum Control. You access reports from the IBM Spectrum Control GUI, and work with the reports in the Cognos BI reporting tool. Tivoli Common Reporting provides the Cognos BI tool for IBM Spectrum Control reports.

The Cognos BI reporting tool to view over 70 predefined reports about the capacity and performance of your resources in Tivoli Common Reporting. Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics.

You can schedule reports and specify to create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use IBM Spectrum Control monitoring jobs, such as probes and performance monitors, to gather comprehensive information and statistics about your storage resources.

Rollup servers

Rollup reporting combines capacity and status information from multiple instances of IBM Spectrum Control for monitoring and reporting. In rollup mode, subordinate servers collect information from managed resources and roll that data up to master servers.

Rollup capability requires two types of servers: a master rollup server running IBM Spectrum Control, which collects capacity and status information from one or more subordinate rollup servers. Subordinate servers also have IBM Spectrum Control or Tivoli Storage Productivity Center V5.2.0 or later installed, and communicate with the master server during server probes.

The master server collects capacity and status information about storage resources that are managed by the subordinate servers. This master/subordinate server configuration is called *rollup mode*. When you enable rollup mode on the master server, you can see capacity and status information in the rollup resource pages for the storage systems, servers, hypervisors, switches, and fabrics that are monitored by your rollup servers.

You can also configure a master rollup server as a subordinate server to other master servers. This allows you to create a comprehensive rollup reporting structure across your enterprise.

Note: If you configure a master rollup server as a subordinate of another master server, the data collected by that master server is not rolled up to the higher

master server. Master rollup servers only collect and display data from their immediate subordinate servers. For example, if you configure master server 2 as a subordinate to master server 1, the data that is rolled up to master server 2 from its subordinate servers is not collected by master server 1. Master server 1 only collects capacity and status information about master server 2, not the subordinate servers monitored by master server 2.

Master rollup servers must have IBM Spectrum Control V5.2.10 or later installed. You can add subordinate servers that are running any version of IBM Spectrum Control or Tivoli Storage Productivity Center V5.2.0 or later.

Restriction: The following restrictions apply when using rollup mode:

- If the subordinate server is running a version of Tivoli Storage Productivity Center or IBM Spectrum Control earlier than V5.2.10, the master server might not be able to determine the correct Web server port to connect to on the subordinate server. In this case, the default port 9569 will be used. The Web server port is used to create links to subordinate servers in the rollup drop-down menu and the rollup resource pages, such as the Switches (Rollup) page.
- The master server must be running IBM Spectrum Control Advanced Edition or Advanced Select Edition V5.2.10 or later. Subordinate servers can have any version of IBM Spectrum Control, or Tivoli Storage Productivity Center V5.2.0 or later, installed and running.
- After you upgrade a master rollup server from Tivoli Storage Productivity Center to IBM Spectrum Control V5.2.10 or later, you must run probes of the subordinate servers that were added to the master server prior to the upgrade. Use the Start Probe action to run probes of the subordinate servers.

Managing subordinate servers

After you add a subordinate server on the **Settings > Rollup Server Connections** page, you can view information about the server and perform actions on it. You must have IBM Spectrum Control Administrator role authorization to add, remove, and manage subordinate servers on this page.

Information about subordinate servers

IBM Spectrum Control provides the following information about the subordinate servers on the **Rollup Server Connections** page:

Managed resources

The number of block storage systems, servers, switches, fabrics, and hypervisors that the subordinate server is managing. You can view information about these resources on the rollup resource list pages for the resources, such as the Switches (Rollup) page. NPV switches are not included in this number on either the master or subordinate Switches (Rollup) resource list pages.

Data collection

The status of the last probe of the subordinate server, and the time of the next probe of the server. Hourly probes are automatically scheduled when a subordinate server is added to a master server.

Actions

You can perform the following actions on the subordinate server. Unless otherwise noted, you must have Administrator role authority to perform these actions.

Open logs

Opens the probe logs window for the subordinate server in a separate browser window. Any IBM Spectrum Control user role can perform this action.

Start probe

Starts a probe of the subordinate server. If a probe is already running for the server, you must wait until the current probe completes before starting a new probe.

Remove server

Removes the subordinate server from the **Rollup Server Connections** page and the **Servers (Rollup)** list. Also removes any associated resources that are managed by the subordinate server. For example, if a subordinate server is the managing instance for a switch and the subordinate server is removed from the master server, the switch will no longer display on the **Switches (Rollup)** resource list page.


Note: Removing a subordinate server on the **Rollup Server Connections** page does not remove that server from the non-rollup **Servers** page. If the server has been added by using the **Add Server** dialog on the non-rollup **Servers** page, IBM Spectrum Control will continue to manage that server unless you remove it from the **Servers** page.

Modify connection

Allows you to change the host name, user name, password, and port of the subordinate server.

Rollup mode

Rollup mode is a special view of your enterprise and is enabled by configuring your IBM Spectrum Control servers to communicate and gather data through server probes. Rollup mode requires a master server configured to collect capacity and status information from one or more subordinate servers. Once you have configured a master server with at least one subordinate server, you can go to the


rollup icon  in the menu bar to enter rollup mode. The rollup icon does not display until you have configured at least one subordinate server for the master (local) server. You might have to refresh your browser window to see the rollup icon after you add the first subordinate server for the master server you are configuring.

Entering rollup mode

To view rollup capacity and status information from your IBM Spectrum Control servers, you must enter rollup mode.

1. On the **Rollup Server Connections** page, add at least one subordinate server. Subordinate servers must have IBM Spectrum Control installed and running.
2. Move the mouse pointer over the rollup icon in the menu bar and select **Enter rollup mode**.

You will see the following changes in the menu bar:

- The color of the rollup icon changes:  .
- The lower edge of the menu bar changes to orange.

- The menu bar options change to Storage, Servers, Network, Rollup Server Connections.

Using rollup mode

Rollup mode provides lists of resources that are managed by subordinate servers. For example, you can see a list of all the network switches that are managed by a master server and its associated subordinate servers. These resource list pages are similar to the pages you work with when not in rollup mode. However, in some cases rollup mode provides different data columns and different actions that you can take on the rollup storage devices.

View rollup information

1. In the rollup menu bar, select the type of storage device you want to view information about. For example, **Servers**.
2. Select one of the available menu options. The Servers menu bar item has two options: **Servers** and **Hypervisors**.
3. On the resource list page, you can see the data about that resource type that has been collected from the master and subordinate servers.
4. Rollup resource lists function similar to non-rollup pages: you can add and reorder data columns, right-click items in the resource list to perform actions on them, and other familiar functions. You can also use Ctrl-click to select multiple list items.

Actions

You can perform actions on items in the resource lists by right-clicking them or by selecting one or more list items and clicking **Actions**.

View properties

View the properties of the selected item in the resource list. This action is not available for multiple selections at the same time.

Export Export the data for one or more list items. You can specify PDF, CSV, or HTML output.

Units of measurement for storage data

IBM Spectrum Control uses decimal and binary units of measurement to express the size of storage data.

Decimal units such as kilobyte (KB), megabyte (MB), and gigabyte (GB) are commonly used to express the size of data. Binary units of measurement include kibibyte (KiB), mebibyte (MiB), and gibibyte (GiB). Table 16 compares the names, symbols, and values of decimal and binary units.

Table 16. Comparison of binary and decimal units and values

Binary			Decimal		
Name	Symbol	Value (base 2)	Name	Symbol	Value (base 10)
kibibyte	KiB	2 ¹⁰	kilobyte	KB	10 ³
mebibyte	MiB	2 ²⁰	megabyte	MB	10 ⁶
gibibyte	GiB	2 ³⁰	gigabyte	GB	10 ⁹

Table 16. Comparison of binary and decimal units and values (continued)

Binary			Decimal		
Name	Symbol	Value (base 2)	Name	Symbol	Value (base 10)
tebibyte	TiB	2 ⁴⁰	terabyte	TB	10 ¹²
pebibyte	PiB	2 ⁵⁰	petabyte	PB	10 ¹⁵
exbibyte	EiB	2 ⁶⁰	exabyte	EB	10 ¹⁸

Binary units of measurement express the size of data more accurately. When you compare the size of 100 KB to 100 KiB, the difference is relatively small, 2.35%. However, this difference grows as the size of the data values increases. When you compare the size of 100 TB to 100 TiB, the difference is 9.06%.

In general, IBM Spectrum Control uses base 2 values for memory and disk space values, and base 10 values for space on physical hard drives.

Table 17 shows the percentage difference between decimal and binary values across a range of data sizes.

Table 17. Percentage difference between decimal and binary units

Decimal value	Binary equivalent of decimal value	Difference
100 kilobytes (KB)	97.65 kibibytes (KiB)	2.35%
100 megabytes (MB)	95.36 mebibytes (MiB)	4.64%
100 gigabytes (GB)	93.13 gibibytes (GiB)	6.87%
100 terabytes (TB)	90.94 tebibytes (TiB)	9.06%
100 petabytes (PB)	88.81 pebibytes (PiB)	11.19%
100 exabytes (EB)	86.73 exbibytes (EiB)	13.27%

Role-based authorization


Roles determine the functions that are available to users of IBM Spectrum Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

In Version 5.2, the roles that were previously defined in IBM Spectrum Control were consolidated into a smaller set of roles. The following table shows how the roles in versions earlier than 5.2 are mapped to the current set of roles:

Table 18. How roles in previous versions of IBM Spectrum Control are mapped to the roles in version 5.2 and later

Roles in previous versions	Roles in version 5.2 and later	Authorization level
Superuser Productivity Center administrator Disk administrator Fabric administrator Data administrator Tape administrator	Administrator	This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role. Note: When IBM Spectrum Control is first installed, the following operating system groups are assigned the Administrator role: <ul style="list-style-type: none"> • Windows: Administrators • UNIX and Linux: root • AIX: system
Disk operator Fabric operator Data operator Tape operator	Monitor	This role has access to the following read-only functions: <ul style="list-style-type: none"> • Viewing and exporting information about monitored resources • Viewing, acknowledging, and removing alerts • Viewing tasks and data collection jobs • Viewing data paths • Opening management GUIs • Opening logs • Viewing chargeback reports Exception: Users with the Monitor role can provision storage if they are granted permission in a service class. A service class is a logical entity that describes storage capabilities and characteristics and can be used to specify requirements for storage provisioning. For more information about service classes, see .
This role did not exist in versions 5.1 or earlier.	External Application	This role enables users of other applications to use the provisioning capability of IBM Spectrum Control to provision storage. For example, a VMware user with this role can provision storage in the vSphere GUI by using the vSphere Web Client extension for IBM Spectrum Control. If you assign the External Application role to the user, you must also assign one or more service classes to the user. This role does not enable users to log in to the IBM Spectrum Control GUI.

Tips:

- To determine the role of the user who is logged in, click the user icon  in the upper-right corner of any page in the GUI.

- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.
- If a user is not a member of a group that is assigned a IBM Spectrum Control role, no access is granted to that user.
- If you have been assigned the Monitor role, you can only open and view logs from the Data Collection page for the selected resource.

Fabrics and zones

Use the IBM Spectrum Control to learn more about fabrics and zones.

Monitoring fabrics and zones in the GUI

You can monitor switches and fabrics in the IBM Spectrum Control GUI. You can also enable automatic zoning so that new zones are created during storage provisioning, if necessary, to connect a server to a storage system.

You can use the GUI to add switches and fabrics for monitoring. You can add the following types of switches:

- Brocade
- Cisco
- Other switches

After switches and fabrics are probed, you can view detailed status and performance information about the switches and fabrics. You can also test the connection to the switch or fabric.

Information that you can view about a fabric includes the following properties:

- The fabric type
- The number of physical and virtual switches in the fabric
- The number of ports that are on all the switches in the fabric
- The number of switch ports that are online and connected to other ports in a fabric
- The name of the principal switch of the fabric

Information that you can view about a switch includes the following properties:

- The number of ports on a switch
- The number of ports that are connected to a storage resource
- The name of the fabric where a switch is a member
- The name of the vendor or manufacturer for a switch

From the Details page of a fabric, you can open the Zone Sets page to view information about the zone sets on the fabric, and to identify which is the active zone set.

You can enable automatic zoning in the GUI. When automatic zoning is enabled, IBM Spectrum Control creates new zones during provisioning if new zones are needed to connect the storage system with the server. When you set the zoning policy, you can specify whether zone changes from automatic zoning are made to the active zone set, or to a new inactive active zone set.

Zones, zone aliases, and zone sets

Zones, zone aliases, and zone sets allow logical grouping of ports and storage devices within a storage area network. This section describes zoning concepts and elements.

In a storage area network a *zone* is a logical grouping of ports to form a virtual private storage network. Zones that belong to a single SAN can be grouped into a *zone set*, which can be activated or deactivated as a single entity across all switches in the fabric. A zone set can contain one or more zones, and a zone can be a member of more than one zone set. Using zoning, you can automatically or dynamically arrange fabric-connected devices into logical groups across a physical fabric.

Ports and devices in a zone are called *zone members*. A zone can contain one or more zone members. Ports that are members of a zone can communicate with each other, but they are isolated from ports in other zones. Devices, however, can belong to more than one zone. A *zone alias* is a collection of zone members. A zone alias can be added to one or more zones.

Note: In some cases, inactive zones might not have any zone members. Activating an inactive zone set that contains empty zones will fail if a switch does not support empty zones in active zone definitions.

Zone membership can be specified by:

- The N_Port_Name of the N_Port connected to the switch (also known as WWN zoning or port name zoning)
- The N_Port address identifier assigned during fabric login (also known as PortId or FCID zoning)
- The Node_Name associated with the N_Port
- The Domain identification (Domain_ID) and physical port ID of the Switch Port to which the N_Port is attached (also known as domain port zoning)
- An alias name

Zoning supports the use of aliases, which are meaningful names assigned to devices. An alias can also be a group of devices that are managed together to make zoning easier.

There are two types of zoning:

Hardware zoning (port zoning)

In hardware zoning (also called *port zoning*), the members of a zone are the physical ports on a fabric switch.

Software zoning (WWN zoning)

Software zoning uses the Simple Name Server (SNS) that runs inside a fabric switch. It is based on the node WWN or port WWN of the zone members to be included. Software zoning lets you create symbolic names for the zones and zone members.

A *default zone* is a group of devices that are not members of the active zone set. These can communicate with each other but not with members of any other zone. Default zoning is enabled by default. You can use a switch element manager to configure the Default Zone option to enable or disable the default zone independently of the active zone set.

Note:

1. If the default zone is disabled, devices that are not members of the active zone set cannot communicate.
2. If the default zone is disabled and no zone set is active, no devices can communicate.
3. If default zoning is enabled, deactivating the active zone set makes all devices members of the default zone. If default zoning is disabled, all communication stops.
4. If you activate one zone set while another zone set is active, the currently active zone set is deactivated.
5. If your EFC Manager manages multiple fabrics, ensure that you have the correct zone set for the fabric you are currently updating.

You can view zones in SANs using the topology viewer console. Fabric Manager lets you view zone membership.

Note: Certain types of switches and the types of agents configured can affect the zoning operations that you can perform.

Switch zoning capabilities

When you select a fabric in which to perform zoning operations, Fabric Manager determines the capabilities of the switches in the SAN and limits the zoning operations based on that information.

To see a list of supported switches and their capabilities, default values, ranges, and possible effects, go to <http://www.ibm.com/support/docview.wss?uid=swg27047050>.

Zoning for Cisco MDS 9000 switches:

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. Fabric Manager provides basic zone discovery for the Cisco MDS 9000 series switches.

Fabric Manager supports these zone member types:

- N_Port WWN
- FC ID (the fibre channel ID of an N_port attached to the switch)

Fabric Manager does not allow zone management of zones that contain unsupported members.

Brocade switches in zones:

Brocade switches support some non-standard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, Fabric Manager preserves them and does not modify them in any way. If they are part of a zone set that is active at some time, the devices that are part of such zones that are also online are displayed in the topology Zone View.

You can create, change, and delete non-standard zones by using the Brocade switch management application.

Note: IBM Spectrum Control allows you to add empty zones to an inactive zone definition. Activation of a zone set containing empty zones may still fail if the switch does not support empty zones in active zone definitions.

Agents

The IBM Spectrum Control uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, and SNMP agents.

Common Information Model agents

Common Information Model (CIM) agents enable communication between storage devices, such as storage systems and fabric switches, and IBM Spectrum Control. Each storage device that you want to manage must have a CIM agent either installed or embedded unless it is a device that uses Native API connection.

CIM agents are provided by the vendor of the storage subsystem or fabric switch. Each vendor provides unique agent code for their family of storage devices. This code implements a Common Information Model Object Manager (CIMOM) that conforms to the Storage Management Initiative Specification (SMI-S) of the Storage Networking Industry Association (SNIA).

The CIM agent enables communication between the storage device and IBM Spectrum Control. Commands and responses are transmitted between IBM Spectrum Control and the CIM agent using an XML transport layer. The CIM agent to storage device layer uses a vendor-specific proprietary interface.

The CIM agent usually must be installed and configured, so that it can identify the storage devices with which it communicates. Some storage devices, such as fabric switches, contain embedded CIM agents and so do not require that CIM agents be installed. In these cases, IBM Spectrum Control must be configured to point directly to the storage devices that contain the embedded CIM agents.

CIM agents can be referred to as SMI-S providers, CIM proxy agents, and CIMOM agents. CIM agents can be embedded in the device or installed on a separate computer.

Note:

- Do not install multiple CIM agents on a single computer because of port conflicts.
- Do not install a CIM agent on the system where a IBM Spectrum Control server component is installed.

Storage Resource agents

Use Storage Resource agents to collect information about storage resources such as servers, virtual machines, workstations, and HBAs.

You must deploy Storage Resource agents on resources where you want to gather the following information:

- Asset information
- File and file system attributes
- Network-attached storage (NAS) device information
- Topology information

Tip: You can also monitor servers without deploying a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates information

about that server with the known host connections on monitored resources. If a match is made between the server and a monitored resource, you can view topology information and the capacity and performance of the storage that is assigned to that server. For more information about when to deploy a Storage Resource agent or add an agentless server, see “Adding servers” on page 125.

SNMP agents

An SNMP agent, also called an out-of-band agent, is software that runs on a switch and can be used to monitor and manage systems and devices in your network. IBM Spectrum Control uses SNMP to retrieve information about specific switches, such as Cisco switches, and the fabrics that the switches are members of.

SNMP agents are necessary for retrieving virtual storage area network (VSAN) information for Cisco switches.

IBM Spectrum Control uses SNMPv3 (preferred) or SNMPv1 to probe Cisco switches and fabrics, and uses SNMPv2 or SNMPv3 to collect performance data. Some switches are configured to use SNMPv3 by default.


IBM Spectrum Control Monitoring Agent

The IBM Spectrum Control Monitoring Agent can be used by the IBM Tivoli Enterprise Monitoring Server to monitor systems in your enterprise. This agent is an optional program you can install and use in your enterprise.

IBM Tivoli Monitoring monitors and manages system and network applications on a variety of operating systems, tracks the availability of your enterprise system, and provides reports to track trends and troubleshoot problems. The IBM Spectrum Control Monitoring Agent participates in that environment.

For performance reasons, install the IBM Spectrum Control server on a separate system from the IBM Tivoli Monitoring server. You can install the IBM Spectrum Control Monitoring agent on any system you want to monitor.

Before using the IBM Spectrum Control Monitoring Agent, you must have IBM Tivoli Enterprise Monitoring Server installed. For information about the IBM Tivoli Enterprise Monitoring Server, see  http://www.ibm.com/support/knowledgecenter/SSTFXA_6.2.1/com.ibm.itm.doc_6.2.1/welcome.htm.

For information about installing IBM Tivoli Monitoring, see the Quick Start Guide at  http://www.ibm.com/support/knowledgecenter/SSTFXA_6.2.0.1/com.ibm.itm.doc/GI11-8058-00.htm.

To use the IBM Spectrum Control Monitoring Agent, install this agent on the system you want to monitor.

To view the type of reports that are displayed by the IBM Spectrum Control Monitoring Agent, see http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.14/com.ibm.spectrum.sc.doc/fqz0_c_using_tpcma.html.

IBM Spectrum Control REST API

You can use the Representational State Transfer (REST) API for IBM Spectrum Control to access information about resources and to generate custom configuration and performance reports.

Using the REST API with a command line utility

You can use a REST command line utility, for example, GNU Wget.

To access information about the GNU Wget utility, go to <https://www.gnu.org/software/wget/>.

The IBM Spectrum Control REST API is hosted here: <https://<hostname>:9569/srm/REST/api/v1/>.

You can enter a command to authenticate with the REST API and store a security token in a file to use in your later queries, for example, `cookies.txt`.

You can use the utility, by entering the following:

```
wget --post-data "j_username=<user name>&j_password=<password>"  
--no-check-certificate --keep-session-cookies --save-cookies cookies.txt  
https://<hostname>:9569/srm/j_security_check
```

You can use the security token in all subsequent commands issued against the REST API, by issuing the following command:

```
wget --no-check-certificate --load-cookies cookies.txt  
https://<hostname>:9569/srm/REST/api/v1/
```

Example, (partial listing) :

```
},  
{  
  "Description": "Provides a list of ports belonging to a parent storage system.  
Format: \StorageSystems\<id>\Ports For a list of ports associated with a  
port owner, specify the owner type and id in the URL.
```

To see a list of all storage systems, you can append the URL:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems
```

To see more information about a resource type, you can further append the URL:

```
https://<hostname>:9569/srm/REST/api/v1/<Name>
```

To see more information about a specific resource, add the ID to the URL:

```
https://<hostname>:9569/srm/REST/api/v1/<Name>/<id>
```

To query for volumes that belong to a specific storage system, enter:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/57909/Volumes
```

Example, (a partial list is shown) :

```
[  
  {  
    "Acknowledged": "No",  
    "Allocated Space": "1.00",  
    "Capacity": "1.00",  
    "Controller": "Node 0",  
    "Copy Relationship": "",  
    "Easy Tier": "Tiered Pools\No",  
    "Encryption": "",  
    "Enterprise HDD Capacity": "1.00",  
    "Format": "FB",  
    "Hosts": "18",  
    "LSS or LCU": "00",  
    "Last Data Collection": "Apr 10, 2016, 15:03:57",  
    "Name": "CET_RA_vol",
```

```

    "Nearline HDD Capacity": "",
    "Physical Allocation": "Fully Allocated",
    "Pool": "General_Use_1",
    "RAID Level": "RAID 5",
    "SSD Capacity": "",
    "Service Class": "",
    "Shortfall": "",
    "Status": "Normal",
    "Storage System": "DS8000-2107-75BLG91-IBM",
    "Storage Virtualizer": "None",
    "Thin Provisioned": "No",
    "Ticket": "",
    "Tier Distribution": "0",
    "Unallocated Space": "0.00",
    "Unused Space": "",
    "Used Allocated Space": "",
    "Used Space": "1.00",
    "Virtualizer Disk": "None",
    "Volume ID": "0000",
    "Volume Number": "0",
    "Volume Unique ID": "75blg91\0000",
    "id": "79852"
  },
  {
    "Acknowledged": "No",
    "Allocated Space": "1.00",
    "Capacity": "1.00",
    "Controller": "Node 0",
    "Copy Relationship": "",
    "Easy Tier": "Tiered Pools\No",
    "Encryption": "",
    "Enterprise HDD Capacity": "1.00",
    "Format": "FB",
    "Hosts": "18",
    "LSS or LCU": "00",
    "Last Data Collection": "Apr 10, 2016, 15:03:57",
    "Name": "CET_RA_vol",
    "Nearline HDD Capacity": "",
    "Physical Allocation": "Fully Allocated",
    "Pool": "General_Use_1",
    "RAID Level": "RAID 5",
    "SSD Capacity": "",
    "Service Class": "",
    "Shortfall": "",
    "Status": "Normal",
    "Storage System": "DS8000-2107-75BLG91-IBM",
    "Storage Virtualizer": "None",
    "Thin Provisioned": "No",
    "Ticket": "",
    "Tier Distribution": "0",
    "Unallocated Space": "0.00",
    "Unused Space": "",
    "Used Allocated Space": "",
    "Used Space": "1.00",
    "Virtualizer Disk": "None",
    "Volume ID": "0001",
    "Volume Number": "1",
    "Volume Unique ID": "75blg91\0001",
    "id": "79851"
  },
  {
    "Acknowledged": "No",
    "Allocated Space": "1.00",
    "Capacity": "1.00",
    "Controller": "Node 0",
    "Copy Relationship": "",
    "Easy Tier": "Tiered Pools\No",

```

```

"Encryption": "",
"Enterprise HDD Capacity": "1.00",
"Format": "FB",
"Hosts": "17",
"LSS or LCU": "00",
>Last Data Collection": "Apr 10, 2016, 15:03:57",
>Name": "CET_RA_vol",
>Nearline HDD Capacity": "",
>Physical Allocation": "Fully Allocated",
>Pool": "General_Use_1",
>RAID Level": "RAID 5",
>SSD Capacity": "",
>Service Class": "",
>Shortfall": "",
>Status": "Normal",
>Storage System": "DS8000-2107-75BLG91-IBM",
>Storage Virtualizer": "None",
>Thin Provisioned": "No",
>Ticket": "",
>Tier Distribution": "0",
>Unallocated Space": "0.00",
>Unused Space": "",
>Used Allocated Space": "",
>Used Space": "1.00",
>Virtualizer Disk": "None",
>Volume ID": "0010",
>Volume Number": "16",
>Volume Unique ID": "75blg91\0010",
"id": "79836"
},
{
"Acknowledged": "No",
"Allocated Space": "9.00",
"Capacity": "9.00",
"Controller": "Node 1",
"Copy Relationship": "",
"Easy Tier": "Tiered Pools\No",
"Encryption": "",
"Enterprise HDD Capacity": "9.00",
"Format": "FB",
"Hosts": "2",
"LSS or LCU": "01",
>Last Data Collection": "Apr 10, 2016, 15:03:58",
>Name": "Livi_L3_528_1",
>Nearline HDD Capacity": "",
>Physical Allocation": "Fully Allocated",
>Pool": "General_Use_2",
>RAID Level": "RAID 5",
>SSD Capacity": "",
>Service Class": "",
>Shortfall": "",
>Status": "Normal",
>Storage System": "DS8000-2107-75BLG91-IBM",
>Storage Virtualizer": "None",
>Thin Provisioned": "No",
>Ticket": "",
>Tier Distribution": "0",
>Unallocated Space": "0.00",
>Unused Space": "",
>Used Allocated Space": "",
>Used Space": "9.00",
>Virtualizer Disk": "None",
>Volume ID": "0100",
>Volume Number": "0",
>Volume Unique ID": "75blg91\0100",
"id": "79742"
},

```

The information provided in the `Format:` field in the example shows what combinations are possible. For example:

`https://<hostname>:9569/srm/REST/api/v1/StorageSystems/<id>/Ports`

Note:

- If the `Format:` field is omitted, you can add the ID of one resource to the URL.
- The name elements in the URL are case sensitive.
- If the URL is not used correctly, you receive an error message:

```
{"result":{"type":"E","msgId":"BPCUI0099E","time":"Apr 4, 2016 16:25:07","text":
"The storage resource is not available."}}
```

Using the REST API with a web browser

A convenient way to use the IBM Spectrum Control REST API is within the interface itself. You can quickly access resource information for your reporting capabilities using a web browser.

To use the REST API with a web browser, use these steps:

1. Log on to the IBM Spectrum Control GUI. To start IBM Spectrum Control, see “Starting IBM Spectrum Control” on page 39.
2. Open a new tab in your web browser.
3. Enter `https://<hostname>:9569/srm/REST/api/v1/`.

For example, to see a list of fabrics, enter `https://<hostname>:9569/srm/REST/api/v1/Fabrics` in a new tab of your web browser (partial listing) :

```
[
  {
    "Acknowledged": "No",
    "Active Zone Set": "",
    "Connected Switch Ports": "0",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Fabric Type": "Cisco",
    "Last Successful Probe": "N/A",
    "Links": "0",
    "Location": "",
    "NPV Switches": "0",
    "Name": "VSAN0002",
    "Parent Fabric": "san_director",
    "Principal Switch of Fabric": "tpc-70swt-csc",
    "Probe Schedule": "N/A",
    "Probe Status": "N/A",
    "Status": "Error",
    "Switch Ports": "0",
    "Switches": "1",
    "Virtual": "Yes",
    "WWN": "2002000DECAC5081",
    "id": "2663"
  },
  {
    "Acknowledged": "No",
    "Active Zone Set": "",
    "Connected Switch Ports": "18",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "2",
    "Fabric Type": "Cisco",
```

```

    "Last Successful Probe": "N\A",
    "Links": "4",
    "Location": "",
    "NPV Switches": "0",
    "Name": "san_director",
    "Parent Fabric": "",
    "Principal Switch of Fabric": "",
    "Probe Schedule": "N\A",
    "Probe Status": "N\A",
    "Status": "Error",
    "Switch Ports": "104",
    "Switches": "2",
    "Virtual": "No",
    "WWN": "",
    "id": "2017"

```

To see the information for a specific fabric in your environment, append the URL with the ID:

<https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017>

Example (partial listing) :

```

{
  "Acknowledged": "No",
  "Active Zone Set": "",
  "Connected Switch Ports": "18",
  "Custom Tag 1": "",
  "Custom Tag 2": "",
  "Custom Tag 3": "",
  "Data Source Count": "2",
  "Fabric Type": "Cisco",
  "Last Successful Probe": "N\A",
  "Links": "4",
  "Location": "",
  "NPV Switches": "0",
  "Name": "san_director",
  "Parent Fabric": "",
  "Principal Switch of Fabric": "",
  "Probe Schedule": "N\A",
  "Probe Status": "N\A",
  "Status": "Error",
  "Switch Ports": "104",
  "Switches": "2",
  "Virtual": "No",
  "WWN": "",
  "id": "2017"

```

To see the switches that belong to the fabric "id": "2017", enter:

<https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017/Switches>

Example (partial listing) :

```

[
  {
    "Acknowledged": "No",
    "Connected Fabrics": "",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Domain ID": "10",
    "Fabric": "unstable_BRCD_1-11",
    "Firmware": "v6.4.3g",
    "IP Address": "9.11.91.241",
    "Last Successful Monitor": "None",
    "Last Successful Probe": "None",

```

```

    "Links": "",
    "Location": "",
    "Mode": "Native",
    "Model": "Brocade 4100",
    "Name": "mdm-y76-swt",
    "Parent Switch": "",
    "Performance Monitor Interval (min)": "",
    "Performance Monitor Status": "Disabled",
    "Ports": "",
    "Principal Switch of Fabric": "cjswitch4",
    "Probe Status": "Never Probed",
    "Serial Number": "1070029",
    "Status": "Normal",
    "Vendor": "IBM",
    "Virtual": "No",
    "WWN": "100000051E347790",
    "id": "60989"
  },
  {
    "Acknowledged": "No",
    "Connected Fabrics": "",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Domain ID": "239",
    "Fabric": "VSAN1 tpc70swt-csc",
    "Firmware": "3.3(1c)",
    "IP Address": "9.11.98.150",
    "Last Successful Monitor": "None",
    "Last Successful Probe": "None",
    "Links": "2",
    "Location": "",
    "Mode": "",
    "Model": "DS-C9509",
    "Name": "director",
    "Parent Switch": "director",
    "Performance Monitor Interval (min)": "",
    "Performance Monitor Status": "Disabled",
    "Ports": "80",
    "Principal Switch of Fabric": "director",
    "Probe Status": "Never Probed",
    "Serial Number": "nwg064000a",
    "Status": "Warning",
    "Vendor": "Cisco Systems",
    "Virtual": "Yes",
    "WWN": "2001000530007CDF",
    "id": "2025"
  },

```

More examples of the IBM Spectrum Control REST API:

Table 19. IBM Spectrum Control REST API URL examples

IBM Spectrum Control REST API	Description
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems">https://<hostname>:9569/srm/REST/api/v1/StorageSystems	A list of all monitored storage systems.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002	Specific information for one storage system (ID is available from this URL: <a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems">https://<hostname>:9569/srm/REST/api/v1/StorageSystems) .
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes	A list of all volumes for one storage system.

Table 19. IBM Spectrum Control REST API URL examples (continued)

IBM Spectrum Control REST API	Description
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance</code>	The available performance metrics for a volume.
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication</code>	The remote mirror replication relationships for storage systems.
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000</code>	Sample performance granularity information and a start and end time for the performance data.

Chapter 5. Scenarios

This section contains scenarios which illustrate how to perform specific tasks using IBM Spectrum Control.

The provided scenarios set out situations that can be solved using IBM Spectrum Control. When using these scenarios you should remember the following:


- The scenarios are intended to be step-by-step instructions that you follow to complete a specific job. Each scenario contains multiple tasks that must be followed in the order given.
- Specific values are given in the scenarios, such as user name, IP Address, probe name, and so forth. These are for illustration purposes only and you must replace them with values appropriate for your system.
- The scenarios do not provide in-depth conceptual information about the tasks. See the appropriate topics in the IBM Spectrum Control Knowledge Center for more information about the concepts behind the tasks.
- You will complete a scenario successfully if you follow the instructions correctly and your system performs as expected.

Optimizing the performance of storage virtualizers

In these scenarios, you balance pools and analyze tiering to help resolve performance hotspots and manage storage resources efficiently.

You can optimize the placement of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified block storage

Learn more: For information about the components that can be used with IBM Spectrum Control Version 5.2 (or later), see  IBM Spectrum Control interoperability matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

To optimize the placement of volumes in storage virtualizer pools, you must complete the following tasks:

- Probe the storage virtualizers.
- Set the tier level of the storage pools that you want to analyze.
- Collect performance data to analyze the pools and volumes.

When you run the analysis to balance pools, you optimize storage performance by redistributing volume workloads across pools on the same tier.

When you run the tiering analysis, you optimize storage performance by specifying thresholds to move volumes to higher and lower tiers.

Related tasks:

“Optimizing storage tiering” on page 319

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

“Optimizing storage pools” on page 325


Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Managing tasks for tiering storage, balancing pools, and transforming storage” on page 346

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Redistributing volume workloads to balance pools

In this scenario, you identify performance hotspots and then run the balance pools analysis to help resolve them.



The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems.

If the activity deviation percentage of the pool exceeds the threshold set for the pool, which is $> 10\%$, the pool is identified as a candidate for balancing. If the activity deviation percentage of the pool lies within the threshold range set for the pool, which is $< -10\%$ to $> +10\%$, the pool is considered to be balanced.

Tip: On the Pools page, the activity level of the pool is shown in the **Activity** column. The difference between the activity level of the pool and all of the other pools on the same tier and storage system is shown in the **Activity Deviation (%)** column.

To balance pools, complete the following steps:

1. From the **Storage** menu, click **Storage Systems**.

2. To group the pools by storage system, click the **Storage System** column heading.

Tip: Examine the activity deviation percentage values for the pools. Select pools with low and high activity levels on each of the tiers that you want to balance so that the volumes can be moved from high-activity pools to low-activity pools.

3. Right-click the pools on the tiers that you want to balance and then click **Balance Pools**.
4. Select the period that you want to use to analyze the performance data.
5. Click **Analyze**.

The Balance Analysis task opens.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule > Analysis**.

Related tasks:

“Optimizing storage pools” on page 325

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Modifying the criteria for balancing pools” on page 326

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

Related reference:

“Criteria for identifying the pools that require balancing” on page 327

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

Analyzing and re-tiering volumes in pools on tier 1

In this scenario, you analyze the current placement of volumes in pools on tier 1 so that you can move volumes with low workloads from tier 1 to tier 2 or tier 3 pools.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 storage systems. The pools on the storage virtualizer are assigned to tier 1, tier 2, and tier 3.

Choosing the performance data that is used in the analysis



1 Analysis period and data collection days

The analysis period is set to 10 days and Saturdays and Sundays are excluded from the analysis.


Mo	Tu	We	Th	Fr	Sa	Su
Mo	Tu	We	Th	Fr	Sa	Su

- ☐ The performance data that is collected is used in the analysis.
- ☒ The performance data that is collected isn't used in the analysis.



2 Data collection period

09:00 is selected as the start time and 18:00 is selected as the end time.



3 Analysis schedule


You create the analysis on Monday and schedule the analysis to run every 14 days.

Mo	Tu	We	Th	Fr	Sa	Su	
Mo	Tu	We	Th	Fr	Sa	Su	Mo

- ☒ The day that you created the analysis is the day that the analysis is run.

For example, you want to analyze the data that is collected over 10 days and when the volumes are most active. So, you set the period for the analysis to 10 days and exclude the data that is collected at the weekend. You also exclude the data that is collected before 9 AM and after 6 PM because you know that the volumes are most active from 9 AM to 6 PM.

Tiering and volume placement criteria



1 Tiering criteria

The source volumes that are selected are in pools on Tier 1, and the target pools that are selected are on Tier 1, Tier 2, and Tier 3. The tiering thresholds for the volumes are >1000 I/O per second for Tier 1 and >500 I/O per second for Tier 2.

Tiering Criteria	Tier 1: Volume Analysis	Tiering recommendations
Tier 1 I/O Rate: >1000	I/O Rate = 900	Place volume in pool on Tier 2
Tier 2 I/O Rate: <1000 and >500	I/O Rate = 400	Place volume in pool on Tier 3
Tier 3 I/O Rate: <500		



2 Volume placement criteria

To ensure that the pools that are selected as target pools can handle the additional workload, you specify the maximum I/O rates for the target pools on each tier.

Target Pools	Below Threshold	Add Volume
Target Pool 1	Yes	Yes
Target Pool 2	No	No

You want volumes with the highest workloads in tier 1 pools and you want volumes with low workloads in tier 2 or tier 3 pools. To tier your storage, you enter two I/O rate thresholds to generate recommendations to re-tier the volumes.

Tip: You can tier volumes by I/O rate or I/O density.

To ensure that the pools that are selected as destination pools for the volumes can handle the additional workload, you specify the maximum I/O rates for pools on each tier of storage. Only the pools with I/O rates below the maximum I/O rate can be selected as destination pools for the volumes.

1. From the **Storage** menu, click **Storage Systems**.
2. Right-click the storage virtualizer and click **View Details**.
3. In the navigation pane, click **Pools**. To organize the pools by tier level, click the **Tier** column heading.
4. Select the pools on tier 1, right-click and then click **Analyze Tiering**.
5. Select all of the pools on tier 2 and tier 3 as target pools, and then click **Next**.
6. Select the period that you want to use to analyze the performance data.
7. Set the threshold for moving volumes from tier 1 to tier 2. The volumes with I/O rates above the threshold that you set remain in the tier 1 pools. The volumes with I/O rates below the threshold that you set are moved down to tier 2 pools.

8. Set the threshold for moving volumes from tier 2 to tier 3. The volumes with I/O rates above the threshold that you set remain in the tier 2 pools. The volumes with I/O rates below the threshold that you set are moved down to tier 3 pools.
9. Set the maximum I/O rate for pools on each tier.
10. Click **Analyze**.

The Analyze Tiering task is shown.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule > Analysis**. Alternatively, you can modify the criteria that you set for analyzing the volumes on the Tasks page and run the analysis again.

Related tasks:

“Modifying the criteria for analyzing tiering” on page 324

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Collocating volumes

In this scenario, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes when you enter the criteria for analyzing tiering and balancing pools.

To ensure that volumes in the same storage pool that are assigned to the same server or hypervisor are kept together, you want to enforce the collocation of volumes. By enforcing the collocation of volumes, you prevent the placement of related volumes in destination pools that might be on multiple back-end storage systems.

Multiple host connections to the same hypervisor or server: If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

1. From the **Storage** menu, click **Pools**.
2. Right-click the pools that you want to analyze, and then click **Analyze Tiering**.
3. Select the target storage pools.
4. On the Optimize the Placement of Volumes page, ensure that **Collocate volumes** is set to **Yes**.
5. Click **Analyze**.

If a volume requires re-tiering, a recommendation is generated to move all of the volumes that are assigned to the same server to the same destination pool. If none of the target pools have sufficient space to accommodate all of the volumes, then recommendations to move the volumes are not generated.

Related tasks:

“Optimizing storage pools” on page 325

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Optimizing storage tiering” on page 319

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

“Modifying the criteria for analyzing tiering” on page 324

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Monitoring capacity usage at different levels of a business hierarchy

The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

The department model comprises these main elements:

- The department
- The subdepartments that the department contains
- The applications, or the application subcomponents, or both that the department uses

In the following scenario, you want to create a department model that monitors the capacity and space usage of the Books Sales department, and these subdepartments:

- Wholesale
- Retail
- Online

To monitor the capacity and space usage of the departments, you add the applications that the department and its subdepartment use to each of the subdepartments.

The department model that you create determines how you can view capacity usage and space information. To view the total capacity and space usage of a department regardless of the subdepartments that the department contains, you create a department and add the applications and subcomponents to the department.

If, however, you want to see the capacity and space usage of the department and its subdepartments, you create a department hierarchy by adding subdepartments to departments. You then add the applications and application subcomponents that each subdepartment uses to the subdepartments. You can then view capacity and space usage for the department and for the individual subdepartments.

You can create more complex models of departments by adding subdepartments to subdepartments. In the following illustration, the department model is extended to include two more layers of subdepartments under the Wholesale subdepartment.



To monitor the capacity and space usage of the departments in your business organization, complete the following tasks:

- Create the Book Sales department.
- Create the Wholesale, Retail, and Online departments to represent the subdepartments.
- Add the Wholesale, Retail, and Online as subdepartments to the Book Sales department.
- Create the Book Sales DB application and the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents that are used by the departments.

Different applications require different storage capabilities, and different levels of performance and uptime. The storage requirements of the departments are constantly growing and they need to modify their behaviors in a way that justifies the cost of their actions concerning storage utilization. The data modeling of storage resources using the application and department concept enables you to plan and implement a chargeback system if necessary.

Related concepts:

“Departments” on page 270

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

Related tasks:

“Creating departments” on page 272

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

“Adding and creating subdepartments” on page 273

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

“Viewing information about departments” on page 275

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchical business structure.

“Creating applications” on page 263

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

“Adding subcomponents” on page 267

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Comparing storage usage in each department


In this scenario, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

To monitor the capacity and space usage of a department, you create the department and subdepartments, and you add the application subcomponents to the subdepartments.

You want to monitor the capacity and space usage of departments in an organization that sells books.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online



Although you can add applications and application subcomponents when you create departments, it is easier and quicker to create the applications and application subcomponents beforehand.

You want to know the total amount of storage space that the Book Sales department uses, and you want to know the amount of space that each subdepartment uses. To know how much space the Book Sales and its subdepartments use, you must associate the department and its subdepartments with the applications that they use. For example, the Book Sales department and its subdepartments use a database application named Book Sales DB. The Book Sales DB application contains these subcomponents:

- Wholesale Transactions
- Retail Transactions

- Online Transactions

To create the department and subdepartments, you complete these tasks:

1. You create the Wholesale department and add the applications and application subcomponents that the subdepartment uses.
2. You create the Retail department and add the applications and application subcomponents that the subdepartment uses.
3. You create the Online department and add the applications and application subcomponents that the subdepartment uses.
4. You create the Book Sales department. You do not have to add applications or subcomponents to the Book Sales department because the Book Sales department inherits the applications and application subcomponents that are added to the subdepartments.
5. You add the Wholesale, Retail, and Online departments as subdepartments to the Book Sales department.

When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

You want to target what department and what storage resources might be responsible for a change or a upward trend in significant storage usage. By correctly classifying the department or departments you can properly plan to archive or switch the storage within the organization.

Related concepts:

“Departments” on page 270

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

Related tasks:

“Creating departments” on page 272

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

“Adding and creating subdepartments” on page 273

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

“Viewing information about departments” on page 275

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal business structure.

“Creating applications” on page 263

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

“Adding subcomponents” on page 267

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another

level to your business hierarchy.

Using applications and subcomponents to monitor capacity and space usage

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchical manner.

You can create a simple application model, which consists of an application and use a filter to associate the storage resources that the application uses with the application. Alternatively, you can create a complex application model that comprises an application and application subcomponents and then add filters to associate the storage resources that each application subcomponent uses with each of the subcomponents.

Besides creating applications to monitor the performance, capacity, and space usage of the applications in your business organization, you also create applications to add them to departments. When you add applications to departments, you can monitor the capacity and space usage of the departments and you can monitor the performance of the storage resources that are associated with the applications that each department uses.

To monitor the capacity of an application, you create an application model that is based on the following elements:

- The application
- The subcomponents that the application contains
- The applications, application subcomponents, or both that are added to the department

In the following scenario, you want to create an application model that monitors the capacity, space usage, and performance of the Book Sales DB application. The application model that you want to create comprises these application subcomponents:

- Wholesale Transactions
- Retail Transactions
- Online Transactions

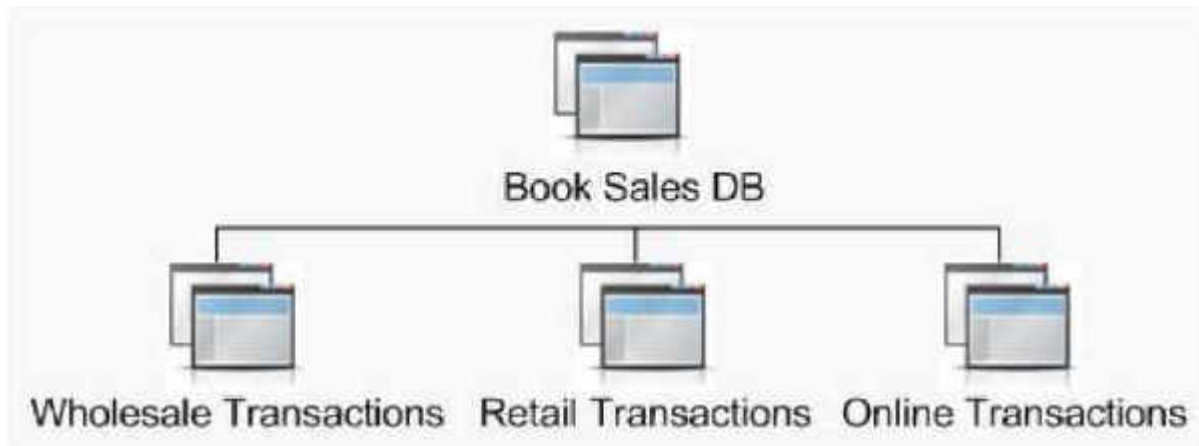


Figure 1. An application with a three subcomponent hierarchy

To understand how capacity is trending for storage in an application and subcomponents, complete these tasks:

1. Create the Book Sales DB application.
2. Create the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents.
3. Add the Wholesale Transactions, Retail Transactions, and Online Transactions as subcomponents to the Book Sales DB application.
4. Create the Book Sales department and the Wholesale, Retail, and Online subdepartments to be associated with the applications.

Related concepts:

“Applications” on page 261

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

Related tasks:

“Creating applications” on page 263

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

“Adding subcomponents” on page 267

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

“Viewing information about applications” on page 268

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Viewing storage capacity and usage trends

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

You want to understand how capacity is trending for storage that is used by the Book Sales DB application and the Wholesale Transaction, Retail Transactions and Online Transactions subcomponents and the association the application has with the Book Sales department and subdepartments.

The storage resources that are used by the application and that are used by each subcomponent are associated with the application and its subcomponents. For example, the storage resources are assigned as follows:

1. Create a database application called Book Sales DB and assign resources that use the volumes on the IBM® Storwize V7000 with names that begin with **bksales** for storing sales transactions.
2. Create the subcomponent applications named Wholesale Transaction, Retail Transactions and Online Transactions.
3. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales_ws** to the Wholesale Transactions subcomponent, which is used for storing wholesale sales transactions.
4. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales_rt** to the Retail Transactions subcomponent, which is used for storing retail sales transactions.
5. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales_ol** to the Online Transactions subcomponent, which is used for storing online sales transactions.

Note: On the Create Filter page, you can specify a name pattern to determine which volumes to include. For volumes, you can specify name patterns to determine from which servers, storage systems, or pools the volumes are selected. You can then click **Preview** to view the volumes that are selected for inclusion in your application.

To monitor space usage for the Book Sales department and its subdepartments, you associate the Book Sales DB application and the Wholesale, Retail, and Online Transactions subcomponents of the application as follows:

- You associate the Book Sales DB application with the Book Sales department.
- You associate the Wholesale subdepartment with the Wholesale Transactions subcomponent.
- You associate the Retail subdepartment with the Retail Transactions subcomponent.
- You associate the Online subdepartment with the Online Transactions subcomponent.




Figure 2. The application, subcomponent, assigned resources and department association

You can view the capacity information that is collected about the application on the Applications page. You can also view capacity information for each application and subcomponent on the details page. When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

Related concepts:

“Applications” on page 261

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

Related tasks:

“Creating applications” on page 263

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

“Adding subcomponents” on page 267

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

“Adding applications to departments” on page 274

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

“Viewing information about applications” on page 268

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Limiting a set of storage to provision in the vSphere Web Client extension

In this scenario, a storage administrator limits the storage that a vSphere user can provision.

You plan to dedicate two storage pools to be used with vSphere. You want to allow a vSphere user to provision storage by using the vSphere Web Client extension for IBM Spectrum Control. You want to define the two storage pools as the resources that the user can use.

1. Create a capacity pool that the vSphere user can provision from. Place the two storage pools that you designated as provisioning targets for vSphere in the capacity pool. To add storage pools to the capacity pool, follow the instructions in “Adding resources to capacity pools” on page 301.
2. Create a service class and specify that the service class can be used to provision storage only from the capacity pool. To create the service class, follow the instructions in “Creating service classes” on page 295.
3. Create a user ID for the vSphere user and assign the External Application role to that ID. To assign the role to the user, follow the instructions that are provided in the product documentation at <http://www.ibm.com/support/>

knowledgecenter/SS5R93_5.2.14/com.ibm.spectrum.sc.doc/
fqz0_t_associate_user_group_to_roles.html.

4. Associate the user with the service class. To associate the user with the service class, follow the instructions in “Modifying service classes” on page 297.

When you complete these tasks, a vSphere user can provision storage by using the vSphere Web Client extension, but they are limited to the service class that was created. Because the service class is associated with the capacity pool that contains your designated vSphere storage pools, you limited where the vSphere user can provision data.

Related concepts:

“Cloud configuration” on page 52

To take advantage of the simplified provisioning and optimization capabilities in IBM Spectrum Control, some configuration is required.

Related tasks:

“Provisioning block storage” on page 316

Use the vSphere Web Client extension for IBM Spectrum Control to provision block storage on IBM Spectrum Control. You can also have a data store created on the storage volume.

“Provisioning file storage” on page 318

Use the vSphere Web Client extension for IBM Spectrum Control to provision file storage on IBM Spectrum Control.

Viewing NPIV connections between server ports and switch ports in a fabric

In this scenario, a storage administrator views the N-Port ID Virtualization (NPIV) connections between the server ports and the switch port in the fabric. With NPIV, multiple node ports can be logically connected to one switch port.

You want to view connectivity information for a switch port that has NPIV connections.

1. In the menu bar, go to **Network > Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.

Tip: Alternatively, in the menu bar, go to **Network > Fabrics**. Double-click the fabric that you want to view, or right-click the fabric and click **View Details**. In the **Internal Resources** section, click **Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.

2. In the **Internal Resources** section, click **Ports**. You can view details about the ports, including the number of NPIV connections.
3. In the **Connected NPIV Ports** column for the port that you want to view, click the number of ports. You can view the following information about the NPIV connections:
 - In the **Connected NPIV Ports** tab, you can view details about the NPIV connections, such as the port type and the name of the connected resource. The port type can be N_Port (end node port) or NP_Port (proxy node port).
 - In the **Connectivity** tab, you can view details about the physically connected end-node port.


Tip: If IBM Spectrum Control does not differentiate between the physical and logical connections for the switch port, the details about the NPIV connections are displayed in the **Connectivity** tab. The **Connected NPIV Ports** tab is not displayed.

Exporting performance data for a SAN Volume Controller system

In this scenario, you export performance data for a SAN Volume Controller to a compressed package. You then send the package to IBM Support.

Some volumes on a SAN Volume Controller system have performance problems. You consult IBM Support who require detailed performance data about the SAN Volume Controller system to diagnose the problem. You are asked to export performance data by using IBM Spectrum Control.

1. Go to **Storage > Block Storage Systems**.



2. Right-click the SAN Volume Controller system, then click **Create Performance Support Package**.



Block Storage Systems

✓ 12 Normal
⚠ 3 Warning
✗ 10 Error

Storage Systems ✗ Alerts Tasks Performance

⋮ Actions + Add Storage System View Performance

Name	Condition	Probe Stat...
2107.7503611	✓ Normal	✗ Failed
DS8000-2107-75FG401-L...		ailed
DS8000-2107-75XC891-L...		cce...
oxide		iled
Storwize V3700-2072-tpcr...		cce...
Storwize V3700-2072-v37...		cce...
Storwize V5000-2078-tpcv...		cce...
tbnet-app		ver...
XIV-2810-XIV_A-IBM		cce...
XIV-2810-XIV_D-IBM		cce...
XIV-2810-XIV_E-IBM		cce...

View Properties
 View Details
 View Performance
 View Capacity
Export Performance Data
 Data Collection
 Edit Alert Definitions
 Edit Alert Notification Settings
 Add to General Group
 Add to Capacity Pool
 Connections
 Remove
 Open Storage System GUI

3. Select a time range for the support package. Typically, the time range includes the time when the performance problem occurred.
4. If you are instructed to do so, select the **Advanced export** check box. Click **Next**.

Export Performance Data

Select a time range to export performance data for ☒ Advanced export

Last: 4 hours **8 hours** 12 hours

03:30 Apr 21, 2017 - 11:30 Apr 21, 2017

Next
Cancel

5. Click **Create** when you are prompted to create the package with the details that you specified.

The package is created and your browser downloads the package.

Related reference:

“Exporting performance data for storage systems and fabrics” on page 237

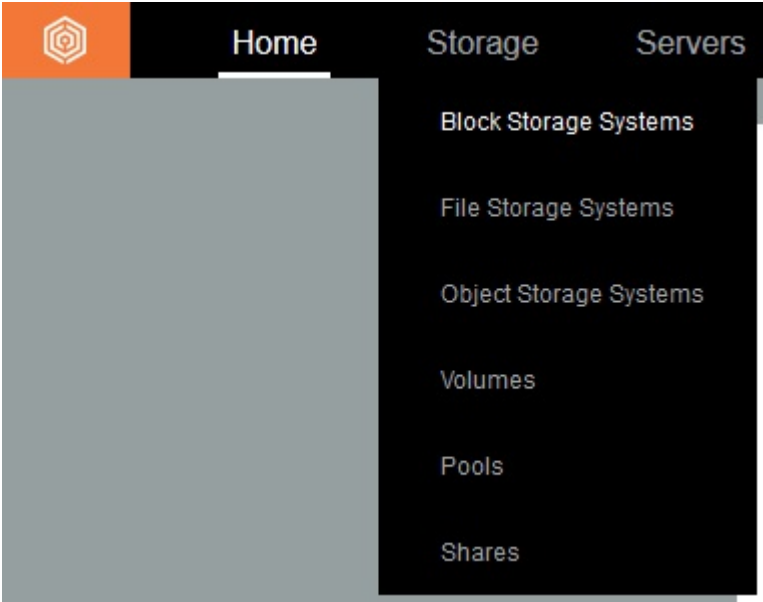
To help resolve performance issues with storage systems and fabrics, you can export performance data for the resources to a compressed file. If you contact IBM Support to help you analyze the problem, you might be asked to send this file.

Comparing the performance of storage systems

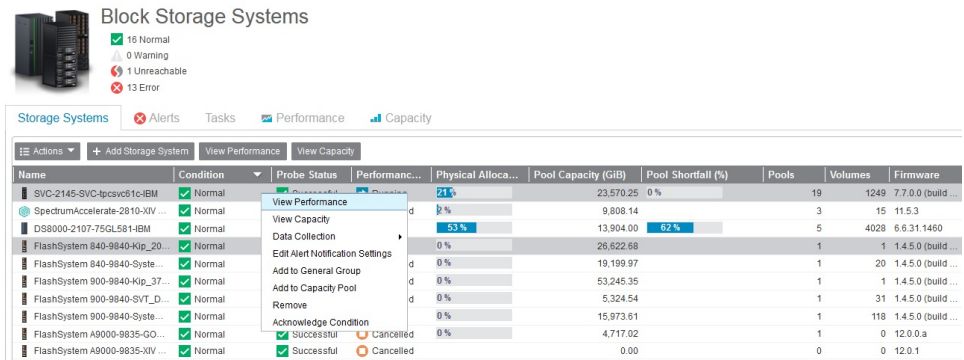
In this scenario, you compare the performance of two storage systems.

Your company recently purchased an IBM FlashSystem 840 storage system. The storage system was added to IBM Spectrum Control and the performance of the storage system is being monitored. Your managers ask you to compare the performance of the FlashSystem 840 storage system to another storage system so that they can assess the performance of their investment. In particular, your managers are interested in the response times and the I/O rates of the FlashSystem 840.

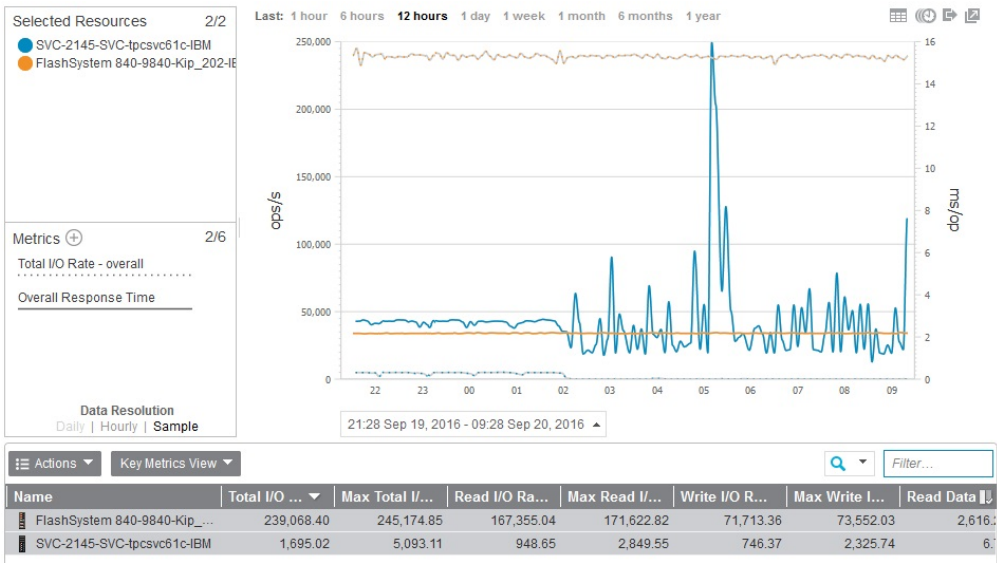
1. Go to **Storage > Block Storage Systems**.



2. Click the FlashSystem 840 storage system and the non-flash storage system that you want to compare the FlashSystem 840 with.
3. Right-click the selected storage systems, then click **View Performance**.



A performance view is displayed. The view shows the performance of the two storage systems on a chart and in a table. By default, the performance view shows the Total I/O Rate - overall and Overall Response Time metrics. The time period for the information is the last 12 hours.



Troubleshooting performance

Bob is a storage administrator. One morning, he receives a ticket: A critical medical application, Epic Database, has a performance problem. Bob must investigate the problem and identify the cause.

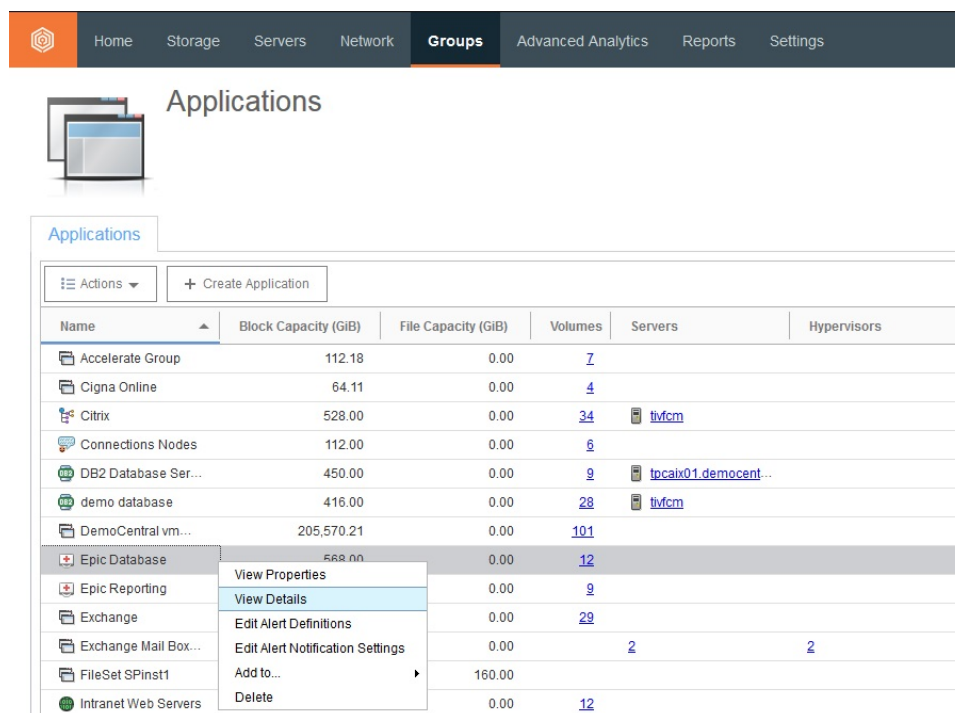
Previously, Bob created an Epic Database application on IBM Spectrum Control and specified the resources that the application uses.

1. Bob clicks **Applications** on the dashboard to troubleshoot only the Epic Database resources.


Applications 17



2. Bob right-clicks the **Epic Database** application on the Applications page, then clicks **View Details**.



Bob immediately notices the high response times of one volume, vsc-epic-db-test-goldp1, in the Most Active Volumes chart. He sees at a glance that this volume is well above 10 ms/op.



Bob is familiar with the response times and other network conditions that are acceptable for the application. He knows that this response time is too high for this application. He also notes that the high response time is sustained, not just spikes. Bob decides to look more closely at the performance of the volume.

- Bob clicks **Volumes** in the Related Resources section, then clicks the **Performance** tab.
- Bob uses the chart and table in the **Performance** tab to further investigate the cause of the problem with the Epic Database application. He uses the chart and table in the **Performance** tab to investigate by using the following techniques:

- Change which metrics are displayed in the performance charts. He clicks a metric in the chart legend to display only that metric. He also clicks the **Select Chart Metrics** icon to add other metrics to the chart.
- Change which resources are displayed in the performance charts. He clicks a resource in the legend to display only that metric. He also clicks resources in the table to add them to the chart.
- View the relationships between resources.
- Compare metrics by time range. To make this comparison, he can select either a predefined time range or a custom time range.
- Compare different metrics within the same time range.
- Compare the same metrics across different time ranges.
- Identify the best or worst values for a performance metric from a large list of resources. For example, Bob might want to see which of the managed disks is handling the most workload, or which volumes have the highest response times.
- Compare performance charts side by side. To compare charts, Bob opens a chart in a new window and arranges the charts beside one another.
- Synchronize all charts so that they all show data for the same time period.
- Export the data from a chart in CSV format.

To see how Bob investigates the performance problem, watch the IBM Spectrum Control Performance troubleshooting techniques (6:49) video on YouTube at <https://youtu.be/4J1jh5hYV-U>.

Monitoring IBM Spectrum Scale performance

You can use IBM Spectrum Control to monitor the performance of the IBM Spectrum Scale clusters in your storage environment.

In the following video, you learn how to use IBM Spectrum Control to compare the performance of multiple IBM Spectrum Scale clusters and then analyze the clusters that have the heaviest I/O loads or the longest response times.△△

Watch the Monitoring IBM Spectrum Scale performance with IBM Spectrum Control (5:38) video on YouTube at <https://youtu.be/9kPbDo9bQSG>.

In the video, you learn how to use the following techniques to identify performance issues and bottlenecks in your storage:△

- Analyze the performance of all your clusters over a period to determine which clusters have the heaviest I/O loads or the longest response times.
- View which of the nodes that mount a file system have the heaviest I/O loads and highest CPU utilization.
- Compare the file systems on a cluster to determine which file systems have the heaviest I/O loads and the longest response times.
- Monitor the performance of SAN attached volumes that are backing the Network Shared Disks △(NSDs) in the cluster.△ If you have block storage devices, such as IBM Spectrum Accelerate or IBM FlashSystem, that are providing storage for your clusters, you can monitor the block volumes that are backing the NSDs.

Chapter 6. Managing resources

Use IBM Spectrum Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

Use the vSphere Web Client extension for IBM Spectrum Control to manage a vSphere virtual storage environment. You can complete provisioning tasks and view information about storage resources in the vSphere Web Client extension.

Resources that you can monitor

















You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Information about object storage systems includes information about their object resources and related resources. Object resources are resources in a GPFS cluster, such as accounts and containers, that enable the GPFS cluster to be used as an object storage system.

Fabrics

A *fabric* is a network of hubs, switches, adapter endpoints, and connecting cables that support a communication protocol between devices. The following table shows the internal resources and related resources that you can monitor for fabrics.

Table 20. Fabric resources that you can monitor.











Fabric icon	Internal resources	Related resources	*Connected resources (to ports)
 Physical fabric  Virtual fabric	 Logical switches  Physical switches  Inter-switch connections  Blades  Ports  Zone sets	 Hypervisors  Servers  Storage systems  NPV switches (Cisco switches in NPV mode and Brocade switches in Access Gateway mode)  Discovered ports	<p>All monitored resources that are connected to ports, plus the following resources:</p> <p> Unmanaged resource. The name is determined from a fabric zone alias.</p> <p> Unmanaged resource. The name is determined from a storage system's mapping to a host.</p> <p> Unmanaged resource. The name is determined from one of a storage system's mappings to a host when multiple mappings are found.</p>
<p>*Connected resources (to ports): The icons for any resources that are managed by IBM Spectrum Control and connected to a port are shown in this column. Icons are also shown for connected resources whose names are known from a fabric zone alias or a storage system's mapping to a host (host connection) but are not currently being managed. If both a fabric zone alias and a storage system's mapping to a host are found for the same unmanaged resource, the fabric zone alias is shown as the name.</p>			

Hypervisors

A *hypervisor* is a layer of software or a physical device that manages multiple, virtual instances of operating systems on the same hardware. Hypervisors that you can monitor include VMWare vSphere data sources such as ESX, ESXi, and vCenter servers. If a hypervisor is part of a cluster, information about that cluster is also available.

The following table shows the internal resources and related resources that you can monitor for hypervisors.











Table 21. Hypervisor resources that you can monitor.

Hypervisor icon	Internal resources	Related resources
	 Virtual machines  Controllers  Disks  Data stores  VMDKs	 Storage systems  Fabrics  Switches  Cluster

Servers

A *server* is a computer or host that is connected to a network and provides an access point to that network. If a server is part of a cluster, information about that cluster is also available. You can monitor the following internal resources and related resources for servers.

Table 22. Server resources that you can monitor.

Server icon	Internal resources	Related resources
	 Controllers  Disks  Volume groups  File systems and logical volumes  Shares and exports	 Storage systems  Fabrics  Switches  VMDKs

Storage systems






















A *storage system* is a device, such as a Redundant Array of Independent Disks (RAID) controller that creates and manages other storage devices. A storage virtualizer is a storage system that virtualizes storage space from internal storage or from another storage system.

Storage systems can be configured for block storage, file storage, a combination of block and file storage, or a combination of file and object storage. The method that a storage system uses for managing data determines the internal resources or object resources that are monitored.

The following tables show the internal resources, object resources, and related resources that you can monitor for block, file, and object storage systems.




















Block storage systems

Table 23. Block storage system resources that you can monitor

Block storage system icon	Internal resources	Related resources
	 Disks  Drives  External disks  Host connections  I/O groups  Managed disks  Modules  ,  Nodes  Pools  Ports  RAID arrays  Volumes	 Back-end storage systems  Fabrics  Hypervisors  Servers  Switches  Virtualizer storage systems  GPFS clusters


















File storage systems

Table 24. File storage system resources that you can monitor

File storage system icon	Internal resources	Related resources
	 Disk Controllers  Clusters  Disks  Shares  Filesets  File systems  Network shared disks  Nodes  Pools  Snapshots  Quotas	 Block storage systems  Fabrics  GPFS clusters  Object storage systems  Hypervisors  Servers  Switches

Object storage systems

Table 25. Object storage system resources that you can monitor

Object storage system icons	Object internal resources	Object resources	Related resources
 Object storage system  IBM Cloud Object Storage	 Access pools  Accesser nodes  Mirrors  Sites  Slicestor nodes  Storage pools  Vaults	 Accounts  Containers	 Back-end storage systems  Fabrics  GPFS clusters  Hypervisors  Servers  Switches

Switches

A *switch* is a network infrastructure component to which multiple nodes attach. In IBM Spectrum Control, you can monitor Fibre Channel (FC) switches. The following table shows the internal resources and related resources that you can monitor for switches.

Table 26. Switch resources that you can monitor.

















Switch icon	Internal resources	Related resources	*Connected resources (to ports)
 Physical switch  Logical switch	 Inter-switch connections  Blades  Logical switches  Physical switches  Ports	 Fabrics  Hypervisors  Servers  Storage systems  Switches  Discovered ports	All monitored resources that are connected to ports, plus the following resources:  Unmanaged resource. The name is determined from a fabric zone alias.  Unmanaged resource. The name is determined from a storage system's mapping to a host.  Unmanaged resource. The name is determined from one of a storage system's mappings to a host when multiple mappings are found.

Table 26. Switch resources that you can monitor (continued).

Switch icon	Internal resources	Related resources	*Connected resources (to ports)
<p>*Connected resources (to ports): The icons for any resources that are managed by IBM Spectrum Control and connected to a port are shown in this column. Icons are also shown for connected resources whose names are known from a fabric zone alias or a storage system's mapping to a host (host connection) but are not currently being managed. If both a fabric zone alias and a storage system's mapping to a host are found for the same unmanaged resource, the fabric zone alias is shown as the name.</p>			

Related tasks:

“Adding ESX and ESXi hypervisors” on page 122

Add ESX and ESXi hypervisors for monitoring by IBM Spectrum Control.

“Adding storage systems” on page 115

To monitor a storage system in IBM Spectrum Control, configure a connection to the storage system and schedule the collection of data.

“Adding fabrics and switches” on page 120

Add fabrics and switches for monitoring.

“Adding servers” on page 125

Add servers for monitoring. You can add a server by deploying a Storage Resource agent, or you can add an agentless server.

Adding resources

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Related concepts:

“Alerting and event processing” on page 148

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected.

Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Related reference:

“Triggering conditions for alerts” on page 370

Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

Required user roles for monitoring resources

When you add storage systems and hypervisors for monitoring, you must provide a user name and password for logging in to those resources. The role or user group that is assigned to the user name determines the data collection and provisioning functions that you can use.

The following roles are associated with the user names that IBM Spectrum Control uses to log in to resources. Specify user names when you add a resource for monitoring. These roles are different from IBM Spectrum Control roles, which are assigned to users that log in to IBM Spectrum Control.

Any roles that are not listed, but include the privileges of the roles that are listed, can also be used for monitoring resources.

Table 27. Required roles for storage system and hypervisor users.

Resource	Required role for performance monitors	Required role for probes	Required role for provisioning
DS6000 DS8000	Monitor role or higher	Monitor role or higher	Administrator, Physical operator, or Logical operator
SAN Volume Controller IBM Storwize family Storwize V7000 Unified (block storage) FlashSystem V840	Administrator	Monitor	Administrator
Storwize V7000 Unified (file storage) IBM SONAS	Performance monitoring is not available	Monitor	Administrator, or Export Administrator + Storage Administrator. For more information about roles, see IBM Knowledge Center for IBM SONAS .
XIV	Monitor role or higher	Any role	Storage administrator
IBM Spectrum Accelerate IBM FlashSystem A9000 IBM FlashSystem A9000R	Monitor role or higher	Any role	Provisioning is not available
IBM FlashSystem 840 IBM FlashSystem 900	Monitor role or higher	Monitor role or higher	Provisioning is not available
IBM Spectrum Scale (file storage)	Any role	Root or non-root*	Provisioning is not available

Table 27. Required roles for storage system and hypervisor users (continued).

Resource	Required role for performance monitors	Required role for probes	Required role for provisioning
IBM Spectrum Scale (object storage)	Performance monitoring is not available	The Keystone admin role. Information is collected only for the object storage accounts and containers that the user has access to. If you want to monitor all accounts and containers, the user must also be assigned the role that is defined in the reseller_admin_role configuration option in the Swift proxy server. The default value for the reseller_admin_role option is ResellerAdmin.	Provisioning is not available
IBM Cloud Object Storage	Performance monitoring is not available	Operator, System Administrator, or Super User role	Provisioning is not available
NetApp devices	Operator role or higher	Operator role or higher	Storage administrator
Cicso switches	network-admin	network-admin	Not applicable
Hypervisors (such as ESX, ESXi, and vCenter Servers)	Performance monitoring is not available	A role that has permission to browse through data stores. For example: Administrator role or Virtual Machine Power User role.	<ul style="list-style-type: none"> Provisioning within the IBM Spectrum Control: any role Provisioning with the vSphere Web Client extension: Storage administrator or Administrator roles.
Rollup servers	Performance monitoring is not available	Administrator	Provisioning is not available
Note: * You can add GPFS clusters and GSS systems as a non-root user, but that user must have privileges to run a set of specified administration commands using the sudo command on the cluster node.			

Adding storage systems


To monitor a storage system in IBM Spectrum Control, configure a connection to the storage system and schedule the collection of data.

To add a storage system, you must have Administrator privileges.

You can add the following types of storage systems for monitoring:

- IBM Cloud Object Storage
- DS4000, DS5000, or DS6000
- DS8000
- ESS
- IBM FlashSystem family
- IBM SONAS
- IBM Spectrum Accelerate
- IBM Spectrum Scale
- SAN Volume Controller
- Storwize family
- Storwize V7000 Unified
- XIV
- EMC VMAX family, EMC VNX family, and EMC VNXe family
- All others (managed by CIM agents)

SAN Volume Controller, IBM Storwize family, IBM FlashSystem V840 Enterprise Performance Solution, and IBM FlashSystem V9000 storage systems run IBM Spectrum Virtualize to virtualize their storage. In this documentation, IBM Spectrum Virtualize is used to refer collectively to SAN Volume Controller, Storwize, FlashSystem V840, and FlashSystem V9000 storage systems.

For a complete list of the versions of storage systems and CIM agents that you can add, see  IBM Spectrum Control interoperability matrix for storage systems at <http://www.ibm.com/support/docview.wss?uid=swg27047049>.

1. To add a block storage system, go to **Storage > Block Storage Systems**. To add a file storage system, go to **Storage > File Storage Systems**. To add an object storage system, go to **Storage > Object Storage Systems**.
2. Click **Add Storage System**.
3. Click the icon for the type of storage system that you want to add.

Tip: Click **Storwize Family** to add the following storage systems:

- IBM® Storwize® V3500
- Storwize V3700
- Storwize V5000
- Storwize V7000

Click **FlashSystem Family** to add the following storage systems:

- FlashSystem V840
- FlashSystem V9000
- FlashSystem 840
- FlashSystem 900
- FlashSystem A9000

- FlashSystem A9000R

Click **EMC** to add the following storage systems:

- EMC VMAX family with support for SMI-S 1.6
- EMC VNX family with support for SMI-S 1.6
- EMC VNXe family with support for SMI-S 1.6
- All other EMC storage systems with support for SMI-S 1.2

4. Complete the connection information for the storage system and schedule the collection of data.

The storage system is added for monitoring. A probe is automatically run to collect status and asset information about the storage system.

Storage systems can be configured for block storage, file storage, a combination of block and file storage, or a combination of file and object storage. The method that a storage system uses for managing data determines the internal resources and object resources that are monitored.

When the collection of data is complete, you can view status information and capacity data about the storage system on the resource list page for the block, file, or object storage system. For example, for block storage systems, you can view the information on the Block Storage Systems page.

Adding and configuring NetApp devices

After adding and configuring a NetApp device, you can use IBM Spectrum Control to administer and monitor the status, configuration, capacity, and information that is collected about the device.

A NetApp device can be added to IBM Spectrum Control as a block storage device, a file storage device, or both. IBM Spectrum Control collects different information from the device, depending on the type of storage system you specify.

- When you add a NetApp device as a block storage system, IBM Spectrum Control collects data by using a CIM agent.
- When you add a NetApp device as a file storage system, IBM Spectrum Control collects data through the native API (NAPI) of the device.

Important:

- For new IBM Spectrum Control installations, you can no longer deploy Storage Resource agents for monitoring NetApp devices. You must now monitor those devices through a CIM agent or a direct connection. However, if you upgrade from a previous version of the product, you can continue to manage NetApp devices that have Storage Resource agents deployed from that previous version.
- When you upgrade from a previous version to IBM Spectrum Control 5.2.10 or later, the existing alert definitions for NetApp filers are migrated to the new version. However, to generate alerts for those filers, you must monitor them through direct connections rather than through Storage Resource agents. To set up direct connections for those NetApp filers, add them again for monitoring and enter the credentials for connecting to them directly. After you add the NetApp filers, their previous entries are automatically removed and alerts will be generated for the alert definitions that were migrated.

To add a NetApp device for monitoring by IBM Spectrum Control, complete the following steps:

1. To add a NetApp device as a block storage system, go to **Storage > Block Storage Systems**. To add the device as a file storage system, go to **Storage > File Storage Systems**.
2. Click **Add Storage System**.
3. Click the icon for the type of storage system that you want to add: **NetApp** for to add the device as a file storage system, **All others** to add the device as a block storage system.
4. Complete the connection information for the storage system and schedule the collection of data.

After the initial discovery has completed, you can view information about the NetApp device in the **File Storage Systems** or **Block Storage Systems** pages in the GUI.

Monitoring IBM Spectrum Scale without requiring root privileges

Enable users that do not have root privileges on a GPFS cluster node to monitor IBM Spectrum Scale™ storage systems.

Before you add the IBM Spectrum Scale storage system for monitoring, complete this task.

To grant access to the user, make the following changes to the sudoers file on the cluster node that is used for authentication:

- Add the set of administration commands that are required to monitor the storage system.
 - Associate the user with the set of administration commands.
1. Log on to the cluster node that is used for authentication with a user name that has root privileges.
 2. To edit the sudoers file, enter the following command:
`visudo -f /etc/sudoers`
 3. Add the following command aliases to the sudoers file. The command aliases contain the commands that the user must be able to issue to monitor the storage system.

Important: Each command alias must be on a single line without line breaks.

```
Cmdnd_Alias TPC_GPFS_MMCMDS = /usr/lpp/mmfs/bin/mmsdrquery, /usr/lpp/mmfs/bin/mmconfig,
/usr/lpp/mmfs/bin/mmgetstate, /usr/lpp/mmfs/bin/mmnodeclass, /usr/lpp/mmfs/bin/mmnodefs,
/usr/lpp/mmfs/bin/mmdf, /usr/lpp/mmfs/bin/mmnodeclass, /usr/lpp/mmfs/bin/mmnodefs,
/usr/lpp/mmfs/bin/, /usr/lpp/mmfs/bin/mmnodeclass, /usr/lpp/mmfs/bin/mmnodefs,
/usr/lpp/mmfs/bin/mmrepquota, /usr/lpp/mmfs/bin/mmnodepolicy, /usr/lpp/mmfs/bin/mmnodepolicy
```

```
Cmdnd_Alias TPC_GPFS_MMDSH = /usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmdsh --version,
/usr/lpp/mmfs/bin/mmdsh -N * /lib/udev/scsi_id --whitelisted *,
/usr/lpp/mmfs/bin/mmdsh -N * /sbin/blockdev --getsize64 *,
/usr/lpp/mmfs/bin/mmdsh -N * /usr/bin/getconf DISK_SIZE *,
/usr/lpp/mmfs/bin/mmdsh -f 20000 -N linuxNodes 'cat /sys/class/fc_host/*',
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces node list,
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces service list -a,
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces address list|grep object_database_node,
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces address list
--by-node|grep object_database_node,
/usr/lpp/mmfs/bin/mmdsh -v -N cesNodes /usr/lpp/mmfs/bin/mmobj config list
--ccrfile object-server.conf --section DEFAULT --property devices,
/usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N * "test -e /opt/IBM/zimon/ZIMonSensors.cfg &&
(grep -w 'collectors' -A 4 /opt/IBM/zimon/ZIMonSensors.cfg | grep -w 'host') || true",
/usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N nonWindowsNodes hostname
```

```
Cmdnd_Alias TPC_GPFS_OTHER = /bin/cat *release, /usr/bin/lsb_release -a
```

```
Cmdnd_Alias TPC_GPFS_CMDS = TPC_GPFS_MMCMDS, TPC_GPFS_MMDSH, TPC_GPFS_OTHER
```

4. To enable the user to issue the commands, add the following lines after the command aliases:

```
Defaults:user_name !requiretty
user_name ALL=(ALL) TPC_GPFS_CMDS
```

where *user_name* is the user name that you enter when you add the storage system for monitoring.

The user that you added to the sudoers file can now monitor the IBM Spectrum Scale storage system.

Verifying that probe data can be collected for IBM Spectrum Scale with object storage

To probe object storage, the IBM Spectrum Control server must be able to connect to the OpenStack Swift and Keystone endpoints that are used to access object services.

To verify that the server can access the object services, complete these steps:

1. List the URLs for the Keystone and Swift services by running commands on a GPFS cluster node that is configured for object storage. Run the following commands with a user name that has root privileges:

```
. ~/openrc
openstack endpoint list
```
2. Ensure that the IBM Spectrum Control server can connect to the IP addresses and host names that are included in the Keystone and Swift services URLs. For example, the URL for the Keystone service might be `http://gpfs420proto1:5000/v3`. The IBM Spectrum Control server must be able to connect to the `gpfs420proto1` host name.

Configuring the collection of performance data for IBM Spectrum Scale

Configure the collection of information about the performance of IBM Spectrum Scale file storage systems.

To collect performance data, you must first configure the IBM Spectrum Scale performance monitoring tool on the GPFS cluster. Then, use the IBM Spectrum Control GUI to schedule the collection of performance data.

1. Configure the sensor component of the IBM Spectrum Scale performance monitoring tool by editing the `/opt/IBM/zimon/ZIMonSensors.cfg` file on the sensor nodes.

Set the host property in the collectors section of the file to the GPFS cluster node where the collector component is running. The host property must be set to one of the following options:


- An IP address that can be reached by the IBM Spectrum Control server
 - A host name that resolves to an IP address that is reachable by the IBM Spectrum Control server
2. Enable the following sensors on each node for which you want to collect performance data: CPU, Memory, GPFSNode, GPFSFileSystem.
 3. Start the collector component of the performance monitoring tool on a single GPFS cluster node. The collector can run on any node.
 4. Start the sensor component of the performance monitoring tool on each node for which you want to collect performance data. By default in IBM Spectrum Scale version 4.1.1, the sensors are started only on protocol nodes.

5. To receive data from the performance monitoring tool, ensure that the IBM Spectrum Control server can connect to port 9084 on the GPFS cluster node where the collector component is running.
6. Schedule the collection of performance data by creating a performance monitor in the IBM Spectrum Control GUI. You can create performance monitors in the GUI when you add resources for monitoring or you can create them later.

A performance monitor is created for the IBM Spectrum Scale storage system. If a successful probe run is completed for the resource, the performance monitor runs according to the defined interval.

To check the progress of a performance monitor, you can complete the following actions:

- View the Performance Monitor Status column on the File Storage Systems page.
- From the File Storage Systems page, right-click a row, and select **Data Collection > Open Performance Monitor Logs**.

To learn more about configuring the performance monitoring tool, see  http://www.ibm.com/support/knowledgecenter/STXKQY_4.1.1/com.ibm.spectrum.scale.v4r11.adv.doc/bl1adv_PMToverview.htm.

Configuring OpenStack access to monitor the object storage system

Configure OpenStack access for the user name that is used to monitor the IBM Spectrum Scale object storage system.

Before you add the IBM Spectrum Scale object storage system for monitoring, ensure that the user name that you use to monitor the storage system has access to object storage accounts.

Tip: The terms "account" and "project" mean the same thing. Swift, the OpenStack object storage service, uses the term "account" and Keystone, the OpenStack identity service, uses the term "project."

1. Set the object storage account and the domain for the user. By default, the domain is set to Default and cannot be modified after you create the user account. Choose one of the following actions:
 - a. To set the account and domain when you create a user account, use the following command:


```
openstack user create --domain domainname --project projectname
--password Password username
```
 - b. To set the account and the domain for an existing user name, use the following command:



```
openstack user set --project projectname username
```
2. Assign the admin role for an object storage account to the user name. Use the following command:


```
openstack role add --user username --project projectname admin
```
3. To monitor all accounts on the object storage system, assign to the user name the role that is defined in the **reseller_admin_role** configuration option in the Swift proxy server. The default value for the **reseller_admin_role** option is ResellerAdmin. For example, use the following command:


```
openstack role add --user username --project projectname ResellerAdmin
```

Restriction: If you do not assign the ResellerAdmin role, information is collected only for the object storage accounts that the user has admin access to.

To learn more about configuring and assigning user roles in OpenStack, see

 http://docs.openstack.org/admin-guide/cli_manage_projects_users_and_roles.html.

Adding fabrics and switches


Add fabrics and switches for monitoring.


For a Cisco switch to successfully receive and respond to queries from IBM Spectrum Control, the following basic requirements must be met:

- IBM Spectrum Control can use SNMPv3 (preferred) or SNMPv1 to probe switches and fabrics. The SNMPv3 protocol is preferred because it provides better security, but switches that use the SNMPv1 protocol are also supported. Some switches are configured to use SNMPv3 by default.
- If the switch uses an SNMP agent, the Fibre Alliance FC Management MIB (FA MIB) and Fibre Channel Fabric Element MIB (FE MIB) must be enabled on the switch.
- When using the SNMPv1 protocol, the community string that is configured in IBM Spectrum Control must match one of the community strings that are configured on the switch with read access. Cisco switches must additionally have a community string match for write access. The default community strings in IBM Spectrum Control are “public” for read access and “private” for write access. Other community strings can be defined on the switches, but are not used. Community strings are not relevant when using the SNMPv3 protocol.
- SNMP access control lists must include the IBM Spectrum Control system. These access control lists are defined and configured on the switches. Some lists automatically include all hosts, while others exclude all by default.
- The Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) protocols must be enabled on the switch. Some switches, such as the Cisco Nexus 5000 series, require you to enable these protocols. Otherwise, IBM Spectrum Control will not recognize the switch when you try to add it using the **Add Switches and Fabrics for Monitoring** dialog. For instructions on how to configure Cisco switches for FCoE enablement, go to the Cisco product website at <http://www.cisco.com> and click **Support**.

In the GUI, you can add switches and fabrics for monitoring at the same time. You can add the following types of switches:

- Brocade
- Cisco
- Other fabric vendors

For a complete list of the switches and their versions that you can add, see the switches and directors information in the  IBM Spectrum Control interoperability matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

Help tips in the GUI: To view descriptions of the information that you must enter for a fabric or switch, hover the mouse pointer over the related help icons  .

1. In the menu bar, go to **Network > Switches**.
2. Click **Add Switch**.
3. Select the vendor: **Brocade**, **Cisco** or **Other**.

4. Specify how you want to discover the switches and the fabrics that are connected to the switches. For Cisco switches and fabrics, you must use SNMP agents. For Brocade and other vendors, you can use the vendor's CIM support as an alternative to SNMP.

Type of switch agent	Steps
SNMP agents	<p>For Cisco, Brocade or other vendors, do the following:</p> <ul style="list-style-type: none"> Specify the version of SNMP: v3 (preferred) or v1. If you choose to use v3, specify the User name and associated Authentication password. Select an Authentication protocol and an Encryption protocol. If you choose to use v1, define the Read community and the Write community. Identify the Host name or the IP address of each switch you want to add. <p>Note: For Cisco fabrics based on SNMPv3, both authentication protocol and authentication password are mandatory. For Brocade and other fabric vendors, authentication is not mandatory when managing or adding switches using SNMPv3.</p>
CIM agents	<p>For Brocade switches and fabrics, do the following:</p> <ul style="list-style-type: none"> Select Monitor with Network Advisor to use a CIM agent on Network Advisor. The switches that are managed by the agent are automatically discovered. After this step, you can also specify to use Simple Network Management Protocol (SNMP) data sources for the switches. You can use SNMP for redundancy. Select Monitor without Network Advisor if you do not want to use a CIM agent as the data source. In this case, you are prompted to add switches and their fabrics through SNMP agents. <p>For other vendors:</p> <ul style="list-style-type: none"> Select Use CIM agent as data source if you want the switches that are managed by the CIM agent to be automatically discovered. After this step, you can also select to use SNMP agents for the switches. Do not select Use CIM agent as data source if you do not want to use a CIM agent as a data source. In this case, you are prompted to add switches and their fabrics through SNMP agents.

5. Specify the display name and location of the switches and fabrics that are discovered. For switches that were discovered through a CIM agent, you can also select to monitor them through SNMP. When you add a Brocade director switch as an SNMP data source, use the IP address or host name of the chassis, rather than the IP address or host name of the primary or secondary core routing blades.
6. Schedule data collection for the switches and fabric. The type of resource determines the data collection jobs that you can schedule:
 - Fabrics: probes
 - Switches: probes, performance monitors

Probes collect status and asset information. Performance monitors collect metrics that measure performance.

The fabric or switch is added for monitoring . A probe is automatically run to collect status and asset information about the resource.

When the collection of data is complete, you can view status information about the fabric or switch on the Fabrics page or Switches page.

Enabling the default alerts for resources: To enable the default alerts for a resource so that you are automatically notified when certain conditions are detected, you must view the alert definitions for that resource. For example, to view the alert definitions for a switch, go to **Network > Switches**, right-click a switch, and select **Edit Alert Definitions**. The default alerts for the switch are automatically enabled after you access the Edit Alert Definitions page.

Agent types for switch and fabric functions

Depending on the functions that you want to enable, you can use CIM or SNMP agents to monitor switches and fabrics.

The vendors of switches can help you determine which agents to use.

Table 28. Agent types for switch and fabric functions

Function	Brocade	Cisco	Other
Monitor performance	CIM agent	SNMP agent	Other fabric vendor switches cannot be used for performance monitoring.
Collect information about switches and switch ports	CIM agent	SNMP agent	CIM agent or SNMP agent
Collect information about topology connectivity	CIM agent	SNMP agent	CIM agent or SNMP agent
Collect information about zoning information and zone control	CIM agent	SNMP agent	CIM agent
Generate alerts	CIM agent	SNMP agent	CIM agent or SNMP agent


Adding hypervisors

You can add hypervisors for monitoring by IBM Spectrum Control. Hypervisors include VMware vSphere data sources such as ESX and ESXi, and vCenter Server systems. When you add vSphere data sources, you can then collect data, generate reports, and manage storage that is related to those resources.

To add a hypervisor, you specify connection information for the ESX or ESXi data source. You can add multiple hypervisors by specifying connection information for a vCenter Server system.

Adding ESX and ESXi hypervisors

Add ESX and ESXi hypervisors for monitoring by IBM Spectrum Control.

You can add an ESX or ESXi hypervisor by specifying connection information for that hypervisor. For a complete list of hypervisors that you can add, see the  IBM Spectrum Control interoperability matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>, and go to the *Agents, Servers and Browsers* section.

1. In the menu bar, go to **Servers > Hypervisors**.
2. Click **Add Hypervisor**.

3. Select **ESX/ESXi** and enter connection information about the hypervisor.
4. Schedule a probe of the hypervisor.
5. Follow the instructions in the GUI to add the hypervisor.

The hypervisor is added for monitoring. A probe is automatically run to collect status and asset information about the hypervisor.

When the collection of data is complete, you can view status information and capacity data about the hypervisor on the Hypervisors page.

Enabling the default alerts for hypervisors: To enable the default alerts for a resource so that you are automatically notified when certain conditions are detected, you must view the alert definitions for that resource. For example, to view the alert definitions for a hypervisor, go to **Servers > Hypervisors**, right-click a hypervisor, and select **Edit Alert Definitions**. The default alerts are automatically enabled for the hypervisor after you access the Edit Alert Definitions page.

Related tasks:

"Adding vCenter Server systems" on page 124

For IBM Spectrum Control to monitor multiple hypervisors, you must first add a vCenter Server system.

Checking permissions to browse data stores:

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

When you add a VMware data source in IBM Spectrum Control, the user name that you specify must have permission to browse through the data stores on VMware. IBM Spectrum Control must browse through the data stores to collect information from the hypervisors. However, the "Read Only" role as defined by VMware does not allow IBM Spectrum Control to browse the data stores. You can use the "Virtual Machine Power User" role if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

To verify that a VMware user is assigned the correct role and privileges to monitor VMware data sources, follow these steps:

1. Ensure that the user role has the required VMware datastore permissions by completing the following steps:
 - a. Connect the vSphere Web Client to the VMware data source. The data source can be an ESX server, a vCenter Server Appliance, or a vCenter Server.
 - b. From the Inventories view, select **Hosts and Clusters**.
 - c. Select a host, and click the **Related Objects** tab.
 - d. View the datastores by clicking the **Datastores** tab.
 - e. Right-click a datastore, and select **File Browser**. If you can view the **Files** tab for the datastore, your browse permission is working correctly.
2. Determine the role that is assigned to the user by logging in to the vSphere Web Client by using the administrator user ID. From the Administration view, select **Roles**. Verify the role name that is assigned to the user.
3. Determine the privileges that are assigned to the role by selecting the user's role and clicking **Privileges**. Expand the privilege groups to view the specific privileges.


4. Optional: If you must edit the privileges for the role, select the role and click the **Edit role action** icon. Select privilege groups or expand to select specific privileges.

For more information about VMware user roles, go to the VMware documentation center and search for *vSphere users and permissions*.

Adding vCenter Server systems

For IBM Spectrum Control to monitor multiple hypervisors, you must first add a vCenter Server system.

When you add vCenter Server version 5.1 or later, you can deploy the vSphere Web Client extension for IBM Spectrum Control and register the latter as a VASA provider. You can then provision storage, view reports, and publish alerts in the vSphere Web Client about storage that is monitored by IBM Spectrum Control.

You can add multiple hypervisors by specifying connection information for a vCenter Server system. For a complete list of hypervisors and vCenter Server systems that you can add, see the  IBM Spectrum Control interoperability matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>, and go to the *Agents, Servers and Browsers* section.

1. In the menu bar, go to **Servers > Hypervisors**.
2. Click **Add Hypervisor**.
3. Select **VMware vCenter** and enter connection information about the vCenter Server system.
4. Optional: To deploy the vSphere Web Client extension for IBM Spectrum Control and register IBM Spectrum Control as a VASA provider, in the **Deploy vSphere extension** page, enter your credentials. For **VMware vCenter Server**, enter a user name and password for the vSphere administrator user. For **IBM Spectrum Control**, enter a user name and password for a IBM Spectrum Control user ID. You must have Administrator, Monitor, or External Application authority to deploy the extension. If you entered a vSphere administrator user name and password for the vCenter Server system, enter only the IBM Spectrum Control credentials.
5. Schedule a probe for the hypervisors that were discovered.
6. Follow the instructions in the wizard to add the vCenter Server system.

After a vCenter Server system is added for monitoring, probes collect status and asset information about its hypervisors.

When the collection of data is complete, you can view status information and capacity data about the hypervisors on the Hypervisors page.

Enabling the default alerts for hypervisors: To enable the default alerts for a resource so that you are automatically notified when certain conditions are detected, you must view the alert definitions for that resource. For example, to view the alert definitions for a hypervisor, go to **Servers > Hypervisors**, right-click a hypervisor, and select **Edit Alert Definitions**. The default alerts are automatically enabled for the hypervisor after you access the Edit Alert Definitions page.

Related tasks:

“Checking permissions to browse data stores” on page 123

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

Adding servers

Add servers for monitoring. You can add a server by deploying a Storage Resource agent, or you can add an agentless server.

Choose which type of server to add and which method you use to add the server.

To determine which type of server to add, follow these guidelines:

Table 29. Choosing the type of server to add

Condition	Add an agentless server	Deploy a Storage Resource agent to a server
You want to view the monitored resources to which a server is connected.	X	X
You want to view the amount of storage that the server is assigned from other monitored resources.	X	X
You want to view the performance of the SAN-related resources that are connected to the server.	X	X
The restrictions within your organization prohibit the deployment of agents to a server.	X	
You want to view information about the resources that are internal to a server. Internal resources include controllers, disks, volume groups, file systems, logical volumes, shares, exports, and directories.		X

You can use several methods to add one or more servers for monitoring. Use the following information to determine the method that best meets the needs of your environment:

Manually

Add one server by manually entering configuration information.

From a File List

Add one or more servers by importing configuration information from a comma-delimited file.

From Discovered Servers

Add one or more servers that are members of a Windows domain. If you select this method, IBM Spectrum Control searches for servers that are in the Windows domain where the product is installed. When you select to add one server from the list of discovered servers, you can configure more information about that server. When you select to add multiple servers, the servers are automatically created without additional configuration information.

Restriction: This method is available only when you add an agentless server. You cannot use this method when you add a server by deploying a Storage Resource agent.

Example: Server-centric monitoring of SAN resources without an agent

In IBM Spectrum Control, you can view the connectivity of all your top-level resources in a SAN environment. These resources include servers, hypervisors, fabrics, switches, and storage systems. In previous releases, to get a complete view of the connectivity between servers and other top-level resources, you were required to deploy Storage Resource agents on all of those servers. However, in some environments it was not always possible to deploy Storage Resource agents, and thus a complete view of all resource relationships was not available. In IBM Spectrum Control V5.2 and later, you can add servers without deploying Storage Resource agents and still view the overall connectivity and performance of SAN storage that is assigned to those servers.

When you add an agentless server for monitoring and specify its worldwide port names, IBM Spectrum Control automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, the following information can be determined:

- The amount of SAN storage that is assigned to the server.
- The storage systems that are mapping volumes to the server.
- The relationships between the server and other monitored resources.

If performance monitors collect data about the related storage systems, you can also view the performance of the internal resources that are directly associated with storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.

Adding servers without Storage Resource agents

Add agentless servers for monitoring. An agentless server is a server or virtual machine that is monitored by IBM Spectrum Control but does not have a Storage Resource agent deployed to it.

When you add an agentless server, enter its IP address or host name and other optional information such as its OS type, location, and ports. IBM Spectrum Control uses this information to correlate the agentless server with the storage systems, fabrics, and switches that are being monitored. For example, the ports that you enter for an agentless server are compared to the ports on host connections for monitored storage systems. If there is a match between the server and a host connection, you can view the connectivity between the server and storage system and the capacity of volumes that are assigned to the server.

Use any of the following methods to add agentless servers for monitoring:

- Add a server by manually entering its information.
- Add a server from a host connection on a monitored storage system.
- Add a server from a discovered port that is connected to a monitored switch.
- Add a server from a discovered port that is connected to a monitored fabric.

Add an agentless server by manually entering its information.

1. In the menu bar, go to **Servers > Servers**.

2. Click **Add Server**.
3. Ensure that the **Deploy an agent for full server monitoring** check box is cleared.
4. Select a method for adding a server and complete the steps.

Add a server from a host connection on a monitored storage system.

Use this method when you want to configure a host connection to represent a physical server. When you select a host connection, you are specifying the internal resources (volumes, ports) that belong to that server from the perspective of the related storage system.

To add an agentless server that is based on a known host connection, complete the following steps:

1. In the menu bar, go to **Storage > Block Storage Systems**.
2. Right-click a storage system and select **View Details**.
3. Click **Host Connections** to view a list of host connections that are mapped to a storage system.
4. Right-click a host connection and select **Add Server**. This action is not available under the following conditions:
 - A Storage Resource agent or agentless server is already associated with the host connection.
 - The host type of the host connection is configured for SAN Volume Controller.
5. Enter information about the server, assign ports and complete any other information that you might need.

Add a server from a discovered port that is connected to a monitored switch.

Discovered ports are ports on resources that are not monitored by IBM Spectrum Control but are visible to a monitored switch. To add an agentless server that is based on a discovered port, complete the following steps:

1. In the menu bar, go to **Network > Switches**.
2. Right-click a switch and select **View Details**.
3. Click **Discovered Ports** to view a list of unmanaged ports that are visible to the switch.
4. Right-click an unknown port and select **Add Server**.
5. Enter information about the server, assign ports and complete any other information that you might need.

After you complete the task, the port is removed from the list of discovered ports and added as a server to the list of related resources for the switch.


Add a server from a discovered port that is connected to a monitored fabric.

Discovered ports are ports on resources that are not monitored by IBM Spectrum Control but are connected to ports on a monitored fabric. To add an agentless server that is based on a discovered port, complete the following steps:

1. In the menu bar, go to **Network > Fabrics**.
2. Right-click a switch and select **View Details**.
3. Click **Discovered Ports** to view a list of unknown ports that are connected to ports in the fabric.
4. Right-click an unknown port and select **Add Server**.

5. Enter information about the server, assign ports and complete any other information that you might need.

After you complete the task, the port is removed from the list of discovered ports and added as a server to the list of related resources for the fabric.

Help tips in the GUI: To view descriptions of the information that you must enter for a server, hover the mouse pointer over the related help icons  .

Restrictions:

- The host name or IP address of an agentless server can be up to 24 characters in length and must contain valid characters. You can enter the following characters when you enter a host name or IP address:
 - A - Z (uppercase characters)
 - a - z (lowercase characters)
 - 0 - 9 (numeric characters)
 - Special characters: - (minus) . (period)
- If the location of an agentless server exceeds 64 characters, it is truncated when the server is added. To avoid the truncation, ensure that the location value for an agentless server is equal to or fewer than 64 characters in length.

File List:

Add one or more agentless servers by importing configuration information from a comma-delimited file. IBM Spectrum Control reads the entries in the file and automatically adds the corresponding servers.

The input file for importing server information can contain entries for one or more servers. Each line in the file represents a server that you want to add. Information about each server must be organized in the following format:

```
host name or IP address,OS type,machine type,location,custom tag 1,custom tag 2,  
custom tag 3, WWP1, WWP2, WWPn
```

where:

- host name or IP address is required for each server entry. An IP address can be in an IPv4 or IPv6 format. You must include a host name or IP address for each server that you want to add. The host name or IP address can be up to 24 characters in length and must contain valid characters. You can enter the following characters when you enter a host name or IP address:
 - A - Z (uppercase characters)
 - a - z (lowercase characters)
 - 0 - 9 (numeric characters)
 - Special characters: - (minus) . (period)
- OS type is optional and represents the operating system of the server. The OS type for a server must be one of the following values:
 - Windows
 - Linux
 - AIX
 - Solaris
 - HP-UX

- Machine type is optional and determines whether the server is a physical device or a virtual machine. The machine type for a server must be one of the following values:
 - Physical
 - Virtual
- Location is optional and represents the physical location of the server. If the location of an agentless server exceeds 64 characters, it is truncated when the server is added. To avoid the truncation, ensure that the location value for an agentless server is equal to or fewer than 64 characters in length.
- custom tag 1, custom tag 2, and custom tag 3 are optional and represent any additional information that you want to provide about the server. The custom tag values can be up to 64 characters in length. If the length exceeds 64 characters, the custom tag value is truncated when the server is added.

Tip: The custom tags can be displayed on the Servers page or can be included as report columns when you generate reports for the server.

- WWPN is optional and represents the worldwide port name that is associated with the server. You can include one or more WWPNs for each server. The value for worldwide port name (WWPN) must be 16 hexadecimal characters in length. If the WWPN for a server entry is already associated with a server that is being monitored, that WWPN is not created for the server.

If a server has no information for an optional field, you must add a comma separator unless there are no further values for that server.

Example

```
host1,Windows,Virtual,San Jose,Building 1,Finance Department,Payments,
2100000056AE45B2,2100000056AE45B3 # Comment
host5,Linux,Physical,London,Marketing Department,,,2100000056EC45A4,
21000000565E4FB6
systemxyz.storage.widgets.com,Solaris,Chicago,,,21000000567E5FB7
198.51.100.22,HPUX,Virtual,,,,21000000566E6FB6,21000014566E6FB7,
2100001566E6FB8
2001:DB8:0:0:0:0:0:0,Windows,Physical,Tokyo
```

Tip: If you add many servers in the same input file, the adding process might require some time to complete. To run the process in the background, click **Close** on the Add Server window. If you remain on the Servers page, a message is displayed when the servers are added. If you leave the Servers page during processing, a message is not displayed. To confirm that the servers were added, you must go back to the Servers page and refresh the list.

Restrictions

Keep in mind the following restrictions when you import server information from an input file:

- Each line in the file represents one server.
- Enter a "#" at the beginning of a line to comment out a line. When you add a server, the server on that line is not added when the list is imported.
- If there are syntax problems in the file, then none of the servers in the file are added.

Adding servers with Storage Resource agents

Add servers by deploying Storage Resource agents. Deploy Storage Resource agents if you want to enable full server monitoring.

To add a server, you must have Administrator privileges.

For each installation of IBM Spectrum Control, you can deploy only one Storage Resource agent on each server. If you attempt to deploy additional Storage Resource agents on a server, the deployments fail.

1. In the menu bar, go to **Servers > Servers**.
2. Click **Add Server**.
3. Select **Deploy an agent for full server monitoring**.
4. Select a method for adding a server. You can choose one of the following methods:
 - Add a server by manually entering information about the server and the Storage Resource agent.
 - Add one or more servers by importing configuration information from a comma-delimited file.
5. On the Deploy Agent page, configure deployment information for the Storage Resource agents.

If you add multiple servers with different operating systems, separate configuration pages are displayed for agents that are deployed on Windows servers and agents that are deployed on UNIX servers.

6. On the Configure page, schedule the deployment of the Storage Resource agents.

If you are deploying agents on multiple servers, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the IBM Spectrum Control server.
7. Schedule the time and frequency that probes are run for the servers.

If you add multiple servers, a time span is calculated during which the servers are probed.
8. To add the servers, click **Finish**.

A probe is automatically run for a server after the agent is successfully deployed.

Use **Agent State** on the Servers page to monitor the status of the agent deployment.

Enabling the default alerts for servers: To enable the default alerts for a resource so that you are automatically notified when certain conditions are detected, you must view the alert definitions for that resource. For example, to view the alert definitions for a server, go to **Servers > Servers**, right-click a server, and select **Edit Alert Definitions**. The default alerts for the server are automatically enabled after you access the Edit Alert Definitions page.

File List:

Add one or more servers with Storage Resource agents by importing the configuration information from a comma-delimited file.

The web-based GUI guides you through the following steps for adding servers:

- Select the input file.
- Configure deployment information.
- Schedule the agent deployment and data collection for the servers.

The comma-delimited file that you use to import the configuration information for servers can contain entries for a single server or multiple servers. Each line in the file represents a server that you want to add. The information about each server must be organized in the following format:

host name or IP address,OS type,location,custom tag 1,custom tag 2,custom tag 3

where:

- host name or IP address is required for each server entry. An IP address can be in an IPv4 or IPv6 format. A host name or IP address can contain the following characters or symbols:
 - A - Z (uppercase characters)
 - a - z (lowercase characters)
 - 0 - 9 (numeric characters)
 - Symbols: - . : _
- OS type is required and represents the operating system of the server. The OS type for a server must be one of the following values:
 - Windows
 - Linux
 - AIX
 - Solaris
 - HP-UX
- location is optional and represents the physical location of the server. The location value can be up to 64 characters in length. If the length exceeds 64 characters, the location value is truncated when the server is added.
- custom tag 1, custom tag 2, and custom tag 3 are optional and represent any additional information that you want to provide about the server. The custom tag values can be up to 64 characters in length. If the length exceeds 64 characters, the custom tag value is truncated when the server is added.

Tip: The custom tags can be displayed on the Servers page or can be included as report columns when you generate reports for the server.

If a server has no information for an optional field, you must add a comma separator unless there are no further values for that server.

Example

```
host1,Windows,San Jose,Accounting department
host5,Linux,London,Finance department
198.51.100.22,HP-UX,,Computing department
2001:DB8:0:0:0:0:0:0,Windows,Tokyo
```

Tips:

- If the comma-delimited file contains entries for multiple servers, it might take some time to add the servers. To confirm that the servers are added, check the Status column on the Servers page.
- To comment out a line, enter a "#" at the beginning of the line. The server on that line is not added when the list is imported.

Example:

```
# host1,Windows,San Jose,Accounting department
```

- If there are syntax problems in the file, none of the servers in the file are added.

Adding rollup servers

Rollup servers gather capacity and status information from IBM Spectrum Control servers in your enterprise. You can add rollup servers for monitoring on the **Settings > Rollup Server Connections** page.

IBM Spectrum Control supports master and subordinate rollup servers. Master rollup servers gather capacity and status information about the storage resources that are managed by subordinate rollup servers. You can view this information by enabling IBM Spectrum Control rollup mode. These managed resources include servers, hypervisors, fabrics, switches, and block storage systems.

Configuring rollup servers


To configure rollup servers, you add one or more subordinate servers to a master server. The master server uses probes to collect capacity and status information from the subordinate servers. Master servers must have IBM Spectrum Control V5.2.10 or later installed and running. Subordinate servers must have Tivoli Storage Productivity Center V5.2 or later installed, or any version of IBM Spectrum Control.

Adding a rollup server

To add and configure subordinate rollup servers, go to **Settings > Rollup Server Connections**. You can add one or more subordinate servers to the master (local) server. You can also add a master rollup server as a subordinate server to another master server.

Viewing rollup information

To view information about the resources managed by your rollup servers, you must be in *rollup mode*. To enter rollup mode, hover

your pointer over the rollup mode icon  on the menu bar and select **Enter rollup mode**. You cannot enter rollup mode unless you have added at least one subordinate server to the master (local) server.

Removing resources

Remove resources that you no longer want to monitor with IBM Spectrum Control.

1. In the menu bar, go to the type of resource that you want to remove.

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics
Rollup subordinate servers	Settings > Rollup Server Connections

2. Right-click the resource and click **Remove**. To remove rollup subordinate servers, click **Actions > Remove Server**.

Note: To remove a Cisco fabric, you must remove all the switches in that fabric. The Cisco fabric is then automatically removed.



3. Click **Remove** to confirm that you want to remove the resource.

When you remove a resource, it is no longer monitored by IBM Spectrum Control. All the data that is associated with the resource, including historical data, is removed from the database. All the scheduled tasks and data collection jobs for the resource are canceled.

Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

Table 30. Data collection jobs, resources, and related tasks

Data collection job	Resources	Related tasks
Probe Use probes to collect status, asset, configuration, and capacity information about resources.	Storage systems Servers Hypervisors Fabrics Switches Rollup servers (capacity and status information only)	<ul style="list-style-type: none">• Schedule a probe when you add a resource for monitoring.• Modify the schedule for a probe at any time on the list and details pages for a resource. For example, to modify the schedule for a block storage system, go to the Block Storage Systems page or the details page for the block storage system.• Use the Rollup server connections page to manage probes for rollup servers.  Learn more...
Performance monitor Use performance monitors to collect metrics that measure the performance of resources.	Block storage systems File storage systems Object storage systems Switches	<ul style="list-style-type: none">• Schedule a performance monitor when you add a resource for monitoring.• Modify the schedule for a performance monitor at any time on the Performance page, and on the list and details pages for the resource.  Learn more...

Related tasks:

“Viewing capacity alerts and violations” on page 254

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

“Defining alerts for attribute and capacity changes” on page 151

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.

Collecting asset and status data

Use probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource list and resource details pages, for example, the Block Storage Systems page, or the Block Storage Systems details page.

Creating probes

Probes are data collection jobs that collect status and asset information about your storage resources. You can create probes when you add resources for monitoring or after you add the resources for monitoring.

You can create probes for one or more resources at the same time. You can also create an automated or manual probe. Schedule an automated probe to manage your probe jobs efficiently and to minimize the load on the IBM Spectrum Control server. If you want to run a probe at a specific time, schedule a manual probe.

1. To schedule a probe for a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics
Rollup subordinate servers	Settings > Rollup server connections

Tip: For switches, probes also collect asset and topology information about the fabrics that include those switches.

2. Locate the resources that you want to probe.
3. Select the resource rows and click **Actions > Data Collection > Schedule**. For Rollup subordinate servers, select **Actions > Start probe** for the server.
4. On the Data Collection Schedule window, set the probe to **Enabled**.
5. From the following table, choose the type of probe that you want. Make your selections for the data collection, and save your selections.

Type of probe	Selections for probe
Automated probe	Select the frequency of the data collection. Click Save .
Manual probe	Select the time and the frequency of the data collection. Click Save .

For automated probes, the start time is determined by an algorithm that minimizes the number of concurrent probe jobs to avoid excessive load on the IBM Spectrum Control server. Automated probe jobs are scheduled by IBM Spectrum Control within the time interval that you defined in the Automated Probe Run Window.

If you schedule manual probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the IBM Spectrum Control server.

Probe schedules are created for the resources that are selected. The probe jobs run according to the modified schedule. If you disabled the probes, IBM Spectrum Control stops collecting asset, status, and performance data for the monitored resources.

You select three storage systems, `storage_system_a`, `storage_system_b`, and `storage_system_c`. A probe schedule was defined for `storage_system_c`, but not for `storage_system_a` and `storage_system_b`. You set the time and frequency fields, and set the probe to **Enabled**. When you click **Save**, probe schedules are created for `storage_system_a` and `storage_system_b`. The probe schedule that was defined for `storage_system_c` remains unchanged.

How automated probes are scheduled:

When you schedule probes, you can manually specify when they run, or you can have IBM Spectrum Control automatically schedule them. Automated probe scheduling helps remove the complexity of manually planning when probes are run and ensures that server resources (CPU, Memory, and so on) required for data collection are distributed over time. Distributing the use of system resources over time helps avoid excessive load on IBM Spectrum Control at any one time.

To schedule probes automatically, an algorithm is used to determine the best times for when the probes are run. To determine this time, the algorithm considers the existing schedules of probes for all types of resources. For example, when an automated probe of a storage system is scheduled, IBM Spectrum Control attempts to minimize the number of probes that are running concurrently by analyzing the existing probes of switches, servers, hypervisors, and other storage systems.

IBM Spectrum Control also analyzes the frequency, start time, and expected duration of existing probe schedules. This analysis includes probes that were scheduled automatically or manually. After the analysis, IBM Spectrum Control creates a unique probe schedule for each day of the week.

Keep in mind the following information about automatic probe scheduling:

- When you schedule a new probe during the week, the schedule of existing probes does not change. Existing probes are considered when the schedule of the new probe is determined, but they continue to run according to the previously defined schedule for that week.
- When a probe is scheduled to run on a frequency of every X days where $X > 1$, IBM Spectrum Control determines the days for that frequency based on a number of conditions. Some of the conditions include whether the probe was run previously as a scheduled probe and how much time is available until the end of the current probe window.
- Automated probes can be rescheduled by IBM Spectrum Control under the following conditions:

- When you change the time range for the automated probe schedule, IBM Spectrum Control reschedules all the automated probes to optimize when probes are run during that new time range.
- On Saturday night (in the time zone of the IBM Spectrum Control server) of every week, all automated probes are rescheduled. During the rescheduling process, the algorithm is again applied to determine the best times to run automated probes on each day of the upcoming week. Because all automated probes are rescheduled weekly, IBM Spectrum Control ensures that the average number of concurrent probes is kept minimized for your environment on an ongoing basis.

Verifying that a probe is running for a resource

Verify that a probe is running for a resource to ensure that IBM Spectrum Control is collecting asset and status information about the resource.

To verify that a probe is running for a resource, use the resource list page in the web-based GUI.

1. From the menu bar, go to a resource list page for a resource type. For example, to verify that a probe is running for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Locate the resource that you want to check.
3. View the Probe Status column. A value of Successful or Running indicates that the probe is collecting asset and status information about the resource. The Last Successful Probe column shows the most recent date and time when the probe successfully collected data.
4. Optional: If the Probe Status column shows a status of Failed or Warning, right-click the resource and click **Data Collection > Open Probe Logs** to view the messages in the log file.
5. Optional: If the Probe Status column shows a status of Never probed, right-click the resource and click **Data Collection > Schedule** to schedule a probe for the resource.

Modifying probes

Probes are data collection jobs that collect status and asset information about monitored resources in your environment. To modify a probe in IBM Spectrum Control, use the resource list page for the monitored resource in the web-based GUI.

You can select one or more devices and modify their schedules.

1. To modify a probe for a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

Tip: For switches, probes also collect asset and topology information about the fabrics that include those switches.

2. Locate the resources that you want to probe.
3. Select the resource rows and click **Actions > Data Collection > Schedule**.
4. On the Data Collection Schedule window, modify the schedule for the probes, and click **Save**.

For automated probes, the start time is determined by an algorithm that minimizes the number of concurrent probe jobs to avoid excessive load on the IBM Spectrum Control server. Automated probe jobs are scheduled by IBM Spectrum Control within the time interval that you defined in the Automated Probe Run Window.

If you schedule manual probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the IBM Spectrum Control server.

The following rules determine the settings for the probe fields when you open the Data Collection Schedule window:

- The Enable Probe check box is checked by default. Uncheck this box to disable the probe operation for all selected resources.
- If all selected resources are configured with automated probe scheduling, then Automatically is selected; if all selected resources are configured with manual probe scheduling, then Manually is selected. If some resources are configured with automatic probe scheduling and some resources are configured with manual probe scheduling, then neither Automatically nor Manually is selected.
- Fields that are blank indicate that more than one value (Automatically or Manually) is used among the selected resources. If the field contains a red asterisk (*), you must set a value. If the field is blank, the existing values are preserved unless you select a new value to be applied to all of the selected resources.

The probe jobs run according to the modified schedule. If you disabled the probes, IBM Spectrum Control stops collecting asset and status data for the monitored resources.

The changes that you make to the probe schedule are applied to all the resources that you select. For example, you select two servers with different probe frequency values. Probes are disabled for both servers. When you open the Data Collection Schedule window, the probe status field is set to **Disabled** and the frequency field is blank. You set the status field to **Enabled** and the frequency field to Every 2 days. When you click **Save**, the probe is enabled with the new frequency value for the two servers that you selected.

Configuring alerts for probes

Configure the alert notifications that are generated if a probe fails to run.

You can define alerts for a resource if you want to be notified when the following conditions occur:



- The status of its probe is error or warning. An error status occurs when a probe did not complete and no data was collected about a resource. A warning status occurs when a probe completes, but might not have collected a complete set of data.

- A specified amount of time has passed since a probe collected data about a resource. You can use this type of alert to be notified when up-to-date configuration and status data is not being collected about a resource and its existing data might be stale.

To define an alert for the probe that collects data about a resource, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. Right-click a resource in the list and click **Edit Alert Definitions**.
3. To enable the **Probe Status** alert for an attribute, click the corresponding switch icon  and select a status to be alerted on. You can use this alert to be notified when a probe is only collecting partial data about a resource or not collecting any data.
4. To enable the **Last Successful Probe** alert for an attribute, click the corresponding switch icon  and specify how much time can pass before an alert is generated. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a resource and its existing data might be stale.
5. Click **Save**.

Starting probes

Use the **Start Probe** action to immediately collect data about resources.

If a probe job is not defined for the resource, use the **Schedule** action to create a probe.

You can start probes for multiple resources at the same time.

Tip: When you start a probe for a switch, the probe gathers statistics about the fabric that the switch is a part of.

1. From the menu bar, go to the resource list page for a resource type. For example, to start a probe for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Locate the resources that you want to probe.
3. Select the resource rows and click **Actions > Data Collection > Start Probe**. The **Start Probe** action is available for the resources that you select if at least one of the resources meets the following criteria:
 - A probe job is defined for the resource.
 - The probe job is enabled.
 - The probe job is not currently running.

The probe status changes to Running.

Probes are started for all the resources that meet the criteria. For example, you select two storage systems, `Storage_System_a` and `Storage_System_b`. Probes are enabled for both of the storage systems and a probe is running for `Storage_System_a`. The probe is started for `Storage_System_b` when you select **Start Probe**. `Storage_System_a` is not affected by the action because a probe is already running for the resource.


Viewing probes for a specific resource

You can view information for specific resources on the resource details pages. For example, use the Storage Systems details page to view the date and time when data was last collected.

1. In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want to view the resource information for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Locate the resource that you want to view information for.
3. Right-click the resource row and select **View Details**.
4. In the General section of the resource details page, click **Data Collection**.

The Data Collection pane is displayed, which contains information about the data collection jobs that are defined for the resource.

Note: To view information about rollup resources, you must be in *rollup mode*. To

enter rollup mode, move the mouse pointer over the rollup icon  in the menu bar and select **Enter rollup mode**.





Viewing probe logs

A probe log file is created for each run of a probe on a storage resource. Probe logs include detailed information about the status, actions, and progress of a probe.

1. In the menu bar in the web-based GUI, go to a resource list page for a resource type. For example, to view the probe logs for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Locate the resource that you want to view.
3. Right-click the resource row, and select **Data Collection > Open Probe Logs**.

Restriction: If you select multiple rows, the **Open Probe Logs** action is not available.

The Logs page opens and shows log entries for the most recent run of the probe.

4. Optional: On the Logs page, to view a previous probe log, select a log from the **Select a log** list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only entries that have the following statuses:
 -  Only error entries
 -  Only warning entries
 -   Error and warning entries

6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

You can view the following information on the Logs page:

- The overall status of a probe. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the job run.
- The status for each probe action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The description of an action.

Collecting performance data

Before you can monitor the performance of the resources in your environment, you must collect data about those resources. IBM Spectrum Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the Performance Monitors page to view and manage performance monitors and to track threshold violations.

You must schedule and run performance monitors before you can complete other tasks, such as optimizing storage tiering, balancing pools, and running performance reports.

Before you can view performance information for resources, you must complete the following tasks:

- Add the storage system or switch for monitoring by IBM Spectrum Control.
- Schedule a performance monitor to collect performance data about the resource.
- Define a performance alert so that you are notified when the performance of a resource might represent a potential problem. For example, you can define an alert threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.

The performance information that is gathered includes metrics that measure the performance of volumes, ports, and disks. You can view performance metrics on resource list and detail pages for switches and storage systems. For example, you can view performance metrics that are related to switches in the **Performance** tab on the Switches page.

Creating performance monitors

Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in the web-based GUI when you add resources for monitoring or at a later time.

To create performance monitors, you must have Administrator privileges.

To create performance monitors for resources in your environment, you can use the list and detail pages for the resources. For example, to create performance monitors for switches, use the Switches page. You can create performance monitors for multiple resources at the same time.

Restrictions:

- You can define only one performance monitor for a resource.
- You can create performance monitors for the following resources only:
 - Storage systems

- Switches that are managed by a CIM agent
1. From the menu bar, go to a resource list page for a resource type. For example, to create performance monitors for storage systems, go to **Storage** and select the type of storage system you want to monitor.
 2. Locate the resources that you want to monitor.
 3. Select the resource rows and click **Actions > Data Collection > Schedule**.
 4. On the Data Collection Schedule window, set the performance monitors to **Enabled**. You can set a performance monitor to **Enabled** only if the probe is set to **Enabled**.
 5. Select the interval for the performance monitors, and click **Save**. The interval represents the number of minutes over which samples of performance data are averaged. If you are creating performance monitors for multiple resources, the interval list displays the performance monitor intervals that are common to all resources. For example, you select two storage systems. `Storage_System_a` specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. `Storage_System_b` specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

Restrictions:

Note the following restrictions when you create performance monitors:

- a. If you specify a 1-minute interval for performance monitors, data collection samples might be inconsistent lengths. The inconsistent lengths occur because it might take longer than 1 minute to collect the data on some systems. The inconsistent lengths can cause the data in the historical line chart to be misleading. This problem can also occur for the 5-minute summary chart when a 1-minute interval was specified for data collection. Check the messages in the job log to determine the length of time that data collection takes for a resource. If the data collection takes longer than 1 minute, consider changing the interval to 5 minutes. You can use a 1-minute interval when you want to troubleshoot a resource.
- b. If you specify a 1-minute interval, the amount of data that is stored in the database repository increases significantly. IBM Spectrum Control stores only 7 days of sample data that is collected at 1-minute intervals.

Performance monitors are created for the resources that are selected, and that meet the criteria. For example, you select three storage systems, `Storage_System_a`, `Storage_System_b`, and `Storage_System_c`. You enable the performance monitors and set the interval field. When you click **Save**, performance monitors are created for `Storage_System_a` and `Storage_System_b`. `Storage_System_c` is not affected by the action because you cannot configure performance monitors for that storage system.

If a successful probe run is completed for a resource, the performance monitor runs according to the defined interval.

To check the progress of a performance monitor, you can complete the following actions:

- From the resource list page, right-click a resource row, and select **Data Collection > Open Performance Monitor Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.

- View the Performance Monitor Status column on the resource list page.

Verifying that a performance monitor is running for a resource

Verify that a performance monitor is running for a resource to ensure that IBM Spectrum Control is collecting performance metrics about the resource. Use performance monitors to collect metrics about the performance of storage systems and switches.

To verify that a performance monitor is running for a resource, use the resource list page in the web-based GUI.

1. From the menu bar, go to a resource list page for a resource type. For example, to verify that a performance monitor is running for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Locate the resource that you want to check.
3. View the Performance Monitor Status column. A value of Completed or Running indicates that the performance monitor is collecting performance metrics about the resource. The Last Successful Monitor column shows the most recent date and time when the performance monitor successfully collected data.
4. Optional: If the Performance Monitor Status column shows a status of Running with problems, Completed with warnings, Failed, or Canceled, you can view the logs to troubleshoot the error or warning messages. To view the performance monitor logs, right-click the resource row and click **Data Collection > Open Performance Monitor Logs**.
5. Optional: If the Performance Monitor Status column shows a status of Disabled, to enable the performance monitor for the resource, right-click the resource row and click **Data Collection > Schedule**. On the Data Collection Schedule page, select **Enabled**.

Tip: You can set the performance monitor to **Enabled** only if the probe is set to **Enabled**. If a probe is set to **Disabled**, the performance monitor cannot run.

Modifying performance monitors

To modify all of the performance monitors that are used by IBM Spectrum Control, use the Performance Monitors page in the web-based GUI. To modify performance monitors for specific resources, use the resource list and resource detail pages.

To modify performance monitors, you must have Administrator privileges.

You can modify the schedule for multiple performance monitors at the same time.

You can enable or disable performance monitors and modify the interval. The interval represents the number of minutes over which samples of performance data are averaged.

1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**. Detailed information about all performance monitors is shown in the **Performance Monitors** tab.
2. Locate the performance monitors that you want to modify.
3. Select the performance monitor rows and click **Actions > Schedule**.
4. On the Data Collection Schedule window, modify the schedule for the performance monitors.

If you select multiple performance monitors, the following rules are used to determine the field settings when you open the Data Collection Schedule window:

- If all of the performance monitors that you select are enabled, the **Enabled** value is displayed. If all of the performance monitors are disabled, the **Disabled** value is displayed. Otherwise, the state field for the performance monitors is blank.
- If the performance monitors have different interval values, the interval field is blank. If the performance monitors have the same interval value, the value is displayed.

Intervals that are common to all performance monitors that you select are displayed in the interval list. For example, you select performance monitors for two storage systems. `Storage_System_a` specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. `Storage_System_b` specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

5. Click **Save**.

The changes are applied to all the performance monitors that are selected. The performance monitors run according to the modified schedule. If you disabled a performance monitor, IBM Spectrum Control stops collecting performance metrics for the storage system or switch.

To check the progress of the performance monitors, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the most recent date and time when performance data was collected about a resource.
- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor job. Use this information to troubleshoot any errors that might occur when the job runs.

Configuring alerts for performance monitors

Configure the alert notifications that are generated if a performance monitor fails to run.



You can define alerts for a storage system or switch if you want to be notified when the following conditions occur:

- The status of its performance monitor is error or warning. An error status occurs when a performance monitor did not complete and no performance data was collected about a resource. A warning status occurs when a performance monitor completes, but might not have collected a complete set of performance data.
- A specified amount of time has passed since a performance monitor collected data about a resource. You can use this type of alert to be notified when up-to-date performance data is not being collected about a resource and its existing data might be stale.

To define an alert for the probe that collects data about a resource, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Switches	Network > Switches

2. Right-click a resource in the list and click **Edit Alert Definitions**.
3. To enable the **Performance Monitor Status** alert for an attribute, click the corresponding switch icon  and select a status to be alerted on. You can use this alert to be notified when a performance monitor is only collecting partial data about a resource or not collecting any data.
4. To enable the **Last Successful Monitor** alert for an attribute, click the corresponding switch icon  and specify how much time can pass before an alert is generated. You can use this alert to be notified when up-to-date performance data is not being collected about a resource and its existing data might be stale.
5. Click **Save**.

Starting and stopping performance monitors

To start or stop performance monitors for storage systems and switches, use the Performance Monitors page in the web-based GUI.

To start or stop performance monitors, you must have Administrator privileges.

You can start or stop performance monitors for multiple resources at the same time.

Restriction: A successful probe must be completed on the resource before a performance monitor can start. If a successful probe is not run, you cannot start the performance monitor for the resource.

1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**.
2. Click the **Performance Monitors** tab.
3. Locate the performance monitors that you want to start or stop.
4. Select the performance monitor rows and click **Actions > Start** or **Actions > Stop**.

Performance monitors are started or stopped for all the resources that meet the criteria. For example, you select two storage systems, `Storage_System_a` and `Storage_System_b`. Performance monitors are enabled for the storage systems and both resources are probed. A performance monitor is running for `Storage_System_a` but not for `Storage_System_b`. The performance monitor is started for `Storage_System_b` when you select **Start**. `Storage_System_a` is not affected by the action because a performance monitor is already running for the resource.

Starting or stopping a performance monitor has the following effects:

- Start** Starting a performance monitor starts the immediate collection of data from the monitored resource. The status of the performance monitor changes to Running.
- Stop** Stopping a performance monitor stops the collection of performance data from the resource. If a performance monitor is running, it stops. Performance data is not collected from the resource until you restart the performance monitor. The status of the performance monitor changes to Canceled.

To check the progress of the performance monitor, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the total number of threshold violations that were detected during the last sampled time of the performance monitor.

- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.

Viewing performance monitors

To view all of the performance monitors that IBM Spectrum Control uses to collect metrics for measuring the performance of resources, use the Performance Monitors page. To view the performance monitors for specific resources, use resource list and resource details pages.

Viewing all performance monitors:

To view all of the performance monitors that IBM Spectrum Control uses to collect metrics for measuring the performance of storage systems and switches, use the Performance Monitors page.

The Performance Monitors page displays one row for each storage system or switch that IBM Spectrum Control is monitoring. If a performance monitor is not scheduled for a resource, the status column is set to Disabled for the performance monitor row.

Before IBM Spectrum Control can collect performance metrics for a resource, the resource must fulfill the following conditions:

- A probe must be run for the resource.
 - For switches, the switch must be managed by a CIM agent.
1. In the menu bar, go to **Home > Performance Monitors**. On the **Performance Monitors** tab, the information about performance monitors is organized into columns. These columns include the status of performance monitors, the names of the monitored resources, and the most recent dates and times when data was successfully collected.
 2. Optional: View the status icons at the top of the page to view a summary of performance monitor statuses. This summary includes the number of performance monitors that have a Running, Running with problems, Failed, or Not Running status.

Viewing performance monitors for specific resources:

To view the performance monitors for specific storage systems and switches, use the resource detail pages in the web-based GUI. For example, use the Storage Systems detail page to view performance monitor information for a specific storage system.

If performance monitor information is not displayed on the resource detail page, this indicates that the resource does not support the collection of performance metrics by IBM Spectrum Control.

1. In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want to view the performance monitor for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Right-click a resource in the list, and select **View Details**. The resource detail page is displayed.

In the **General** section, the number next to **Data Collection** represents the number of data collection jobs that are associated with the resource. For example, **Data Collection (2)** indicates that you can view a performance monitor and a probe for the resource.

The icon next to the **Data Collection** link represents the most critical status of the data collection jobs that are associated with the resource. For example, if a job failed, the following icon is shown:

Data Collection (2) 

3. In the **General** section, click **Data Collection** to view information about the performance monitor for the switch or storage system. On the **Data Collection** pane, you can view information such as the status of the most recent run of a performance monitor, the performance monitor interval, and the most recent data and time that data was successfully collected.

Tip: You can use the columns on the resource list pages to view information about the performance monitors that are defined for storage systems or switches. For example, use the Storage Systems page to view performance monitor information for storage systems.





Viewing performance monitors logs:

Use performance monitor logs to view detailed information about the status, actions, and progress of a performance monitor. You can use this information to troubleshoot any errors that might occur when a performance monitor is running.

Performance monitor logs contain informational, warning, and error messages that are related to each action that is taken during the processing of a performance monitor.

1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**.
2. In the **Performance Monitors** tab, locate the performance monitor that you want to view.
3. Right-click the performance monitor row, and select **Open Logs**.

Restriction: If you select multiple rows, the **Open Logs** action is not available. The Logs page opens and shows log entries for the most recent performance monitor log. A new log file is created when you stop and restart a performance monitor.

4. Optional: On the Logs page, to view a previous performance monitor log, select a log from the **Select a log** list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only entries that have the following statuses:
 -  Only error entries
 -  Only warning entries
 -   Error and warning entries
6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a performance monitor. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the job run.
- The status for each performance monitor action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Collecting information about shares on storage systems

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Set the **saveNonRoot** parameter to **saveNonRoot=1** in the TPCD.config file to view the information that is collected about shares in the GUI.

1. Open the TPCD.config file.

For Windows operating systems:

By default, the TPCD.config file is in the *installation_dir\data\config* directory.

For Linux or UNIX operating systems:

By default, the TPCD.config file is in the *installation_dir/data/config/* directory.

Where *installation_dir* is where IBM Spectrum Control is installed.

2. In the [server] section of the TPCD.config file, change the value of the **saveNonRoot** parameter to 1 as shown in the following example:

```
saveNonRoot=1
```

3. Save the TPCD.config file.
4. Restart the Data server.

When you provision new shares, you can view information about the shares on the Provisioning task page in the GUI. On the Shares page, you can view more information about shares, for example, the servers that have access to the share.

Collecting information about the sizes of GPFS™ snapshots

By default, IBM Spectrum Control does not collect information about the size of GPFS snapshots. To view that information, you must enable IBM Spectrum Control to collect information about the size of GPFS snapshots.

Restriction: When you enable data to be collected about the size of GPFS snapshots, the collection might put a significant load on the GPFS cluster during the probe. This load might affect other applications that use the same cluster.

1. Run the following command to set the *Probe.GetGPFSSnapshotSize* property for the Device server to true:

```
tpctool setdscfg -user user_ID -pwd password -url localhost:9550  
-property Probe.GetGPFSSnapshotSize true
```

Where *user_ID* is an IBM Spectrum Control user ID and *password* is the password for the IBM Spectrum Control user ID.

2. Restart the IBM Spectrum Control Device server and the web server.

Alerting and event processing

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Tip: You can define alerts for individual resources or for the set of resources that are included in an application or general group.

Conditions that generate alerts are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource.

Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem V840, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.


The conditions that trigger alert notifications depend on the type of resource that you are monitoring. In general, the following types of conditions can trigger alerts:

- An attribute or configuration of a resource changed
- The capacity of a resource fell outside a specified range
- The performance of a resource fell outside a specified range
- The storage infrastructure was changed, such as a new or removed resource
- Data is not being collected for a resource

For a list of conditions that can trigger alerts, see “Triggering conditions for alerts” on page 370.

Determining which type of alert to use

To determine whether to define alerts for individual resources or for the set of resources that are included in an application or general group, follow these guidelines:

Which type of alerts to use?	Scenario	 Learn more
Resource alerts	<p>You want to receive alert notifications about changes for a specific resource, or its internal resources. For example, for a switch, you can alert on the attributes of the switch itself, and on the attributes of its blades, inter-switch connections, and ports.</p> <p>If you define an alert for a resource, for example, a performance alert for the ports on a storage system, the alert threshold value applies to all of the ports on the storage system. You cannot apply different alert thresholds to internal resources of the same type on a resource.</p>	<ul style="list-style-type: none"> • Resource alerts for attribute and capacity changes • Resource alerts for performance metrics • Custom resource alerts
Application alerts	<p>Use application alerts in the following scenarios:</p> <ul style="list-style-type: none"> • You want to receive alert notifications for all the resources of a certain type in an application. For example, if your application uses multiple servers, you can define the server alerts once for the application and the alerts apply to all the servers. If you later add more servers to the application, the existing application alerts apply to those servers also. • You want to apply different thresholds to internal resources of the same type on a storage system. For example, you have production applications and test applications that use volumes on a SAN Volume Controller. The production applications require response times of 6 milliseconds or less while the test applications can tolerate response times up to 30 milliseconds. You can use application alerts to set separate response time thresholds for volumes used by the different applications, depending on the needs of that application. 	<ul style="list-style-type: none"> • Application alerts for attribute and capacity changes • Application alerts for performance metrics
General group alerts	<p>Use general group alerts in the following scenarios:</p> <ul style="list-style-type: none"> • You want to receive alert notifications about changes for a subset of the resources of a particular type. For example, you can detect when the ports that are used for replication on your SAN Volume Controller have insufficient buffer-to-buffer credit. Alert notifications are not generated for ports that are not used for replication. • You want to receive alert notifications about changes for a group of resources that are logically related. You can group all the storage systems at a specific location or all the servers that use a particular operating system. For example, you can receive alert notifications when the used space of any of your Linux servers exceeds 80%. 	<ul style="list-style-type: none"> • General group alerts for attribute and capacity changes • General group alerts for performance metrics

Related tasks:

“Adding resources” on page 112

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

“Viewing capacity alerts and violations” on page 254

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

“Defining alerts for attribute and capacity changes” on page 151

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.

Related reference:







“Triggering conditions for alerts” on page 370

Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

Alert severities

IBM Spectrum Control determines the severity of alert conditions that it detects on monitored resources. Use the severity level to help determine the priority in which you resolve alerts.

Each severity is represented by an icon in the web-based GUI.

Icon	Status	Description
	Critical	An alert with a Critical severity represents a serious problem on a resource or on its internal resources. Resolve these problems as soon as possible. Review the condition that triggered an alert for more information about the problem.
	Critical - Acknowledged	An alert with a Critical severity was acknowledged. A Critical - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Warning	An alert with a Warning severity represents potential problems on a resource or on its internal resources. Resolve these problems after you fix any critical alerts. Review the condition that triggered an alert for more information about the problem.
	Warning - Acknowledged	An alert with a Warning severity was acknowledged. A Warning - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Informational	An alert with an Informational severity does not represent a problem, but is intended to provide information about actions related to a resource.
	Informational - Acknowledged	An alert with an Informational severity was acknowledged. An Informational - Acknowledged severity indicates that an alert was reviewed and can be ignored.

Defining alerts

Define how and when you are alerted to changes in your storage environment. You can define alerts for individual resources or for the set of resources that are included in an application or general group.

Defining alerts for resources

Define alerts for changes in the configuration, attributes, and performance of monitored resources.

Defining alerts for attribute and capacity changes:


You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.


1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. Right-click a resource in the list and click **Edit Alert Definitions**.
3. Click the type of resource that you want to alert on. You can alert on one or more attributes of the resource itself, its internal resources, or a combination of both. For example, for a switch, you can alert on the attributes of the switch itself, and on the attributes of its blades, inter-switch connections, and ports.
4. Click the category of the attributes that you want to alert on.

Category	Description
General	Attributes for the key properties of a resource, such as status, version changes, removals, discoveries, state, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.
Performance	Attributes for the performance of storage systems and switches, such as I/O rates, data rates, response times, and port frame rates.




5. To enable the alert for an attribute, click the corresponding switch icon .



Tip: To disable the alert for an attribute, click the switch icon  again. When you disable an alert, notifications are not generated for that attribute.

6. Specify the criteria for generating an alert. Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts. For example, for a capacity attribute such as Available Disk Space, you can specify that an alert is generated when the amount of available space on a resource's disks is less than or equal to 50 GiB. The operator (less than or equal to) + a specified amount of space (50) + the unit of measurement (GiB) is the criteria that determines if an alert is generated.

Tips:

- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
 - Some attributes can use the operators such as *is*, *is not*, and *contains*. For example, for the Firmware attribute on a DS8000 storage system, select the operator **contains**. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3. This alert definition might be useful if you want to be notified when the firmware for a storage system was reverted to a previous version for some reason.
7. Assign a severity to an alert. Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new pool is added to a storage system.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the amount of available space on a file system falls below a specified threshold.


8.  Specify how you are notified of alert violations and if you want a script to be run.
9.  Specify how often you are notified of alerts.



Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.

Option	Description
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

Restrictions:

- For some attributes, not all the suppression options are available. Specifically, you cannot suppress the alert notifications for the following attributes: New *resource*, Removed or Deleted *resource*, Last Successful Probe, and Last Successful Monitor.
- For attributes that use the **changes** operator (for example, the Firmware attribute), only the **Do not suppress alerts** and **Only generate alerts every** suppression options are available.

10. Optional:  Duplicate an alert. Use this action when you want to define another alert for the same attribute but with different criteria and settings. Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.
For example, for a capacity attribute such as Available Disk Space, you might want to define the following alerts:
 - Define a warning alert  to be generated when the amount of available space on a resource's disks is less than or equal to 50 GiB.
 - Duplicate the alert, but this time, specify a critical severity  when the amount of available space on a resource's disks is less than or equal to 10 GiB.
- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Available Disk Space attribute, you can configure the notification settings  so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

11. Click **Save**.

Conditions that generate alerts are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource.

You can view information about alert violations at the following locations in the GUI:

Table 31. Locations where you can view alert violations for attribute and capacity alerts.

Task	Location in GUI
View all alerts	In the menu bar, select Home > Alerts .
View all the alerts that were generated for storage systems	<ol style="list-style-type: none"> 1. In the menu bar, select one of the following options: <ul style="list-style-type: none"> • Storage > Block Storage Systems • Storage > File Storage Systems • Storage > Object Storage Systems 2. Click the Alerts tab.
View the alerts for a specific storage system	<ol style="list-style-type: none"> 1. In the menu bar, choose one of the following options: <ul style="list-style-type: none"> • Storage > Block Storage Systems • Storage > File Storage Systems 2. Double-click a storage system. 3. Click Alerts in the General section.
View all the alerts that were generated for switches	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Switches. 2. Click the Alerts tab.
View the alerts for a specific switch	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Switches. 2. Double-click a switch. 3. Click Alerts in the General section.
View all the alerts that were generated for fabrics	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Fabrics. 2. Click the Alerts tab.
View the alerts for a specific fabric	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Fabrics. 2. Double-click a fabric. 3. Click Alerts in the General section.
View all the alerts that were generated for hypervisors	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Hypervisors. 2. Click the Alerts tab.
View the alerts for a specific hypervisor	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Hypervisors. 2. Double-click a hypervisor. 3. Click Alerts in the General section.
View all the alerts that were generated for servers	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Servers. 2. Click the Alerts tab.
View the alerts for a specific server	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Servers. 2. Double-click a server. 3. Click Alerts in the General section.

Related concepts:

“Collecting data” on page 133

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

“Alerting and event processing” on page 148

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Related tasks:

“Viewing capacity alerts and violations” on page 254

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

Defining alerts for performance metrics:


You can define alerts that are triggered when the performance of a storage system or switch falls outside a specified threshold.


To define a performance alert, select a metric that you want to measure and specify a threshold value. When the performance of that resource falls outside the threshold, an alert is generated. For example, you can define an alert that notifies you when the back-end response times for managed disks on a SAN Volume Controller exceed 35 milliseconds per read operation. The Overall Back-end Response Time is a metric that measures the average number of milliseconds that it takes to service each read operation on a managed disk.

Tip: You cannot apply different thresholds to internal resources of the same type. For example, a storage system has two ports, Port A and Port B. You cannot set different threshold values for each port on the storage system. When you define a performance alert for ports, its threshold value is applied to all of the ports on the storage system.

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Switches	Network > Switches

2. Right-click a resource in the list and click **Edit Alert Definitions**.
3. Click the type of resource that you want to alert on. You can alert on the performance of the resource itself, its internal resources, or a combination of both. For example, for a switch, you can alert on the performance of the switch itself, and on the performance of its inter-switch connections and ports.
4. Click **Performance**.
5. Click **Add Metrics**.
6. Select one or more metrics to alert on and click **OK**. Metrics are available for measuring and alerting on the performance of volumes, ports, and disks. The type of storage system or switch determines the metrics that you can select.
7. To enable the alert for a performance metric, click the corresponding switch icon .

Tip: To disable the alert for a performance metric, click the switch icon  again. When you disable an alert, notifications are not generated for that metric.




8. Specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a. Select an operator. An operator determines if an alert is triggered when the performance of a resource is greater than or equal to or less than or equal to the specified threshold value.
 - b. Enter a threshold value. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.




Tips for threshold values:

- IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.





However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.

- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
9. Assign a severity to an alert. Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%.

10.  Evaluate the current and historical performance of a resource to help determine the threshold value for an alert. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
 - a. To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
 - b. To filter any violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
11. Click **Done**.
12.  Specify how you are notified of alert violations and if you want a script to be run.
13.  Specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

14. Optional:  Duplicate an alert. Use this action when you want to define another alert for the same metric but with different criteria and settings. Duplicating alerts can be helpful in the following situations:
 - When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.
For example, for the CRC Error Rate metric for ports, you might want to define the following alerts:
 - Define a warning alert  to be generated when the number of frames per second that are received with cyclic redundancy check (CRC) errors is greater than or equal to .01 counts per second.
 - Duplicate the alert, but this time, specify a critical severity  when the CRC error rate is greater than or equal to .033 counts per second.
 - When you want to send alert notifications to different people based on the severity of an alert.
In the previous example for the CRC Error Rate metric, you can configure the notification settings  so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

15. Click **Save**.

To determine whether alerts are generated, the performance metrics that are collected for resources are compared with the threshold values that you defined for those alerts. You can view information about alert violations at the following locations in the GUI:

Table 32. Locations where you can view alert violations for performance metrics.

Task	Location in GUI
View all the alerts that were generated for storage systems	<ol style="list-style-type: none">1. In the menu bar, choose one of the following options:<ul style="list-style-type: none">• Storage > Block Storage Systems• Storage > File Storage Systems2. Click the Alerts tab.
View all the alerts that were generated for switches	<ol style="list-style-type: none">1. In the menu bar, select Network > Switches.2. Click the Alerts tab.
View the alerts for a specific storage system	<ol style="list-style-type: none">1. In the menu bar, choose one of the following options:<ul style="list-style-type: none">• Storage > Block Storage Systems• Storage > File Storage Systems2. Double-click a storage system.3. Click Alerts in the General section.
View the alerts for a specific switch	<ol style="list-style-type: none">1. In the menu bar, select Network > Switches.2. Double-click a switch.3. Click Alerts in the General section.
View all performance alerts	<ol style="list-style-type: none">1. In the menu bar, select Home > Alerts.2. Click the performance icon on the summary chart.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a performance alert for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.

Defining custom alerts:

You can define alerts that are triggered when two or more changes occur in the attributes, capacity, and performance of resources.

To define a custom alert, select the attributes, capacity, and performance metrics that you want to combine to trigger an alert and specify their conditions and threshold values. You can combine conditions for the resource and its internal resources into a custom alert. The alert is triggered when the conditions for the attributes and capacity of the resource are met, and the performance of the resource falls outside the threshold values.




For example, you can create a custom alert that notifies you when the overall response time for the volumes on a SAN Volume Controller system is worse than 20 milliseconds per operation and the system CPU utilization on the nodes on the system is greater than 70%. The Overall Response Time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a

volume. The System CPU Utilization is a metric that measures the average percentage of time that the processors on nodes are busy doing system I/O tasks.

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. Right-click a resource in the list and click **Edit Alert Definitions**.
3. Click **Custom**.
4. Click **Create Alert**.
5. Enter a name for the alert.
6. Assign a severity to the alert. Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.

7. To add the first triggering condition for the alert, select the resource itself or an internal resource that you want to alert on.
8. Select the category of the attribute that you want to alert on.

Category	Description
General	Attributes for the key properties of a resource, such as status, version changes, removals, discoveries, state, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.

Category	Description
Performance	Attributes for the performance of storage systems and switches, such as I/O rates, data rates, response times, and port frame rates.

9. To add a general or capacity attribute, specify the criteria for generating an alert. Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts. For example, for a capacity attribute such as Available Disk Space, you can specify that an alert is generated when the amount of available space on a resource's disks is less than or equal to 50 GiB. The operator (less than or equal to) + a specified amount of space (50) + the unit of measurement (GiB) is the criteria that determines if an alert is generated.

Tips:





- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
 - Some attributes can use the operators such as *is*, *is not*, and *contains*. For example, for the Firmware attribute on a DS8000 storage system, select the operator **contains**. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3. This alert definition might be useful if you want to be notified when the firmware for a storage system was reverted to a previous version for some reason.
10. To add a performance attribute, specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a. Select an operator. An operator determines if an alert is triggered when the performance of a resource is greater than or equal to or less than or equal to the specified threshold value.
 - b. Enter a threshold value. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.

However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.

- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.

11.  For performance attributes evaluate the current and historical performance of a resource to help determine the threshold value for an alert. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
 - a. To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
 - b. To filter any violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
12.  Duplicate an alert. Use this action to add a second and subsequent condition to the alert.
13. Repeat steps 7 - 12 to add more conditions to the alert.
14.  Specify how you are notified of alert violations and if you want a script to be run.
15.  Specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

16. Click **Save**.

You can view information about custom alert violations at the following locations in the GUI:

Table 33. Locations where you can view alert violations for custom alerts

Task	Location in GUI
View all alerts	In the menu bar, select Home > Alerts .

Table 33. Locations where you can view alert violations for custom alerts (continued)

Task	Location in GUI
View all the alerts that were generated for storage systems	<ol style="list-style-type: none"> 1. In the menu bar, select one of the following options: <ul style="list-style-type: none"> • Storage > Block Storage Systems • Storage > File Storage Systems • Storage > Object Storage Systems 2. Click the Alerts tab.
View the alerts for a specific storage system	<ol style="list-style-type: none"> 1. In the menu bar, choose one of the following options: <ul style="list-style-type: none"> • Storage > Block Storage Systems • Storage > File Storage Systems 2. Double-click a storage system. 3. Click Alerts in the General section.
View all the alerts that were generated for switches	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Switches. 2. Click the Alerts tab.
View the alerts for a specific switch	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Switches. 2. Double-click a switch. 3. Click Alerts in the General section.
View all the alerts that were generated for fabrics	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Fabrics. 2. Click the Alerts tab.
View the alerts for a specific fabric	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Fabrics. 2. Double-click a fabric. 3. Click Alerts in the General section.
View all the alerts that were generated for hypervisors	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Hypervisors. 2. Click the Alerts tab.
View the alerts for a specific hypervisor	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Hypervisors. 2. Double-click a hypervisor. 3. Click Alerts in the General section.
View all the alerts that were generated for servers	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Servers. 2. Click the Alerts tab.
View the alerts for a specific server	<ol style="list-style-type: none"> 1. In the menu bar, select Servers > Servers. 2. Double-click a server. 3. Click Alerts in the General section.

Scenarios for custom alerts:

Use custom alerts to combine attributes and conditions from multiple resources into a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred in your storage.

The storage requirements of your environment determine the custom alerts that you create. For example, your storage systems might run critical production applications on tier 1 storage. In this case, you don't want the performance of the tier 1 storage to fall below a certain threshold. To be notified when that situation occurs, you can create a custom alert that checks if the overall response time of tier 1 storage is too high.

Tip: You can define alerts for storage systems, servers, hypervisors, fabrics, and switches and their internal resources. However, the resources in an alert must be along the same data path. For example, if you create an alert for a storage system volume, resources along the same data path include the pool that volume is part of, the switch ports that are used to access that volume, and the host that maps the volume.

Use the following example scenarios as guides to help you create custom alerts for your environment.

Receive alerts when the response time of storage on a specific tier is too high

Your storage systems run both critical production applications and noncritical test applications. The production applications use tier 1 storage, while the test applications use storage on tiers 2 and 3.

To ensure consistent, top performance for tier 1 storage, you want to be notified when its response time is higher than 6ms/op so that you can resolve the bottleneck. However, to avoid too many alerts, you do not want to receive notifications when the response time of tier 2 or 3 storage exceeds 6ms/op.

Solution

Define a custom alert that checks if the volumes used by an application are in Tier 1 pools and if their Overall Response Time is higher than 6 ms/op.

For a storage system, set up a custom alert with the following attributes and conditions:

High response time for tier 1 storz

Pools > General > Tier is 1

Volumes > Performance > Overall Response Time >= 6 ms/op

Attribute	Condition
Overall Response time (Storage System > Volumes > Performance)	Greater than or equal to (>=) 6 ms/op
Tier (Storage System > Pools > General)	Tier 1

Receive alerts when the response time of volumes is too high, but do not generate these alerts during the time when batch and backup jobs are running

To ensure the consistent, top performance of your volumes, you want to be notified when their response times are becoming too high. However, to avoid too many alerts, you do not want to receive notifications when batch and backup jobs are running on your storage. You understand that these jobs can cause an expected spike in response times and do not require action on your part.

Solution

Define a custom alert that checks if the Read Response Time of volumes exceeds an amount that is more than expected in your environment and the Read Transfer Size is less than 256 KiB/op. Typically, read transfer sizes greater than 256KiB/op indicate that batch or backup jobs are running in the background.

For a storage system, set up a custom alert with the following attributes and conditions:

☒ Slow performance of volumes

Volumes

Performance

Read Response Time

>=

20

ms/op

Volumes

Performance

Read Transfer Size

<=

256

KiB/op

Attribute	Condition
Read Response Time (Storage System > Volumes > Performance)	Greater than or equal to (>=) 20 ms/op
Read Transfer Size (Storage System > Volumes > Performance)	Greater than or equal to (<=) 256 KiB/op

Receive alerts if a port is being used for both inter-node communication and host I/O exchanges

You want to avoid potential bottlenecks by ensuring that storage system ports aren't being used for both inter-node communication in the local cluster and for I/O exchanges to host computers. You can also use this custom alert to check for adherence to best practices that are related to configuring ports for nodes with 8 or more ports. It does not apply to nodes that contain only 4 ports.

Solution

Define a custom alert that checks if the I/O rate for ports indicates exchanges between local nodes and hosts. For a storage system, set up a custom alert with the following attributes and conditions:

☒ Dual use ports

Ports

Performance

Total Port to Host I/O Rate

>=

0.01

ops/s

Ports

Performance

Total Port to Local Node I/O Rate

>=

0.01

ops/s

Attribute	Condition
Total Port-to-Host I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s
Total Port-to-Local Node I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s

Tip: Optionally, you can define other custom alerts to be notified of this situation, depending on your storage requirements. For example:

Attribute	Condition
Total Port-to-Disk I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s
Total Port-to-Local Node I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s
Total Port-to-Remote I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s

Receive alerts for link resets that are not associated with link initialization

You want to identify link resets that are generated in response to hardware failures or link congestion. Link Resets generated by link initialization are ignored.

Solution

Define a custom alert that checks if link resets occur and if those resets are not associated with a link initialization. For a switch, set up a custom alert with the following attributes and conditions:

The screenshot shows a configuration window for an alert titled "Link reset not initialized". It features a green toggle switch on the left and a red 'X' icon on the right. Below the title bar, there are three rows of configuration fields. Each row starts with a "Ports" dropdown menu, followed by a "Performance" dropdown menu, then a specific attribute name, a comparison operator, a numerical value, and a unit "cnt/s". The first row is for "Link Reset Received Rate" with a greater-than-or-equal-to operator and a value of 0.01. The second row is for "Loss of Sync Rate" with a less-than-or-equal-to operator and a value of 0. The third row is for "Loss of Signal Rate" with a less-than-or-equal-to operator and a value of 0.

Attribute	Condition
Link Reset Received Rate (Storage System > Ports > Performance) OR Link Transmitted Received Rate (Storage System > Ports > Performance)	Greater than or equal to (\geq) .01 cnt/s
Sync Loss (Storage System > Ports > Performance)	Less than or equal to (\leq) 0 cnt/s
Signal Loss (Storage System > Ports > Performance)	Less than or equal to (\leq) 0 cnt/s

Receive alerts for invalid word transmissions that are not associated with link initialization

You want to identify invalid transmission words that are generated because of poor link quality. Poor or marginal link quality can be caused by a bad SFP, HBA, or cable. Invalid transmission words that are generated by link initialization are ignored.

Solution

Define a custom alert that checks if invalid word transmissions occur and if those resets are not associated with a link initialization. For a switch, set up a custom alert with the following attributes and conditions:

The screenshot shows a configuration window for an alert titled "Invalid word transmission not initiz". It features a green toggle switch on the left and a red 'X' icon on the right. Below the title bar, there are three rows of configuration fields. Each row starts with a "Ports" dropdown menu, followed by a "Performance" dropdown menu, then a specific attribute name, a comparison operator, a numerical value, and a unit "cnt/s". The first row is for "Invalid Transmission Word Rate" with a greater-than-or-equal-to operator and a value of 0.01. The second row is for "Loss of Sync Rate" with a less-than-or-equal-to operator and a value of 0. The third row is for "Loss of Signal Rate" with a less-than-or-equal-to operator and a value of 0.

Attribute	Condition
Invalid Transmission Word Rate (Storage System > Ports > Performance)	Greater than or equal to (\geq) .01 cnt/s
Sync Loss (Storage System > Ports > Performance)	Less than or equal to (\leq) 0 cnt/s

Attribute	Condition
Signal Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s

Suppressing alerts:


Use suppression options to determine how often you receive alerts.

To specify a suppression option for an alert definition, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. Right-click a resource in the list and click **Edit Alert Definitions**.

3. To suppress notifications for an alert, click the suppression icon . Specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

4. Click **Done**.
5. Click **Save**.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a performance alert definition for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.

Defining alerts for applications

Define alerts for changes in the configuration, attributes, and performance of the servers, volumes, filesets, and shares in your application.

You can quickly identify and address the critical conditions that are detected on the servers, volumes, filesets, and shares in your application by using alerts.

Application alerts can be helpful in the following situations:

- When you want to apply different thresholds to internal resources of the same type on a storage system.

For example, you have production applications and test applications that use volumes on a SAN Volume Controller. The production applications require response times of 6 milliseconds or less while the test applications can tolerate response times up to 30 milliseconds. You can use application alerts to set separate response time thresholds for volumes used by the different applications, depending on the needs of that application.

- When you want to quickly define alerts for multiple resources of the same type. You can define alerts once for the application and the alerts apply to all the resources of that type in the application.

For example, if your application uses multiple servers, you can define the server alerts once for the application and the alerts apply to all the servers. If you later add more servers to the application, the existing application alerts apply to those servers also.

Tips for working with alerts in application hierarchies:

- When you define an application alert for a resource such as a volume, the alert applies to all the resources of that type that belong to that application and all child applications.
- If you define an application alert for a resource such as a server, and the server also belongs to child applications, the alert is generated once at the parent application level. Separate alerts are not generated for each of the child applications that contain the server.
- If a child application has the same alert as a parent application but with different conditions, separate alerts are generated for the child application for the different alert conditions.

Related concepts:

“Applications” on page 261

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

Defining application alerts for attribute and capacity changes:

You can define alerts that are triggered when the attributes or capacity of the servers, volumes, filesets, or shares in an application change.


Attributes represent the key properties and configuration of a resource, such as status, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used file system space, total disk space, and volume used space.

To define an application alert for attribute and capacity changes, complete these steps:

1. In the menu bar, select **Groups > Applications**.
2. Right-click an application in the list and click **Edit Alert Definitions**.
3. Click the type of resource that you want to set an alert for. You can set an alert for one or more of the attributes of servers, volumes, filesets, or shares.
4. Click the category of the attributes that you want to set an alert for.

Category	Description
General	Attributes for the key properties of a resource, such as status, removals, discoveries, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as used space, available file system space, total disk space, and volume used space.
Performance	Attributes for the performance of storage system volumes, such as I/O rates, data rates, and response times.

5. Enable the alert for an attribute by clicking the corresponding switch icon  .



Tip: To disable the alert for an attribute, click the switch icon  again. When you disable an alert, notifications are not generated for that attribute.


6. Specify the criteria for generating an alert. Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts.


For example, for a capacity attribute such as Available Disk Space, you can specify that an alert is generated when the amount of available space on a resource's disks is less than or equal to 50 GiB. The operator (less than or equal to) + a specified amount of space (50) + the unit of measurement (GiB) is the criteria that determines if an alert is generated.


Tip: Not all attributes require criteria for generating alerts, for example, the New Fileset attribute for an IBM Spectrum Scale storage system. The category and type of an attribute determine whether you can specify criteria and the options that you can select.

7. Assign a severity to an alert to help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.




Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new fileset is added to a storage system.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.


Option	Description
	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the amount of available space on a file system falls below a specified threshold.

8. Click the icon  to specify how you are notified of alert violations and whether you want a script to be run.

9. Click the icon  to specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive only one notification for an alert violation, even if the condition is violated multiple times. This option is the default.
Only generate alerts every	Receive one notification when an alert violation is initially detected. The alert is suppressed and no notifications are sent until a specified time was exceeded and the alert is triggered again.
Do not alert until the condition is violated for more than	Receive notifications for an alert only after a certain time limit was exceeded and if the violation was not cleared.

10. Optional: Click the icon  to duplicate an alert. Use this action when you want to define another alert for the same attribute but with different criteria and settings. Duplicating alerts can be helpful in the following situations:
- When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.
For example, for a capacity attribute such as Available Disk Space, you might want to define the following alerts:
 - Define a warning alert  to be generated when the amount of available space on a resource's disks is less than or equal to 50 GiB.
 - Duplicate the alert, but this time, specify a critical severity  when the amount of available space on a resource's disks is less than or equal to 10 GiB.
 - When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Available Disk Space attribute, you can configure the notification settings  so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

11. Click **Save**.

Conditions that generate alerts are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource. For IBM Spectrum Scale,

status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource.

You can view information about application alert violations at the following locations in the GUI:

Table 34. Viewing application alert violations for attribute and capacity alerts.


Task	Location in GUI
View all application alerts	In the menu bar, select Home > Alerts .
View alerts for a specific application	<ol style="list-style-type: none"> 1. Go to Groups > Applications. 2. Right-click an application and select View Details. 3. Click Alerts in the General section.
View alerts for a specific group	<ol style="list-style-type: none"> 1. Go to Groups > General Groups. 2. Right-click an application and select View Details. 3. Click Alerts in the General section.


Defining application alerts for performance metrics:

You can define alerts that are triggered when the performance of the volumes that belong to an application fall outside a specified threshold.

To define an application alert for performance metrics, select the metric that you want to measure for the volumes in an application and specify a threshold value. When the performance of a volume falls outside the threshold, an alert is generated. For example, by using the Read Response Time metric you can define an alert that notifies you when the response time for a volume exceeds 20 milliseconds per read operation.

To define an application alert for performance metrics, complete these steps:




1. In the menu bar, select **Groups > Applications**.
2. Right-click an application in the list and click **Edit Alert Definitions**.
3. Click **Volumes** and then click **Performance**.
4. Click **Add Metrics**.
5. Select one or more metrics to alert on and click **OK**.
6. Enable the alert for a performance metric by clicking the corresponding switch icon .




Tip: To disable the alert for a performance metric, click the switch icon  again. When you disable an alert, notifications are not generated for that metric.

7. Specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a. Select an operator. An operator determines whether an alert is triggered when the performance of a resource is greater than or equal to or less than or equal to the specified threshold value.
 - b. Enter a threshold value. For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:


- Thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. To help determine threshold values for volumes, collect performance data over time to establish a baseline of the normal and expected performance.
 - For some metrics, lower values might indicate that a volume has performance problems and higher values might indicate that a volume is performing efficiently. For example, low cache-hit percentages can increase response time and cause performance problems. Therefore, you might define an alert that is triggered when the cache-hit percentage falls below a threshold value.
8. Assign a severity to an alert to help you more quickly identify and address the critical performance conditions that are detected on volumes. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, you might assign a warning severity to alerts that notify you when the average Read Response Time for volumes is greater than 10 ms/op but less than 20 ms/op.
	Assign this severity to alerts that are critical and must be resolved. For example, you might assign a critical severity to alerts that notify you when the average Read Response Time for volumes is greater or equal to 20 ms/op.

9. Click the icon  to evaluate the current and historical performance of a volume to help determine the threshold value for an alert. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts. A red horizontal line represents the threshold value for critical alerts.
- a. To customize the chart, click **Top 10** or **Bottom 10** to show volumes according to their performance, click a time period, and change the start and end dates for the data that is displayed.
 - b. To filter any violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
10. Click **Done**.
11. Click the icon  to specify how you are notified of alert violations and if you want a script to be run.
12. Click the icon  to specify how often you are notified of alerts.



Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.

Option	Description
Only alert once until problem clears	Receive only one notification for an alert violation, even if the condition is violated multiple times. This option is the default.
Only generate alerts every	Receive one notification when an alert violation is initially detected. The alert is suppressed and no notifications are sent until a specified time was exceeded and the alert is triggered again.
Do not alert until the condition is violated for more than	Receive notifications for an alert only after a certain time limit was exceeded and if the violation was not cleared.

13. Optional: Click the icon  to duplicate an alert. Use this action when you want to define another alert for the same metric but with different criteria and settings. Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.

For example, for the Read Response Time metric for volumes, you might want to define the following alerts:

- Define a warning alert  to be generated when the average time that it takes to service each read operation is greater than or equal to 10 milliseconds.
- Duplicate the alert, but this time, specify a critical severity  when the average time is greater than or equal to 20 milliseconds.
- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Read Response Time metric, you can

configure the notification settings  so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

14. Click **Save**.

To determine whether alerts are generated, the performance metrics that are collected for volumes are compared with the threshold values that you defined for those alerts.

You can view information about application alert violations at the following locations in the GUI:

Table 35. Viewing application alert violations for performance metrics

Task	Location in GUI
View all performance alerts for applications	<ol style="list-style-type: none"> 1. In the menu bar, select Home > Alerts. 2. Click the performance icon on the summary chart. 3. To view only applications alerts, in the filter fields select Alert Category and Application Group.
View alerts for a specific application	<ol style="list-style-type: none"> 1. Go to Groups > Applications. 2. Right-click an application and select View Details. 3. Click Alerts in the General section.

Table 35. Viewing application alert violations for performance metrics (continued)

Task	Location in GUI
View all application alerts that were generated for storage system volumes	<ol style="list-style-type: none"> 1. In the menu bar, select Storage > Block Storage Systems. 2. Click the Alerts tab.
View the application alerts for volumes on a specific storage system	<ol style="list-style-type: none"> 1. In the menu bar, select Storage > Block Storage Systems. 2. Double-click a storage system. 3. Click Alerts in the General section.

Tip: If a performance monitor is already collecting data about a volume when you add, modify, or remove a performance alert for that volume, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.



Defining custom alerts for applications:


Define a custom alert to combine multiple conditions from multiple resources for an application in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the servers, volumes, filesets, and shares in an application.

For example, you can create a custom alert that notifies you when the used space on both volumes and filesets is greater than a specified amount.

Tip: In a custom alert, the resource types that you specify conditions for must be associated with the same type of top-level resource. For example, if you include conditions for storage system volumes, you can also include additional conditions for filesets and shares, but you cannot include conditions for servers.

1. In the menu bar, select **Groups > Applications**.
2. Right-click an application in the list and click **Edit Alert Definitions**.
3. Click **Custom**.
4. Click **Create Alert**.
5. Enter a name for the alert.
6. Assign a severity to the alert. Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Used Space for a server is greater than or equal to 75%.

Option	Description
	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Used Space for a server is greater than or equal to 85%. The default severity for custom alerts is critical.

7. Select the type of resource that you want to alert on. For example, if one of the conditions that you want to check for in the custom alert is Used Space for volumes, select **Volumes**.
8. Select the category of the attribute that you want to alert on. For example, for Used Space for volumes, select **Capacity**.

Category	Description
General	Attributes for the key properties of a resource, such as status, version changes, removals, discoveries, and states.
Capacity	Attributes for capacity statistics of a server, volume, or fileset, such as available and used space, total disk space, assigned SAN space, and capacity.
Performance	Attributes for the performance of volumes, such as I/O rates, data rates, response times, and cache hits.





9. For general and capacity attributes, specify the criteria for generating an alert. Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts. For example, for a capacity attribute such as Used Space, you can specify that an alert is generated when the amount of used space on a resource is more than or equal to 75%. The operator (less than or equal to) + a specified amount of space (75) + the unit of measurement (%) is the criteria that determines if an alert is generated.

Tips:

- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
 - Some attributes can use the operators such as *is*, *is not*, and *changes*. For example, for the Used Inodes attribute for a fileset, select the operator **changes** to be notified if the current number of used inodes for a fileset changes.
10. For performance attributes, specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a. Select an operator. An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.
 - b. Enter a threshold value. For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- For metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.

11.  For performance attributes evaluate the current and historical performance of a resource to help determine the threshold value for an alert. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
 - a. To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
 - b. To filter any violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
12.  Add another condition to the custom alert and repeat steps 7 - 11. If you do not want to add another condition, continue to the next step.
13.  Specify how you are notified of alert violations and if you want a script to be run.
14.  Specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

15. Click **Save**.

You can view information about custom alert violations for applications at the following locations:

Table 36. Viewing information about custom alert violations for applications.

Task	Location in GUI
View all application alerts	In the menu bar, select Home > Alerts . Tip: To view only applications alerts on the Alerts page, in the filter fields select Alert Category and Application Group .
View alerts for a specific application	1. Go to Groups > Applications . 2. Right-click an application and select View Details . 3. Click Alerts in the General section.
View alerts for a specific group	1. Go to Groups > General Groups . 2. Right-click an application and select View Details . 3. Click Alerts in the General section.

Defining alerts for general groups

Define alerts for changes in the configuration, attributes, and performance of the resources in your general groups.

Before you can define general group alerts for a resource type, you must first add resources of that type to the group or one of its subgroups. For example, if you want to set an alert for the attributes of a switch, you must first add one or more switches to your general group or one of its subgroups.

Tips for working with alerts in general group hierarchies:

- When you edit the alert definitions for a general group, you can edit the alert definitions for its subgroups at the same time. The resources for all subgroups of the general group are shown when you edit the alert definitions for a parent general group.
- If you define a general group alert for a resource such as a server, and the server also belongs to subgroups, the alert is generated once at the parent group level. Separate alerts are not generated for each of the subgroups that contain the server.
- If a subgroup has the same alert as a parent group but with different conditions, separate alerts are generated for the subgroup for the different alert conditions.

Related reference:

“General groups” on page 279

View and administer the general groups that are used to alert on logically related storage resources. You can define alerts for a general group to notify you about changes in the configuration, attributes, and performance of the resources in the group.

Defining general group alerts for attribute and capacity changes:

You can define alerts that are triggered when the attributes or capacity of the resources that belong to a general group change.


Attributes represent the key properties and configuration of a resource, such as status, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used file system space, total disk space, and volume used space.

To define a general group alert for attribute and capacity changes, complete these steps:

1. In the menu bar, select **Groups > General Groups**.
2. Right-click a group in the list and click **Edit Alert Definitions**.
3. Click the type of resource that you want to set an alert for. You can set an alert for one or more of the attributes of the resource itself or its internal resources. For example, for a switch, you can set an alert for the attributes of the switch itself and for the attributes of the blades, inter-switch connections, and ports.
4. Click the category of the attributes that you want to alert for.

Category	Description
General	Attributes for the key properties of a resource, such as status, removals, discoveries, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.
Performance	Attributes for the performance of storage systems and switches, such as I/O rates, data rates, response times, and port frame rates.

5. Enable the alert for an attribute by clicking the corresponding switch icon .




Tip: To disable the alert for an attribute, click the switch icon  again. When you disable an alert, notifications are not generated for that attribute.


6. Specify the criteria for generating an alert. Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts.


For example, for a capacity attribute such as Available Disk Space, you can specify that an alert is generated when the amount of available space on a resource's disks is less than or equal to 50 GiB. The operator (less than or equal to) + a specified amount of space (50) + the unit of measurement (GiB) is the criteria that determines whether an alert is generated.

Tips:


- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
 - Some attributes can use the operators such as *is*, *is not*, and *contains*. For example, for the Firmware attribute on a DS8000 storage system, select the operator **contains**. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3.
7. Assign a severity to an alert to help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new pool is added to a storage system.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the amount of available space on a file system falls below a specified threshold.

8. Click the icon  to specify how you are notified of alert violations and whether you want a script to be run.



9. Click the icon  to specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive only one notification for an alert violation, even if the condition is violated multiple times. This option is the default.
Only generate alerts every	Receive one notification when an alert violation is initially detected. The alert is suppressed and no notifications are sent until a specified time was exceeded and the alert is triggered again.
Do not alert until the condition is violated for more than	Receive notifications for an alert only after a certain time limit was exceeded and if the violation was not cleared.


10. Optional: Click the icon  to duplicate an alert. Use this action when you want to define another alert for the same attribute but with different criteria and settings. Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.

For example, for a capacity attribute such as Available Disk Space, you might want to define the following alerts:

- Define a warning alert  to be generated when the amount of available space on a resource's disks is less than or equal to 50 GiB.
- Duplicate the alert, but this time, specify a critical severity  when the amount of available space on a resource's disks is less than or equal to 10 GiB.

- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Available Disk Space attribute, you can configure the notification settings  so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

11. Click **Save**.

Conditions that generate alerts are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource.

You can view information about alert violations for general groups at the following locations in the GUI:

Table 37. Viewing general group alert violations for attribute and capacity alerts


Task	Location in GUI
View all general group alerts	In the menu bar, select Home > Alerts .
View alerts for a specific general group	<ol style="list-style-type: none"> 1. Go to Groups > General Groups. 2. Right-click a general group and select View Details. 3. Click Alerts in the General section.


Defining general group alerts for performance metrics:

You can define alerts that are triggered when the performance of the storage systems or switches that belong to a general group fall outside a specified threshold.

To define a general group alert for performance metrics, select the metric that you want to measure and specify a threshold value. When the performance of that resource falls outside the threshold, an alert is generated. For example, you can define an alert that notifies you when the back-end response times for managed disks on a SAN Volume Controller exceed 35 milliseconds per read operation. The Overall Back-end Response Time is a metric that measures the average number of milliseconds that it takes to service each read operation on a managed disk.

To define a general group alert for performance metrics, complete these steps:

1. In the menu bar, select **Groups > General Groups**.
2. Right-click a general group in the list and click **Edit Alert Definitions**.
3. Click the type of resource that you want to alert on.
4. Click **Performance** and then click **Add Metrics**.
5. Select one or more metrics to alert on and click **OK**.
6. Enable the alert for a performance metric by clicking the corresponding switch icon .

Tip: To disable the alert for a performance metric, click the switch icon  again. When you disable an alert, notifications are not generated for that metric.




7. Specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a. Select an operator. An operator determines whether an alert is triggered when the performance of a resource is greater than or equal to or less than or equal to the specified threshold value.
 - b. Enter a threshold value. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.


Tips for threshold values:

- IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.

However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, no recommended values are provided. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance.


- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
8. Assign a severity to an alert to help you more quickly identify and address the critical performance conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.


Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%.

9. Click the icon  to evaluate the current and historical performance of a resource to help determine the threshold value for an alert. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts. A red horizontal line represents the threshold value for critical alerts.


- a. To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
- b. To filter any violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.

10. Click **Done**.

11.  Specify how you are notified of alert violations and if you want a script to be run.



12. Click the icon  to specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive only one notification for an alert violation, even if the condition is violated multiple times. This option is the default.
Only generate alerts every	Receive one notification when an alert violation is initially detected. The alert is suppressed and no notifications are sent until a specified time was exceeded and the alert is triggered again.
Do not alert until the condition is violated for more than	Receive notifications for an alert only after a certain time limit was exceeded and if the violation was not cleared.


13. Optional: Click the icon  to duplicate an alert. Use this action when you want to define another alert for the same metric but with different criteria and settings. Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.

For example, for the CRC Error Rate metric for ports, you might want to define the following alerts:

- Define a warning alert  to be generated when the number of frames per second that are received with cyclic redundancy check (CRC) errors is greater than or equal to 0.01 counts per second.
- Duplicate the alert, but this time, specify a critical severity  when the CRC error rate is greater than or equal to 0.033 counts per second.

- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the CRC Error Rate metric, you can configure the notification settings  so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

14. Click **Save**.

To determine whether alerts are generated, the performance metrics that are collected for resources are compared with the threshold values that you defined for those alerts. You can view information about general group alert violations at the following locations in the GUI:

Table 38. Viewing general group alert violations for performance metrics

Task	Location in GUI
View all performance alerts for general groups	<ol style="list-style-type: none"> 1. In the menu bar, select Home > Alerts. 2. Click the performance icon on the summary chart. 3. To view only general groups alerts, in the filter fields select Alert Category and General Group.
View alerts for a specific general group	<ol style="list-style-type: none"> 1. Go to Groups > General Groups. 2. Right-click a general group and select View Details. 3. Click Alerts in the General section.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a performance alert for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.


Defining custom alerts for general groups:



Define a custom alert to combine multiple conditions from multiple resources for a general group in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the resources in a group.

For example, you can create a custom alert that notifies you when the response times for volumes and pools in a group exceed a certain threshold and the performance monitors for storage systems are running without problems.

Tip: In a custom alert, the resource types that you specify conditions for must be associated with the same type of top-level resource. For example, if you include conditions for storage systems, you can also include additional conditions for pools and volumes, but you cannot include conditions for servers.

1. In the menu bar, select **Groups > General Groups**.
2. Right-click a group in the list and click **Edit Alert Definitions**.
3. Click **Custom**.
4. Click **Create Alert**.
5. Enter a name for the alert.
6. Assign a severity to the alert. Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

Option	Description
	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Used Space for a server is greater than or equal to 75%.
	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Used Space for a server is greater than or equal to 85%. The default severity for custom alerts is critical.

7. Select the type of resource that you want to alert on. For example, if one of the conditions that you want to check for in the custom alert is Overall Response Time for volumes, select **Volumes**.
8. Select the category of the attribute that you want to alert on. For example, for Overall Response Time for volumes, select **Performance**.

Category	Description
General	Attributes for the key properties of a resource, such as status, version changes, removals, discoveries, and states.
Capacity	Attributes for capacity statistics of resources, such as available and used space, total disk space, assigned SAN space, and capacity.
Performance	Attributes for resources, such as I/O rates, data rates, response times, and cache hits.

9. For general and capacity attributes, specify the criteria for generating an alert. Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts. For example, for a capacity attribute such as Used Space, you can specify that an alert is generated when the amount of used space on a resource is more than or equal to 75%. The operator (less than or equal to) + a specified amount of space (75) + the unit of measurement (%) is the criteria that determines if an alert is generated.





Tips:

- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
- Some attributes can use the operators such as *is*, *is not*, and *changes*. For example, for the Firmware attribute on a DS8000 storage system, select the operator **contains**. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3. This alert definition might be useful if you want to be notified when the firmware for a storage system was reverted to a previous version for some reason.

10. For performance attributes, specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a. Select an operator. An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.

- b. Enter a threshold value. For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- For metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
 - For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
11.  For performance attributes, evaluate the current and historical performance of a resource to help determine the threshold value. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
- a. To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
 - b. To filter any violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
12.  Add another condition to the custom alert and repeat steps 7 - 11. If you do not want to add another condition, continue to the next step.
13.  Specify how you are notified of alert violations and if you want a script to be run.
14.  Specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

15. Click **Save**.

You can view information about custom alert violations for general groups at the following locations:

Table 39. Viewing information about custom alert violations for general groups.

Task	Location in GUI
View all group alerts	In the menu bar, select Home > Alerts . Tip: To view only group alerts on the Alerts page, in the filter fields select Alert Category and General Group .
View alerts for a specific group	1. Go to Groups > General Groups . 2. Right-click a group and select View Details . 3. Click Alerts in the General section.

Viewing and modifying alert definitions

View and modify the attributes and conditions that trigger alerts for resources and the notification settings for those conditions.




Viewing alert definitions for a resource

View the attributes, triggering conditions, notification settings, suppression settings, and threshold values that are defined in alerts.

To view an alert definition, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:


Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics


2. Right-click a resource in the list and click **Edit Alert Definitions**.
3. Click a category to view the alert definitions for that category. For example, click **Performance** or **Capacity** to view the performance and capacity alerts that are defined for a resource. Click a resource category, such as **Volumes**, to view all the alerts that are defined for that resource. To view the alerts with multiple conditions that are defined for a resource, click the category, **Custom**.
4. Optional:  View how you are notified of alert violations and if a script is specified.
5. Optional:  Evaluate the current and historical performance of a resource. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
6. Optional:  View how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

Disabling or enabling alert definitions

Disable an alert definition so that IBM Spectrum Control does not generate notifications for its specified condition. You can re-enable alert definitions at any time.

1. In the menu bar, select the resource that contains the alert definition to disable. For example, to disable the alert definition for a block storage system, select **Storage > Block Storage Systems**.
2. Right-click the resource that contains the alert definition and select **Edit Alert Definitions**.
3. Locate the alert definition that you want to disable.
4. Click the switch icon next to the definition so that it appears like this: .

Tip: To enable the definition again, click the icon so that it appears like this: .

5. Click **Save**.

Modifying alert definitions






Modify an alert definition to change its triggering conditions, notification settings, threshold values, and suppression settings.

To edit an alert definition for a resource, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. Right-click the resource that contains the alert definition and select **Edit Alert Definitions**.
3. Locate the alert definition that you want to edit.

4. To enable or disable an alert definition, click the corresponding switch icon:
 -  The alert is enabled.
 -  The alert is disabled.
5. Edit the criteria for generating an alert. Criteria includes operators such as greater than or equal to, less than or equal to, is, is not, and contains, and it can include values for storage space, performance metrics, and time measurements.
6. Change the severity for the alert. Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.
7.  For performance alerts, edit the threshold value by evaluating the current and historical performance of a resource. In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
8.  Change how you are notified of alerts and if you want a script to be run.
9.  Specify how often you are notified of alerts.

Option	Description
Do not suppress alerts	Receive alert notifications whenever an alert violation is detected.
Only alert once until problem clears	Receive one notification for an alert violation, even if the condition is violated multiple times.
Only generate alerts every	Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.
Do not alert until the condition is violated for more than	Do not receive notifications for an alert unless the violation is not cleared for longer than the specified time.

10. Click **Save** when you are done editing the definition.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a performance alert for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.

Viewing and modifying alert notification settings

View and modify the alert notification settings to determine the actions that are taken when alert conditions are detected for resources. The settings are applied to all of the conditions that are specified for the resources. You can view and modify the alert notification settings for a single resource or for multiple resources.

Viewing the alert notification settings for a resource

View the notification settings that determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the alert definitions that are enabled for the resource.

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. To view the alert notification settings for a resource, choose one of the following options:

Option	Description
Right-click the resource in the list and click Edit Alert Notification Settings .	Use this option to view the notification settings for an alert from the resource list page.
Right-click the resource and click View Details . The resource page is displayed. In the General section, click Alerts , and then click the Notification Settings tab.	Use this option to view the notification settings for an alert from the resource page.

3. View the notification settings that apply to all of the alert definitions that are enabled for the resource.
4. To close the **Notification Settings** page, click **Cancel**.

Viewing the alert notification settings for multiple resources

View the notification settings that determine the actions that are taken when alert conditions are detected for multiple resources. The settings are applied to all of the alert definitions that are enabled for the resources.

When you select multiple resources, the check box choice for each notification setting indicates how the setting is applied to the resources:

Icon	Status of the alert notification
<input checked="" type="checkbox"/>	The notification setting applies to all of the selected resources.
<input type="checkbox"/>	The notification setting does not apply to any of the selected resources.
<input type="checkbox"/>	The notification setting has different values for the selected resources.

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems

Resource	Menu bar
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. To view the alert notification settings for multiple resources, press Ctrl and right-click each resource. Click **Edit Alert Notification Settings**.
3. To close the **Notification Settings** page, click **Cancel**.

Modifying the alert notification settings for a resource

You can modify the alert notification settings to determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the alert definitions that are specified for the resource.

Configure IBM Spectrum Control to send alert notifications by using email, SNMP traps, or a Tivoli Netcool/OMNIbus server.

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

2. To modify the alert notification settings for a resource, choose one of the following options:

Option	Description
Right-click the resource in the list and click Edit Alert Notification Settings .	Use this option to edit the notification settings for an alert from the resource list page.
Right-click the resource and click View Details . The resource page is displayed. In the General section, click Alerts , and then click the Notification Settings tab. Click Edit .	Use this option to edit the notification settings for an alert from the resource page.

3. Specify the email addresses that you want to notify when an alert is generated. The email addresses that you specify override the addresses that were defined at **Settings > Alert Notifications**.

Tip: The email addresses that you define are applied to all the alert definitions for a resource, unless overridden. To override these settings, edit the notification settings for a specific alert definition.

4. Optional: Send alert notifications to a Tivoli Netcool/OMNIbus server.

- Optional: Send SNMP trap messages when an alert is generated.
- Optional: Write alert messages to the OS log and select their level of severity.
- Click **Save**.

Modifying the alert notification settings for multiple resources

Modify the alert notification settings to determine the actions that are taken when alert conditions are detected for multiple resources. The settings are applied to all of the alert definitions that are enabled for the selected resources.

Configure IBM Spectrum Control to send alert notifications by using email, SNMP traps, or a Tivoli Netcool/OMNIbus server.

When you select multiple resources, the check box choice for each notification setting indicates how the setting is applied to the resources:

Icon	Status of the alert notification
<input checked="" type="checkbox"/>	The notification setting applies to all of the selected resources.
<input type="checkbox"/>	The notification setting does not apply to any of the selected resources.
<input type="checkbox"/>	The notification setting has different values for the selected resources.

- To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

- To modify the alert notification settings for multiple resources, press Ctrl and right-click each resource. Click **Edit Alert Notification Settings**.
- Specify the email addresses that you want to notify when an alert is generated. The email addresses that you specify override the addresses that were defined at **Settings > Alert Notifications**.

Tip: The email addresses that you define are applied to all the alert definitions for the selected resources, unless overridden. To override these settings, edit the notification settings for a specific alert definition.

- Optional: Send alert notifications to a Tivoli Netcool/OMNIbus server.
- Optional: Send SNMP trap messages when an alert is generated.
- Optional: Write alert messages to the OS log and select their level of severity.
- Click **Save**.

Viewing alerts

View the alerts that were detected on monitored resources.

Viewing all alerts

Use the Alerts page in the GUI as a central location to view and manage all the alerts that were detected by IBM Spectrum Control.

The Alerts page is organized into two sections:

Alerts summary chart



The summary chart at the top of the page is a visual representation of the alerts that were detected on monitored resources. The chart consists of stacked bars with the number and severity of current alerts on the y-axis and the resource type on the x-axis. The y-axis labeling scales dynamically according to the highest number of alerts in any of the categories.



Alerts list


The alerts list shows information about all the alerts that were detected on monitored resources. Information in the list is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected.

The version of IBM Spectrum Control that detected an alert determines where information about the alert is displayed. Alerts that were generated in versions of the product before Tivoli Storage Productivity Center version 5.1 are only displayed on the Alerts page. Alerts that were generated in Tivoli Storage Productivity Center version 5.1 or later are displayed on the Alerts page and in the **Alerts** tab on the pages for the related resources.

1. In the menu bar in the GUI, go to **Home > Alerts**. Information about alerts is shown in the list of alerts.

Tip: Hover the mouse pointer over a condition to view a summary of the event that triggered the alert.

2. Optional: To view the number of alerts for a specific type of resource, hover the mouse pointer over the related icon in the summary chart. For example, if you want to view the number of Critical, Warning, and Informational alerts that were triggered by hypervisors, hover the mouse pointer over the **Hypervisors** icon .
3. Optional: To filter the list of alerts for a specific type of resource, click the icon for that resource in the summary chart. For example, if you only want to view a list of alerts for hypervisors, click the **Hypervisors** icon .

Tip: To view descriptions of the icons and columns on the Alerts page, click the **Help** icon  in the upper-right corner of the page.

Viewing alerts for a type of resource

Use resource list pages in the IBM Spectrum Control GUI to view only the alerts for a specific type of top-level resource. Each resource list page has an **Alerts** tab that lists only the alerts for monitored storage systems, servers, hypervisors, switches, or fabrics. For example, you can use the list page for block storage systems to view only the alerts that were detected on monitored block storage systems.

1. In the menu bar in the GUI, go to a type of top-level resource. For example, if you want to view the alerts that are related to block storage systems, go to **Storage > Block Storage Systems**.

2. Click the **Alerts** tab to view a list of all the alerts for the monitored resources of that type.


The information about alerts is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected. Hover the mouse pointer over a condition to view a summary of the event that triggered the alert.

Tip: The icon on the **Alerts** tab represents the most severe alert condition that was detected on the monitored resources. For example, if one of the jobs for storage systems failed, the following icon is shown on the tab:



Alerts.

3. Optional: View the icons at the top of the **Alerts** tab for a summary of alert severities. This summary includes the number of alerts that have a Critical, Warning, Informational, or Acknowledged status.

Tip: To view descriptions of the icons and columns on the **Alerts** tab, click the **Help** icon  in the upper-right corner of the page.

Viewing alerts for a specific resource

Use resource detail pages in the IBM Spectrum Control GUI to view only the alerts for a specific top-level resource. Each resource details page has an **Alerts** link that lists only the alerts for a monitored storage system, server, hypervisor, switch, or fabric. For example, you can use the details page for a storage system to view only the alerts that were detected for that storage system.


1. In the menu bar in the GUI, go to a type of top-level resource. For example, if you want to view the alerts that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor.
2. Right-click a resource in the list and select **View Details**. In the **General** section, the number next to **Alerts** represents the number of alert conditions that were detected on the resource. For example, **Alerts (25)** indicates that 25 alerts were detected on the resource.

The icon next to the Alerts link represents the most severe alert condition that was detected on the resource. For example, if alerts with a Warning severity were detected the following icon is shown:

Alerts (25) 

3. Click the **Alerts** link in the **General** section to view a complete list of the alerts that were detected on the resource.

The information about alerts is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected. Hover the mouse pointer over a condition to view a summary of the event that triggered the alert.

Tip: To view descriptions of the icons and columns that are shown for alerts, click the **Help** icon  in the upper-right corner of the page.

4. Optional: View the icons at the top of the page for a summary of alert severities that were detected. This summary includes the number of alerts that have a Critical, Warning, Informational, or Acknowledged status.

Acknowledging alerts

Some alerts in IBM Spectrum Control are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

Tip: Acknowledging alerts does not affect the status of resources. The status of a resource is determined by its consolidated status and the propagated status of its internal resources, not by alerts that are generated. For information about how to acknowledge the status of a resource, see “Acknowledging the condition and status of resources” on page 210.

To acknowledge an alert, complete the following steps:

1. In the IBM Spectrum Control GUI, go to a list of alerts that includes the alerts that you want to acknowledge. For example, if you want to acknowledge the alerts for a specific server, complete the following steps.
 - a. In the menu bar, go to **Servers > Servers..**
 - b. Right-click the server that has alerts you want to acknowledge and select **View Details**.
 - c. Click the **Alerts** link.
2. Right-click the alert that you want to acknowledge and select **Acknowledge alerts**. The severity of the alert is changed to Acknowledged on the Alerts page, resource list page, and resource details page.

For example, if you acknowledge a server alert that has a Warning severity, it is shown on the Alerts page, Servers page, and Server details page with the following severity:



Warning - Acknowledged

3. Optional: To acknowledge multiple alerts at the same time, press Ctrl and click each alert. To select a series of alerts, select the first alert, and then press Shift and click the last alert. When the alerts are selected, right-click a row in the list and select **Acknowledge alerts**.
4. Optional: To acknowledge all the alerts that are shown in a list, select **Acknowledge all alerts** from the **Actions** menu.
5. Optional: To acknowledge all the alerts with an Informational severity in a list, select **Acknowledged informational alerts** from the **Actions** menu. You might want to acknowledge all Informational alerts when you want to view only alerts that represent potential or serious problems that require action to resolve.

Removing alerts

Remove alerts when you no longer require to view them.

By default, alerts are automatically removed based on retention settings that are defined on the History Retention page in the IBM Spectrum Control GUI. However, you can use the following procedure to manually remove alerts at any time.

1. In the GUI, go to a list of alerts that includes the alerts that you want to remove. For example, if you want to remove alerts for all different types of top-level resources, go to **Home > Alerts**. If you want to remove alerts for a specific type of resource such as storage systems, go to **Storage** and select the type of storage system you want to monitor, and click the **Alerts** tab.

2. Right-click the alert that you want to remove and select **Remove alerts**.
3. Click **Remove** to remove the alert. When you remove an alert, it is no longer shown on the Alerts page or in the **Alerts** tabs on the related resource list and resource detail pages.

For example, if you remove an alert for a fabric, that alert is removed from the Alerts page, Fabrics page, and Fabric details page.
4. Optional: To remove all the alerts that are shown in a list, select **Remove all alerts** from the **Actions** menu.
5. Optional: To remove all the acknowledged alerts that are shown in a list, select **Remove acknowledged alerts** from the **Actions** menu.

Viewing information about resources

You can view detailed information about the resources that are monitored by IBM Spectrum Control. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

Remember: Data collection jobs retrieve information about monitored resources. To ensure that the latest information about resources is available in the GUI, run data collection jobs regularly. To learn about how to collect data, see “Collecting data” on page 133.

How information is organized

Information about monitored resources is organized into different pages in the GUI. You can use the pages to view summary information about all types of monitored resources, such as storage systems, fabrics, and switches. You can see what resources are consuming storage and the alerts that are being detected on those resources. You can also view a chart of the most active storage systems in your environment.

Dashboard view

The dashboard is view displayed when you start a new session of IBM Spectrum Control or when you go to **Home > Dashboard** in the menu bar. Use the dashboard to view the following summary information about a storage environment:

- Condition of storage: block, file, and object storage systems
- Condition of storage consumers: hypervisors and servers
- Condition of network resources: fabrics and switches
- Number error, warning, and information alerts that are generated on resources
- Most active storage systems in your environment.

You can click on resource icons in the dashboard to add more resources for monitoring and to see detailed information about resources that IBM Spectrum Control is currently monitoring. Position the mouse pointer over the lines in the Most Active Storage Systems chart to see performance information about specific resources. For more information about the dashboard, see “Dashboard view” on page 197.

Resource list pages

Use a resource list page as a central location for viewing information about the top-level resources of a specific type. For example, use the File Storage Systems

page to view a list of the file storage systems that are monitored by IBM Spectrum Control. Each resource list page corresponds to a specific type of resource, which includes block storage systems, file storage systems, object storage systems, volumes, pools, servers, hypervisors, switches, and fabrics. You can view the following information on resource list pages:

- Overall condition of the resources
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

To learn about how to access the list pages for resources, see “Viewing information about top-level resources” on page 200.

Note: Volumes and pools are internal resources of storage systems, but have their own resource list pages so that you can view them all in a central location.

Resource detail pages

Use resource details pages to view information about resources that are internal or related to a top-level resource. For example, use the File Storage System details page to view detailed information about a specific file storage system. You can view the following information on resource details pages:

- The overall capacity and storage usage for a resource
- The status of resources
- Key information about a resource
- Alerts that were generated for a resource
- Data collection that was run for a resource
- Information about components that exist in top-level resources
- Information about the resources that are related to another resource

To learn about how to access detail pages for resources, see “Viewing information about internal, object, and related resources” on page 201.

Alerts page

Use the Alerts page to view and manage the alert conditions that were detected on monitored resources. Specifically, you can complete the following actions:

- View the overall number of alerts that were detected on monitored resources. This information is presented in a bar chart for the following types of alerts:
 - Database
 - Fabric
 - Hypervisor
 - Other
 - Performance
 - Server
 - Storage system
 - IBM Spectrum Control
- View information about specific alerts, including severity, last occurrence, and the resources where they were detected.
- Acknowledge alerts that were reviewed but are not yet resolved.
- Unacknowledge alerts.

- Remove alerts from the page.

To access the Alerts page, go to **Home > Alerts** in the menu bar.

Performance Monitors page

IBM Spectrum Control can collect information about the performance of monitored storage systems and switches. This information includes key performance metrics. On the Performance Monitors page, you can view detailed information about performance monitors that are collecting data. Use this information to identify performance monitors that are encountering problems when they run. You can also view the metrics that are collected by a performance monitor for a resource.

To access the Performance Monitors page, select **Home > Performance Monitors**.

Tasks page

Use the Tasks page to view and manage the tasks that IBM Spectrum Control uses to optimize resources and provision storage. Tasks are created when you complete any of the following actions:

- Provisioning storage by assigning storage to servers, hypervisors, or a cluster.
- Analyzing tiering by moving volumes to higher or lower tiers.
- Transforming storage by moving or converting volumes.
- Balancing pools by distributing the workload of volumes across pools on the same tier.

To access the Tasks page, go to **Home > Tasks**.

System Management page

Use the System Management page to view information about the overall condition of IBM Spectrum Control. You can view information about the servers on which IBM Spectrum Control is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.

To access the System Management page, go to **Home > System Management** in the menu bar.

Reporting

Use the optional Cognos Business Intelligence reporting tool in the GUI to view detailed reports about resources. The following reports are available:

Predefined reports

These reports are included with IBM Spectrum Control. The following types of predefined reports are available:

- Capacity and relationships of resources
- Performance
- Historical space
- Storage tiering

Custom reports

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information

about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To learn about how to use the Cognos Business Intelligence reporting tool, see

 [Reporting in IBM Spectrum Control](#).

Dashboard view


Use the dashboard to monitor the overall condition of monitored resources and identify potential problem areas in a storage environment. The dashboard also displays alerts, as well as a chart showing the most active storage systems.

Viewing the status of monitored resources

The dashboard provides status information for the resources that IBM Spectrum Control is monitoring.

The overall status for each resource type

A status symbol and number are shown below each resource icon on the dashboard. This symbol represents the most critical status that was detected on the internal resources for a resource type. For example, if 20 storage systems are monitored, and an error was detected on a volume for

one of those storage systems, a  1 status is shown below the storage systems icon or label in the dashboard. If no errors, warnings, or unreachable statuses were detected on the monitored resources, then a green symbol is shown without a number.

Use the status to quickly determine the condition of your monitored resources and if any actions must be taken. On the dashboard, the following statuses might be shown for a resource type:



Error (red)

An error status was detected on one or more of the monitored resources for a resource type. Error statuses represent serious problems on a resource. Resolve these problems as soon as possible. View the details of a resource to learn more about the status of its internal resources.



Unreachable (orange)

One or more of the monitored resources for a resource type are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.



Warning (yellow)

A warning status was detected on one or more of the monitored resources for a resource type. Warning conditions are not critical, but represent potential problems. View the details of a resource to learn more about the status of its internal resources.



Normal (green)

No warnings or errors were detected on the monitored resources for a resource type.



Unknown status (gray)

One or more of the monitored resources for a resource type have an unknown status. This status might occur if no data was collected about a resource. To change an unknown status, use a data collection job to collect status information about the resource.

For information about how status is determined, see “How the condition of a resource is determined” on page 204.

The number of statuses that were detected and acknowledged for each resource type

To view the total number of statuses for a resource type, move the mouse pointer over the related status symbol. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a IBM Spectrum Control user. An acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.

Viewing the status of monitored storage systems

Storage systems can be configured for block storage, file storage, object storage, a combination of block and file storage, or a combination of file and object storage. The top left section of the dashboard shows the type, number, and status of the storage systems that IBM Spectrum Control is monitoring. The number in the circle next to the storage system type shows how many storage systems of that type are being monitored. The status icon below each storage system type shows the status of the monitored systems. Click the storage system type, for example Block Storage Systems, to see more detailed status and information about those storage systems. Use the information in the horizontal bar charts to view details about the used and free space of the monitored storage systems in your environment:

Storage systems that are configured for block data

The following information is shown for monitored storage systems that are configured for block data:

Used Block Storage Space

Amount of space that is being used for block data on all the storage systems that are configured for block data.

Available Block Storage Space

Amount of unused space on all the storage systems that are configured for block data.

Storage systems that are configured for file data

The following information is shown for monitored storage systems that are configured for file data:

Used File System Space

Amount of file system space that is being used by files and directories on the storage systems and filers that are configured for file data.

Available File System Space

Amount of unused storage space on the storage systems and filers that are configured for file data.

Storage systems that are configured for object data

The following information is shown for monitored storage systems that are configured for object data:

Used Object Storage Space

Amount of storage space that is used on all the object storage systems.

Available Object Storage Space

Amount of storage space that is available on all the object storage systems.

Viewing the status of monitored resources that are consuming storage

The diagram in the top middle section of the dashboard provides a visual representation of the monitored resources that are consuming storage. You can click the icons in this section to add resources of that type or to see detailed information about the resources you have already added. The number next to each icon shows the number of those resources that are being monitored by IBM Spectrum Control. If IBM Spectrum Control is monitoring resources of a particular type, such as servers, you will see a status icon that shows the overall status of those types of resources.

The types of resources that are consuming storage

Each type of resource in the diagram is represented by an icon. To learn more detailed information about the individual resources that are monitored, click the icon in the diagram. For example, if you click the departments icon, a new page shows detailed information about the individual departments that are being monitored.

The number of resources that are monitored for each resource type

A number in a circle is shown next to each resource icon. This number represents the number of monitored resources for each resource type. For example, if 4 file storage systems are monitored, the value ④ is shown. This number excludes resources that are not monitored by IBM Spectrum Control.

Tip: If no resources are being monitored for a resource type, click its icon on the dashboard to add a resource of that type. For example, if the dashboard shows a 0 next to the departments icon, you can click the icon or the circle for departments to add a department. The GUI guides you through the process for adding the resource for monitoring. For more information about adding resources, see “Adding resources” on page 112.

Viewing the status of monitored network fabrics and switches

The top right section of the dashboard shows the number and status of the fabrics and switches that IBM Spectrum Control is monitoring. The number in the circle next to the fabric or switch icon shows how many fabrics and switches are being monitored. The status icon below the fabric or switch shows the status of those entities. Click the icon to see more detailed status and information about those fabrics and switches. The dashboard shows the following information for the monitored fabrics and switches in your environment:

Fabrics

- The number of fabrics that are being monitored.
- The most critical status of the monitored fabrics. If more than one fabric is monitored, the dashboard shows the most critical status level of those fabrics.

Switches

- The number of switches that are being monitored.
- The most critical status of the monitored switches. If more than one fabric is monitored, the dashboard shows the most critical status level of those switches.

Viewing the status of alerts that were detected on monitored resources

Alerts are generated when IBM Spectrum Control detects certain conditions or events on monitored resources. The Alerts section of the dashboard shows a status summary of the alerts that were detected in your environment during a specified time period. This summary includes the number of alerts that have the following statuses: Critical, Warning, and Informational.

Use the following actions to manage the display of statuses for alerts:

- To select the time period in which you want to view alert statuses, click the down arrow and select a time period. You can view the status of all alerts that were detected in the last hour, last day, or last week.
- To view more detailed information about alerts, click **View all alerts** to access the Alerts page.

Viewing the most active storage systems in your environment

The Performance chart provides information about the most active storage systems in your storage environment. Position your cursor over a data point in the chart to see the total I/O rate for a particular storage system. The I/O rate is displayed in operations per second (ops/s) on the left vertical axis, and in milliseconds per operation (ms/ops) on the right vertical axis. The horizontal axis shows the time of the performance data in hourly increments.

Viewing information about top-level resources


Use a resource list page as a central location for viewing information about all the monitored resources of a specific type. Resources that have a list page include storage systems, volumes, pools, shares, servers, hypervisors, switches, and fabrics.


1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view. You can access the following types of resources:
 - **Storage:** block storage systems, file storage systems, object storage systems, volumes, pools, shares
 - **Servers:** servers, hypervisors
 - **Network:** switches, fabrics
2. Select the type of resource that you want to view. For example, if you want to view information about servers, go to **Servers > Servers**. A list of monitored servers is displayed on Servers page.
3. Locate the resource that you want to view in the list. Each column in the list shows information about the resource.
4. Optional: To view key attributes of a specific top-level resource without leaving the page, right-click the resource and select **View Properties**. A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.

5. Optional: To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**. A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, and related resources.

A resource list page shows information about the monitored resources for the selected type. You can view the following information on this page:

- The overall status of the resources for a specific type
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

Tip: To view descriptions of the information that is available on a resource list page, click the **Help** icon  in the upper-right corner of the page.

To view information about rollup resources, you must be in *rollup mode*. To enter rollup mode, move the mouse pointer over the rollup icon  in the menu bar and select **Enter rollup mode**.

Related reference:

“Resources that you can monitor” on page 107

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Viewing information about internal, object, and related resources


Use resource details pages to view information about a specific top-level resource and its internal, object, and related resources.

1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view. You can access the following types of resources:
 - **Storage:** block storage systems, file storage systems, object storage systems, volumes, pools, shares
 - **Servers:** servers, hypervisors
 - **Network:** switches, fabrics
2. Select the type of the resource that you want to view. A resource list page shows the monitored resources for the selected type. For example, if you want to view a block storage system, go to **Storage > Block Storage Systems**. A list of monitored block storage systems is displayed.
3. In the list of resources, right-click the resource that you want to view and select **View Details**. A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, object resources, and related resources.
4. To view information about the internal resources of the top-level resource, click the name of the internal resource in the **Internal Resources** section. A list of the internal resources is displayed. Each column in the list shows information about the resources. For example, on the details page for a server, click **Controllers** to view a list of the controllers that are associated with the server.

5. Optional: To view key attributes of a specific internal resource without leaving the page, right-click the resource and select **View Properties**.
6. To view information about the object resources in an object storage system, click the name of the object resource in the **Object Resources** section. Object resources are associated only with object storage systems, such as IBM Spectrum Scale that is configured for object storage.
7. To view information about resources that are related to the top-level resource, click the name of the related resource in the **Related Resources** section. A list of the related resources is displayed. For example, on the server details page, click **Block Storage Systems** to view a list of the block storage systems that have at least one disk (volume) assigned to the server.
8. Optional: To view key attributes of a specific related resource without leaving the page, right-click the resource and select **View Properties**.
9. Optional: To view the resource details page for a related resource, right-click the resource and select **View Details**.

A resource details page shows information about a top-level resource. The information is available in the following links and sections on the page:

- **Overview:** The overall capacity and storage usage for a resource.
- **Properties:** Key information about a resource.
- **Alerts:** Alerts that were generated for a resource.
- **Tasks:** Jobs that were run for a resource.
- **Internal Resources:** Information about the components that exist in a higher-level resource.
- **Object Resources:** Information about the resources in a GPFS cluster, such as accounts and containers, that enable the GPFS cluster to be used as an object storage system.
- **Related Resources:** Information about the resources that are related to another resource. A related resource is external to a higher-level resource, but is related to it through assigned storage, a network connection, or virtual hosting.

Tip: To view descriptions of the information that is available on a resource details page, click the **Help** icon  in the upper-right corner of the page.

Related reference:




“Resources that you can monitor” on page 107

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Monitoring the status and condition of resources

Monitor the operational condition of storage systems, servers, hypervisors, fabrics, and switches and the status of their internal resources. Use this information to identify potential problem areas in a storage environment.

Table 40. Monitoring the status and condition of resources

	Explanation	Steps to view status or condition
Status	<p>The status of a resource that is reported by its own hardware to IBM Spectrum Control. For example:</p> <ul style="list-style-type: none">• If the cooling fans in a storage system are stopped and the internal temperature is too high, an error status  is reported by that storage system.• If a disk on a storage system is starting, a warning status  is reported for that disk by the storage system. <p>Tip: The status of internal resources is used to determine the operational condition of the associated top-level resources.</p>	<ol style="list-style-type: none">1. In the menu bar, go to the resource type that you want to view. For example, if you want to view the status of switches, go to Network > Switches.2. Right-click a resource and select View Details. A status icon is shown next the image of the resource and its internal resources.
Condition	<p>The overall operational condition of a storage system, server, hypervisor, fabric, or switch. This condition represents the most critical status that was detected on the resource itself and on its internal resources.</p> <p>For example, if an error status was detected on a storage system pool, an error icon  is shown for the overall condition of the storage system. If no errors, warnings, or unreachable statuses were detected on a resource or on its internal resources, then a green symbol is shown for the condition of the storage system.</p>	<p>Dashboard view</p> <p>In the menu bar, go to Home > Dashboard. The icons that show overall condition are shown next to the image of a resource type.</p> <p>Resource list pages</p> <p>In the menu bar, go to the resource type that you want to view. For example, if you want to view the condition of switches, go to Network > Switches.</p> <p>The overall condition of a resource is shown in the Condition column, and aggregated in the condition icons at the top of the page.</p>

IBM Spectrum Control provides a number of different icons to help you quickly determine the health of resources.

Table 41. Possible statuses and conditions of resources


Icon	Health	Explanation
	Error	A serious problem was detected on a resource or on its internal resources. Resolve these problems as soon as possible.

Table 41. Possible statuses and conditions of resources (continued)

Icon	Health	Explanation
	Error - Acknowledged	<p>An Error status was detected and acknowledged. An Error - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>An acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a volume is Error, the condition of the associated storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall condition of the storage system.</p>
	Unreachable	A resource is not responding to requests from the IBM Spectrum Control server. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.
	Unreachable - Acknowledged	<p>An Unreachable status was detected and acknowledged. An Unreachable - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>An acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a controller is Unreachable, the condition of the associated server is also unreachable. If the Unreachable status of the controller is acknowledged, its status is not used to determine the overall condition of the server.</p>
	Warning	A Warning status represents potential problems on a resource or on its internal resources. This status is not critical.
	Warning - Acknowledged	<p>A Warning status was detected and acknowledged. An Warning - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>An Warning - Acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a controller is Warning, the condition of the associated server is also Warning. If the Warning status of the controller is acknowledged, its status is not used to determine the overall condition of the server.</p>
	Normal	No warnings or errors were detected on a monitored resource.
	Unknown	A resource is known to IBM Spectrum Control but is not monitored. To change an Unknown status, run a probe to collect status information about the resource.

How the condition of a resource is determined

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

To determine the condition of top-level resources (storage systems, servers, hypervisors, fabrics, and switches), IBM Spectrum Control uses the status of those resources and the status of their internal resources. The statuses of the following internal resources are used to calculate the overall condition of a top-level resource.

Table 42. Internal resources that are used to determine the condition of top-level resources




































Top-level resource	Internal resources that are used to determine the condition of a top-level resource
Fabric 	 Ports  Switches
Switch 	 Blades  Inter-switch connections  Logical switches  Ports
Hypervisor 	 Disks  Virtual machines  VMDKs
Server 	 Controllers  Disks

Table 42. Internal resources that are used to determine the condition of top-level resources (continued)

Top-level resource	Internal resources that are used to determine the condition of a top-level resource
<p>Block storage system</p> 	<ul style="list-style-type: none">  Disks  Drives  External disks  Host connections  I/O groups  Managed disks  Modules  ,  Nodes  Pools  Ports  RAID arrays  Volumes
<p>File storage system</p> 	<ul style="list-style-type: none">  Network shared disks  Nodes
<p>Object storage system</p> 	<ul style="list-style-type: none">  Network shared disks  Nodes

The following statuses of internal resources are used to help calculate the condition of top-level resources:

- Normal
- Warning
- Error

Tip: Statuses that are acknowledged are not used to calculate the overall condition of a top-level resource.






Internal resources for a top-level resource might have different statuses. IBM Spectrum Control uses the most critical status of an internal resource to help determine the overall condition of a top-level resource. For example, in a storage system, a port might have an Error status, a pool might have a Warning status, and multiple controllers might have an Unknown status. In this case, if the storage system itself has a normal status, its overall condition is Error because it is the most critical status that was detected on internal resources.

The following table shows some of the possible combinations of statuses and the resulting, overall condition for a top-level resource.

Table 43. Propagation of the statuses for resources

Error 	Unreachable ¹ 	Warning 	Normal 	Unknown ² 	Resulting condition for a top-level resource
				X	 Unknown
			X		 Normal
			X	X	 Normal
		X			 Warning
		X		X	 Warning
		X	X	X	 Warning
	X				 Unreachable
	X			X	 Unreachable
	X		X	X	 Unreachable
	X	X	X	X	 Unreachable
X					 Error
X				X	 Error
X			X	X	 Error
X		X	X	X	 Error
X	X	X	X	X	 Error

Table 43. Propagation of the statuses for resources (continued)

Error 	Unreachable ¹ 	Warning 	Normal 	Unknown ² 	Resulting condition for a top-level resource
Note: 1. The Unreachable status applies only to top-level resources. 2. The Unknown status of an internal resources is not used to determine the condition of a top-level resource.					

Viewing the overall condition of resources

Use the dashboard to view the overall, operational condition of resources, and identify potential problem areas in a storage environment. The dashboard shows the total number of different conditions on storage systems, servers, hypervisors, fabrics, switches, and on the resources in departments and applications.

The condition of a resource represents the most critical status that was detected on that resource or on its internal resources. For example, if an error status was detected on a storage system pool, the overall condition of the storage system is

error .


1. In the menu bar, go to **Home > Dashboard**.
2. View the condition icons that are shown with each resource type.

For example, if 20 block storage systems are monitored, but 10 storage systems have internal resources with an error status, 2 storage systems are unreachable, and 4 storage systems have internal resources with a warning status, the following conditions icons are shown for block storage systems:

Block Storage Systems

 10  2  4

If no Error or Warning statuses were detected on the internal resources for

block storage systems, the normal condition icon  is shown.

Related reference:

“How the condition of a resource is determined” on page 204

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

Viewing the condition of specific types of resources

Use resource list pages as a central location for viewing the condition of storage systems, servers, hypervisors, switches, and fabrics.

1. In the menu bar, go to the type of top-level resource that you want to view. For example, if you want to view the condition of switches, go to **Network > Switches**.
2. To view the total number of different conditions that were detected for the type of resource, check the condition icons at the top of the resource list page. The number next to the icon represents how many occurrences of each condition were detected. For example, if 15 switches are monitored, but 5 switches have internal resources with a warning status, 2 switches have internal resources

with an unreachable status, and 3 switches have internal resources with an error status, and the following condition information is displayed at the top of the page:



5 Normal





5 Warning



2 Unreachable





3 Error

3. To view the condition for specific resources, check the **Condition** column in the list of resources.
4. Optional: Click the **Search** icon  to filter the list so that only resources with a specified condition are shown. You can use filtering to identify the resources with the most serious conditions in your environment.
5. Optional: Select **Condition**.
6. Optional: In the **Choose Value** list, select the condition on which you want to filter. The list of resources is automatically refreshed and only resources with the filtered status are shown. For example, select **Warning**  to view only the resources with a Warning condition.

Viewing the status of resources

Use detail pages to view the status of storage systems, servers, hypervisors, switches, and fabrics and the status of their internal resources. The status of a resource is reported to IBM Spectrum Control by its hardware.

The status of a resource is different than the operational condition that is calculated by IBM Spectrum Control for top-level resources:

- *Condition* is determined by propagating the statuses of resources and is only shown for top-level resources.
- *Status* represents the status of a resource as reported by its hardware and is shown for top-level resources and internal resources. The following examples illustrate statuses that might be reported by a resource:
 - If the cooling fans in a storage system are stopped and the internal temperature is too high, an error status  is reported by that storage system.
 - If a disk on a storage system is starting, a warning status  is reported for that disk by the storage system.

Tip: View condition to identify which top-level resources might be encountering problems. View status to identify the specific resources that are causing the problems in a top-level resource.

1. In the menu bar, go to the resource type that you want to view. For example, if you want to view the status of block storage systems and their internal resources, go to **Storage > Block Storage Systems**.
2. Right-click a resource and select **View Details**.
3. To view the status of the top-level resource, view the status icon that is displayed next to its image in the upper-left corner of the page.

4. To view the aggregated status of an internal resource, view the status icon that is displayed next to its image in the **Internal Resources** section. If a resource has a status other than normal, an icon is displayed for the most critical status. For example, if a server has three controllers with an error status and two controllers with a warning status, the error status icon is shown for controllers in the **Internal Resources** section:


 **Controllers (5)** 

5. To view the status of a specific internal resource, click the name of that internal resource in the **Internal Resources** section. A list of the monitored internal resources is displayed. For example, on the details page for a server, click **Controllers (5)** to view a list of the five controllers that are associated with the server.
6. Check the **Status** column to view the status for each internal resource.

Tip: When you view information about some internal resources, the Status column might show values that are more specific than Error, Unreachable, Warning, Unreachable, Unknown, and Normal. Use this additional status to determine the cause of the problem for an internal resource.

For example, in the **Internal Resources** section on the details page for a storage system, an error status might be shown for managed disks:

 **Managed Disks (5)** 

When you view the list of managed disks, the Status column might show an Offline status for a specific managed disk. When a managed disk is reported as offline, IBM Spectrum Control uses an error status icon  to reflect that status.

Acknowledging the condition and status of resources

Sometimes the status of resources might represent problems that commonly occur but can be ignored. In such cases, you can acknowledge those statuses so that they are not used when determining the overall condition of storage systems, servers, hypervisors, fabrics, and switches.

Acknowledging the condition of top-level resources

A summary of the conditions for storage systems, servers, hypervisors, fabrics, and switches is available in the Dashboard view and on resource list pages. You can acknowledge the condition of a top-level resource so that its condition is not shown as Error, Unreachable, or Warning on those pages.

1. In the menu bar, go to the type of resource that has a condition you want to acknowledge. For example, if you want to acknowledge the condition of a switch, go to **Network > Switches**.
2. In the list of resources, locate the resource with the Error, Unreachable, or Warning condition that you want to acknowledge.
3. Right-click the resource and select **Acknowledge condition**. A window shows that the condition for the specified resource was acknowledged.
4. Optional: To acknowledge the condition of multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge condition**.

5. Click **Close**.

When you acknowledge the condition of a top-level resource, the following actions are taken:

Dashboard view

- The condition of the acknowledged resource is not used to determine the condition icon that is displayed for that resource type on the Dashboard view. For example, if you acknowledge the unreachable condition for a switch, the condition of that switch is not used to determine the number of unreachable conditions that is shown for switches on the Dashboard view.
- The acknowledged condition is shown as **Acknowledged** for a resource on the Dashboard view. For example, if 10 switches have Unreachable conditions and you acknowledge the condition for one of them, the following information is shown for switches:




9 Unreachable (1 Acknowledged)

Resource details page

- The condition icons at the top of the page are updated. For example, if 10 switches have Unreachable conditions, and you acknowledge one of them, the total number Unreachable conditions decreases by one and the number of Unreachable - Acknowledged conditions increases by one. The following information is updated at the top of the Switches page:



9 Unreachable  **1 Unreachable - Acknowledged**

- The value in the **Condition** column for the resource is updated to show an acknowledged condition. For example, if you acknowledge the Unreachable condition for a switch, the following value is shown in the **Condition** column:



Unreachable

Related reference:

“Resources that you can monitor” on page 107

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

“How the condition of a resource is determined” on page 204

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

Acknowledging the status of internal resources

The statuses of internal resources are used to determine the condition of the associated storage systems, servers, hypervisors, fabrics, and switches. Sometimes these statuses might represent problems that commonly occur or are known, but can be ignored. In such cases, you can acknowledge the status of an internal resource so that it is not used to determine the condition of a top-level resource.

For example, if the status of a volume is Error, the condition of the associated storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall condition of its storage system. In this

case, if the other internal resources of the storage system are Normal, then the condition of the storage system is also Normal.

1. In the menu bar, go to the type of resource in which the internal resource is located. For example, if you want to acknowledge the status of internal resources for a switch, go to **Network > Switches**.
2. Right-click a resource and select **View Details**. For example, if you want to acknowledge the status of a port on a switch, right-click that switch and select **View Details**.
3. In the **Internal Resources** section of the resource details page, click the name of the internal resource. For example, if you want to acknowledge the status of a port on a switch, click **Ports**.
4. Right-click the resource and select **Acknowledge status**. A window shows that the status for the specified resource was acknowledged. For example, if you want to acknowledge the status of a port, right-click the port and select **Acknowledge status**.
5. Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge status**.
6. Click **Close**. The status of the related, top-level resource might take several minutes to update.

When you acknowledge the status of an internal resource, the following actions occur:

- The status of the internal resource is no longer used to determine the condition of the associated top-level resource.
 - The status icons at the top of the page for the internal resource are updated.
- For example, if you acknowledge the Error status for a controller on a server, the total number of Error statuses decreases by one and the number of Error - Acknowledged statuses increases by one on the Controllers page. If originally there were five controllers with Error statuses, the following updated status information is shown:

 4 Error  1 Error - Acknowledged

Related reference:

“Resources that you can monitor” on page 107

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

“How the condition of a resource is determined” on page 204

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

Monitoring vaults in IBM Cloud Object Storage

Monitor access risk and storage risk for vaults in IBM Cloud Object Storage.

You can use this information to complete the following tasks:

- Identify the vaults that cannot be accessed and the vaults that are at risk of access failure.

- Identify the vaults that are not available for write operations and the vaults that are at risk of storage failure.

Monitoring access to vaults in IBM Cloud Object Storage

In these scenarios, you use IBM Spectrum Control to identify the vaults that cannot be accessed and the vaults that are at risk of access failure. Then, you identify the access pools and the COS Accesser nodes that the vaults depend on.

How the accessibility value for a vault is determined

The accessibility value that is calculated for each vault comprises the number of COS Accesser nodes that are available to access the vault. For example, if five COS Accesser nodes are configured for a vault and four of the nodes are in a failed state, then the accessibility value is one.

On the details page for IBM Cloud Object Storage, the Vaults by Accessers chart shows the vaults that are inaccessible and the accessibility values for the remaining vaults in the storage system.


Investigating vaults that cannot be accessed:

In this scenario, you use IBM Spectrum Control to identify the vaults that cannot be accessed because of COS Accesser node failures.

You have a backup application that uses IBM Cloud Object Storage vaults. Because the application is critical to your business, you want to quickly identify vaults that cannot be accessed because of COS Accesser node failures. You also want to know the access pools that the vaults depend on and the COS Accesser nodes that failed in those pools.

1. To determine whether any of your vaults cannot be accessed, click **Storage > Object Storage Systems**.
2. Right-click the storage system that is used by your backup application and click **View Details**.
3. In the navigation pane, click **Overview**. The first line of the Vaults by Accessers chart shows the number of vaults that cannot be accessed, for example, Inaccessible 4.
4. To identify the vaults that are inaccessible, click **Vaults** in the navigation pane. To sort the data, click the **Accessibility** column.









A value of 0 in the Accessibility column indicates that the vault is inaccessible.

The error status  indicates that the vault is inaccessible because none of the COS Accesser nodes that are configured for the vault are available. All the inaccessible vaults are configured to use just one access pool, AccessPool2.



Vaults

✓ 104 Normal
! 2 Warning
✗ 5 Error

Actions ▾		
Name	Access Pools	Accessibility ▾
 accessertest1	 AccessPool2	✗ 0
 demovault7	 AccessPool2	✗ 0
 myvault	 AccessPool2	✗ 0
 vault2	 AccessPool2	✗ 0

5. To identify the COS Accesser nodes that failed in the access pool, which in this scenario is called AccessPool2, click **Accesser Nodes** in the navigation pane. To sort the data, click the **Access Pool** column.

To restore access to the vaults, fix the COS Accesser nodes that failed or add COS Accesser nodes to the access pool.

Monitoring vaults that are at risk of access failure:

In this scenario, you use IBM Spectrum Control to check which vaults are at risk because of COS Accesser node failures.

You have a backup application that uses IBM Cloud Object Storage vaults. Because the application is critical to your business, you want to monitor the tolerance of the vaults to COS Accesser node failures and identify the vaults that are at risk of access failure.















For the vaults with a high risk of failure, you want to identify the access pools that the vaults depend on and the status of the COS Accesser nodes in the access pools.

1. Click **Storage > Object Storage Systems**.
2. Right-click the storage system that is used by your backup application and click **View Details**.
3. In the navigation pane, click **Overview**. The Vaults by Accessers chart shows the tolerance of the vaults to COS Accesser node failures. For example, a line on the chart that shows a value of 1 Node Failure 57 Vaults at Risk indicates that 57 vaults rely on a single COS Accesser node for access. If the node fails, the vault cannot be accessed.
4. To identify the vaults with an accessibility value of 1, click **Vaults** in the navigation pane. To sort the data, click the **Accessibility** column.



Vaults

✓ 104 Normal
! 2 Warning
✗ 5 Error


Actions ▾		
Name	Access Pools	Accessibility ▾
 failovertest	2	 1
 test2	2	 1
 DemoVault	 AccessPool1	1
 arun878172498111...	 AccessPool1	1
 arunagiri123	 AccessPool1	1
 arunagiri123.meta	 AccessPool1	1
 containermodetest	 AccessPool1	1

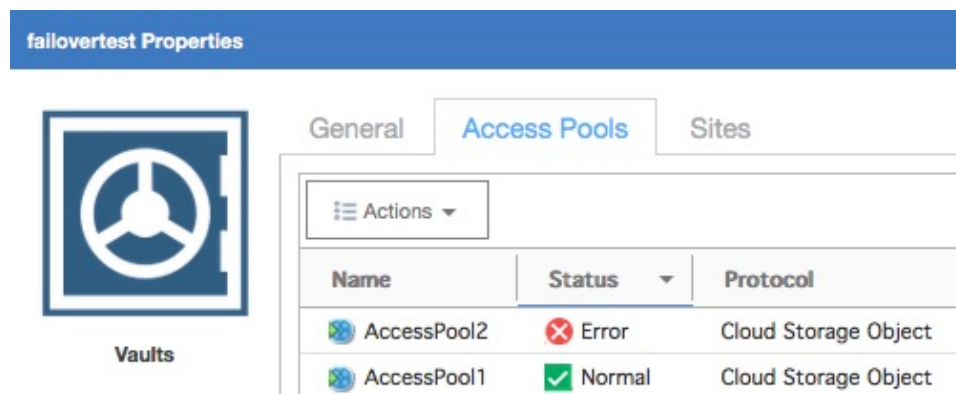
In this scenario, vaults such as containermodetest depend on only one access pool, AccessPool1. The absence of a warning status in the Accessibility column indicates that the access pool is healthy.

Tip: To reduce the risk of access failure for the vaults that depend on AccessPool1, add COS Accesser nodes to the access pool or deploy the vaults to other pools.

For the test2 and failovertest vaults, you see that two access pools are

configured. The warning status  indicates that some of the COS Accesser nodes in the pools are not available.

- To investigate the access issue, click the number in the **Access Pools** column. To sort the data, click the **Name** column. In this scenario, AccessPool2 is shown with an error status , which indicates that the COS Accesser nodes in the pool are not available.



- To identify the COS Accesser nodes that failed in the access pool, click **Accesser Nodes** in the navigation pane. To sort the data, click the **Access Pool** column.

Fix the COS Accesser nodes that failed or add COS Accesser nodes to the access pool.

Calculating the failure tolerance for vaults

The failure tolerance value for vaults represents the tolerance of a vault to drive failures across all of the COS Slicestor nodes in its storage pool. The lower the value for drive failure tolerance, the greater the risk that write operations can no longer be completed for the vault.

The calculation of drive failure tolerance is based on the following values:

- The write threshold value in the IDA that is configured for the vault. The write threshold is the minimum number of COS Slicestor nodes that must be available to complete write operations for the vault.
- The number of available COS Slicestor nodes in the storage pool that the vault belongs to.
- The value for drive failure tolerance for each of the available COS Slicestor nodes in the storage pool. This value represents the tolerance of the node to drive failures and is based on the number of failed drives in the node and the drive error threshold. A value of 2 means that the node will be unavailable to store vault data if 2 drives fail in the node.

For example, a vault in your IBM Cloud Object Storage environment is configured with an IDA of 8-4-6. The IDA determines that the vault data is stored across 8 COS Slicestor nodes and a minimum of 6 nodes must be available to complete write operations for the vault.

The nodes have the following values for drive failure tolerance:

Node	Drive Failure Tolerance
node_1	3
node_2	4
node_3	0*
node_4	2
node_5	4
node_6	1
node_7	3

Node	Drive Failure Tolerance
node_8	3

*A failure tolerance value of 0 means that node_3 is not available to store vault data.

7 COS Slicestor nodes are available in the storage pool and a minimum of 6 nodes must be available to complete write operations. Therefore, if 2 more nodes fail, the write threshold value is exceeded and write operations can no longer be completed for the vault.

Use the following steps to calculate the value for drive failure tolerance for a vault:

1. Use the IDA column on the Vaults page to identify the write threshold value for the vault. In the example, the write threshold value is 6.
2. Use the Slicestor Nodes page to determine the number of available COS Slicestor nodes in the storage pool that the vault belongs to. A value > 0 in the Drive Failure Tolerance column means that the node is available. In the example, 7 nodes are available.
3. Calculate how many nodes must fail for the number of available nodes to fall below the write threshold value. In the example, $(7 - 6) + 1 = 2$. When 2 nodes fail, the vault is not available for write operations.
4. Use the Drive Failure Tolerance column on the Slicestor Nodes page to identify the 2 available nodes in the storage pool that have the lowest tolerance to drive failures.

In the example, the 2 nodes are as follows:

Node	Drive Failure Tolerance
node_4	2
node_6	1

5. To calculate the overall failure tolerance for the vault, sum the failure tolerance values for the 2 nodes. $2 + 1 = 3$.
The drive failure tolerance for the vault is 3. If 3 drives fail, write operations can no longer be completed for the vault.

Monitoring the performance of resources

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.


You can use the performance monitoring to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the alerts that were triggered when the performance of a resource fell outside of a specific range.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.

- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including performance information.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view performance information for resources, you must complete the following tasks:


Task	 Learn more
Add storage systems and switches for monitoring by IBM Spectrum Control.	For information about how to add these resources, see the following topics: <ul style="list-style-type: none"> • “Adding storage systems” on page 115 • “Adding fabrics and switches” on page 120
Ensure that performance data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	For information about verifying data collection and scheduling performance monitors, see the following topics: <ul style="list-style-type: none"> • “Verifying that a performance monitor is running for a resource” on page 142 • “Creating performance monitors” on page 140
Define performance alerts to be notified if the performance of a resource falls outside a specified range and might represent a potential problem. When you define a performance alert for a resource, select a specific metric that you want to measure. For example, you can define a threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.	For information about how to define performance alerts, see “Defining alerts for performance metrics” on page 155.

Related reference:

“Performance metrics” on page 403

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can view performance information or specify alerts for when the performance of a resource falls outside a specified range.

Related information:


 [Video] IBM Spectrum Control Performance Troubleshooting Techniques
Learn about the performance troubleshooting techniques that you can use in IBM Spectrum Control. Watch Bob, the storage administrator, as he investigates a problem with a critical medical database application.


Viewing performance information

View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

Before you view the performance of resources, ensure that performance data was collected for those resources during the time ranges that you want to analyze. IBM Spectrum Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. For information about how to verify that performance monitors were run, see “Verifying that a performance monitor is running for a resource” on page 142.

You can access the performance view for a resource from different locations in the web-based GUI:

Task	Accessing the performance view in the web-based GUI
View the performance of specific storage systems, switches, and their internal resources.	<ol style="list-style-type: none"> 1. In the menu bar, go to the type of resource that you want to view. For storage systems, depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. For switches, go to Network > Switches. 2. Right-click a resource and select View Performance. By default, the following information is displayed in a chart for each type of resource: <ul style="list-style-type: none"> • For storage systems, the top 5 storage systems with the highest total I/O rate are shown. • For switches, the top 5 switches with the highest total port data rate are shown. 3. To view the performance of an internal resource, right-click the top-level resource in the chart legend and select the option for the internal resource that you want to view. For example, for a SAN Volume Controller, right-click its name in the chart legend and select MDisk Performance to view the performance of its managed disks. For managed disks, the top 5 contributors to the storage system workload are shown on the chart. The default metrics are total back-end I/O rate and total back-end response time. <p>Tips:</p> <ul style="list-style-type: none"> • To add or change the metrics that are displayed for a resource, click the select metrics icon  and select the metrics to display. • You can also view the performance of internal resources on the details page for the related, top-level resource. For more information, see Table 44 on page 222. • To view the performance of multiple resources at the same time, press Shift or Ctrl and click the resources that you want to view. Press Shift and click to select consecutive rows in the list of resources; press Ctrl and click to select non-consecutive rows. Then, right-click any of the selected resources and select View Performance. Each resource is represented by a separate line in the performance chart and separate rows in the table view.

Task	Accessing the performance view in the web-based GUI
View the overall performance of resources of the same type.	<ol style="list-style-type: none"> 1. In the menu bar, go to the type of resource that you want to view. For storage systems, depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. For switches, go to Network > Switches. 2. Click the Performance tab. By default, the following information is displayed in a chart: <ul style="list-style-type: none"> • For storage systems, the top 5 storage systems with the highest total I/O rate are shown. • For switches, the top 5 switches with the highest total port data rate are shown. <p>Tip: To add or change the metrics that are displayed for a resource, click the select metrics icon  and select the metrics to display.</p>
View the real-time performance of a SAN Volume Controller or Storwize storage system.	<ol style="list-style-type: none"> 1. Depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. 2. Right-click a storage system and select View Real-Time Performance. A graphical representation of the processor usage of the system, and the bandwidth of the volumes, interfaces, and MDisk of the system is displayed. <p>You can view the real-time performance of a storage system when the following conditions apply:</p> <ul style="list-style-type: none"> • The storage system is a SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V5000, Storwize V7000, or Storwize V7000 Unified system. • The storage system must be at microcode level 7.3 or later. • The storage system is running. To test the connection, right-click the storage system, and select Connections > Test Connection. • The storage system must be available to the system where the browser is running, through the local area network. To check if the storage system is available, right-click the storage system and click Open Storage System GUI.

Task	Accessing the performance view in the web-based GUI
View the alerts that were generated by threshold violations and the performance of the resources where the violation occurred.	For information about how to view performance alerts, see “Viewing performance alerts” on page 235.
View the performance of a resource that is included in a specific performance monitor.	<ol style="list-style-type: none"> 1. In the menu bar in the web-based GUI, go to Home > Performance Monitors. 2. In the Name column, locate the name of the resource that you want to view. 3. Right-click the performance monitor and select View Performance.
View predefined reports and create custom reports that show the performance of resources. In the optional Cognos Business Intelligence reporting tool, you can view reports about storage systems, components of storage systems, switches, and switch ports.	<ol style="list-style-type: none"> 1. In the menu bar in the web-based GUI, go to Cognos > View predefined reports. 2. To view predefined reports, navigate to the performance report that you want to view and click the name of the report. To create a custom report about the performance of resources, in the menu bar, go to Cognos > Create performance reports.

In the web-based GUI, you can also view the performance of internal resources on the details pages for the related, top-level resources. For example, you can view the performance of the controllers for a storage system on the details page for that storage system. To access the performance view for an internal resource on a details page, select the type of internal resource and click the **Performance** tab.

The following information summarizes the internal resources and locations where you can access a performance view.

Table 44. Internal resources and locations where you access the performance view

Internal resource	Locations in the GUI where you can access the performance view
Pools	<ul style="list-style-type: none"> • Storage > Pools • Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Pools <p>Available only for pools that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, IBM Spectrum Accelerate, and XIV systems.</p>
Volumes	<ul style="list-style-type: none"> • Storage > Volumes • Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Volumes

Table 44. Internal resources and locations where you access the performance view (continued)

Internal resource	Locations in the GUI where you can access the performance view
Controllers	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Controllers
Disks	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Disks Available only for disks that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.
Host Connections	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Host Connections Available only for host connections that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, and XIV systems.
I/O Group	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click I/O Groups
Modules	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Modules
Nodes	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Nodes
Managed Disks	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Managed Disks Available only for nodes that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.
Ports on storage systems	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Ports
Ports on switches	Network > Switches , and click Ports

Table 44. Internal resources and locations where you access the performance view (continued)

Internal resource	Locations in the GUI where you can access the performance view
RAID Arrays	<p>Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click RAID Arrays</p> <p>Available only for RAID arrays that are associated with System Storage DS series and ESS storage systems.</p>

Related information:



Performance statistics (for SAN Volume Controller systems)

Real-time performance statistics provide short-term status information for the SAN Volume Controller system. The statistics are shown as graphs in the management GUI.



Performance statistics (for Storwize V7000 Unified systems)

Real-time performance statistics provide short-term status information for the Storwize V7000 Unified system. The statistics are shown as graphs in the management GUI.

How performance information is displayed

The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Performance chart

The top section of the performance view shows information about the selected resources. You can view this information in the following formats:



Chart

The chart shows a visual representation of how the performance of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.



Table

The table shows performance and asset information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Performance chart legend

The bottom section of the performance view shows more information about the resources in the chart. The information is formatted into rows

and columns. Each row represents a resource that was selected for the view. Each column provides asset and performance information about a resource.

Each metric on the chart is represented by two columns in the legend. The first column includes the weighted average of the metric for each resource in the selected time range. The second column represents the peak data point of the metric for each resource in the selected time range. For example, if a metric in the chart is Total I/O Rate, the following columns are displayed in the legend: Avg Total I/O Rate, Worst Total I/O Rate. The worst value can be the minimum or the maximum number, depending on the metric. For example, the worst value for an I/O rate metric is a maximum number, while the worst value for the Cache Hit Percentage metric is a minimum number.

You can also view information about the resources that are related to the resources in the legend. For example, if a SAN Volume Controller is listed in the legend, right-click it to view the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage. For more information about the resources that you can view in the chart legend, see “Resources in the chart legend” on page 229.

Tip:

- If you hide a resource in the performance chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

For example, if the current date is January 1, and you view a volume with the time range set to December 1 to December 7, the legend shows the capacity of that volume on December 7. If the capacity of the volume was changed between December 7 and January 1, this historical capacity is different from the current capacity of the volume on January 1.

Related concepts:

“Customizing lists of resources, tasks, and alerts” on page 337

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

Controls:

Each performance view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



View chart

View performance information in a chart format. The chart shows a visual representation of how resource performance trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of

measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of performance information at a specific time.



View table

View performance information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that is available.

For information how to filter, sort, and customize the columns in a table, see “Customizing lists of resources, tasks, and alerts” on page 337.



Metrics

Add metrics to the performance chart. On the Select Chart Metrics window, metrics are organized into three tabs: Volume Metrics, Disk Metrics, and Port Metrics. The metrics and tabs that are available depend on the type of resource that is being shown in the chart. The number next to the name of the tab represents the number of metrics that are currently selected from that tab.

You can select multiple metrics at the same time and from different tabs, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see “Performance metrics” on page 403.



Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.



Specify granularity

Determine the granularity of the data that is shown in a performance view. Granularity determines the points that are shown on the y-axis of a chart, the points on the lines in a chart, and the rows in a table. When you first

view performance information, the default time range is the last 12 hours and the granularity matches the frequency of data collection by the performance monitor (sample).

To change the granularity of chart or table, select one of the following options from the granularity menu:



(daily)

Set the granularity to show 1-day increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 day.



(hourly)

Set the granularity to show one-hour increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 hour.



(Sample)

Set the granularity to match the frequency of data collection by a performance monitor. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments.

This granularity is only available if the time range of the chart is less than 2 weeks.

If the performance monitor collects sample data at 1-minute intervals, data is displayed as 1-minute intervals in the chart only when data is available for all of the selected time range. If data is not available at 1-minute intervals for any portion of the selected time range, the data is displayed in 5-minute intervals.

For example, the performance monitor for a resource collects sample data every 5 minutes. Six hours ago, the performance monitor was changed to start collecting data every minute. You want to view sample data for the resource for the last 12 hours. However, the performance monitor did not collect sample data every minute for 6 of the last 12 hours. Therefore, the data is displayed in 5-minute intervals rather than 1-minute intervals.

Restriction: Each granularity has a maximum time range for displaying data. This time range is dynamically calculated based on the resources and metrics that are shown on the chart, and the data points per line. When the maximum range is exceeded, the granularity icon is disabled and a message indicates that there are too many lines on the chart. To avoid this problem, reduce the time range or the number of resources or metrics that are selected.

Specify a time range

The time range of the performance information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the performance view, the default time range is the last 12 hours.

Ensure that performance data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



Hide and show chart controls

The performance view lists the names of resources that are being shown in the chart and the controls for customizing that chart. Click the left arrow icon to hide the controls and display the chart in the entire window. Click the right arrow icon again to show the controls.



Export information about the chart to a file

Export information on a performance view to a CSV file.



Open the performance view in a separate web browser window

Open a duplicate of the current performance view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original performance view for comparison.



Synchronize the time range across all the open performance views

Synchronize the time range across all the performance views that are displayed in separate browser windows. Use this action when you change the time range in a performance view and want to apply the same time range to the other performance views.

Restriction: This action affects only the performance views that are displayed in separate browser windows. The time range of the performance view that is shown as part of the main window for the GUI is not affected.

For example, you can view the storage system volumes that are assigned to a server. You can then open separate browser windows for the host connections, pools, and managed disks that are related to one of the volumes. If you change the time range in the performance view for host connections, click the synchronize button to apply the same time range to the views for pools and managed disks. The performance view of the original volume is not changed.

Actions for managing the resources in the chart legend

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the **Actions** menu:

View Properties

View key details about a resource, including asset, status, configuration, capacity, and performance information.

View resource Performance

View the performance of resources that are internal or related to a resource in the chart legend. For example, when you view the performance of a SAN Volume Controller, you can right-click it and view the performance of its internal resources, such as disks, volumes, ports, managed disks, and back-end storage.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the performance view in the main window of the GUI.

Restriction: If performance data was not collected for an internal or related resource, that resource is not shown in the view. For example, if you right-click a storage system and select **Volume performance**, only the volumes for which performance data was collected are shown.

For more information about the resources that you can view in the chart legend, see “Resources in the chart legend.”

Key Metrics View

View a set of predefined metrics for resources in the chart legend.

Custom View

Customize the set of metrics that are shown for resources in the chart legend.

Resources in the chart legend:

In the performance view, performance metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, you can open separate performance views for their internal or related resources. The separate performance views use the same time range as the performance view in the main window of the GUI.

For example, when you view the performance of a SAN Volume Controller, you can right-click it in the chart legend to open a performance view for its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.

The type of resource determines the internal and related resources that you can view. The following tables show the internal and related resources that are available when you right-click a resource in the chart legend.

Table 45. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, RAID array, pool, controller, host connection, port	Storage virtualizer, switch

Table 45. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend (continued)

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Volume	Storage system, RAID array, pool, controller, host connection	Managed disk
RAID array	Storage system, volume, pool, controller, host connection	None
Pool	Storage system, volume, RAID array, controller, host connection	None
Controller	Storage system, volume, RAID array, pool, host connection	None
Host connection	Storage system, volume, RAID array, pool, controller	None
Storage system port	Storage system	Switch port

Table 46. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, pool, node, I/O group, disk, managed disk, host connection, port	Back-end storage system, storage virtualizer, switch
Volume	Storage system, pool, node, I/O group, disk, managed disk, host connection	None
Pool	Storage system, volume, node, I/O group, disk, managed disk, host connection	None
Node	Storage system, volume, pool, I/O group, disk, managed disk, host connection, port	None
I/O group	Storage system, volume, pool, node, disk, managed disk, host connection, port	None
Disk	Storage system, volume, pool, I/O group, node, managed disk, host connection	None
Managed disk	Storage system, volume, pool, node, I/O group, disk, host connection	Back-end volume

Table 46. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend (continued)

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Host connection	Storage system, volume, pool, node, I/O group, disk, host connection	None
Storage system port	Storage system, I/O group, node	Switch port

Table 47. Resources that you can view for IBM Spectrum Scale in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	File system, node	None
File system	Storage system, node	None
Node	Storage system, file system	None

Table 48. Resources that you can view for IBM Spectrum Accelerate in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 49. Resources that you can view for XIV systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 49. Resources that you can view for XIV systems in the chart legend (continued)

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Host connection	Storage system, module, pool, volume	None
Storage system port	Storage system, module	Switch port

Table 50. Resources that you can view for IBM Spectrum Accelerate in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 51. Resources that you can view for SMI-S BSP systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, port	Storage virtualizer, switch
Volume	Storage system	Managed disk
Storage system port	Storage system	Switch port

Table 52. Resources that you can view for switches in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Switch	Port	Storage system, switch
Port	Switch	Storage system port, switch port

Table 53. Resources that you can view for servers in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Server	None	Storage system, volume, pool, controller, module, node, I/O group, disk, managed disk, host connection, switch, storage system port, switch port

Saving URLs for performance views

You can create a bookmark or a favorite for a performance view.

The link in the bookmark or favorite stores details of the performance view, such as the resources, metrics, interval, and other details. The link also stores the time period that is specified in the time selector at the top of the chart.

You can also copy the URL from the address bar if you want to share the URL with colleagues.

Do one of the following steps:

- In Firefox, right-click in the performance view and click **Bookmark This Page**.
- In Internet Explorer, right-click in the performance view and click **Add to favorites**.

Identifying performance issues for IBM Spectrum Virtualize storage systems

You can identify key metrics that are outside of a standard range for storage systems that run IBM Spectrum Virtualize by using the performance guidelines.

Before you view the performance of your storage systems that run IBM Spectrum Virtualize, ensure that performance data was collected for those systems during the time ranges that you want to analyze. IBM Spectrum Control uses performance monitors to collect data so that you can measure the performance of storage systems and switches. To verify that performance monitors were run, follow the instructions in Verifying that a performance monitor is running for a resource.

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM System Storage SAN Volume Controller, IBM Storwize, IBM FlashSystem V840 Enterprise Performance Solution, and IBM FlashSystem V9000 storage systems.

Restriction: Performance guidelines are not provided for FlashSystem storage systems or for file storage in Storwize V7000 Unified storage systems.

1. To select a storage system, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems

2. From the list of storage systems that are displayed, right-click the resource that you want to view and click **View Details**. For example, select a SAN Volume

Controller or a Storwize storage system on the Block Storage Systems page, or select a Storwize V7000 Unified storage system on the File Storage Systems page.

3. In the General section of the resource details page, click **Performance**. The key performance charts are displayed with data that was collected from monitored resources over the last 24 hours. The performance of each monitored resource is charted against the best practice guidelines for IBM Spectrum Virtualize storage systems.
4. Optional: If you want to view the performance charts for a period other than the default 24-hour period, change the date for all of the charts. Click **Last 24 hours** and use the calendar to specify a date.
5. Optional: Drill down to view detailed information about a specific I/O group, port, or node. Click a resource from the chart's list.

For example, select the name of an I/O group in the **Read Response Time by I/O Group** chart to view the average number of milliseconds taken by the I/O group to complete a read operation.

By default, metrics are shown on the charts for the following time periods:

- The last 24 hours

- The same 24-hour period for yesterday

- The same 24-hour period for the day seven days before today

If you selected a date from the calendar, the metrics are shown on the charts for the following time periods:

- The selected date

- The day before the selected date

- The day seven days before the selected date

Best practice performance guidelines

You can view the key metrics that are outside of a standard range for storage systems that run IBM Spectrum Virtualize by using the performance guidelines. The guidelines were established by a historical analysis of storage environments.

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM System Storage SAN Volume Controller, IBM Storwize family, IBM FlashSystem V840 Enterprise Performance Solution, and IBM FlashSystem V9000 storage systems.

Restriction: Performance guidelines are not provided for FlashSystem storage systems.

Most of the performance charts show an orange line that indicates the best practice value for the metric. These best practice guidelines are established as the levels that allow for a diverse set of workload characteristics while maintaining a stable performance profile. The other lines on each chart represent the measured values for the metric for the resources on your storage system: I/O groups, ports, or nodes. You can use the lines to compare how close to potentially becoming overloaded your resources are. If your storage system is responding poorly and the charts indicate overloaded resources, you might have to better balance the workload. You can balance the workload between the nodes of the cluster, potentially adding more nodes to the cluster, or move some workload to other storage systems.

The charts show the hourly performance data measured for each resource on the selected day. Use the following charts to compare the workloads on your storage system with the best practice guidelines:

Node Utilization Percentage by Node

Compare the guideline value for this metric, for example, 60% utilization, with the measured value from your system.

Overall Port Bandwidth Percentage by Port

Compare the guideline value for this metric, for example, 50%, with the measured value from your system. Because a cluster can have many ports, the chart shows only the eight ports with the highest average bandwidth over the selected day.

Port-to-Local Node Send Response Time by Node

Compare the guideline value for this metric, for example, 0.6 ms/op, with the measured value from your system.

Port-to-Remote Node Send Response Time by Node

Because latencies for copy-services operations can vary widely, a guideline is not established for this metric. Use this chart to identify any discrepancies between the data rates of different nodes.

Read Response Time by I/O Group

Compare the guideline value for this metric, for example, 15 ms/op, with the measured value from your system.

System CPU Utilization by Node

Compare the guideline value for this metric, for example, 70% utilization, with the measured value from your system.

Total Data Rate by I/O Group

Because data rates can vary widely, a guideline is not established for this metric. Use this chart to identify any significant discrepancies between the data rates of different I/O groups because these discrepancies indicate that the workload is not balanced.

Write Response Time by I/O Group

Compare the guideline value for this metric, for example, 5 ms/op, with the measured value from your system.

Zero Buffer Credit Percentage by Node

Compare the guideline value for this metric, for example, 20%, with the measured value from your system.

Tip: The guidelines are not thresholds, and they are not related to the alerting feature in IBM Spectrum Control. To create performance alerts that use the guidelines as thresholds, go to a resource detail page in the web-based GUI, click **Alerts** in the General section, and then click **Definitions**.

Viewing performance alerts

View the alerts that were generated when the measured value of a performance metric falls outside of the specified threshold for a resource.

To generate performance alerts, you must collect performance data and define performance alerts for that resource. For information about how to collect performance data, see “Creating performance monitors” on page 140. For information about how to define performance alerts, see “Defining alerts for performance metrics” on page 155.

1. Complete one of the following tasks to view performance alerts:

Task	Steps
View the performance alerts for a specific storage system	<ol style="list-style-type: none"> 1. In the menu bar, choose one of the following options: <ul style="list-style-type: none"> • Storage > Block Storage Systems • Storage > File Storage Systems 2. Double-click a storage system. 3. Click Alerts in the General section.
View all the performance alerts that were generated for switches	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Switches. 2. Click the Alerts tab.
View the performance alerts for a specific switch	<ol style="list-style-type: none"> 1. In the menu bar, select Network > Switches. 2. Double-click a switch. 3. Click Alerts in the General section.


2. Right-click the performance alert that you want to view and select **View Alert**. In this view, a chart shows the performance of the resource that violated the threshold and generated the alert. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:

- The blue line represents the performance of the resource for the selected metric.
- The red line represents the threshold that was violated.
- The dotted line represents the time when the violation occurred.

For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:

- The performance threshold was violated on an internal resource other than a port.
- The performance threshold was violated on any of the following storage systems:
 - DS6000
 - DS8000
 - XIV
 - IBM Spectrum Accelerate
 - ESS
 - SAN Volume Controller
 - Storwize V3500
 - Storwize V3700
 - Storwize V7000
 - Storwize V7000 Unified

You can also view the servers or hosts to which those volumes are assigned.

3. Optional: To access the full performance view of a resource that violated the threshold, click the open performance view icon  on the chart.

Exporting performance data for storage systems and fabrics

To help resolve performance issues with storage systems and fabrics, you can export performance data for the resources to a compressed file. If you contact IBM Support to help you analyze the problem, you might be asked to send this file.

You can export performance data for the following types of resource: storage systems, nodes, modules, I/O groups, host connections, pools, RAID arrays, managed disks, disks, volumes, file systems, switches, ports, and inter-switch connections.

The summarization type for the exported data can be sample, hourly, or daily.

The performance data is created as CSV files, one file for each resource type. By default, the files are compressed as .zip files.

Related tasks:

“Exporting performance data for a SAN Volume Controller system” on page 100
In this scenario, you export performance data for a SAN Volume Controller to a compressed package. You then send the package to IBM Support.

Exporting performance data by using the GUI

You can export performance data for a managed resource. If you contact IBM Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

The performance data might be large, especially if the data is for storage systems that have many volumes, or the performance monitors are running with a 1-minute sampling frequency. If the time range for the data is greater than 12 hours, volume data and 1-minute sample data is automatically excluded from the performance data even if it is available. To include volume data and 1-minute sample data, select the **Advanced package** option when you export performance data.

1. In the menu bar, select the type of storage system. For example, to create a compressed file for a block storage system, go to **Storage > Block Storage Systems**. To create a compressed file for a fabric, go to **Network > Fabrics**.
2. Right-click the storage resource, and then click **Export Performance data**.

Exporting performance data by using a script

Use the `exportPerformanceData` script to export performance data for managed resources. If you contact IBM Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

Running the `exportPerformanceData` script

You can run the `exportPerformanceData` script on the server where IBM Spectrum Control is installed. The script is run from the `installation_directory/scripts` directory on the server. You can export the performance data for one or more managed resources.

The `exportPerformanceData` script can be run on Windows and AIX/Linux servers:

- Windows default location: `C:\Program Files\IBM\TPC\scripts\exportPerformanceData.bat`
- AIX/Linux default location: `/opt/IBM/TPC/scripts/exportPerformanceData.sh`
or `/usr/IBM/TPC/scripts/exportPerformanceData.sh`

The following arguments are required when you run the **exportPerformanceData** script:

-user *user_name*

Specifies an IBM Spectrum Control user ID.

-pwd *password*

Specifies the password for the IBM Spectrum Control user ID.

-resNames *resource_name1 resource_name2 resource_name3...*

Specifies the names of storage system, switch, or fabric names as displayed in the UI, separated by spaces. If the resources have spaces in the names, then you must enclose the resource names in double quotation marks ("").

To view the optional parameters that the script uses, run the script with the **-help** option.

Examples

To collect and export performance data to a compressed file for a storage system:

```
>cd c:\Program Files\IBM\TPC\scripts\  
>exportPerformanceData -user scUser1 -pwd scPwd1  
-out "c:\web data\performance reports\Array1.20160402"  
-resTypes storageSystem,Nodes,Pools,Ports  
-start "2016-02-04 00:00:00" -end "2016-02-04 23:59:59"  
-summType hourly -advPkg yes -zip no  
-resNames "Production Array 1"
```

This command creates a directory Array1.20160402 in c:\web data\performance reports\ if the directory does not exist already. The command writes to that directory the CSV files that contain the hourly performance data for the whole storage system, nodes, pools, and ports. The data is for the storage system named Production Array 1. The time range of the data is from Apr 2, 2016 0:00 AM to Apr 2, 2016 11:59 PM. One minute sample and volume data are included where they are available in the specified time range.

Monitoring the capacity of resources

IBM Spectrum Control can collect information about the capacity and space usage of block, file, and object storage resources. This information includes key metrics that can help you measure, identify, and troubleshoot capacity and space usage issues in your storage. You can view capacity metrics for storage resources, including tiers.


You can use the capacity and space usage information to complete the following tasks:

- Measure, compare, and troubleshoot the capacity and space usage of resources such as:
 - Block storage systems and pools
 - Compressed, thin-provisioned, and Easy Tier volumes
 - File systems, pools, and filesets in file storage systems
 - Containers in object storage systems
 - Tiers
- Review the alerts that were triggered when the capacity or space usage of a resource fell outside a specific range.

- Customize views of capacity and space usage so that you can analyze specific resources and metrics during time ranges that you specify.
- View capacity and space usage information in a chart or table format to help you quickly identify where and when issues are occurring. The chart is a visual representation of how the capacity and space usage of resources trend over time.
- Drill down into resources to view detailed information about the capacity and space usage of internal resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the capacity and space usage of its pools.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including capacity and space usage information.
- Export capacity and space usage information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view capacity and space usage information for resources, you must complete the following tasks:

Task	 Learn more
Add storage systems for monitoring by IBM Spectrum Control.	For information about how to add these resources, see “Adding storage systems” on page 115.
Ensure that capacity data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	For information about verifying data collection, see “Collecting data” on page 133.
Define capacity and space usage alerts to be notified if the capacity or space usage of a resource falls outside a specified range and might represent a potential problem. When you define a capacity or space usage alert for a resource, select a specific metric that you want to measure. For example, you can define an alert that notifies you when the total used space for a storage system falls outside a specified range.	For information about how to define capacity and space usage alerts, see “Defining alerts for attribute and capacity changes” on page 151.

Related reference:

“Capacity metrics for block storage systems” on page 469

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

“Capacity metrics for file storage systems” on page 482

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

“Capacity metrics for object storage systems” on page 484

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

“Capacity metrics for tiers” on page 485

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

Viewing capacity information

View the capacity and space usage information of resources such as block storage systems, volumes, and pools. You can view similar information for filesets, file systems, and file system pools as internal resources of file storage systems. You can view information for containers as internal resources of object storage systems. Additionally, you can view capacity and space usage information for resources that violated a specified capacity range and generated an alert.

Before you view the capacity and space usage of resources, ensure that capacity data was collected for those resources during the time ranges that you want to analyze. Data collection is usually scheduled when resources are added for monitoring by IBM Spectrum Control.

You can view the following capacity information:

- “Capacity charts for block storage”
- “Capacity charts for file storage” on page 241
- “Capacity charts for object storage” on page 241
- “Capacity charts for tiers” on page 242

Capacity charts for block storage

Table 54. View capacity charts for block storage

Actions	Navigation
View capacity charts for block storage systems.	<ol style="list-style-type: none">1. From the Storage menu, click Block Storage Systems.2. Right-click one or more storage systems and click View Capacity.
View the capacity charts for all block storage pools.	<ol style="list-style-type: none">1. From the Storage menu, click Pools.2. Right-click one or more pools and click View Capacity.
View the capacity chart for the pools in a storage system.	<ol style="list-style-type: none">1. From the Storage menu, click Block Storage Systems.2. Right-click the storage system and click View Details.3. In the navigation menu, click Pools.4. Right-click one or more pools and click View Capacity.
View the capacity chart for all of the compressed, thin-provisioned, and Easy Tier volumes.	<ol style="list-style-type: none">1. From the Storage menu, click Volumes.2. Right-click one or more volumes that are compressed, or thin-provisioned, or enabled for Easy Tier, and then click View Capacity.

Table 54. View capacity charts for block storage (continued)

Actions	Navigation
View the capacity chart for all of the compressed, thin-provisioned, and Easy Tier volumes in the storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click Block Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click Volumes. 4. Right-click one or more volumes that are compressed, or thin-provisioned, or enabled for Easy Tier, and then click View Capacity.

Capacity charts for file storage

Table 55. View capacity charts for file storage

Actions	Navigation
View capacity charts for all of the filesets in a file storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click File Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click Filesets. 4. Right-click one or more filesets and click View Capacity.
View the capacity charts for all of the file systems in a file storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click File Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click File Systems. 4. Right-click one or more file systems and click View Capacity.
View the capacity charts for all of the file system pools in a file storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click File Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click File System Pools. 4. Right-click one or more file system pools and click View Capacity.

Capacity charts for object storage

Table 56. View capacity charts for object storage

Actions	Navigation
View the capacity charts for all of the containers in an object storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click Object Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click Containers. 4. Right-click one or more containers and click View Capacity.

Capacity charts for tiers

Table 57. View capacity charts for tiers

Actions	Navigation
View the capacity charts for tiers.	<ol style="list-style-type: none">1. From the Groups menu, click Tiers.2. Right-click one or more tiers and click View Capacity.

Related reference:

“Capacity metrics for block storage systems” on page 469

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

“Capacity metrics for file storage systems” on page 482

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

“Capacity metrics for object storage systems” on page 484

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

“Capacity metrics for tiers” on page 485

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

How capacity information is displayed

The capacity view is displayed when you view the capacity of resources, such as block storage systems or pools, including the capacity of resources in alert violations. Information in the capacity view is organized into a chart or table and a legend.

Capacity chart

The top section of the capacity view shows information about the selected resources, such as block storage systems or pools. You can view this information in the following formats:



Chart

The chart shows a visual representation of how the capacity of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.



Table

The table shows capacity information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Capacity chart legend

The bottom section of the capacity view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides capacity and space usage information about a resource.

You can also view information about the resources that are related to the resources in the legend. For example, in the capacity view for block storage systems, if a SAN Volume Controller is listed in the legend, right-click it to view the capacity of its internal resources, such as pools.

Tip:

- If you hide a resource in the capacity chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

For example, if the current date is January 1, and you view a pool with the time range set to December 1 to December 7, the legend shows the capacity of that pool on December 7. If the capacity of the pool was changed between December 7 and January 1, this historical capacity is different from the current capacity of the pool on January 1.

Related reference:

“Capacity data” on page 469

Use the IBM Spectrum Control GUI to view the capacity data that is collected about the storage systems in your environment.

Controls for capacity views:

Each capacity view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



View chart

View capacity information in a chart format. The chart shows a visual representation of how resource capacity trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of capacity information at a specific time.



View table

View capacity information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that are available.

For information how to filter, sort, and customize the columns in a table, see “Customizing lists of resources, tasks, and alerts” on page 337.



Metrics

Add metrics to the capacity chart. On the Select Chart Metrics dialog, the metrics that are available depend on the type of resource that is being shown in the chart. The number next to the name of the metrics type represents the number of metrics that are currently selected.

You can select multiple metrics at the same time, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see “Capacity data” on page 469.



Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.

Specify a time range

The time range of the capacity information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the capacity view, the default time range is the last month.

Ensure that capacity data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



Export information about the chart to a file

Export information on a capacity view to a CSV file.



Open the capacity view in a separate web browser window

Open a duplicate of the current capacity view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original capacity view for comparison.

Actions for managing the resources in the chart legend

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the **Actions** menu:

View Properties

View key details about a resource, including asset, status, configuration, and capacity information.

View resource Capacity

View the capacity of resources that are internal or related to a resource in the chart legend. For example, when you view the capacity of a SAN Volume Controller, you can right-click it and view the capacity of its internal resources, such as pools.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the capacity view in the main window of the GUI.

Resources in the capacity chart legend:

In the capacity view, capacity metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, such as block storage systems, you can open separate capacity views for their internal resources, such as pools or volumes. The separate capacity views use the same time range as the capacity view in the main window of the GUI.

For example, when you view the capacity of a SAN Volume Controller, you can right-click it in the chart legend to open a capacity view for its internal resources, such as pools or volumes.

You can view capacity metrics for the following resources:

- Block storage systems
- Internal resources of block storage systems: volumes and pools
- Internal resources of file storage systems: filesets, file systems, and file system pools
- Internal resources of object storage systems: containers
- Tiers

The following table summarizes the resources and the internal resources for which you can view capacity metrics.

Table 58. Internal resources for which you can view capacity metrics

Resource	Internal resources
Block storage system	<ul style="list-style-type: none">• Volumes• Pools

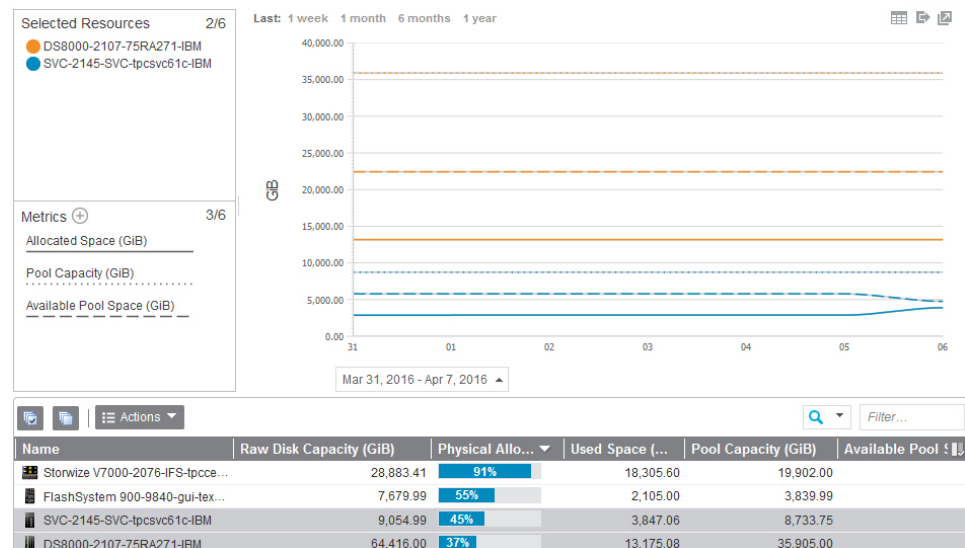
Table 58. Internal resources for which you can view capacity metrics (continued)

Resource	Internal resources
File storage system	<ul style="list-style-type: none"> • Filesets • File systems • File system pools
Object storage system	<ul style="list-style-type: none"> • Containers

Investigating capacity trends for block storage systems

Use the capacity charts and the information that is shown in the tables for the block storage systems to check which storage systems have the highest growth rates and which storage systems might need more capacity.

1. From the **Storage** menu, click **Block Storage Systems**.
2. Click **Capacity** and select the storage systems and the capacity metrics that you want to review.




Tip: To check which pools consume the most storage, right-click the storage systems and click **Pool Capacity**.

Investigating capacity trends for block storage pools

Use the capacity charts and the information that is shown in the tables for the block storage pools to check which storage pools have the highest growth rates, which storage pools might need more capacity, and which storage pools are at risk because too much capacity is committed to the thin-provisioned volumes.

1. From the **Storage** menu, click **Block Storage Systems**. Alternatively, you can click **Storage > Pools**.
2. Right-click the storage system and click **View Details**.
3. In the navigation pane, click **Pools**.
4. Click **Capacity** and select the pools and the capacity metrics that you want to review.




Tip: To investigate the pools with the highest shortfall rates, right-click the pools in the table and click **Volume Capacity**.

Investigating capacity trends for volumes

Use the capacity charts and the information that is shown in the tables to check the rate at which the thin-provisioned volumes consume capacity, the compression savings for the volumes that were converted to compressed volumes, and the changes in the distribution of Easy Tier volumes across the SSD, Enterprise HDD, and Nearline HDD drives.

1. From the **Resources** menu, click **Block Storage Systems**.
2. Right-click the storage system and click **View Details**.
3. In the navigation pane, click **Volumes**.
4. Click **Capacity** and select the volumes and the capacity metrics that you want to review.



Tip: To check the distribution of Easy Tier volumes, click **Select Chart Metrics**




and add the Easy Tier metrics to the chart.

Investigating capacity trends for file systems

Use the capacity charts and the information that is shown in the tables for the file systems to check which file systems have the highest growth rates and which file systems might require additional capacity.

1. From the **Storage** menu, click **File Storage Systems**.
2. Right-click the storage system and click **View Details**.
3. In the navigation menu, click **File Systems**.
4. Click **Capacity** and select the file systems and the capacity metrics that you want to review.



Tip: To check which pools consume the most storage, right-click the file systems and click **File System Pool Capacity**.

Investigating capacity trends for file system pools

Use the capacity charts and the information that is shown in the tables for the file system pools in your storage systems to check which storage pools have the highest growth rates so that you can plan future storage requirements.


1. From the **Storage** menu, click **File Storage Systems**.
2. Right-click the storage system and click **View Details**.
3. In the navigation menu, click **File System Pools**.
4. Click **Capacity** and select the pools and the capacity metrics that you want to review.



Investigating capacity trends for filesets

Use the capacity charts to check which filesets have the highest growth rates in your file storage systems.


1. From the **Storage** menu, click **File Storage Systems**.
2. Right-click the storage system and click **View Details**.
3. In the navigation menu, click **Filesets**.
4. Click **Capacity** and select the filesets that you want to review.



Investigating capacity trends for containers

Use the capacity charts to check which containers have the highest growth rates in your object storage systems.

1. From the **Storage** menu, click **Object Storage Systems**.
2. Right-click the storage system and click **View Details**.
3. In the navigation menu, click **Objects**.
4. Click **Capacity** and select the objects that you want to review.



Investigating capacity trends for tiers

Use the capacity charts and the information that is shown in the tables to check which tiers have the highest growth rates and which tiers might require more capacity.

1. From the **Groups** menu, click **Tiers**.
2. Click **Capacity**.

Creating bookmarks for URLs of capacity views

You can create a bookmark for the URL of a capacity view. You can also open a duplicate of the capacity view to change and compare views.

The bookmark stores details of the capacity view, such as the resources, metrics, interval, and other details. It also stores the time period that is specified in the time selector at the top of the chart.

To create a bookmark for a capacity view, follow these steps:

1. Open the capacity view for which you want to create a bookmark.
2. Right-click and click the relevant option for creating a bookmark.

Tip: To open a duplicate of the current capacity view, click the icon for opening a new browser window. You can change the information that is displayed in the new window while retaining the original capacity view for comparison.

Viewing the capacity of external storage

View the used space and capacity of external pool storage that is used by file systems in IBM Spectrum Scale. External pools can include storage that is provided by IBM Cloud Object Storage, Amazon Simple Storage Service (S3), OpenStack Swift, IBM Spectrum Archive™, IBM Spectrum Protect, and other storage providers.


For all external storage that is being used by IBM Spectrum Scale file systems, you can complete the following tasks:

- View used space to understand how much data is being migrated from those file systems to external storage.
- Charge back storage costs to the departments within your organization.
- Gauge shortfall in case you need to recall data from external pools.

For external storage that is provided by IBM Cloud Object Storage, you can also view capacity information, including the percentage of space (active and inactive) that is being used. Use this information to complete the following tasks:

- View how storage space is distributed between internal and external storage tiers.
- Identify the external pools that are running out of space.
- Determine when you need to reconcile files between an IBM Spectrum Scale and an external storage tier. For example, when you remove files from a file system, no policy is available to automatically remove the cloud objects that are created. This situation might cause orphan objects (inactive data) on the cloud. To clean up inactive data and keep the cloud in sync with the file system, you can reconcile files.

The following figure illustrates the relationship between IBM Spectrum Scale and the IBM Cloud Object Storage system that provides storage for external pools:



To view the used space and capacity of external storage, complete the following steps:

1. From the **Storage** menu, click **File Storage Systems**.
2. Right-click a IBM Spectrum Scale storage system and click **View Details**.
3. In the navigation pane, click **Pools**.
4. View the following columns for information about the used space and capacity of external pools:

Capacity (%)

The percentage of used space in the file system pool. Hover over the bar for a specific pool to view statistics about its used and available space.

For external pools that are connected to cloud services or storage providers, only the amount of used space is shown. Cloud services and other storage providers might include IBM Spectrum Protect, IBM Spectrum Archive, Amazon Simple Storage Service (S3), or OpenStack Swift.

For external pools that are connected to vaults in IBM Cloud Object Storage, the percentage of used space includes active and inactive data. Space values are determined from the usable vault capacity that is associated with the cloud account and the hard quota value (if a hard quota is configured). Hover over the bar for a specific pool to view statistics about its used active space, used inactive space, and available space.

Active and inactive data in external pools

Active data is data that has corresponding stub files on the GPFS file system. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system. Data is classified as inactive under these conditions:

- When you migrate a file off a GPFS file system cluster, its data is still visible and accessible through its stub file. When the stub file is deleted, the data in the external pool is no longer accessible and is considered inactive.
- When you modify a migrated file, it's recalled from the external pool and modified locally. The previous data for the file remains in the external pool but is no longer accessible through the GPFS file system and is considered inactive.

Tip: To keep the files in sync between an external pool and a GPFS file system and clean up inactive data, you can run the IBM Spectrum Scale `mmcloudgateway files reconcile` command. For more information, see [Reconciling files](#) between IBM Spectrum Scale file system and cloud storage tier.

External Used Space

The amount of used space on an external pool. This value includes active data and inactive data. Active data is data that has corresponding stub files on the GPFS file system. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system cluster.

Available for: External pools that are provided by IBM Cloud Object Storage.

Inactive Used Space

The amount of inactive used space on an external pool. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system cluster. Available for external pools that are provided by Cloud Object Storage only.

Available for: External pools that are provided by IBM Cloud Object Storage.

Used Space

The amount of space in the file system pool that is being used. For external pools, used space includes active data only. Active data is data that has corresponding stub files on the GPFS file system.

Available for: All storage systems.

Related information:

[Introduction to Transparent Cloud Tiering for IBM Spectrum Scale](#)

Identifying shortfall before data is recalled from external storage

View the percentage of migrated data in external pools that does not fit into the available space on an IBM Spectrum Scale file system.

This percentage, called *shortfall*, represents the relative risk of not having enough internal file system space if you recall all data from the associated external pools. The higher the percentage, the higher the amount of migrated data that does not fit.

To view the potential shortfall for a file system, complete the following steps:

1. From the **Storage** menu, click **File Storage Systems**.
2. Right-click a IBM Spectrum Scale storage system and click **View Details**.
3. Under **Internal Resources**, click **File Systems**.
4. Locate the file system that you want to check and view the value in the Shortfall column.

To calculate shortfall, the following formula is used:

$$[(\text{External Pool Migrated Data} - \text{Available Space}) \div \text{External Pool Migrated Data}] \times 100$$

For example, if the available space on a file system is 200 GiB, but the used space of migrated data on external pools is 500 GiB, the shortfall percentage is 60% (300 GiB). In this case, consider adding capacity to the file system before recalling data.

$$[(500 - 200) \div 500] \times 100 = 60\%$$

If the file system has sufficient available space to contain all the migrated data in external pools, no shortfall exists. For example, if the available space on a file system is 500 GiB, and the used space in external pools is only 200 GiB, the shortfall is 0%.

Tip: Only the space that is used by migrated data in external pools is included in the calculation for shortfall. Pre-migrated data is not included because it exists in both the internal and external pools.

Viewing capacity alerts and violations

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

To view alert violations for the capacity of a resource, you must collect capacity data and define capacity alerts for that resource.

1. For block storage systems, go to **Storage > Block Storage Systems**. For file storage systems, go to **Storage > File Storage Systems**.
2. Right-click a storage system and click **View Details**.
3. In the navigation menu, click **Alerts**.
4. Right-click the resource alert that you want to view, such as an alert which was defined for a capacity metric for the resource, and click **View Alert**.

Related concepts:

“Collecting data” on page 133

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

“Alerting and event processing” on page 148

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Related tasks:

“Defining alerts for attribute and capacity changes” on page 151

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used space, total space for disk, file systems, volumes, and pools, and reserved space.

Monitoring and administering applications, departments, and general groups

You create applications to monitor the storage capacity, space usage, and performance of applications, and you create departments to monitor the space usage of the applications in the department. You can structure your applications and departments hierarchically to match the structure of your business organization.

You can map the application to the storage resources that the application uses to do capacity trending, health monitoring, and performance troubleshooting tasks. For example, a department might use 15 applications and be part of another five departments. Also a department might share storage resources with another department, subdepartment or application even if they do not belong in the same business hierarchy. Another example, might be if you are contacted because an application within a department is experiencing a performance issue. You carry out tasks like capacity reporting or trending on the behalf of the application.

Create general groups to quickly view information about storage resources that have common characteristics. For example, you might group the subset of ports on a SAN Volume Controller that are used for inter-node communication, or all the storage systems with lease agreements that end in the current year, or the storage systems that are used by a critical business application.

Application administration

Go to the pages for creating and administering the applications that you want to monitor.

Table 59. Create and administer applications and subcomponents

Actions	Navigation
Create applications and assign storage resources directly to the application.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. On the Applications page, click Create Application. 3. On the Select a Method page, click Assign storage resources to the application.
Create subcomponents that can have its own storage resources.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. On the Applications page, click Create Application. 3. On the Select a Method page, click Create a subcomponent that can have its own set of storage resources.
Create a filter to use in assigning resources.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. On the Applications page, right-click a row in the Applications list and click View Details. 3. In the Members section, click Filters. 4. On the Filters page, click Create Filter.
Add applications as subcomponents to an application.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. Right-click the applications that you want to add as subcomponents, and then click Add to.. > Application. 3. On the Add to Application page, select the applications and click Save.

Table 59. Create and administer applications and subcomponents (continued)

Actions	Navigation
Add applications to departments.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. On the Applications page, select one or more applications, right-click and then click Add to.. > Department. 3. On the Add to Department page, select the applications that you want to add as members and click Save.
Add applications to subdepartments.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartment, and click View Details. 5. In the General section, click Applications. 6. Click Create Application.
Remove applications and subcomponents that you do not want to monitor.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. On the Applications page, select one or more applications, and click Delete. 3. If you select Delete subcomponents, the subcomponents that belong only to that application are removed.
Remove subcomponents within direct context of an application that you do not want to monitor.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. Right-click the application and click View Details. 3. In the General section, click Subcomponents. 4. Right-click the subcomponent, and click Remove from Application.
Remove selected resources from applications or subcomponents	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. Right-click the application and click View Details. 3. In the Related Resources section, click an eligible resource. 4. Optional: In the Members section, click the storage resource that is part of the application. 5. On the resource page, right-click the resource, and click Remove from Application.

Table 59. Create and administer applications and subcomponents (continued)

Actions	Navigation
Troubleshoot applications	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. Right-click the application, and then click View Details. <p>Investigate the performance of the storage resources that are associated with the application by viewing the charts on the Overview page. You can investigate the performance of related resources by clicking the resource and clicking the Performance tab.</p>
Troubleshoot subcomponents	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Applications. 2. Right-click the application, and then click View Details. 3. In the General section, click Subcomponents. 4. Right-click the subcomponent, and click View Details. <p>Investigate the performance of the storage resources that are associated with the subcomponent by viewing the charts on the Overview page. You can investigate the performance of related resources by clicking the resource and clicking the Performance tab.</p>

Department administration

Go to the pages for creating and administering the departments that you want to monitor.

Table 60. Create and administer departments and subdepartments

Actions	Navigation
Create departments and associate members (applications and departments) to create a hierarchal structure.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. On the Departments page, click Create Department. 3. On the Add Members page, click Add one or more applications or Create subdepartment.
Add departments as subdepartments to a department.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. On the Departments page, right-click the departments that you want to add as subdepartments, and then click Add to Department.

Table 60. Create and administer departments and subdepartments (continued)

Actions	Navigation
Add applications to departments.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Applications. 4. Click Create Application.
Add applications to subdepartments.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartment, and click View Details. 5. In the General section, click Applications. 6. Click Create Application.
Remove departments, subdepartments, and applications, that are members of the department.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. On the Departments page, right-click one or more departments and click Delete. 3. If you select Delete subdepartments and applications, the subdepartments and applications that belong only to the department, that is being removed, are removed.
Remove applications from a department or subdepartment.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartment, and click View Details. 5. In the General section, click Applications. 6. Select and right-click the applications, and then click Delete or Remove from Department.
Remove subdepartments within direct context of a department that you do not want to monitor.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartments that you want to remove, and then click Remove from Department.

General group administration

Administer the general groups that you use to organize storage resources with common characteristics.

Table 61. Administering general groups

Actions	Navigation
Create a general group and add resources.	<ol style="list-style-type: none"> 1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add file storage systems, click Storage > File Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Click Add to new group, specify a name and description for the new group, and click Save.
Add resources to general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add block storage systems, click Storage > Block Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the resources to one or more existing groups, click Add to existing groups, click the appropriate groups in the list, and click Save. b. To add the resources to a new group, click Add to new group, specify a name and description for the new group, and click Save.
Define and edit the alert definitions for a general group. Learn more	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click Edit Alert Definitions.
Edit the alert notification settings for general groups. Learn more	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Edit Alert Notification Settings.
Delete general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Delete. 3. To delete subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>

Table 61. Administering general groups (continued)

Actions	Navigation
Remove resources from a general group.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click View Details. 3. In the Members section, click the type of resource that you want to remove. For example, if you want to remove some of the volumes from the group, click Volumes. 4. Right-click the resources that you want to remove and click Remove from General Group. 5. Click OK to confirm that you want to remove the resources from the general group and from all of its subgroups.
Add general groups as subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the groups as subgroups of existing groups, click Add to existing groups, click one or more groups in the list, and click Save. b. To add the groups as subgroups of a new group, click Add to new group, specify a name and description for the new group, and click Save.
Remove subgroups from a group hierarchy.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click View Details. 3. Click Subgroups in the General section. 4. Right-click the general groups that you want to remove and click Remove from General Group. <p>When you remove a subgroup from its parent group, the subgroup is moved to the same level in the hierarchy as the parent group. The subgroup is still monitored by the product.</p>
Create a subgroup.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click the general group that you want to create a subgroup for and click View Details. 3. Click Subgroups in the General section. 4. Click Create Subgroup. 5. Specify a name and description for the new group. 6. Customize the icon for the group by clicking the existing icon and selecting another icon. 7. Click Create.

Table 61. Administering general groups (continued)

Actions	Navigation
Delete subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click the general group that contains the subgroups that you want to delete and click View Details. 3. Click Subgroups in the General section. 4. Right-click one or more subgroups and click Delete. 5. To delete further subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>

Related tasks:

“Monitoring capacity usage at different levels of a business hierarchy” on page 91
The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

“Comparing storage usage in each department” on page 93

In this scenario, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

“Using applications and subcomponents to monitor capacity and space usage” on page 95

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchal manner.

“Viewing storage capacity and usage trends” on page 96

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

Applications

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

An application can be part of a department and have subcomponents that can contain a five level deep hierarchy. Applications that are grouped together can range from large line of business systems to specialized software, in a department that runs on either client computers or servers. For example, an application might be an automated billing system in the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

Another example might be an application named *Database* and the database contains logs and data, and a best practice is to place the logs and data on separate


volumes. You might create subcomponents for the logs and data and these subcomponents become part of the *Database* application.

A storage administrator can use applications for planning purposes to diagnose how much storage a particular application is consuming in relation to the other applications in the business hierarchy. Another applicable usage for applications might be troubleshooting performance issues. The administrator checks the related resources that make up the application to see if there is a health status issue or can examine the performance of volumes that are used by the application, and, then troubleshoot along a particular data path.

The Applications page shows capacity information about the applications that are monitored and the actions that you can use to view and manage them. To see more detailed information, view properties and details for the application. If the application is associated with a department, information about that department is also shown.

The following information shows the resources and objects that can be added to the application.

Table 62. Resources and objects that can be added to the application

Object	Resources or objects that can be added to the application
 Application	<ul style="list-style-type: none">• Applications• Servers• Hypervisors• Volumes• Data Stores• Shares• File Systems (Only server file systems that are monitored through a Storage Resource Agent.)• Filesets• Volume Groups• Virtual Machines (Can only be added with the CLI modifyappgroup command)

Related tasks:

“Defining alerts for applications” on page 167

Define alerts for changes in the configuration, attributes, and performance of the servers, volumes, filesets, and shares in your application.

Creating an application hierarchy

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

In this scenario, you want to monitor the capacity, space usage, and performance of a DB2 application.

The DB2 application has two subcomponents, which you want to monitor separately:

- logs
- data

To monitor the DB2 applications and the data and logs subcomponents, complete these tasks:

- Create an application that is named DB2.
- Create applications that are named data and logs
- Associate storage resources with the data and logs applications. For example, in this case, associate the volumes that the data application uses, type `data_vol*`. For the logs application, type `logs_vol*` in the **Volumes** field of **A selection of volumes that match a pattern** filter. Then click **Belonging to** and click **Storage System**. Type the name of the storage system or the pattern that matches the name of the storage system.
- Add the data application and the logs application as subcomponents to the DB2 application.

Tip:

When you create complex application models for applications, it is not necessary to create a filter to associate storage resources with the application because the application inherits the storage resources that are associated with its subcomponents. However, if you want to monitor storage resources that are not associated with the subcomponents, you must create a filter that associates the additional storage resources that you want to monitor with the application.

You can now view information about the capacity, space usage, and performance of the storage resources on the Overview page for the application and subcomponents.

You can create departments and add applications to the departments.

Related concepts:

“Application and department hierarchies” on page 277

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

Creating applications:

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

Use the GUI to create applications, assign resources, and subcomponents for monitoring by IBM Spectrum Control.

To determine the storage capacity usage by the storage resources that the applications use, you create a data model of the applications for your storage environment.

To create an application, complete these steps:

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, click **Create Application**.
3. On the Create Application page, enter information about the application.

Note: You cannot have a duplicate application name on the same level in your application hierarchy. For example, you create an application named *Finance* and create a subcomponent named *Business*. You can create another application

named *Business* as a top level application, but you cannot add the new *Business* application as a subcomponent of the *Finance* application.

You can customize the icon for your application by clicking the existing icon on the Create Application page and selecting another icon from the Customize Icon page.

4. Click **Create**.
5. Optional: On the Select a Method page, select **Assign storage resources to the application that was created** and follow the instructions to add the resources directly to the application.
6. Click **Preview** to see a list of filter matches before you finalize your filtering criteria, otherwise, click **Save**.
7. Optional: Select **Create a subcomponent that can have its own set of storage resources** and follow the instructions to create a subcomponent of the application.
8. Click **Create**.

When the application creation is created, you can view status information, block and file capacity and additional data on the Applications page and in properties. For a detailed view of the application and related resources, on the Applications page, right-click the application and select **View Details**. To see the subcomponents for the application, click **Subcomponent** on the application details page.

Related tasks:

“Adding subcomponents” on page 267

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

“Using filters to assign resources to applications”

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

“Viewing information about applications” on page 268

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Using filters to assign resources to applications:

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

When you create a filter to assign storage resources to applications, the resources that match the criteria that are defined in the filter are automatically added to the application or subcomponent (member application).

The use of wildcard characters enables you to customize and expand your search to add related resources. The asterisk (*) and the question mark (?) can be used as wildcard characters. The asterisk matches one or more characters. The question mark matches a single character. Use a comma-separated list to add multiple resources. As new resources are added or modified and they match the existing filter, they automatically become part of the application. For example, a new server is added, a volume is renamed or a share is mounted on a server.

If you edit a filter, the resources that no longer match the filter are removed from the application and the resources that now match the filter are added to the application.

If you remove the filter, all storage resources that match the filter criteria are removed from the application.

Filtering is extremely useful in storage environments where resources are predefined and named in a consistent format. Instead of adding each resource separately, you create filtering criteria that matches your naming format in your environment. You can use filtering when you create a new application or directly from the details page for an existing application.

You can only select one resource type filter at a time during the creation workflow, but you can create multiple filters for applications.

The filters are divided into the following categories:

All storage that belongs to a set of servers/hypervisors

You can add servers or hypervisors to the application or subcomponent (member application).

For example, in the **Define Name Patterns** section, enter *blade, in the **Server/Hypervisor** field and click **Save**. The filter adds all known servers or hypervisors that have a name that ends with blade to the application.

If you enter server???.abc.com, in the **Server/Hypervisor** field, the filter adds all servers, in the domain with names that begin with server followed by three characters and ends with .abc.com to the application. If you enter, accounting*.*, the filter adds all servers with names that begin with accounting regardless of the domain the system belongs to.

If you enter cs-blade47.lab.abc.com, cs-blade49.lab.abc.com or cs-blade47*, cs-blade49*, the filter adds multiple servers to the application.

A selection of volumes that match a pattern

You can add volumes of a certain name format, volumes that belong to a server or hypervisor, volumes from a specific storage system or volumes from a pool to the application or subcomponent. You can only select one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter *vol, in the **Volumes** field, mark the **Belonging to** box, click **Server/Hypervisor**, enter *.lab.abc.com and click **Save**. The filter adds all volumes that have a name that ends with vol and are mapped to servers that have a name that ends with .lab.abc.com to the application.

If you enter *vol, in the **Volumes** field, mark the **Belonging to** box, click **Storage System**, and enter ds8000. The filter adds all volumes that have a name that ends with vol that belong to the ds8000 storage system.

A selection of shares that match a pattern

You can add shares and exports that are mounted on a server or hypervisor, or shares and exports that are exported from a specific storage system to the application or subcomponent. You can select only one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter *share in the **Shares** field, mark the **Belonging to** box, click **Server/Hypervisor**, enter

*.lab.abc.com and click **Save**. The filter adds shares that have a name that ends with share that are mounted on servers that have a name that ends with .lab.abc.com .

If you enter share3 in the **Shares** field, mark the **Belonging to** box, click **Storage System** and enter SONAS, the filter adds all the shares that are named share3 and are exported from the SONAS storage system.

A selection of filesets that match a pattern

You can add filesets that are associated with file systems, from a storage system to the application or subcomponent. You can select only one resource option per filter creation.


For example, in the **Define Name Patterns** section, if you enter * in the **Filesets** field, mark the **Belonging to** box, click **Storage System**, enter SONAS* and click **Save**, the filter adds all the filesets that are exposed by storage systems that have names that begin with SONAS to the application.

If you enter * in the **Filesets** field, mark the **Belonging to** box, click **File System** and enter fs1, the filter adds all the filesets that are associated with the file systems named fs1 to the application.

To create a filter, complete these steps:

1. In the menu bar, go to **Groups>Applications**.
2. On the Applications page, right-click a row in the Applications list and click **View Details**.
3. In the **General** section, click **Filters**.
4. On the Filters page, click **Create Filter**.
5. On the Create Filter page, select a resource type filter and enter the necessary information.
6. Click **Preview** to see a list of filter matches before you finalize your filtering criteria.
7. On the Filter Matches page, click **Save** to create the filter.

The **General** section of the details page shows the number of filters that you created, for example, **Filters (6)**. Click **Filters** to access the Filters page and see the filters that were created. For example, The number of matches: 26, which indicates there are 26 matches for the filter criteria. Click the link to see the resources for the filter matches that are available.

Tip: To view a description of a field on the Create Filter page, hover the mouse pointer over the help icon  for that field.

Related tasks:

“Creating applications” on page 263

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

“Adding subcomponents” on page 267

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Adding subcomponents:

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

To create a subcomponent for an existing application, complete these steps:

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, select one or more applications, right-click and then, click **Add to.. > Application**.
3. On the Add to Application page, choose the applications that you want to become subcomponents of your selected applications. The **Hierarchy** column shows the path of only the ancestor applications; it does not show departments.
4. Click **Save**.
5. Optional: On the details page for the application, click **Subcomponents**.
6. On the Subcomponents page, click **Create Subcomponent**.
7. On the Create Subcomponent page, enter information including name, description, type, and subtype.
8. Click **Create**.

After the subcomponent is created, you can view the results on the Subcomponent page and in the properties. For a more detailed view of the subcomponent, select the subcomponent from the Subcomponent page and click **Actions > View Details**.

Adding resources manually to applications

You can add resources directly from a resource page to your applications.

Only the following resources can be added directly to applications:

- Servers
- Hypervisors
- Volumes
- Volume groups
- Data stores
- Shares
- File systems (Only server file systems that are monitored through a Storage Resource Agent.)
- Filesets

To add a resource directly to an application, complete these steps:

1. In the menu bar, go to **Groups > Servers**.
2. On the Servers page, select a server from the Servers list, right-click and click **Add to Application**.
3. On the Add to Application page, select one or more applications to assign the server to.
4. Click **Save**.

The resource is added to the application and is displayed on the Applications page.

Related tasks:

“Using filters to assign resources to applications” on page 264

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name

pattern matching.

Adding resources with the command line interface

You can use the **modifyappgroupviafile** command-line interface (CLI) command to automate your bulk assignment of eligible storage resources to applications. To add a single resource to an application, use the **modifyappgroup** command.

Related information:

 [modifyappgroupviafile command](#)

 [modifyappgroup command](#)

Viewing information about applications


You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

To view the details of an application, use these steps:

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, right-click an application and select **View Details**.

The details page shows information about the specific application. The information is available from the following sections on the page:

- **Overview:** The overall capacity trending about the application and performance information. An interactive breadcrumb that shows the hierarchy of applications and departments and a list of all the top level applications.
- **Properties:** The key information about the application.
- **Subcomponents:** The number of subcomponents that are a members of the selected application. If there is more than one subcomponent, the number of subcomponents is displayed.
- **Members:** The information on the resource types that were added directly to the application or subcomponent. The number next to the resource type, for example, Servers (2), shows the number of resources of that type in the application.
- **Filters:** The number of filters created and the filtering criteria used for the resources that were added to the application.
- **Related Resources:** The information about the resources that are related to the application.

Tip: To view descriptions of the information that is available on an application details page, click the **Help** icon  in the upper-right corner of the page.

Related concepts:

“Application and department hierarchies” on page 277

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

Related tasks:

“Creating an application hierarchy” on page 262

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

Removing applications and subcomponents

You can remove applications and their subcomponents. The data that was collected to monitor the applications is also removed from the database repository.

To remove applications and subcomponents from a level of the business hierarchy, complete the following steps:

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, select one or more applications, and click **Delete**.
 - a. If you select **Delete Subcomponents**, the subcomponents that belong only to the application are removed.
 - b. If you do not select **Delete Subcomponents**, any subcomponents of the application, that is being removed, then become members of the parent application or department. If the application is a top level application, and, the subcomponents are not members of any other application, the subcomponents become top level applications.

The applications and subcomponents are removed and are no longer displayed on the Applications page.

To remove subcomponents from within the direct context of the application, use these steps:

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, select an application from the Application list and right-click **View Details**.
3. Optional: Select a application from the Applications page and right-click **View Properties**.
 - a. Click the **Subcomponents** tab and continue with steps 5-6.
4. On the application details page, click **Subcomponents** in the General section.
5. On the Subcomponents page, select the subcomponent and right-click **Remove from Application**.
6. Confirm that you want to remove the subcomponent and click **OK**.

Removing resources from applications

Related resources within the context of an application can be removed from applications. The action is only available when viewing the related resources that belong to an application on the details page.

The following resources can be removed from an application:

- Volumes
- Shares
- Servers
- File systems (Only server file systems that are monitored through a Storage Resource Agent.)
- Filesets
- Volume Groups
- Hypervisors
- Data Stores

Resources that were added through a filter match are removed if the filter is modified and no longer includes the resource that was removed from the application, the filter is removed from IBM Spectrum Control, or the resource is removed from IBM Spectrum Control. Resources can be added back to the application if the filter is modified and includes the resource that was originally removed, a new filter is created that matches the resource, the resources are added to IBM Spectrum Control, or you use the **Add to Application** action.

For example, a server named, abc.xyz.com is added to an application named, ApplicationABC because it matched the server name filter, abc*. You decide to remove the abc.xyz.com server from the ApplicationABC application and remove the abc.xyz.com server from IBM Spectrum Control. Later you want to add the abc.xyz.com server back to IBM Spectrum Control. The abc.xyz.com server is added to the ApplicationABC application because the name matches the server resource filter that you originally created for the ApplicationABC application. If you do not want the resource added back to the application, you can further refine your filters. Click the **View Filters** link on the confirmation dialog to make any adjustments to the filtering criteria before creation.

Resources that are automatically added to an application because they are related to resources that were explicitly added to the application *are not* removed from the application. The resources are automatically removed when the corresponding related resource that was added to the application, is also removed from the application.

For example, you add a server to the application that has one volume from a storage system. The volume is added as a related resource of that server in the application. If you attempt to remove the volume from the application, the volume is *not* removed since it is a related resource. The volume is only removed when the server, that was explicitly added to the application, is removed.

To remove resources directly from an application, use these steps

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, select an application from the Application list and right-click **View Details**.
3. On the application details page, click an eligible resource in the **Related Resource** section.
4. On the resource page, select the resource from the list, and right-click **Remove from Application**.
5. Confirm that you want to remove the resource.
6. Click **Remove**.

The resources that were added to the application are removed from the application.

Related tasks:

“Using filters to assign resources to applications” on page 264

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

Departments

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.


A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or application even if they do not belong in the same hierarchy. For

example, in a collaboration scenario, a single IBM SAN Volume Controller might be shared by multiple departments. You can provide reports on the storage capacity that is used at different levels of your business hierarchy to regulate and compare the capacity storage, and in turn, track the trending of storage usage by specific departments.

The Departments page shows storage capacity information about the top level departments, the subdepartments and any applications that belong to the department that are monitored by IBM Spectrum Control and the actions that you can use to view and manage the resources that are defined to them. To see more detailed information, view properties and details for the department.

The following information shows the objects that you can add using the Add to Department function.

Table 63. Objects that can be added to the department

Object	Objects that can be added to the department
 Department	<ul style="list-style-type: none"> • Applications • Departments

Creating a department hierarchy

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

In this scenario, you want to monitor the capacity and space usage of departments in a business organization that sells books. You also want to monitor the performance of the storage resources that are associated with the applications or application subcomponents that are used by the departments.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online

To monitor the storage resources that the Books Sales department uses and that each of its subdepartments use, complete these tasks:

-
- Create a department for Book Sales.
- Create the Wholesale, Retail, and Online subdepartments that you want to add to the department:
 - Create the subdepartments as departments.
 - Add the departments as subdepartments to the department.
 - Add the applications to the subdepartments.

You can now view charts and information about the capacity and space usage of the department and each of the subdepartments on the Departments page and the Subdepartments page. You can also view charts and information about the performance of the storage resources that are associated with the applications or application subcomponents that the departments and subdepartments use.

Tip: In the **Subdepartments** column for the department on the Departments page, the name of the subdepartment is shown. If you add two or more departments to a subdepartment, the number of subdepartments is shown. To view information about the subdepartments, click the name of the subdepartment or the number.

You can view the total capacity and space usage for the department and the capacity and space usage for each of the subdepartments. You can add to your department hierarchy by adding more applications and departments.

Creating departments:

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

Use the web-based GUI to create departments, subdepartments (member departments) and add applications for monitoring different levels of a business hierarchy.

To create a department, complete these steps:

1. In the menu bar, go to **Groups > Departments**.
2. On the Departments page, click **Create Department**.
3. On the Create Department page, enter information about the department. You cannot have a duplicate department name on the same level in your department hierarchy. For example, you create a department named *Finance* and create a subdepartment named *Business*. You can create another department named *Business* as a top level department, but if you want to make your new *Business* department a subdepartment of the *Finance* department, you see a message that indicates there is already a department named *Business* that is a member of the *Finance* department; your department creation is denied.
You can customize the icon for your department by clicking the existing icon on the Create Department page and selecting another icon from the Customize Icon page.
4. Click **Create** and then **Close**.
5. Optional: On the Add Members page, click **Add one or more applications**.
 - a. On the populate page, select the applications that to add to your department. If you do not want to add any of the applications that are listed, click **Create Application** to create a new application to add to the department.
 - b. Click **Finish**.
6. Optional: On the Add Members page, click **Create subdepartment**.
 - a. On the Create Subdepartment page, enter your information.
 - b. Click **Create**.

To view the newly created department, go to the Departments page.

To see the applications that you added to your department, see the details page and click **Applications**.

To see the subdepartment you created, see the Subdepartments column on the Departments page or see the details page and click **Subdepartments**.

Continue to add to your department hierarchy by adding applications and departments.

Related tasks:

“Adding and creating subdepartments”

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

“Adding applications to departments” on page 274

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

Adding and creating subdepartments:

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

To add a subdepartment to the department level hierarchy, use these steps:

1. In the menu bar, go to **Groups > Departments**.
2. On the Departments page, select one or more departments, right-click and click **Add to Department**.
3. On the Add to Department page, select one or more departments from the list and click **Save**. The department or departments you select from the Add to Department page become parents of the selected department. The Hierarchy column shows the longest ancestry path of the department. You can also view the applications that are members of the department on the Departments page.
4. Optional: Create a new subdepartment:
 - a. Click **Create Department** to open the Create Department page.
 - b. Enter information about the department.
 - c. Click **Create**.
 - d. On the Add Members page, click **Create a subdepartment**.
 - e. On the Create Subdepartment page, enter your information.
 - f. Click **Create**.
5. Optional: Create a new subdepartment from the details page of the existing department:
 - a. In the **General** section, click **Subdepartments** to open the Subdepartments page.
 - a. Click **Create Subdepartment** to open the Create Subdepartment page.
 - b. Enter information about the subdepartment.
 - c. Click **Create**.

To see the subdepartment you created, see the Subdepartments column on the Departments page or click **Subdepartments** on the details page.

You can continue to create or add subdepartments to the department or add applications to the department.

Related tasks:

“Creating departments” on page 272

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

“Adding applications to departments”

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

Adding applications to departments:

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

To add applications to departments, complete these steps:

1. In the menu bar, go to **Groups > Applications**.
2. On the Applications page, select one or more applications, right-click and then, click **Add to.. > Department**.
3. On the Add to Department page, select one or more departments to add to the selected application or applications. The Hierarchy column shows the longest ancestry path of the department.
4. Click **Save**.
5. Optional: On the Create Department page, enter information about the department.
6. Click **Create**.
7. On the Add Members page, click **Add one or more applications**.
8. On the populate page, select the applications that to add to your department. If you do not want to add any of the applications that are listed, click **Create Application** to create a new application to add to the department.
9. Click **Finish**.

View the added department or departments in the Departments column on the Applications page or the Departments tab of the application properties.

Related tasks:

“Creating applications” on page 263

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

“Creating departments” on page 272

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

Viewing information about departments


You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchical business structure.

To view the details of a department, use these steps:

1. In the menu bar, go to **Groups > Departments**.
2. On the Departments page, right-click a department in the list and select **View Details**.

The department details page shows information about the specific department. The information is available in the following sections on the page:

- **Overview:** The overall capacity trending and performance information about the department. An interactive breadcrumb that shows the hierarchy of the departments and applications and a list of all the top level departments.
- **Properties:** The key information about the department.
- **Applications:** The applications that belong to the department. If there is more than one application, the number of applications is displayed.
- **Subdepartments:** The subdepartments that belong to the department. If there is more than one subdepartment, the number of subdepartments is displayed.
- **Related Resources:** Information about related resources in the storage environment that the applications, which belong to the department, are using or are connected to.

Tip: To view descriptions of the information that is available on a department details page, click the **Help** icon  in the upper-right corner of the page.

Related concepts:

“Application and department hierarchies” on page 277

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

Related tasks:

“Creating a department hierarchy” on page 271

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

Removing departments and subdepartments

You can remove departments, subdepartments and applications that are members of the department. The data that is collected about the departments, subdepartments and applications is removed from the database repository.

To remove departments, subdepartments, and applications, that are members of the department, from a level of the business hierarchy, complete these steps:

1. In the menu bar, go to **Groups > Departments**.
2. On the Departments page, select one or more departments, right-click and click **Delete**.
3. Confirm the information and click **Delete**.

- a. If you select **Delete subdepartments and applications**, the subdepartments and applications that belong only to the department, that is being removed, are removed.
- b. If you do not select **Delete subdepartments and applications**, the subdepartments and applications of the department that are being removed become members of the parent department. If the department is a top level department, and the subdepartments are not members of any other department, they become top level departments.

The departments, subdepartments and applications are removed from the business level hierarchy and are no longer displayed on the Departments page.

To remove subdepartments from within the direct context of a department, use these steps:

1. In the menu bar, go to **Groups > Departments**.
2. On the Departments page, select a department from the list and right-click **View Details**.
3. Optional: Select a department from the Departments page and right-click **View Properties**.
 - a. Click the **Subdepartments** tab and continue with steps 5 and 7.
4. On the department details page, click **Subdepartments** in the General section.
5. On the Subdepartments page, select the subdepartment and right-click **Remove from Department**.
6. Optional: On the Subdepartments page, select the subdepartment and click **Delete**.
7. Reply to the confirmation message and click **OK**.

Removing applications from a department

You can remove an application from within a department. This action is only available from the department detail page or the properties for the department.

To remove applications from within the context of a department, use these steps:

1. In the menu bar, go to **Groups > Departments**.
2. On the Departments page, select a department from the list and right-click **View Details**.
3. Optional: Select a department from the Departments page and right-click **View Properties**.
 - a. Click the **Applications** tab and continue with steps 5-7.
4. On the department details page, click **Applications** from the General section.
5. On the Applications page, select the application and right-click **Remove from Department**.
6. Confirm the information.
7. Click **OK**.

The applications that were added within the context of a department are eligible for removal from within department.

Related tasks:

“Viewing information about departments” on page 275

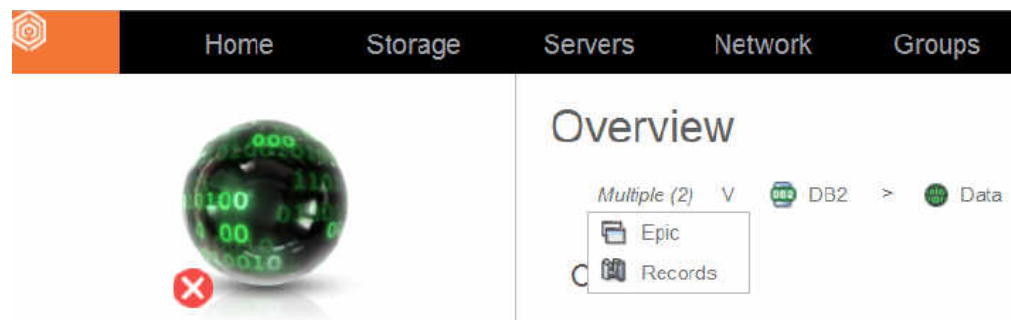
You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal

business structure.

Application and department hierarchies

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

The breadcrumb uses default or customized icons to represent the related applications and departments. The greater-than sign (>) serves as the hierarchy separator and ends with the name of the selected application or department. Only the ancestor names of the selected application are shown.



When you move your mouse pointer over an icon in the breadcrumb, you see the siblings for that application or department. If you move your mouse pointer, for example, over the Data subcomponent, the related application and department elements are shown.

The visual representation of the hierarchy stops when either there are no additional parents or multiple parents are encountered from the selected application or department. When multiple parents are encountered, the number of parents are shown in parenthesis. For example, *Multiple (2)*. Moving the mouse pointer over *Multiple* shows a list of the parents. The siblings of multiple parents are *not* shown.

Move your mouse pointer over a sibling in the breadcrumb and it becomes a name link. Click the application or department name link to navigate to the details for that application or department and view the breadcrumb if it exists.

You can add resources directly on the Overview page for applications from the breadcrumb. For example, in this case, you create an application named DB2 and create two subcomponents, Logs and Data. On the Overview page for the Logs subcomponent, click the Logs subcomponent name link in the breadcrumb and click **Filters** to create a resource filter for that subcomponent. Then, click the Data subcomponent name link and create a resource filter for that application subcomponent. To see a visual representation of the hierarchy for applications, subcomponents, departments, and subdepartments click **View Details**.

Related tasks:

“Creating an application hierarchy” on page 262

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

“Creating a department hierarchy” on page 271

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

Exporting information about applications and departments

You can save information about applications and departments to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of applications and departments.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the GUI has a column or set of columns that determines the order of its rows.
- Filter the list to show only the applications and departments that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the application and department list and detail pages.

Applications and departments list and detail pages

You can export information about the top level applications and departments and the related resources that are shown on application and department list and detail pages.

For example, you can export information about monitored applications and their subcomponents and departments on the Applications page. When you select the export action, the generated file includes the column values for each monitored application in the list.

For example, you can export information about the servers that are associated with an application on the details page.

1. In the GUI, go to the Applications or Departments page and select applications or departments for which you want to export information.
2. Click **Actions** and select **Export** > *file_format*, where *file_format* represents the format that you want to use for a file. The following formats are available:

CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name","Column_name","Column_name","Column_name","Column_name","Column_name"
"data","data","data","data","data","data"
```

where *Column_name* represents the name of a column in a table and *data* represents the data that is associated with a column.

PDF You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

Related concepts:

“Customizing lists of resources, tasks, and alerts” on page 337

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

General groups

View and administer the general groups that are used to alert on logically related storage resources. You can define alerts for a general group to notify you about changes in the configuration, attributes, and performance of the resources in the group.

Organizing your resources into general groups can be helpful in the following situations:

- When you want to receive alert notifications about changes for a subset of the resources of a particular type. For example, you want to detect when the ports that are used for replication on your SAN Volume Controller have insufficient buffer-to-buffer credit. You do not want the alerts to apply to other ports on your SAN Volume Controller. You can group the ports that are used for replication and then define alerts for the group.
- When you want to receive alert notifications about changes for a group of resources that are logically related. Examples include all the servers that use a particular operating system or all the storage systems at a specific location.

General group hierarchies

Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.

Only top-level groups in the general group hierarchy are shown on the General Groups page. If the general group has a subgroup, the name of the subgroup is shown in the Subgroups column. If the general group has more than one subgroup, click the number of groups in the Subgroups column to view the subgroups.

Related tasks:

“Defining alerts for general groups” on page 176

Define alerts for changes in the configuration, attributes, and performance of the resources in your general groups.

Creating a general group hierarchy

Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.

In this scenario, you want to monitor all the ports on your SAN Volume Controller but you also want to separately monitor the following subsets of ports on the SAN Volume Controller:

- Ports that are used for inter-node communication
- Ports that are used for host I/O exchanges

To add the ports to a general group hierarchy, complete the following steps:

1. Add the ports that are used for inter-node communication to a general group, SVC Inter Node Ports.
2. Add the ports that are used for host I/O exchanges to another general group, SVC Host I/O Ports.
3. Add the groups as subgroups of a parent general group, SVC All Ports.

You can now quickly view information about the resources in the hierarchy on the details page for the general groups.

When you view the SVC Inter Node Ports or SVC Host I/O Ports groups, you see information about the specific ports that are members of that subgroup.

When you view the parent group, SVC All Ports, you see information about the ports that are members of the parent group and also the ports in the subgroups.

Creating general groups and adding resources:

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

You can add the following resources and their internal resources to a general group:

- Storage systems
- Servers
- Hypervisors
- Switches
- Fabrics
- Other general groups

Only resources that you specifically add to a group are included as members of the group. For example, if you add a SAN Volume Controller, the internal resources of the SAN Volume Controller, such as volumes and pools, are not automatically added to the group. To add the volumes and pools, you must specifically select those resources and add them to the group.

To create a general group and add resources to the group, complete these steps:

1. Take one of the following actions to go to the list page for the resource or group that you want to add:

- To add top-level resources such as storage systems, servers, hypervisors, switches, or fabrics, go to the appropriate resource page. For example, to add servers, in the menu bar, click **Servers > Servers**.
 - To add internal resources of top-level resources, complete the following steps:
 - a. Go to the resource list page for the top-level resource. For example, to add internal resources of a SAN Volume Controller, in the menu bar, click **Storage > Block Storage Systems**.
 - b. Right-click the resource and click **View Details**.
 - c. In the Internal Resources section, click the type of resource that you want to add, for example, volumes or pools.
 - To add general groups, in the menu bar, click **Groups > General Groups**.
2. Right-click one or more resources or groups and click **Add to General Group**.
 3. Click **Add to new group** and specify a name and description for the new group.
 4. Optional: Customize the icon for a new group by clicking the existing icon and selecting another icon.
 5. Click **Save**.

You can now quickly view information about the resources in the general group hierarchy on the details page for the general group.

Related tasks:

“Adding resources to general groups” on page 282

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

Related reference:

“Viewing and administering general groups” on page 283

View and administer the general groups that you use to alert on logically related storage resources.

Adding general groups as subgroups:

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

1. In the menu bar, click **Groups > General Groups**.
2. Right-click one or more general groups and click **Add to General Group**.
3. Take one of the following actions:
 - To add the groups as subgroups of existing groups, click **Add to existing groups**, click one or more groups in the list, and click **Save**.
 - To add the groups as subgroups of a new group, click **Add to new group**, specify a name and description for the new group, and click **Save**.

Related tasks:

“Adding resources to general groups” on page 282

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

Related reference:

“Viewing and administering general groups” on page 283

View and administer the general groups that you use to alert on logically related

storage resources.

Adding resources to general groups

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

The following resources and their internal resources can be added to a general group:

- Storage systems*
- Servers
- Hypervisors
- Switches
- Fabrics
- Other general groups

Tip: * For IBM Cloud Object Storage, only vaults can be added to general groups.

Only resources that you specifically add to a group are included as members of the group. For example, if you add a SAN Volume Controller, the internal resources of the SAN Volume Controller, such as volumes and pools, are not automatically added to the group. To add the volumes and pools, you must specifically select those resources and add them to the group.

You can create a general group hierarchy by adding one or more general groups as subgroups of a general group. When you define a general group alert for a resource such as a volume, the alert applies to all the volumes that belong to that group and all its subgroups.

To add resources to a general group, complete these steps:

1. Take one of the following actions to go to the list page for the resource or group that you want to add:
 - To add internal resources of top-level resources, complete the following steps:
 - a. Go to the resource list page for the top-level resource. For example, to add internal resources of a SAN Volume Controller, in the menu bar, click **Storage > Block Storage Systems**.
 - b. Right-click the resource and click **View Details**.
 - c. In the Internal Resources section, click the type of resource that you want to add, for example, volumes or pools.
 - To add general groups, in the menu bar, click **Groups > General Groups**.
2. Right-click one or more resources or groups and click **Add to General Group**.
3. Take one of the following actions:
 - To add the resources or groups to a new group, click **Add to new group** and specify a name and description for the new group.
 - To add the resources or groups to one or more existing groups, click **Add to existing groups** and click the appropriate groups in the list.
4. Optional: Customize the icon for a new group by clicking the existing icon and selecting another icon.
5. Click **Save**.

To view the resources and subgroups that are members of the group, complete the following steps:

1. Go to **Groups > General Groups**.
2. Right-click the general group and select **View Details**.
On the general group details page, all the different resource types, such as volumes or pools, that are members of the group or members of its subgroups are shown in the Members section.
3. Click the resource type to view the specific resources. For example, click **Volumes** to view all the volumes that are in the general group or in its subgroups.
4. Click **Subgroups** to view the child groups.

To define alerts for the general group, use the **Edit Alert Definitions** action on the General Groups page.

Related tasks:

“Creating general groups and adding resources” on page 280

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

“Adding general groups as subgroups” on page 281

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

Related reference:

“Viewing and administering general groups”

View and administer the general groups that you use to alert on logically related storage resources.

Viewing and administering general groups

View and administer the general groups that you use to alert on logically related storage resources.

Administering general groups

Table 64. Administering general groups

Actions	Navigation
Create a general group and add resources.	<ol style="list-style-type: none">1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add file storage systems, click Storage > File Storage Systems.2. Right-click one or more resources and click Add to General Group.3. Click Add to new group, specify a name and description for the new group, and click Save.

Table 64. Administering general groups (continued)

Actions	Navigation
Add resources to general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add block storage systems, click Storage > Block Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the resources to one or more existing groups, click Add to existing groups, click the appropriate groups in the list, and click Save. b. To add the resources to a new group, click Add to new group, specify a name and description for the new group, and click Save.
Define and edit the alert definitions for a general group. Learn more	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click Edit Alert Definitions.
Edit the alert notification settings for general groups. Learn more	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Edit Alert Notification Settings.
Delete general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Delete. 3. To delete subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>
Remove resources from a general group.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click View Details. 3. In the Members section, click the type of resource that you want to remove. For example, if you want to remove some of the volumes from the group, click Volumes. 4. Right-click the resources that you want to remove and click Remove from General Group. 5. Click OK to confirm that you want to remove the resources from the general group and from all of its subgroups.

Table 64. Administering general groups (continued)

Actions	Navigation
Add general groups as subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the groups as subgroups of existing groups, click Add to existing groups, click one or more groups in the list, and click Save. b. To add the groups as subgroups of a new group, click Add to new group, specify a name and description for the new group, and click Save.
Remove subgroups from a group hierarchy.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click View Details. 3. Click Subgroups in the General section. 4. Right-click the general groups that you want to remove and click Remove from General Group. <p>When you remove a subgroup from its parent group, the subgroup is moved to the same level in the hierarchy as the parent group. The subgroup is still monitored by the product.</p>
Create a subgroup.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click the general group that you want to create a subgroup for and click View Details. 3. Click Subgroups in the General section. 4. Click Create Subgroup. 5. Specify a name and description for the new group. 6. Customize the icon for the group by clicking the existing icon and selecting another icon. 7. Click Create.
Delete subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click the general group that contains the subgroups that you want to delete and click View Details. 3. Click Subgroups in the General section. 4. Right-click one or more subgroups and click Delete. 5. To delete further subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>

Viewing general groups

Table 65. Viewing general groups

Actions	Navigation
View a list of general groups.	In the menu bar, click Groups > General Groups . The following general groups are shown: <ul style="list-style-type: none">• The top-level groups in general group hierarchies• All general groups that are not part of a general group hierarchy
View subgroups in a general group hierarchy.	<ol style="list-style-type: none">1. In the menu bar, click Groups > General Groups.2. Click the number of groups or the group name that is shown in the subgroups column.3. Click the Subgroups tab.
View details about a general group, including the alert notifications, the resources that are members of the group, and the subgroups.	<ol style="list-style-type: none">1. In the menu bar, click Groups > General Groups.2. Right-click a general group and select View Details.
View the alert definitions or alert notification settings for a general group.	<ol style="list-style-type: none">1. In the menu bar, click Groups > General Groups.2. Right-click a general group and select View Details.3. Click Alerts in the General section.4. Click the Definitions or Notification Settings tabs.
View the alert notifications for a general group.	<ol style="list-style-type: none">1. In the menu bar, click Groups > General Groups.2. Right-click a general group and select View Details.3. Click Alerts in the General section.
View reports about general groups.	In the menu bar, click Cognos . The Cognos BI reporting tool is included in Tivoli Common Reporting. To view reports about general groups, view reports about storage resource groups in the Cognos BI reporting tool.

Edit Alert Notification Settings

View and edit the settings that determine how you are notified when alert conditions are detected for general groups. These notification settings are applied globally to all the alert definitions that are specified for the selected general groups. You can override the global settings by specifying different settings for each alert definition.

For information about the notification settings and actions that you can specify, see “Triggering actions and notifications for alerts” on page 401.

Tips for editing the notification settings for multiple general groups:

- To edit the notification settings for multiple general groups at the same time, press Ctrl + click or Shift + click to select the groups, right-click the selection, and select **Edit Alert Notification Settings**.
- When you edit the notification settings for multiple groups, some of the existing settings might not apply to all of the groups. In these cases of mixed values, the alert notification settings can have the following states:
 - If a setting is blank, that setting does not apply to any of the selected groups.

- If a setting is checked, it applies to all of the selected groups.
- If a setting contains a dot, it does not apply to all of the selected groups, or the setting contains different values for each of the groups.

Related tasks:

“Creating general groups and adding resources” on page 280

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

“Adding resources to general groups” on page 282

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

“Adding general groups as subgroups” on page 281

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

Viewing the data paths of resources

A data path shows the connectivity between host connections, servers, hypervisors, virtual machines, storage systems, and the fabrics through which they share storage. You can use the data path view to monitor status and pinpoint problem areas in selected data paths. This view includes graphical and tabular representations of the top-level resources in a data path.

If a resource does not map storage to other resources, or the resource is not assigned storage from other resources, the data path view is not available.

To view the data path for a resource, complete these steps:

1. In the menu bar, select the type of resource that you want to view. For example, if you want to view the data path for a server, go to **Servers > Servers**.
2. The type of resource that you select determines how to view a data path:
 - On the Servers page, right-click a server and select **View Data Path**.
 - On the Block Storage Systems page, File Storage Systems page, or Object Storage Systems page, right-click a storage system and select **View Details**. On the details page, click **Data Path** in the **General** section.

Depending on how resources are configured, a data path can show between three to six types of top-level resources. For example, a data path with three types of top-level resources might show the following relationships:

server <> fabric <> storage system

A data path with six types of top-level resources might show the following relationships:

virtual machine <> hypervisor <> fabric <>
storage virtualizer <> fabric <> storage system

Before you can view a data path, you must complete the following tasks:

- Add the resource that has the data path you want view. For information about how to add a resource, see “Adding resources” on page 112.

Tip: Resources that are not monitored by IBM Spectrum Control are also shown in a data path view. However, detailed information about their status and internal resources is not available.

- Schedule a probe to collect data about your monitored resources to ensure that the most recent information is shown for a data path. For information about probes, see “Collecting asset and status data” on page 134.

You can use the data path view to complete the following actions:

- View the path of data that is shared between resources. The data path view shows only the resources that have storage mapped between them. For example, if you view the data path for a storage system, only the host connections that are assigned volume space from that storage system are shown. Host connections that are connected to the storage system, but do not have storage assigned storage to them, are not shown.
- View the fabrics through which resources in a data path are communicating.
- View the propagated status of the top-level resources that are in a data path.
- View the status of the internal resources for top-level resources that are in a data path.
- Customize the appearance of the data path view to suit the needs of your environment.
- Export the data paths view as an image or CSV file.

Information about a data path is organized into two views: a topology view and a table view. The topology view shows icons for connected resources and uses lines to represent the path of data through them. The table view includes a row for each resource in a data path. The row is divided into columns that contain information about the resource.

Resources that you can view in a data path

The **Topology View** tab in the data path view includes a visual representation of the data path for a top-level resource. In a data path, icons represent the connected resources and lines represent the path of data through them. For example, resources that are storage systems are shown with an icon that resembles a storage system.

The following table shows the icons for each type of resource that can be shown in a data path.

Table 66. Resource icons in a data path



Resource	Icon
Server (with agent)	
Server (without agent) ¹	

Table 66. Resource icons in a data path (continued)









Resource	Icon
Hypervisor	
Virtual machine	
Fabric	
Virtual fabric	
DS8000	
DS6000	
SAN Volume Controller	
Storwize V7000 ²	

Table 66. Resource icons in a data path (continued)






Resource	Icon
XIV	 A tall, black, vertical storage unit with a single door and a small orange indicator light near the bottom.
ESS	 A tall, black, vertical storage unit with multiple vertical slots and a central control panel.
EMC	 A tall, black, vertical storage unit with multiple horizontal slots and a central control panel.
Other storage systems	 A black, vertical storage unit with multiple horizontal slots and a central control panel.
Unmanaged host connections	 A blue storage unit with a yellow cable connected to a grey server rack.

Table 66. Resource icons in a data path (continued)

Resource	Icon
<p>Tips and restrictions:</p> <ul style="list-style-type: none"> ¹ When agentless servers and servers with agents are grouped together, the icon for servers with agents is used to represent that group. ² If a Storwize V7000 is not monitored by IBM Spectrum Control, but is connected to a monitored hypervisor, it is shown in the data path with the icon for a SAN Volume Controller. <p>For example, when a virtual machine is monitored, a probe collects information about its disks. If the disks are attached to a SAN-attached storage system that is not monitored, IBM Spectrum Control uses the probe information to determine the type of storage system. If the storage system is a Storwize V7000, the icon for SAN Volume Controller is shown instead of the icon for a Storwize V7000.</p> <ul style="list-style-type: none"> Only the resources that are sharing storage in a data path are shown. Servers and hypervisors that are not monitored by IBM Spectrum Control, but are assigned storage from a monitored storage system, are shown as unmanaged host connections. For example, if a storage system has a volume that is mapped to an unknown server, the server is represented by an unmanaged host connection in the view. For storage systems with multiple host connections, the host connections are included in the same resource box if they communicate with the storage system through the same fabric. For example, if a storage system has volumes that are mapped to 20 host connections that use the same fabric, a single resource box represents the 20 host connections. Storage systems can have volumes that are provisioned to more than one server in a clustered environment. Because all the clustered servers can be part of the data path, each server is shown in the view. 	

The status of resources

Each resource in a data path is represented by a resource box. A resource box contains an icon that represents the overall status of a resource. This status is determined by the propagated status of the internal resources for the resource. For example, in a storage system, a port might have an Error status and a pool might have a Warning status. In this case, the Error icon is shown on the resource box for the storage system because the Error condition is the most severe status that was detected on its internal resources.



Some resources in a data path might be shown with an unknown status. Resources can have an unknown status if they are not monitored by IBM Spectrum Control, but are connected to other, monitored resources. IBM Spectrum Control uses the information that it collects about monitored resources to determine basic information about the connected resources that are not monitored.

For example, if a host connection is assigned volumes, but IBM Spectrum Control is not monitoring the fabric through which the host connection communicates with the storage system, the fabric is shown with an unknown status.



Navigating in a data path

A data path might contain more resources that can be fit in the view at the same time. In these cases, you can scroll the view to show different parts of a data path.

1. In the **Topology View** tab, click the **Pointer** button .
2. Click an empty area in the view and drag the mouse pointer to scroll to different areas of a data path.
3. Optional: To view where you are in the overall topology of a data path, click the **Topology overview** button  in the upper-right corner of the topology view. The red box shows the visible area within the overall topology view, and status icons are displayed for the resources in the data path.

Viewing information about top-level resources

You can view detailed information about the top-level resources in a data path. Top-level resources that can be shown in a data path include storage systems, servers, hypervisors, and fabrics.

1. To view key attributes of a specific top-level resource without leaving the **Topology View** tab, right-click the resource and select **View Properties**. A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
2. To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**. A resource details page is displayed and includes information about storage statistics, alerts, jobs, internal resources, and related resources.



Viewing information about internal resources


You can view information about the internal resources that are associated with the top-level resources in a data path. Internal resources are components that exist in a top-level resource.

1. In the **Topology View** tab, hover the mouse pointer over a top-level resource in the data path.
2. View the following information about internal resources:
 - The number of internal resources that are associated with a top-level resource. For example, if a server has 20 disks, the value **Disks (20)** is shown.
 - The overall status of the internal resources. For information about how the overall status of internal resources is determined, see “How the condition of a resource is determined” on page 204.

Zooming in



You can use the zoom buttons in the toolbar or the mouse to magnify areas in a data path.

1. To zoom in on a specific area in a data path, click the **Zoom to selected** button  and click and drag the mouse pointer over the area.
2. Optional: To use the mouse to zoom in on a specific area in a data path, move the mouse pointer over the area where you want to zoom in and scroll the middle mouse button forward.
3. Optional: To zoom in on the area that is in the center of the view, click the **Zoom in** button .

- Optional: To adjust the zoom level of the view so that all the resources in a data path are shown at the same time, click the **Fit to view** button .

Zooming out

You can use the zoom out button in the toolbar or the mouse to zoom out in the view.

- To zoom out from the area that is in the center of the view, click the zoom out button .
- Optional: To use the mouse to zoom out from a specific area in a data path, move the mouse pointer over the area where you want to zoom out and scroll the middle mouse button backward.
- Optional: To adjust the zoom level of the view so that all the resources in a data path are shown at the same time, click the **Fit to view** button .

Changing the layout

You can change the layout of the resources in a data path. The layout determines how resources are arranged in the view.

- Click the **Layout** button in the toolbar. The icon that is displayed for the **Layout** button represents the current layout of the data path.
- Select one of the following layouts:



Hierarchical

Organize resources in a vertical format. Resources that use storage are shown at the top of the view. Resources that provide storage are shown at the bottom of the view.



Tree

Organize resources in a vertical and horizontal layout.




Customize

Customize the layout of a data path. When you customize a layout, use the hierarchical or tree layouts as the base and specify values for the associated options.

- Optional: If you selected **Customize**, click **OK** to apply your settings.


Saving the display settings

You can save the display settings that are defined for a data path view. The next time that you view the data path for the resource, the same display settings are applied.

- In the **Topology View** tab, click the **Save** button .
- Optional: To use the keyboard to save the display settings, press Ctrl+S.
- Click **Save Customizations**.

Printing a view

You can print the current view of a data path in the **Topology View** tab. The printing capability is not available in the **Table View** tab.

- Click the **Print** button .
- Specify the following options when printing:

- In the **Page margin** field, enter the amount of space between the edges of the page and the view. The default setting is 1 inch.
 - In the **Maximum pages in width or height** field, specify the number of pages across which the content of a view is printed. This number determines the maximum number of pages that are printed. For example, if you enter **1**, the content of a view is printed on one page. If you enter **2**, the view is printed across four pages: two pages horizontally and two pages vertically.
3. Click **OK**.
 4. Select the printer settings and click **OK** to print the view.

Exporting a data path to a file

You can save the contents of a data path to an image file or CSV file for offline viewing.

1. In the **Topology View** tab, right-click on an empty area in the data path.
2. Complete one of the following actions to select an output format for the file:
 - Select **Export as Image** to save the data path as a JPEG or PNG file.
 - Select **Export** to save the data path to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information for a data path is organized into the following column in the CSV files: Name, Type, From, RelType, To.
3. Specify the location and name of the file.
4. Save the file.

Provisioning storage

You can assign storage to servers, hypervisors, and clusters on servers and hypervisors in the GUI. You can assign storage to hypervisors with the vSphere Web Client extension for IBM Spectrum Control.

You can assign the following types of storage to servers, hypervisors, and clusters on servers and hypervisors:

- Volumes
- NAS file shares

When you assign storage, you must associate the storage with a service class. You can associate the storage with a default service class such as gold, silver, or bronze, or you can create a user-defined service class. You can also create capacity pools to restrict the provisioning of storage to a defined set of storage resources such as pools on a storage system in a particular location.

Configuring IBM Spectrum Control for provisioning

Before you can use the simplified provisioning capabilities that are provided with IBM Spectrum Control, you must configure your environment.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, IBM Spectrum Control can analyze and optimize volume performance.

To configure IBM Spectrum Control for provisioning, complete the following tasks:

1. Assign storage pools to tiers by setting the tier level of each storage pool. The tier levels are number tags that can reflect any tier structure in the environment.
You must assign storage pools to tiers to later provision volumes that require a certain tier level and to optimize the placement of volumes. Storage pools do not need to be assigned to tiers to provision volumes that do not require a tier level, or to provision file shares.
2. Create service classes. A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.
3. Optional: Create capacity pools. You can restrict provisioning requests to a capacity pool. Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business, or to a particular application.

Configuring service classes

Configure service classes to represent the different levels or types of storage quality that you want to provision. When you are configuring service classes for your installation, consider whether certain storage requests must always be satisfied from a particular set of resources. Also, consider whether you want to allow users without administrator privileges to provision storage, and whether you want their provisioning requests to require administrator approval.

Creating service classes:

Create service classes to specify storage requirements for future provisioning requests. A service class describes attributes and capabilities of block storage or file storage resources, and typically describes particular quality of service.

You must have Administrator privileges to create a service class.

The IBM Spectrum Control GUI guides you through the steps to create a service class. When you provision a volume or share, you define storage requirements by specifying a service class. A block-storage service class specifies the requirements for the block-storage resource from which a new volume is allocated. A file-storage service class defines requirements for the file storage resources from which a new share is allocated.

You can associate a service class with one or more capacity pools. If you do, future provisioning requests can be satisfied only by storage resources in the specified capacity pools.

You can selectively grant users who do not have administrator privileges permission to provision storage by using the service class.

Restrictions:

- To create a block-storage service class, you must have the IBM Spectrum Control Advanced Edition license.

- To associate the service class with a capacity pool, one or more capacity pools must be defined in the IBM Spectrum Control database. If no capacity pool is defined, the Specify Capacity Pool page is not displayed.

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics > Cloud Configuration > Work With Service Classes**. To create a service class, complete the following steps:

1. From the Service Classes page, click **Create Service Class**.
2. Select the type of service class that you want to create:
 - To create a block-storage service class, click the **Block** icon.
 - To create a file-storage service class, click the **File** icon.
3. In the Define Properties page, complete the following steps:
 - a. Specify a name for the service class.
 - b. Optional: Specify a service class description. Descriptions are displayed in the table of service classes on the Service Classes page.
 - c. Specify the general attributes of the service class. Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.
4. Optional: Specify or modify the advanced properties of the service class.
 - a. From the Define Properties page, click **Advanced**.
 - b. If you are creating a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
 - c. If you are creating a block-storage service class, select how the multipath driver on the host is configured. For hypervisors, the multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.
 - d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
5. Optional: In the Specify Capacity Pool page, associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools.
6. Optional: Grant permission to individual users, who do not have administrator privileges, to provision storage by using the service class. By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.
7. Click **Finish** The new service class is listed in the table of service classes on the Service Classes page, and can be specified in future provisioning requests.

Viewing the details of service classes:

You can view the provisioning requirements and storage constraints that a service class defines. You can also view the volumes or shares that were provisioned by using the service class.

From the Service Classes page, you can open a properties page to view the following information for a service class:

- The properties of a service class, such as the required RAID level or storage tier for a block-storage service class.
- The storage constraints that are imposed by the service class. The service class can allow provisioning from all available storage, or can limit provisioning to one or more capacity pools.
- The users who can provision storage by using the service class.
- The volumes or shares that were provisioned by using the service class.
- The candidate storage for the service class. The *candidate storage* is the set of storage resources that can provide the service class.

To open the service class properties window from the Service Classes page, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Service Classes**. The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions > View / Modify**. The service class properties window is displayed.

Modifying service classes:

You can modify the attributes of a service class to change the provisioning requirements, capacity pool storage constraints, user permissions, and service class name.

You must have Administrator privileges to modify a service class.

Important: When volumes or shares are provisioned by specifying a service class, they are associated with the service class. The Volumes page and the Shares page in the web-based GUI can display the name of the service class that was used to provision the volume or share. If you modify a service class, the volumes or shares are still associated with the service class, but might be on storage resources that no longer satisfy the requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the new volumes or shares have attributes that they do not possess.

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics > Cloud Configuration > Work With Service Classes**. To modify a service class, complete the following steps:

1. From the list of service classes, select a service class and then select **Actions > View / Modify**. The service class properties window is displayed.
2. Optional: Click the **General** tab to modify the general attributes of the service class. Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.

3. Optional: Specify or modify the advanced properties of the service class.
 - a. From the **General** tab, click **Advanced**.
 - b. If you are modifying a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
 - c. If you are modifying a block-storage service class, select how the multipath driver on the host is configured. For hypervisors, the multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.
 - d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
4. Optional: Modify storage constraints. If one or more capacity pools are defined in the IBM Spectrum Control database, you can associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools. If a capacity pool was previously associated with the service class, you can disassociate it from the service class.
5. Optional: Modify user permissions. By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.

If users in the Monitor or External Application roles have permission to provision by using the service class, you can revoke that permission. You can specify that only administrators can provision by using the service class, or you can selectively remove users. To remove a user, right-click the user name in the table and then select **Remove** from the context-sensitive menu.

Deleting service classes:

You can delete a service class that is no longer needed.

You must have Administrator privileges to delete a service class.

Deleting a service class does not affect the storage volumes or NAS shares that are already provisioned by using the service class. However, the storage volumes or NAS shares are no longer associated with any service class. When a volume that is not associated with a service class is optimized, the volume can be migrated to any available storage pool. The storage pool to which the volume is migrated might not satisfy the requirements of the service class that was specified when the volume was provisioned.

To delete a service class from the Service Classes page, complete these steps:

1. In the menu bar, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Service Classes**. The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions > Delete**.

Tagging resources to satisfy custom requirements:

A service class can specify tags to create custom requirements for provisioning. Add the same tags to storage resources that satisfy the custom requirements. When the service class is specified during provisioning, only the tagged resources are candidates for provisioning.

You must have Administrator privileges to tag storage resources.

When you provision storage, you indicate your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have to be a candidate for provisioning. In addition to the default properties, a service class can define custom requirements by specifying up to three custom tags. To provide the service class, a storage resource must have all the same tags as the service class. The ability to define custom requirements by tagging service classes and storage resources gives you more control over the selection of storage resources during provisioning. By using tags, you can filter the candidates for provisioning by using any criteria that you choose.

When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a storage pool for the volume. If the service class specifies tags, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a file system or Network Shared Disk (NSD) for the share. If the service class specifies tags, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

Tag only the storage resources that satisfy the intended custom requirements of the service class. During provisioning, only the tagged resources can provide the service class.

1. Open the resource list page for the type of storage resource that you want to tag. To satisfy the custom tag requirements of a block-storage service class, you can tag a storage pool or its containing block storage system. To satisfy the custom tag requirements of a file-storage service class, you can tag a file system or NSD, or the containing file storage system.

Important: The tags on a containing storage system apply to an internal pool, file system, or NSD resource only if the internal resource is not tagged.

- If you are tagging a storage pool, from the menu bar, go to **Storage > Pools**.
- If you are tagging the block storage system that contains the pools, complete the following steps:
 - a. In the menu bar, go to **Storage > Block Storage Systems**.
 - b. Click the **Block Storage** tab.
- If you are tagging a file system or NSD, complete the following steps:
 - a. In the menu bar, go to **Storage > Block Storage Systems**.
 - b. Click the **File Storage** tab.

- c. From the list of storage systems, select the storage system that contains the file system or NSD, and then select **View Details**. Alternatively, you can select **Actions > View Details**.
- d. In the Internal Resources area, click **File Systems** or **Network Shared Disks**.
- If you are tagging the file storage system that contains the file systems or NSDs, complete the following steps:
 - a. In the menu bar, go to **Storage > Block Storage Systems**.
 - b. Click the **File Storage** tab.
2. In the table of storage systems, pools, file systems, or NSDs, select the storage resource that you want to tag, and then select **View Properties**. Alternatively, you can select **Actions > View Properties**.
3. In the properties notebook for the storage system, pool, file system, or NSD, click **Edit**.
4. In the custom tag fields of the properties notebook, specify tags that match the tags that are specified in the service class. To provide the service class, the resource must have all of the same tags.
5. In the properties notebook, click **Save**.

Setting access path host replacements:

If you are going to use the VMware vSphere Web Client extension to provision file shares, verify whether you must provide IBM Spectrum Control with an access path host replacement.

You need to provide an access path host replacement if the cluster name of an IBM SONAS or Storwize V7000 Unified is not a valid host name that is registered with a Domain Name System (DNS) server.

Complete the following steps to verify cluster host names and to provide access host name replacements where necessary:

1. In the menu bar, go to **Storage > Block Storage Systems**.
2. Click the **File Storage** tab.
3. For each IBM SONAS and Storwize V7000 Unified system, complete the following steps to determine whether you must specify a replacement host name.
 - a. From the **File System** tab of the Storage System page, select the storage system, and then select **Actions > View Details**.
 - b. From the list of File Internal Resources, click **Clusters**.
 - c. Determine whether the cluster name is a valid host name that is registered with a DNS server. If the cluster name is registered with a DNS server, the access path that is returned by the IBM Spectrum Control will be valid. If the cluster name is not valid, the VMware vSphere Web Client extension will not be able to access the share unless you specify a host name replacement for the access path.
4. For each IBM SONAS and Storwize V7000 Unified system that requires a host name replacement, complete the following steps:
 - a. From the **File Storage** tab of the Storage Systems page, select the storage system, and then select **Actions > View Properties**.
 - b. In the properties notebook for the storage system, click **Edit**.

- c. In either the Custom Tag 1, Custom Tag 2, or Custom Tag 3 field, enter a fully qualified host name that is mapped to one or more interface nodes of the storage system.

Requirement: For all storage systems that require a replacement host name, ensure that you set the replacement host name by using the same custom tag. For all IBM SONAS and Storwize V7000 Unified systems that do not require a replacement host name, ensure that the custom tag is empty.

5. When you are creating or modifying file-storage service classes that might be used by the VMware vSphere Web Client extension, set the access host replacement to the custom tag that is set on the file storage systems. To set the access host replacement, complete these steps:
 - a. From the list of general properties in the service class properties window, click **Advanced**.
 - b. On the **General** tab of the Advanced Properties window for the file-storage service class, set the access path host replacement to the custom tag that specifies the replacement host name. When the share is provisioned, the access path that is returned by IBM Spectrum Control includes the replacement host name if one is specified on the custom tag for the storage system. If the storage system does not specify a value for the custom tag, the cluster name is used in the access path.

Configuring capacity pools

Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Adding resources to capacity pools:

Divide storage resources into capacity pools to later restrict provisioning or optimization to a specific set of resources. You can also use capacity pools to track the used and available space for block and file storage on any set of storage resources. You can add resources to an existing capacity pool, or you can create a new capacity pool.

You must have Administrator privileges to perform this task.

You can add storage resources to a capacity pool from resource list pages such as the Storage Systems page and the Pools page.

You can add the following resource types to a capacity pool:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM SONAS Network Shared Disks (NSDs)

When you add a storage system to a capacity pool, any storage pool, file system, or IBM SONAS NSD that is an internal resource of that storage system is also indirectly assigned to the capacity pool.

Restriction: A resource can be assigned to only one capacity pool. When you attempt to add one or more resources to a capacity pool that are already assigned to different capacity pool, IBM Spectrum Control displays a message that lists the

resources and their capacity pool assignments. When a message is displayed, ensure that you want to change the capacity pool assignments for the listed resources.

To add one or more resources to a capacity pool, complete these steps:

1. Open the appropriate list page for the resource you want to add. You can add an entire storage system, or you can add individual pools, file systems, or IBM SONAS NSDs.
 - To add a storage system, in the menu bar, go to **Storage > Block Storage Systems** or **Storage > File Storage Systems**.
 - To add a storage pool, in the menu bar, go to **Storage > Pools**.
 - To add a file system or IBM SONAS NSD, complete the following steps:
 - a. In the menu bar, go to **Storage > File Storage Systems**.
 - b. Select the storage system that contains the file system or IBM SONAS NSD, and then select **View Details** from the context-sensitive menu.
 - c. From the Internal Resources list, click **File Systems** or **Network Shared Disks**.
2. To add one or more resource to a capacity pool, select the resources in the list, and then select **Actions > Add to Capacity Pool**. If one or more capacity pools are defined, a list of capacity pools is displayed. If no capacity pools are defined, the Create Capacity Pool window is displayed so that you can add resources to a new capacity pool.

Tip: The list of capacity pools and the Create Capacity Pool window shows capacity information. For the list of capacity pools, the Capacity column shows a bar that summarizes the used and free space in each capacity pool. In the Create Capacity Pool window, a capacity bar summarizes the used and free space of the resources that are being added to the new capacity pool. Move the mouse pointer over the bar to display a text summary and details of the capacity usage.

3. Take one of the following actions:
 - If the Create Capacity Pool window is displayed, go to step 4.
 - If a list of capacity pools is displayed, add the selected resources to an existing capacity pool or to a new capacity pool.
 - To add the selected resource or resources to an existing capacity pool, select the appropriate row in the list of capacity pools, and click **Save**.
 - To add the selected resource to a new capacity pool, click **Create Capacity Pool**. The Create Capacity Pool window is displayed.
4. In the Create Capacity Pool window, specify a name for the capacity pool.
5. Optional: Specify a capacity pool description. Capacity pool descriptions appear in the list of capacity pools on the Capacity Pools page.
6. Optional: Specify custom tags to associate any significant information with the capacity pool. You can specify up to three custom tags. The custom tags can be displayed in the list of capacity pools on the Capacity Pools page.
7. Save the new capacity pool by clicking **OK**.

Viewing the details of capacity pools:

You can view details of a capacity pool, such as its total capacity, used and available space, and member resources.

The Capacity Pools page lists all of the capacity pools that are defined in IBM Spectrum Control. For any capacity pool in the list, you can open a window to view more details about the capacity pool and its storage resources.

To view details of a capacity pool, complete these steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions > View / Modify**. A window is displayed, showing the following details:
 - A bar summarizing the used and free space in the capacity pool in gibibytes. The bar represents the total capacity of the capacity pool. The bar is divided into sections that represent the following measurements:
 - The amount of space that is used for file storage
 - The amount of available space for file storage
 - The amount of space that is used for block storage
 - The amount of block storage space that is not used for volumes, and is not reserved by pending or scheduled provisioning tasks.
 - The amount of block storage space that is reserved by pending or scheduled provisioning tasks

Move the mouse pointer over the bar to display a text summary and details of the capacity usage, and a legend for interpreting the colors on the bar.

- A table that lists the storage systems, storage pools, file systems, and IBM SONAS Network Shared Disks (NSDs) that are in the capacity pool. To view details of any of these resources, select the resource, and then select **Actions > View Properties**.

Modifying capacity pools:

You can modify the attributes of a capacity pool such as its name and description, and you can remove storage resources from a capacity pool.

You must have Administrator privileges to perform this task.

To modify a capacity pool, complete these steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions > View / Modify**. A window that shows details of the capacity pool is displayed. A table lists the storage systems, storage pools, file systems, and IBM SONAS Network Shared Disks (NSDs) that are in the capacity pool.
3. Optional: Modify attributes of the capacity pool.
4. Optional: Remove one or more resources from the capacity pool. Removing a storage resource from a capacity pool does not delete the storage resource, but does remove its capacity from the capacity pool. To remove resources, select the appropriate resources in the list, and then select **Actions > Remove**. The resources are removed from the list, but are not removed from the capacity pool until you save your changes.
5. Click **Save**.

Deleting capacity pools:

You can delete a capacity pool that is no longer needed.

You must have Administrator privileges to perform this task.

When you delete a capacity pool, the storage resources that were added to the capacity pool are no longer members of any capacity pool. Deleting a capacity pool does not affect any volumes or NAS shares that were provisioned from the capacity pool. However, the volumes or NAS shares are no longer associated with a capacity pool.

To delete a capacity pool, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions > Delete**.
3. In the confirmation window, click **Delete**.

Block storage: Calculating available space and determining the placement of volumes

For each block-storage service class, IBM Spectrum Control calculates the amount of available space for volumes of that service class. When you request volumes of a particular service class, IBM Spectrum Control recommends the best location for the volumes from the available space.

The Service Classes page in the web-based GUI displays the available space for each block-storage service class. The available space is the amount of space that is available across storage pools for new volumes of the service class. Typically, not all pools are used in the available-space calculation. Some pools might be eliminated because of the consolidated status of the pool or its storage system. Other pools might be eliminated because they do not satisfy the requirements of the service class. The set of storage pools that are used to calculate the available space are also the candidates for provisioning by using the service class.

When you request storage volumes by using the Provision Storage wizard, you specify the capacity that is required for one or more volumes, and you specify a service class. You can also specify a capacity pool from which the volumes must be provisioned. Based on the volume size that is requested and if a capacity pool is specified, the candidates for provisioning are further refined to include only the pools that have enough available space and are in the specified capacity pool. From the remaining candidates for provisioning, IBM Spectrum Control creates a provisioning task that identifies the best location for the new volumes.

How available space for a block-storage service class is calculated

For each block-storage service class, IBM Spectrum Control calculates the available space for volumes of that service class. Only storage pools that are candidates for provisioning by using the service class are used to calculate the available space.

Identifying candidates for provisioning is a filtering process. For example, if the service class specifies that RAID 10 is required, then any storage pool that is not configured at RAID 10 is eliminated as a candidate for provisioning. IBM Spectrum Control filters the set of available storage resources against certain properties of the

service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning. Only certain properties of a service class are used to identify candidates for provisioning. More advanced properties contain configuration instructions, and are not used to identify candidates for provisioning.

The candidates for provisioning are identified by eliminating storage pools that cannot provide the service class, as follows:

- If a storage pool is a primordial pool or if it is formatted for count key data (CKD) volumes, it is not a candidate for provisioning.
- If a storage system has a condition of Error or Unreachable, none of its internal storage pools are candidates for provisioning. If a storage pool has a status of Error, it is not a candidate for provisioning.

Remember: *Status* is different from *condition*. Status is the status reported by a storage resource without considering the status of its internal resources. Condition is the combined status of the monitored, internal resources for the higher-level resource, and the status of the higher-level resource itself. The condition of a storage system is shown on the Storage Systems page and the condition for a pool is shown on the Pools page. The condition of a resource might not have the same value as the status, which is used to eliminate pools as candidates for provisioning.

Although the condition of a resource might not have the same value as its status, the condition is always at least as severe as the status. The condition reflects the most severe status of any internal resource or the status of the resource itself. If the condition of a resource is Normal, then its status is also Normal.

Tip: The Service Classes page in the web-based GUI displays both the available space and the unavailable space for a service class. The Unavailable Space column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

- If a service class allows provisioning only from certain capacity pools, any storage pool that is not in a specified capacity pool is eliminated as a candidate for provisioning.

Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.

- A storage pool must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect whether the storage pool is a candidate for provisioning.

Table 67. Filtering candidates for provisioning based on service class properties

Service class property:	Candidates for provisioning:
Storage tier	If specified in the service class, storage pools that are assigned to any storage tier within the specified range are candidates for provisioning.

Table 67. Filtering candidates for provisioning based on service class properties (continued)

Service class property:	Candidates for provisioning:
RAID level	<p>If a RAID level is specified, only storage pools that are configured at the specified RAID level are candidates for provisioning.</p> <p>Exception: The IBM XIV Storage System uses a nontraditional approach to data redundancy and does not support any of conventional RAID levels. IBM Spectrum Control always considers XIV storage pools to be configured at RAID 10.</p> <p>Restriction: Storage pools that support multiple RAID levels are not candidates for provisioning when a service class specifies a RAID level. Storage pools that support multiple RAID levels are not candidates even when the specified RAID level is one of the RAID levels that are supported by the storage pool. Storage pools that support multiple RAID levels are only candidates for provisioning when the service class specifies that any RAID level is allowed.</p>
Virtualization	<p>If virtualization is required by the service class, only storage pools on IBM storage virtualizers are candidates for provisioning. The IBM storage virtualizers are SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified</p> <p>If virtualization is not allowed by the service class, only storage pools that are not on IBM storage virtualizers are candidates for provisioning.</p>
Encryption	<p>If encryption is required by the service class, only storage pools on DS8000 storage systems that are encrypted and configured to use the encryption group that is specified in the service class are candidates for provisioning.</p> <p>If encryption is not allowed by the service class, only pools that are not encrypted are candidates for provisioning.</p>
Thin provisioning	<p>If thin provisioning is required by the service class, only storage pools that are configured for thin provisioning are candidates for provisioning. Thin provisioning is supported by the XIV, DS8000, SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified. During provisioning, a thin-provisioned volume is created.</p> <p>If thin provisioning is not allowed by the service class, any storage pool that is not thin provisioned is a candidate for provisioning. In addition, any DS8000, SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified storage pool is a candidate for provisioning, regardless of whether it is configured for thin provisioning. During provisioning, the volume that is created is not thin provisioned.</p>
Compression	<p>If compression is required by the service class, at least one I/O group of an IBM storage virtualizer must be enabled for compression. Only pools of the IBM storage virtualizer are candidates for provisioning.</p>
Overallocation limit	<p>If thin provisioning is required by the service class, IBM Spectrum Control identifies the storage pools that can thin provision volumes as candidates for provisioning. If an overallocation limit is specified, IBM Spectrum Control refines that list to include only the storage pools that have an overallocation percentage that does not exceed the overallocation limit.</p> <p>The overallocation percentage for a pool is determined by using the following measurements, as shown on the Pools page:</p> <ul style="list-style-type: none"> • For an XIV storage pool, the Soft Space divided by the Capacity. • For other storage system pools, the Virtual Allocation.

Table 67. Filtering candidates for provisioning based on service class properties (continued)

Service class property:	Candidates for provisioning:
Resource Tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

After the storage pools that can provide the service class are identified, IBM Spectrum Control calculates the amount of space that is available for provisioning new volumes of the service class. The calculated volume capacity is shown in the Available Space column on the Service Classes page.

The available-space calculation differs depending on whether the block-storage service class requires thin provisioning, as follows:

- If thin provisioning is not allowed, the available space is the sum of the following measurements, as shown on the Pools page:
 - For each XIV System pool, the Available Soft Space minus any Reserved Pool Space
 - For other storage system pools, the Available Pool Space minus any Reserved Pool Space.
- If thin provisioning is required, the available space is the total of the following measurements, as shown on the Pools page, and calculations:
 - For each XIV System pool, the Available Soft Space minus any Reserved Pool Space
 - For other storage system pools, the pool Capacity multiplied by the overallocation limit that is specified in the service class, minus the Total Volume Capacity. Any Reserved Pool Space is subtracted from the result. For DS8000 storage pools, any space that is used by Track Space Efficient (TSE) volumes is subtracted from the calculated available space.

How the recommended placement of storage volumes is determined

When you request one or more new volumes by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new volumes. The same filtering that was used to determine the available space for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of storage pools from which the storage can be provisioned. This set can be all the storage pools that are known to IBM Spectrum Control, or can be limited to the storage pools in a particular capacity pool. Only the storage pools in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available space for the service class.
- Any storage pool that cannot provide the requested volume capacity is eliminated as a candidate for provisioning.
- If virtualization and VDisk mirroring are specified by the service class, then a VDisk copy of the volume must be made in a second storage pool on the same IBM storage virtualizer. If the storage pools on the IBM storage virtualizer cannot provide the requested capacity for both the volume and its mirrored copy, the pools are eliminated as candidates for provisioning.

- If an overallocation limit is specified in the service class, only the storage pools that will not exceed the limit after the new volume or volumes are provisioned are candidates for provisioning.

Remember: Storage pools that have an overallocation percentage that does not exceed the overallocation limit, are included in the available-space calculation. However, if the overallocation percentage would exceed the overallocation limit after the requested volume space is included, the pools are not candidates for provisioning.

After the storage pools that can satisfy the provisioning request are identified, IBM Spectrum Control finds the best location for the storage. The best location for the storage is based on the unallocated volume space in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is given to systems that have available performance data.

File storage: Calculating available space and determining the placement of shares

For each file-storage service class, IBM Spectrum Control calculates the amount of available space for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available space.

For each file-storage service class, IBM Spectrum Control calculates the amount of available space for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available space.

The Service Classes page displays the available space for each file-storage service class. The available space is the amount of space that is available across all IBM SONAS and IBM Storwize V7000 Unified file systems and unused IBM SONAS system NSDs that are known to IBM Spectrum Control, and that can satisfy the requirements of the service class. The set of file storage resources that are used to calculate the available space are also the candidates for provisioning by using the service class.

When you request a file share by using the Provision Storage pages, you specify the capacity that is required for the share, and you specify a service class. You can also specify a capacity pool from which shares must be provisioned. Based on the requested size of the share and whether a capacity pool is specified, the candidates for provisioning are further refined to include only the file storage resources that have enough available space and are in the specified capacity pool. From the remaining candidates for provisioning, IBM Spectrum Control creates a provisioning task that identifies the best location for the new share.

How available space for a file-storage service class is calculated

For each file-storage service class, IBM Spectrum Control calculates the available space for file shares of the service class. Only the file storage resources that are candidates for provisioning by using the service class are used to calculate the available space.

Identifying candidates for provisioning is a filtering process. For example, if the service class requires dedicated storage, then the share must be created on an unused NSD within an IBM SONAS system. File systems that could be candidates

for provisioning if shared storage is allowed are eliminated as candidates for provisioning when dedicated storage is required. IBM Spectrum Control filters the set of available storage resources against each attribute of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning.

The candidates for provisioning are identified by eliminating storage resources that cannot provide the service class, as follows:

- Only IBM SONAS and IBM Storwize V7000 Unified systems are supported for provisioning file shares. No other storage system type is a candidate for provisioning file shares.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal NSDs or file systems are included in the available-space calculation.
- If an NSD or file system has a consolidated status of Error, it is not included in the available-space calculation.

Tip: The status of file systems is not displayed on the GUI. View the status on the NSD on which a file system resides to determine the file system status.

- If the storage volume or NSD on which a file system resides has a consolidate status of Error, the file system is not included in the available-space calculation.

Remember: *Status* is different from *condition*. Status is the status reported by a storage resource without considering the status of its internal resources. Condition is the combined status of the monitored, internal resources for the higher-level resource, and the status of the higher-level resource itself. The condition of a storage system is shown on the Storage Systems page, the condition of a volume is shown on the Volumes page, and the condition of an NSD is shown on the storage system details page. The condition might not have the same value as the status, which is used to eliminate file storage resources as candidates for provisioning.

Although the condition of a resource might not have the same value as its status, the condition is always at least as severe as the status. The condition reflects the most severe status of any internal resource or the status of the resource itself. If the condition of a resource is Normal, then its status is also Normal.

Tip: The Service Classes page displays both the available space and the unavailable space for a service class. The Unavailable Space column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

- If a service class allows provisioning only from certain capacity pools, any storage resource that is not in a specified capacity pool is eliminated as a candidate for provisioning.

Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.

- A file system or NSD must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect the storage placement determination.

Table 68. Filtering candidates for provisioning based on service class properties

Service class attribute:	Candidates for provisioning:
Shared storage / Dedicated storage	If the service class specifies dedicated storage, the share can be created only on an unused NSD within an IBM SONAS system. File systems and other file storage resources are eliminated as candidates for provisioning.
Resource tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

After the file storage resources that can provide the service class are identified, IBM Spectrum Control calculates the amount of space that is available for providing new file shares of the service class.

The available-space calculation differs depending on whether the file-storage service class requires dedicated storage or allows shared storage, as follows:

- If the file-storage service class requires dedicated storage, the available space is the total available space on unused NSDs within each IBM SONAS system. You can verify that an NSD is unused, and view its available space, from the details page of an IBM SONAS system. To view information about the NSDs that are associated with an IBM SONAS system, click **Network Shared Disks**. If the NSD is unused, the File Systems column shows 0. The Available Space column shows the amount of space on the NSD, which is the measurement that is used in the available-space calculation for the service class.
- If the file-storage service class allows shared storage, the file system in which the NAS share is provisioned can contain other NAS shares. In this case the available space is the total of the following measurements:
 - The available space on unused NSDs within each IBM Scale Out Network Attached Storage (IBM SONAS) system.
 - The available space for each file system within each IBM SONAS and IBM Storwize V7000 Unified storage system. You can view the file system available space from the details page of a file storage system. To view information about the file systems that are in a storage system, click **File Systems**. The Available Space column shows the amount of available space on each file system, which is the measurement that is used in the available-space calculation for the service class.

How the recommended placement of file shares is determined

When you request a file share by using the Provision Storage pages, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new share. The same filtering that was used to determine the available space for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of file storage resources from which the storage can be provisioned. This set can be all the storage resources that are known to IBM Spectrum Control, or can be limited to the storage resources in a particular capacity pool. Only the storage resources in the specified set are

candidates for provisioning. This set might be more restrictive than the set used to calculate the available space for the service class.

- Any storage resource that cannot provide the requested capacity for the share is eliminated as a candidate for provisioning.

The shares are created on file systems or IBM SONAS NSDs. After the file systems and NSDs that can satisfy the provisioning request are identified, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available space on the file system or NSD.

Changing the default host definition for provisioned storage

Use the **setdscfg** command to change the default host definition for fibre-channel ports when provisioning storage from IBM Spectrum Virtualize devices.

You must have Administrator authority to use the **setdscfg** command.

When you provision storage from a SAN Volume Controller, IBM Storwize, or IBM FlashSystem family V840 or V9000 device, which does not have an existing host definition for the fibre-channel ports, IBM Spectrum Control creates a new host definition during the provisioning task. The new host definition can access only one I/O group on the source device that you specified when you created the provisioning task. To allow the new host definition to access any I/O group on the source device, use the management application for the source device or use the IBM Spectrum Control **setdscfg** command.

To change the host definition to allow access to any I/O groups on the source device:

1. Start the IBM Spectrum Control CLI by issuing the **tpctool** command.
2. Issue the **setdscfg** command with the following values for the **-property** and **-context** parameters:
 - a. Specify *Configuration.RestrictNewHostsToIogroup* as the property value.
 - b. Specify *DiskManager* as the context value and specify **false** as the context parameter.

For example, to allow the new host definition to access any I/O groups on the source device of the provisioning task:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****  
-property Configuration.RestrictNewHostsToIogroup -context DiskManager false
```

The value of *Configuration.RestrictNewHostsToIogroup* is set to **false** in the IBM Spectrum Control database. This removes the default restriction on the new host definition and allows access to any I/O groups on the provisioning source device.

Tip: To verify that the host definition restriction is changed, enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd *****  
-property Configuration.RestrictNewHostsToIogroup -context DiskManager
```

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

The prerequisites for provisioning storage are as follows:

- You must have the required license. “Functions that are available based on role and license” on page 23

- You must have Administrator privileges or permission to use one or more of the service classes that are created for block-storage or file-service provisioning.
- You must add and probe block-storage or file-storage resources, or both.

The following restrictions apply to volume provisioning:

- If the service class that is associated with the provisioning request requires multipathing or if automatic zoning is enabled, the fabrics and switches must be managed by IBM Spectrum Control. If the fabric is not managed by IBM Spectrum Control, the fabric-related configuration options are ignored when the provisioning task is executed.
- If you provision volumes to physical servers and servers that are virtual machines, you must create separate provisioning tasks using the GUI.
- If you provision multiple servers with storage volumes, the servers must all run on the same operating system. If the servers are virtual machines, the hypervisors that manage the servers must all run on the same operating system.
- If the provisioning request involves assigning multiple volumes to a resource, the volumes are sourced from pools in the same storage system.

NAS file shares can be provisioned only from IBM SONAS or Storwize V7000 Unified storage systems.

Specify the capacity that is required and the service class that you want to associate with the provisioning task. For shares, you must also specify at least one file access protocol. IBM Spectrum Control identifies the storage resources that can provide the storage capacity and that meet the criteria that are specified for the service class. From the set of storage resources that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage and creates a provisioning task. Depending on the privileges that are associated with your role, you can save, execute, or schedule the provisioning task.

Provisioning volumes to servers

You can provision volumes to one or more servers. The GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.

You can provision volumes to one or more servers. If a server is a virtual machine, the storage is provisioned to the hypervisor that is managing the server.

1. In the menu bar, go to **Servers > Servers..**

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Servers**.

2. Select one or more servers, and then select **Actions > Provision Storage**.
3. Follow the instructions to assign the storage volumes.

The GUI enables you to create a provisioning task.

Related tasks:

“Provisioning storage with the IBM Spectrum Control GUI” on page 311
You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning volumes to hypervisors

You can provision volumes to one or more hypervisors. To provision volumes, the GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.

Restriction: The multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.

You can provision volumes to one or more hypervisors.

1. **Servers > Hypervisors**

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Hypervisors**.

2. Select one or more hypervisors, and then click **Actions > Provision Storage**.
3. Follow the instructions to assign the storage volumes.

The GUI enables you to create a provisioning task.

Related tasks:

“Provisioning block storage” on page 316

Use the vSphere Web Client extension for IBM Spectrum Control to provision block storage on IBM Spectrum Control. You can also have a data store created on the storage volume.

“Provisioning storage with the IBM Spectrum Control GUI” on page 311

You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning volumes to a server cluster

You can assign volumes to a cluster of servers with the GUI.

If the cluster is on servers that are virtual machines, the storage is assigned to the hypervisors that manage the servers. The volumes are assigned to all of the nodes in the cluster.

1. In the menu bar, go to **Servers > Servers**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Servers**, and click **Clusters**.

2. Select a cluster, and then click **Actions > Provision Storage**.
3. Follow the instructions to assign the storage volumes.

The GUI enabled you to create a provisioning task.

Related tasks:

“Provisioning storage with the IBM Spectrum Control GUI” on page 311

You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning volumes to a hypervisor cluster

You can assign volumes to a cluster of hypervisors.

Restriction: The multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.

The volumes are assigned to all of the nodes in the hypervisor cluster.

1. In the menu bar, go to **Servers > Hypervisors**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Hypervisors** and click **Clusters**.

2. Select a cluster, and then click **Actions > Provision Storage**.
3. Follow the instructions to assign the storage volumes.

The GUI enabled you to create a provisioning task.

Related tasks:

“Provisioning block storage” on page 316

Use the vSphere Web Client extension for IBM Spectrum Control to provision block storage on IBM Spectrum Control. You can also have a data store created on the storage volume.

“Provisioning storage with the IBM Spectrum Control GUI” on page 311

You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning shares to servers

You can provision a network-attached storage (NAS) file share to one or more servers. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You can provision a NAS file share to one or more servers. The share is made available to the selected servers, but is not automatically mounted from the selected servers.

1. In the menu bar, go to **Servers > Servers..**

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Servers**.

2. Select one or more servers, and then select **Actions > Provision Storage**.
3. Follow the instructions to assign the share.

The GUI enabled you to create a provisioning task.

Related tasks:

“Provisioning storage with the IBM Spectrum Control GUI” on page 311

You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning shares to hypervisors

You can provision a network-attached storage (NAS) file share to one or more hypervisors. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You can provision a NAS file share to one or more hypervisors. The share is made available to the selected hypervisors, but is not automatically mounted from the selected hypervisors.

1. In the menu bar, go to **Servers > Hypervisors**.

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Hypervisors**.

2. Select one or more hypervisors, and then select **Actions > Provision Storage**.
3. Follow the instructions to assign the share.

The GUI enabled you to create a provisioning task.

Related tasks:

“Provisioning file storage” on page 318

Use the vSphere Web Client extension for IBM Spectrum Control to provision file storage on IBM Spectrum Control.

“Provisioning storage with the IBM Spectrum Control GUI” on page 311
You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning shares to a server cluster

You can assign a network-attached storage (NAS) file share to a server cluster using the GUI.

You can provision a NAS file share to a server cluster. The share is made available to the servers in the cluster, but is not automatically mounted from the selected servers.

To provision a NAS file share to a server cluster, follow these steps:

1. In the menu bar, go to **Servers > Servers**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Servers**, and click **Clusters**.

2. Select a cluster, and then select **Actions > Provision Storage**.
3. Follow the instructions to assign the share.

A provisioning task was created for a server cluster.

Related tasks:

“Provisioning storage with the IBM Spectrum Control GUI” on page 311
You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning shares to a hypervisor cluster

You can assign a network-attached storage (NAS) file share to a hypervisor cluster.

You can provision a NAS file share to a hypervisor cluster. The share is made available to the hypervisors in the cluster, but is not automatically mounted from the selected hypervisors.

To provision a NAS file share to a hypervisor cluster, follow these steps:

1. In the menu bar, go to **Servers > Hypervisors**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Hypervisors** and click **Clusters**.

2. Select a cluster, and then select **Actions > Provision Storage**.
3. Follow the instructions to assign the share.

A provisioning task was created for a hypervisor cluster.

Related tasks:

“Provisioning storage with the IBM Spectrum Control GUI” on page 311
You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

Provisioning with zone control

When you provision storage, you can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, IBM Spectrum Control can create zones during storage provisioning to connect a server to a storage system.

When IBM Spectrum Control creates a provisioning task, it identifies the best location for the new storage that satisfies the requirements of the service class. If

automatic zoning is enabled, then, during provisioning, existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more zones are created between a host initiator port and a controller, node, or module port.

IBM Spectrum Control provides the support for zone control and management of Brocade switches and fabrics, and Cisco switches and fabrics.

The Cisco support includes:

- Standard Cisco zone control capabilities
- Cisco Enhanced Zoning features.

To set or modify the zoning policy, complete the following steps:

1. In the menu bar, go to **Advanced Analytics > Provisioning**, and click **Set Zoning Policy**.
2. Enable or disable automatic zoning:
 - Enable automatic zoning to allow IBM Spectrum Control to create zones during storage provisioning. The naming convention for new zones is *host_storage-system_suffix-number*. When you enable automatic zoning, you can specify the following options:

Zone name prefix

If you specify a zone name prefix, all zones that are automatically created by IBM Spectrum Control are prefixed with a string you specify. This prefix can help you identify which zones were automatically created, and which were created manually.

Make changes to the active zone set

Specify whether changes are made to the active zone set or to a new inactive zone set. If this check box is selected, one or more new zones are added to the active zone set. If the check box is cleared, a new inactive zone set is created and the new zone or zones are added to this new inactive zone set. The new inactive zone set will contain only the new zones.

- Disable automatic zoning if you want IBM Spectrum Control to use only existing zones. If you disable automatic zoning, no changes will be made to the zoning for the fabric.
3. Click **Save**.

Provisioning storage with the vSphere Web Client extension

You can use the vSphere Web Client extension for IBM Spectrum Control to manage and monitor your virtual environment and storage devices.

You can provision block and file storage to hypervisors from storage devices that are monitored by IBM Spectrum Control. You can monitor the progress of your storage provisioning activities in task information in the vSphere Web Client.

Provisioning block storage

Use the vSphere Web Client extension for IBM Spectrum Control to provision block storage on IBM Spectrum Control. You can also have a data store created on the storage volume.


Ensure that you complete the following tasks in IBM Spectrum Control:

- If you want to provision storage by using service classes to describe performance and storage requirements and to implement multipathing policies,

ensure that the service classes are created in IBM Spectrum Control. IBM Spectrum Control provides a number of predefined service classes, which you can use, modify, or delete.

- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in IBM Spectrum Control.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into IBM Spectrum Control. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by IBM Spectrum Control.
- Ensure that the hypervisor has fabric connectivity. If automatic zoning is enabled, the fabrics and switches must be managed by IBM Spectrum Control.

For more information about configuring for storage cloud provisioning, see

 http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.14/com.ibm.spectrum.sc.doc/fqz0_c_cloud_config.html.

Start the vSphere Web Client.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select the hypervisor for which you want to provision storage.
3. Click **Actions** and select **All TPC Actions**.
4. Select **Provision Block Storage....**
5. Enter a size for the total usable storage space on the volume, and select the unit of measurement from the list.
6. Select the service class for the block storage that you want to associate with the volume. You can select only service classes with available capacity that you are authorized to use.
7. Select the capacity pool from which you want the volume to be provisioned. You can select only capacity pools with available capacity.
8. Optional: If you want to create a data store on the storage volume when it is provisioned, select the **Create datastore** check box, and enter a name for the data store. The minimum volume size for creating a data store is 2 GiB.
9. Click **OK** to provision the volume.
10. To monitor the progress of the provisioning task, click the **Monitor** tab, and click the **Tasks** tab.
11. You can monitor the progress of the creation of the data store by clicking the **Monitor** tab and then clicking the **Tasks** tab. The data store creation task is called **Create VMFS datastore**.
12. You can monitor the progress of the multipathing task by clicking the **Monitor** tab and clicking the **Tasks** tab. The multipathing task is called **Set logical unit policy**.

Related concepts:

“Service classes” on page 52

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

“Capacity pools” on page 56

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can

also be restricted to resources in a capacity pool.

Related tasks:

“Provisioning file storage”

Use the vSphere Web Client extension for IBM Spectrum Control to provision file storage on IBM Spectrum Control.

Provisioning file storage

Use the vSphere Web Client extension for IBM Spectrum Control to provision file storage on IBM Spectrum Control.

Ensure that you complete the following tasks in IBM Spectrum Control:

- If you want to provision storage by using service classes to describe performance and storage requirements, ensure that the service classes are created in IBM Spectrum Control. IBM Spectrum Control provides a number of predefined service classes, which you can use, modify, or delete.
- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in IBM Spectrum Control.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into IBM Spectrum Control. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by IBM Spectrum Control.
- Determine whether you must set access path host name replacements. For more information, see “Special consideration for provisioning by using the VMware vSphere Web Client extension” on page 56.

For more information, see “Cloud configuration” on page 52.

Start the vSphere Web Client.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select the hypervisor for which you want to provision storage.
3. Click the **Actions** menu item, and select **All TPC Actions**.
4. To provision file storage, select **Provision File Storage...**
5. On the **Share Requirements** page, enter a name for the share.
6. Enter a size for the total usable storage space on the share, and select the unit of measurement from the list.
7. Select the service class for the file storage that you want to associate with the share. You can select only service classes with available capacity that you are authorized to use.
8. Select the capacity pool from which you want the share to be provisioned. You can select only capacity pools with available capacity.
9. Enter the name of the data store that you want to associate with the file system.
10. Click **Next** to continue.
11. On the **Share Permissions** page, specify the user and group with permission to mount the share on the hypervisor.
12. Select the **Read-only** check box if you want to limit the user or group to viewing files on the share.
13. Select the **Secure** check box to create the share with the secure flag set.
14. Click **OK** to provision the share.
15. To monitor the progress of the provisioning task, click the **Monitor** tab, and click the **Tasks** tab.

16. You can also monitor the progress of the creation of the data store by clicking the **Monitor** tab and then clicking the **Tasks** tab. The data store creation task is called **Create NAS datastore**.

Related concepts:

“Service classes” on page 52

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

“Capacity pools” on page 56

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Related tasks:

“Setting access path host replacements” on page 300

If you are going to use the VMware vSphere Web Client extension to provision file shares, verify whether you must provide IBM Spectrum Control with an access path host replacement.

“Provisioning block storage” on page 316

Use the vSphere Web Client extension for IBM Spectrum Control to provision block storage on IBM Spectrum Control. You can also have a data store created on the storage volume.

Optimizing storage tiering

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

Before you analyze tiering, complete the following tasks:

- Probe the storage virtualizers
- Collect performance data
- Set the tier level of the storage pools that you analyze and that you select as target pools

Analyzing tiering		
	1 Collect data Probe the storage virtualizers to collect information about the logical devices, such as the volumes and virtual disks, that are associated with the storage virtualizer.	
	2 Collect performance data Collect performance data to gauge the average workload activity of the volumes that you want to analyze.	
	3 Assign tier levels Assign tier levels to the pools that contain the volumes that you want to analyze and to the target pools.	
	4 Choose the source storage resources Select the resources that you want to analyze: ↳ Servers ↳ Hypervisors ↳ Storage virtualizers ↳ Storage pools ↳ Volumes	
	5 Choose the target pools You select the target pools to place volumes that meet the tiering criteria that you specify.	
	6 Choose the performance data Choose the analysis period, the collection days, and time period that you want to use to analyze the tiering of the volumes.	
	7 Specify the tiering criteria Specify the tiering criteria and the pool saturation rates for the target pools.	

The purpose of analyzing tiering is to tier volumes based on the I/O rate or I/O density thresholds that you set. Depending on the conditions that are set, recommendations are generated.

Storage costs can be reduced, for example, by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers. Storage performance can be improved and used more efficiently by moving volumes with heavy workloads to the tiers that meet the workload requirements of the volumes. When you implement the recommendations, you ensure that the volumes are placed on the tiers that best match the workload requirements of the volumes.

You can optimize the placement of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified block storage
- FlashSystem V9000 and FlashSystem V840

Learn more: For information about the components that can be used with IBM Spectrum Control Version 5.2 (or later), see [IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446](http://www.ibm.com/support/docview.wss?uid=swg21386446).

Volumes can be moved to tiered storage pools on the same storage virtualizer, but volumes cannot be moved from one storage virtualizer pool to another storage virtualizer pool.

Optimizing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

Related tasks:

“Analyzing and re-tiering volumes in pools on tier 1” on page 87

In this scenario, you analyze the current placement of volumes in pools on tier 1 so that you can move volumes with low workloads from tier 1 to tier 2 or tier 3 pools.

“Collocating volumes” on page 90

In this scenario, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes when you enter the criteria for analyzing tiering and balancing pools.

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Add the pools in your storage environment to the tiers that best meet the performance requirements of your storage data.

You can view the total capacity by tier and the capacity that is consumed by the pools that are not tiered. You can also view information about recent storage growth for each tier, including the weekly and average storage growth, and the date when the storage space of the tier is expected to run out.

1. Click **Groups > Tiers**.
2. Review the information that is shown about the distribution of capacity across each tier.
3. Optional: To place pools on lower or higher tiers, complete these steps:
 - a. Click the number in the **Pools** column.
 - b. Select the pools that you want to re-tier.
 - c. Right-click the pools, click **Set Tier**, and click the tier level.
4. Optional: To assign untiered pools to tiers, click **Untiered Pools**.

If the storage space for tier 1 pools is running out, you can down-tier the volumes in the tier-1 pools to lower tiers. Alternatively, you can add more storage to meet the requirements of your tier-1 storage data.

Related tasks:

“Setting the tier level of storage pools”

To analyze tiering and to balance pools, you must set the tier level of storage pools.

“Analyzing tiering by storage virtualizers” on page 323

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

“Analyzing tiering by volumes” on page 323

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

“Tiering volumes by I/O density and I/O rate” on page 323

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Before you analyze tiering or balance pools, you must set the tier level of pools.

To analyze tiering, you must set the tier level of the source pools and the target pools. The source pools are the pools that you select for analysis. The target pools are the pools that are used to relocate volumes to lower or higher tiers of storage.

Tip: You can also set or modify the tier level of pools on the **General** tab of the View Properties page for the storage pool.

Child pools have the same tier level as the parent pool. So if you change the tier level of the parent pool or of the child pool, all of the pools in the parent-child relationship are set to the same tier level.

You can assign tier levels to storage systems that are not storage virtualizers. However, you cannot balance the pools or analyze the tiering of pools other than storage virtualizer pools.

1. From the **Storage** menu, click **Pools**.
2. On the Pools page, select one or more storage pools.
3. Click **Set > Tier**, and then select a tier level. To remove the tier level, click **None**.

The tier level that you selected for the storage pools is shown in the **Tier** column.

Related tasks:

“Investigating the capacity of tiered storage” on page 321

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Analyzing tiering by servers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to servers.

The servers that you select must be connected to one or more storage virtualizers.

1. From the **Servers** menu, click **Servers**.
2. Right-click one or more servers, and then click **Analyze Tiering**. Alternatively, you can click **Actions > Analyze Tiering**.

Analyzing tiering by hypervisors

Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to hypervisors.

The hypervisors that you select must be connected to one or more storage virtualizers.

1. From the **Servers** menu, click **Hypervisors**.
2. Right-click one or more hypervisors, and then click **Analyze Tiering**. Alternatively, you can click **Actions > Analyze Tiering**.

Analyzing tiering by storage virtualizers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

1. From the **Storage** menu, click **Storage Systems**.
2. Right-click one or more storage virtualizers, and then click **Analyze Tiering**. Alternatively, you can click **Actions > Analyze Tiering**.

Related tasks:

“Investigating the capacity of tiered storage” on page 321

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Analyzing tiering by storage pools

Optimize tiering by analyzing the tiering of volumes in storage pools.

1. From the **Storage > Pools** menu, click **Pools**.
2. Right-click one or more storage pools, and then click **Analyze Tiering**. Alternatively, you can click **Actions > Analyze Tiering**.

Analyzing tiering by volumes

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

1. From the **Storage** menu, click **Volumes**.
2. Right-click one or more volumes, and then click **Analyze Tiering**. Alternatively, you can click **Actions > Analyze Tiering**.

Related tasks:

“Investigating the capacity of tiered storage” on page 321

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Tiering volumes by I/O density and I/O rate

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

1. Select the resources that you want to analyze. The source storage pools that are related to the resources that you selected are analyzed to determine whether they meet the workload requirements of the volumes. If the workload requirements of the volume in its current tier are not met, the volume is a candidate for relocation.
2. Select the target storage pools.
3. Include or exclude volumes in mirrored volume relationships from the analysis.

Mirrored volumes: If you want to optimize mirrored volumes, select one or more target pools that do not contain either the primary volume or the secondary volume. Primary volumes and secondary volumes cannot be placed in the same destination pool.

4. Select the period that you want to use to analyze the performance data. You can choose the number of days, days of the week, and hours of the day for the analysis to occur.
5. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool. Recommendations are only generated to move one or more of the related volumes when all of the related volumes can be placed in the same destination storage pool.

Multiple host connections to the same hypervisor or server: If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

6. Specify the type of threshold and the threshold values for moving volumes to higher and lower tiers.
7. Click **Analyze**.

When you complete entering the information to analyze the tiering of the volumes, recommendations for tiering are shown on the Tiering Analysis page. You can create a schedule to run the analysis that is based on the analysis criteria that you entered.

Related tasks:

“Investigating the capacity of tiered storage” on page 321

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Modifying the criteria for analyzing tiering

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

You must run the tiering analysis to create a tiering analysis task.

The criteria that you modify are used to analyze the source volumes and target volumes that were selected when you created the tiering analysis.

1. From the **Home** menu, click **Tasks**.
2. Right-click a tiering analysis task, and click **Edit**.

3. Modify the tiering thresholds, and then click **Analyze**.

The source volumes are analyzed to determine whether one or more of the volumes require re-tiering based on the criteria that you specified. If one or more of the volumes require re-tiering, recommendations are generated to move the volumes to the destination pools that meet the criteria that you specified.

Optimizing storage pools


Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

Before you balance pools, complete the following tasks:

- Probe the storage virtualizer. You must probe the storage virtualizer to collect information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.
- Collect performance data. To gauge the average activity level of the volumes and pools, you must collect sufficient performance data for the storage virtualizer.
- Assign a tier level to the pools that you want to balance.

You can balance the workload of volumes in pools on the same tier on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified block storage
- FlashSystem V9000 and FlashSystem V840

Learn more: For information about the components that can be used with IBM Spectrum Control Version 5.2 (or later), see  IBM Spectrum Control interoperability matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

To balance pools, you must select two or more pools on the same tier level and on the same storage virtualizer. For example, you select three tier 1 pools, and four tier 3 pools. The pools on tier 1 are analyzed and balanced, and the pools on tier 3 are analyzed and balanced.

Optimizing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

Related tasks:

“Setting the tier level of storage pools” on page 322

To analyze tiering and to balance pools, you must set the tier level of storage pools.

“Redistributing volume workloads to balance pools” on page 86

In this scenario, you identify performance hotspots and then run the balance pools analysis to help resolve them.

Balancing pools

You balance pools to distribute the workload of volumes across pools on the same tier and on the same storage virtualizer. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier and on the same storage virtualizer.

1. From the **Storage** menu, click **Pools**.
2. Right-click two or more pools on the same tier level and on the same storage virtualizer, and then click **Balance Pools**.

Tip: If you want to analyze pools that contain volumes in mirrored volume relationships, select one or more pools that do not contain either the primary volume or the secondary volume. The primary volume and the secondary volume cannot be placed in the same destination pool.

3. Choose whether to analyze volumes in mirrored volume relationships. These options are only available if you select one or more volumes with mirrored volumes.
4. Select the period that you want to use to analyze the performance data.
5. Specify the placement of optimized volumes in capacity pools. These options are only available if you select one or more volumes in capacity pools.
6. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool.
7. Click **Analyze**.

When you finish entering the information for balancing the pools, a task is created that shows the recommendations on the Balance Analysis page. You can implement the recommended changes, or you can create a schedule to implement the changes later.

Analyzing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class unless the destination pool can meet the requirements of the service class.

Related tasks:

“Redistributing volume workloads to balance pools” on page 86

In this scenario, you identify performance hotspots and then run the balance pools analysis to help resolve them.

Modifying the criteria for balancing pools

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

To create a balance analysis task, you must run the analysis for balancing the pools.

The criteria that you modify are used to analyze the storage pools that you selected when you entered the criteria for balancing the pools.

1. From the **Home** menu, click **Tasks**.
2. Right-click a balance analysis task, and click **Edit**.
3. Make your changes, and then click **Analyze**.

Based on the criteria that you specified, the pools are analyzed to determine whether the pools require balancing. If one or more pools require balancing, recommendations are generated to redistribute the volumes to pools on the same tier of storage.

Criteria for identifying the pools that require balancing

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

To determine the relative activity of pools on the same tier and on the same back-end storage system, the following values are calculated:

- The current activity level of the pool is calculated.
- The activity deviation percentage of the pool is calculated.

Tip: On the Pools page, the activity level of the pool is shown in the **Activity** column. The difference between the activity level of the pool and all of the other pools on the same tier and storage system is shown in the **Activity Deviation (%)** column.

The following formulas are used to calculate the activity level and the activity deviation percentage for pools.

Table 69. Formulas for calculating the activity level and activity deviation percentage values for pools

Value calculated	Formula that is used in the calculation
Activity level of the pool	$[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%}) + \text{Write I/O Rate}] \div \text{Total Pool Capacity}$
Activity level of the pool on XIV systems	$(\text{Total I/O Rate} \div \text{Total Capacity})$
Activity deviation percentage of the pool	$[(\text{Pool Activity Level} - \text{Average Activity Level}) \div \text{Average Activity Level}] \times 100$ A positive value indicates that the relative activity level of the pool is above average. A negative value indicates that the relative activity level of the pool is below average.

If the activity deviation percentage of the pool lies within the range of <-10% to >10%, the pool is considered to be balanced. If the activity deviation percentage of the pool exceeds the deviation threshold of 10%, the pool is a candidate for balancing.

For example, you want to identify the tier-1 pools that require balancing on a storage environment that consists of a Storwize V7000 Unified storage virtualizer with one DS8000 storage system. You open the Pools page and view the activity deviation percentage values for the tier-1 pools.

The Pools page shows that three tier-1 pools have the following activity levels and activity deviation percentages.

Table 70. Tier-1 pools activity and activity deviation percentages

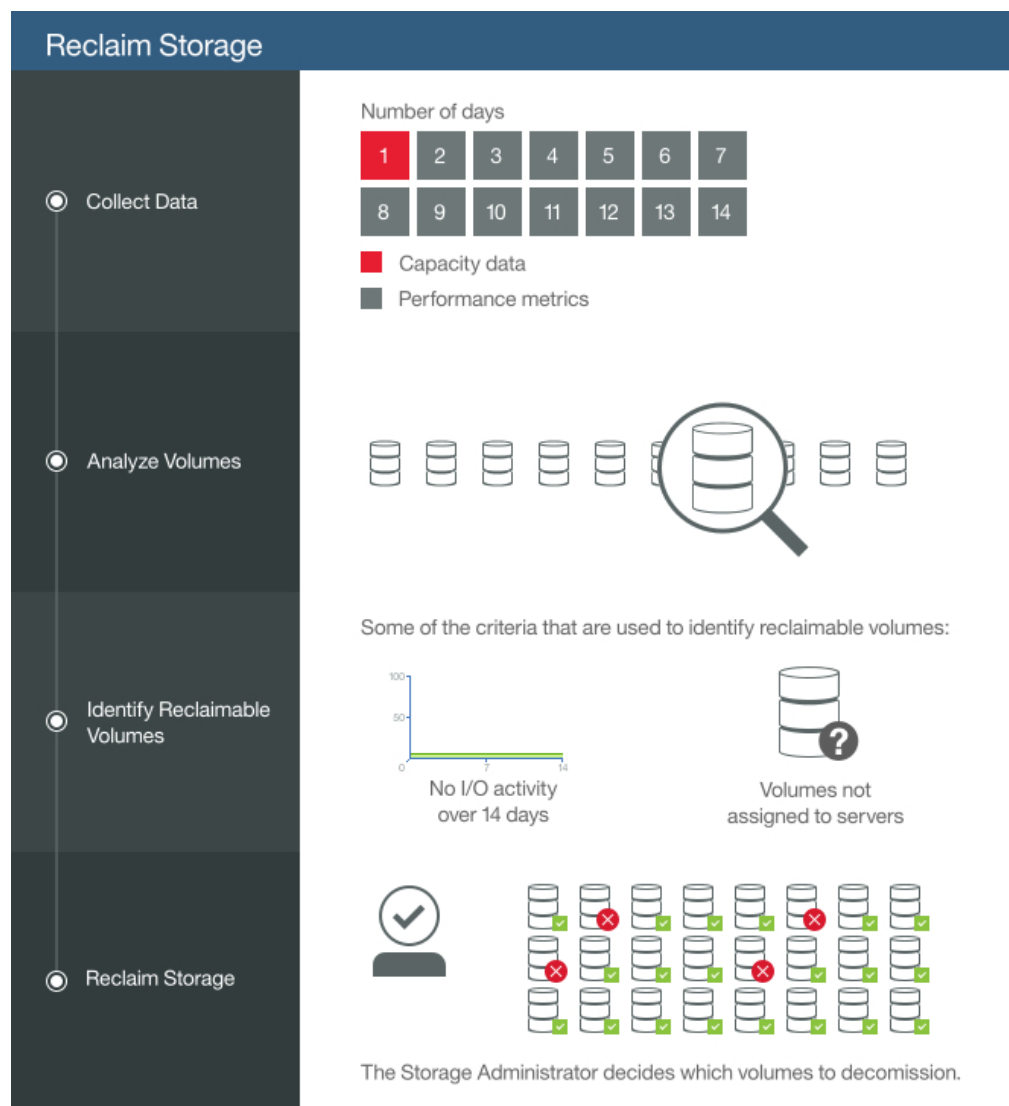
Pool	Activity	Activity deviation percentage
pool_1	110	-3%
pool_2	90	-20%
pool_3	140	24%

The pool_3 pool is a candidate for balancing because the activity deviation value for the pool is 14% higher than the deviation threshold of 10%.

The pool_1 pool is considered to be balanced because the activity deviation percentage of the pool lies within the range for balanced pools. The pool_1 pool and the pool_2 pool, are potential target pools for the volumes that must be moved from the pool_3 pool.

Reclaiming storage

You can view the amount of storage that can be reclaimed in your data center and for each of the block storage systems that are monitored. Based on volume reclamation recommendations, you can reclaim the unused space and more effectively allocate volume space on your storage systems.



To reclaim storage, you must enable performance monitoring to collect the performance metrics about the I/O activity of the volumes. Only block storage systems that support performance monitoring can be analyzed for I/O activity.

By default, capacity data is collected for 1 day and performance metrics are collected for 14 days to identify the block storage that can be reclaimed.

Supported storage systems: Go to IBM Spectrum Control interoperability matrix, and then click the release number in the **Storage** column.

When the data is collected, all of the block storage systems that you monitor in your data center are analyzed. Such criteria as whether the volumes have no I/O activity over 14 days or whether the volumes are assigned to servers is considered to determine whether the volumes are reclaimable.

Restrictions:

- Volumes are identified as reclaimable in IBM Spectrum Accelerate storage systems based on I/O activity. Information about the assignment of volumes to servers is not available.
- Private volumes in EMC storage systems are not included in reclamation analysis.
- Volumes in IBM FlashSystem A9000 storage systems are identified as reclaimable, but the total reclaimable capacity for each volume is unknown and is displayed as zero. This restriction is a current limitation of IBM FlashSystem A9000 storage.

The process for identifying the replica, VDisk mirrored, or FlashCopy® volumes that are reclaimable is as follows:

Replica volumes

The source volumes of volumes with replicas, such as volumes that use Metro Mirror or Global Mirror copy services, are analyzed to determine whether the volumes are reclaimable.

VDisk mirrored volumes

Both copies of the VDisk mirrored volume are analyzed to determine whether the volumes are reclaimable.

FlashCopy volumes

The target volumes of volumes in FlashCopy relationships are analyzed if the source volumes are identified as reclaimable.

You review the information about the volumes that can be reclaimed and decide whether you want to decommission the volumes that are not being used.

Viewing storage reclamation

Use the storage reclamation analysis to see information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view a list of the reclaimable volumes.

Before you view recommendations for reclaiming storage, complete the following tasks:

- Add the storage systems for monitoring and schedule a storage systems probe.
- If your storage environment is configured for storage tiers, ensure that the amount of unused space on tier one is minimal.

By default, the period for analyzing performance data, which is used for generating reclamation recommendations, is set to 14 days. For a more thorough analysis, you can change this period.

View information about volumes that are analyzed and recommended for reclamation.

1. From the **Advanced Analytics** menu, click **Reclamation**.
2. On the View by Reclaimable Capacity page, view the charts and tables for information about reclaimable storage space, volumes that are recommended for reclamation, and volumes that are excluded from analysis.

On the reclamation donut chart, you can view the amount of storage space that is used and the amount of storage space that can be reclaimed. You can also see an estimate of the storage space that can be saved when the volumes that are listed in the table are reclaimed.

On the reclamation by tier bar charts, view the amount of storage space that can be reclaimed on each tier of storage that is defined in your data center.

On the **Recommendations** tab, view the volumes that are identified as potential candidates for reclamation.

On the **Excluded** tab, view the volumes that are excluded from the analysis for reclamation recommendations.

Viewing storage reclamation by storage system

Use the reclamation view of storage systems for detailed information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view information about the storage systems that contain the reclaimable volumes.

Before you view recommendations for reclaiming storage, complete the following tasks:

- Add the storage systems for monitoring and schedule a storage systems probe.
- If your storage environment is configured for storage tiers, ensure that the amount of unused space on tier one is minimal.

By default, the period for analyzing performance data, which is used for generating reclamation recommendations, is set to 14 days. For a more thorough analysis, you can change this period.

View reclamation information about storage systems that contain volumes that are analyzed and recommended for reclamation.

1. From the **Advanced Analytics** menu, click **Reclamation**.
2. On the View by Storage Systems page, view the charts and table for information about reclaimable storage space and for information about the storage systems that contain the reclaimable volumes.

On the reclamation donut chart, you can view the amount of reclaimable space and the non-reclaimable allocated space. If you defined storage tiers, you can view the amount of storage that can be saved for each tier. You can also see an estimate of the storage space that can be saved when the volumes of the storage systems that are listed in the table are reclaimed.

On the reclamation bar chart, view the amount of storage space that can be reclaimed on each storage system in your data center.

In the table, view space, capacity, and reclamation information about each storage system.

Excluding volumes from reclamation analysis

You can select volumes to not include them in the analysis for reclamation recommendations. The volumes are excluded from the reclamation charts.

On the View by Reclaimable Capacity page, you can select a volume or multiple volumes for exclusion from the analysis for reclamation recommendations.

1. From the **Advanced Analytics** menu, click **Reclamation**.
2. On the View by Reclaimable Capacity page, to exclude a volume from the analysis, right-click the volume on the **Recommendations** table, and select **Exclude from Analysis**.

The volume is removed from the **Recommendations** table and added to the **Excluded** table, and the charts are refreshed.

Tip: To include a volume in the analysis, on the **Excluded** table, right-click the volume and select **Include in Analysis**. The volume is removed from the **Excluded** table and added to the **Recommendations** table, and the charts are refreshed.

Transforming and migrating volumes

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Before you transform storage, you must probe the storage virtualizers. The probe collects information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.

You can convert or move volumes in pools on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- FlashSystem V9000 and FlashSystem V840

For volumes in mirrored volume relationships, you can convert the primary volume, the secondary volume, or both volumes.

Restrictions: The restrictions for converting or moving volumes are as follows:

- To move volumes, the destination pool must be on the same storage virtualizer. You cannot move a volume from a pool on one storage virtualizer to a pool on another storage virtualizer.
- To convert or move volumes, image mode volumes must be converted to managed mode volumes.
- To convert fully allocated volumes to compressed volumes, you must have the IBM Real-time Compression™ license.

Migrating volumes with service classes: A warning is generated when you move one or more volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

1. In the menu bar in the web-based GUI, go to **Storage > Volumes**.
2. Click one or more storage volumes, and then select **Transform Storage** from the context-sensitive menu. Alternatively, you can click **Actions > Transform Storage**.

When you complete entering criteria for transforming storage, the Transform Plan page is opened. A task is created which you can run to implement the recommendations.

Related tasks:

“Modifying the period for analyzing performance data”

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. You can extend or shorten the default period for a more accurate analysis.

Modifying the period for analyzing performance data

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. You can extend or shorten the default period for a more accurate analysis.

You must have Administrator authority to use the **setdscfg** command.

By default, the period for analyzing performance data is set to 14 days. When you change the number of days, the period for analyzing performance data is modified for the advanced analytics about storage, that is, tiering, balancing, transforming, and reclamation.

When you transform storage, the performance data that is collected about the pools on storage virtualizers is analyzed. The period that is set for analyzing performance data is used to calculate the average workload activity of the volumes that are selected for analysis. To provide a more accurate analysis of the workload activity of the volumes in your storage environment, you can extend or shorten the period.

For reclamation, the performance data that is collected for the analysis period about all block storage pools is analyzed, regardless of whether the storage is virtualized.

To change the analysis period, complete the following steps:

1. To start the IBM Spectrum Control CLI, issue the **tpctool** command.
2. Issue the **setdscfg** command, and then set the following values by using the **-property** and the **-context** parameters:
 - a. Specify *tiering_pm_days* as the property value.
 - b. Specify *tiering* as the context value, and then specify the number of days for the analysis period.

For example, to change the analysis period from the default value of 14 days to 28 days, issue the following command:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****  
-property tiering_pm_days -context tiering 28
```

The default value is changed to the number of days that you specified, for example 28.

Tip: To verify that the analysis period is changed, you enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd ***** -context tiering
```

Modifying the properties of resources


You can modify the properties of the resources that are monitored by IBM Spectrum Control in the GUI.

The following information summarizes the resources and properties that you can modify, and the steps for modifying each type of resource.

Resource	Properties that can be modified	Steps
Storage systems	<ul style="list-style-type: none"> Name (label) Location Custom tags (3) 	<ol style="list-style-type: none"> Depending on the type of storage system that you want to modify, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. Right-click the storage system that you want to modify and select View Properties. On the properties notebook, click Edit. Modify the properties and click Save.
Pools	<ul style="list-style-type: none"> Tier Back-end storage system type Back-end storage RAID level Back-end storage disk type Back-end storage disks <p>Tip: You can edit the back-end values if the back-end storage system for the related storage virtualizer was not probed and the values for the pool are unknown. If you manually enter values for these properties, IBM Spectrum Control uses the values to help calculate the approximate read I/O capability of the pool.</p>	<ol style="list-style-type: none"> In the menu bar, go to Storage > Pools. Right-click the pool that you want to modify and select View Properties. On the properties notebook, click Edit. Modify the properties and click Save.

Resource	Properties that can be modified	Steps
Servers	<p>Servers with Storage Resource agents</p> <p>You can modify the following properties for servers where agents are deployed:</p> <ul style="list-style-type: none"> • Location • Custom tags (3) • Fabric functions • Agent trace • Trace level • Number of trace files • Trace file size • Run scripts on agent <p>Servers without Storage Resource agents (agentless)</p> <p>You can modify the following properties for agentless servers:</p> <ul style="list-style-type: none"> • Name (label) • OS type • IP address • Virtual machine • Location • Custom tags (3) 	<ol style="list-style-type: none"> 1. In the menu bar, go to Servers > Servers. 2. Right-click the server that you want to modify and select View Properties. 3. On the properties notebook, click Edit. 4. Modify the properties and click Save.
Hypervisors	<ul style="list-style-type: none"> • Location • Custom tags (3) 	<ol style="list-style-type: none"> 1. In the menu bar, go to Servers > Hypervisors. 2. Right-click the hypervisor that you want to modify and select View Properties. 3. On the properties notebook, click Edit. 4. Modify the properties and click Save.
Switches	<ul style="list-style-type: none"> • Name (label) • Location • Custom tags (3) 	<ol style="list-style-type: none"> 1. In the menu bar, go to Network > Switches. 2. Right-click the switch that you want to modify and select View Properties. 3. On the properties notebook, click Edit. 4. Modify the properties and click Save.

Resource	Properties that can be modified	Steps
Fabrics	<ul style="list-style-type: none"> • Name (label) • Location • Custom tags (3) 	<ol style="list-style-type: none"> 1. In the menu bar, go to Network > Fabrics. 2. Right-click the fabric that you want to modify and select View Properties. 3. On the properties notebook, click Edit. 4. Modify the properties and click Save.

Tip: To view the description of a property that you can modify for a resource, hover the mouse pointer over the hover icon  for that property.

Opening the management GUI for a resource

Resources such as storage systems and switches can have their own management GUIs or element managers. In IBM Spectrum Control, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by IBM Spectrum Control.
 - The resource has a management GUI that is web-based
1. In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want open the GUI for a block storage system, go to **Storage > Block Storage Systems**. If you want open the GUI for a switch, go to **Network > Switches**.
 2. Right-click the resource and select the action to open the GUI. For storage systems, select **Open Storage System GUI**. For a switch, select **Open Switch GUI**. The start page for the management GUI is opened in a separate web browser window.

Opening the DS Storage Manager GUI for a DS8000

If the firmware level for a DS8000 is earlier than R6.2, you must complete extra configuration steps before you can open its management GUI. To determine the firmware version of a DS8000 and enable the ability to open its management GUI, complete the following steps:

1. Start IBM Spectrum Control.
2. In the menu bar, go to **Storage > Block Storage Systems**.
3. Right-click the DS8000 and select **View Properties**.
4. Check the version in the **Firmware** field. If the firmware version is earlier than R6.2, continue to the next step. If the version is R6.2 or later, you do not need to complete this procedure.
5. Download and start the PuTTY utility. PuTTY is a free implementation of Telnet and SSH for Windows and UNIX. You can download PuTTY from <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>.
6. On the Session page in the PuTTY tool, complete the following actions:

- a. In the **Host Name (or IP address field)**, enter the IP address of the Hardware Management Console (HMC) for the DS8000. You can determine this IP address by viewing the properties of the DS8000 in the IBM Spectrum Controlweb-based GUI.
 - b. Select **SSH** for a **Connection type**.
 - c. Click **Open**.
7. Enter the user name and password for logging in to the HMC.
8. Run the following command: `touch /opt/esshmc/data/config/ntvbrsen`
9. In the web-based GUI, right-click the DS8000 and select **Open Storage System GUI**.

Exporting information to a file

You can save information about resources, jobs, or alerts to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of resources, jobs, or alerts.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the GUI has a column or set of columns that determine the order of its rows.
- Filter the list to show only the resources, jobs, or alerts that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the following pages:

Resource list and resource details pages

You can export information about the top-level resources, internal resources, or related resources that are shown on resource list pages and resource details pages. For top-level resources, you can also export information about related jobs or alerts.

For example, you can export information about monitored switches on the Switches page. When you select the export action, the generated file includes the column values for each monitored switch in the list. In the **Jobs** or **Alerts** tabs on the Switches page, you can export information about all the jobs or alerts that are related to switches.

The Alerts page

On the **Home > Alerts** page, you can export information about all the alerts that were detected for monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of alerts.

The Jobs page

On the **Home > Jobs** page, you can export information about all the jobs that are associated with monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of jobs.

1. In the GUI, go to a list of resources, jobs, or alerts for which you want to export information.
2. Right click any row in the table and select **Export** > **<file_format>**, where **<file_format>** represents the format that you want to use for a file. The following formats are available:

CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name","Column_name","Column_name","Column_name","Column_name","Column_name"
"data","data","data","data","data","data"
```

where `Column_name` represents the name of a column in a table and `data` represents the data that is associated with a column.

PDF You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

Related concepts:

“Viewing information about resources” on page 194

You can view detailed information about the resources that are monitored by IBM Spectrum Control. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

“Customizing lists of resources, tasks, and alerts”

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

Customizing lists of resources, tasks, and alerts

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.


Related reference:

“How performance information is displayed” on page 224

The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Filtering lists

You can filter the items that are shown in a list. Use filtering when you want to limit the rows that are shown in a list based on values in a specific column or set of columns. For example, when viewing alerts, you can filter the list so that only alerts with a value of Warning in the Severity column are shown.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to filter.
2. Click the **Filter** icon  to determine how to filter the list. The following options are available:

Filter by Default Columns


Select this option to apply the filter to the set of the default columns in a list. Each list contains a set of default columns on which you can filter.

For example, in the lists of alerts on the **Home > Alerts** page, select this option to apply filter text to the following columns at the same time: Condition, Severity, Alert Category, Resource, Internal Resource, Total Occurrences, Alert Name.

Filter by a Specific Column

Select this option to apply filter text to a specific column in a list. For example, on the **Home > Alerts** page, select the **Condition** column if you want to apply filter text to values in that column only.


Restriction: Some columns in lists cannot be used for filtering and are not shown under the **Filter by a Specific Column** option.

3. In **Filter** text box , type the text for filtering a list. You can type letters, whole words, partial words, and numbers. When filtering on certain columns, you can select from a list of values rather than typing in filter text.

For example, when filtering on the **Status** column, you can select one of the following values as the filter text: Error, Unreachable, Warning, Normal, Unknown.

4. Press Enter to apply the filter to the list. Only rows with values that match or partially match the filter text are shown.

For example, on the **Home > Alerts** page, if you select to filter on the Severity column and type critical as the filter text, only alerts that have a severity of critical and critical - acknowledged are shown in the list.


Tip: The name of a column to which a filter applies is shown next to the **Filter** icon .

5. Optional: To remove the filtering for a list, click **Reset**.

Sorting lists

Each list of resources in the web-based GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.

For example, the list of ports on a fabric is ordered by switch name, blade slot, and port number, in that order. You can change the sort order of these columns to organize a list of resources according to your requirements.

Tip: A visual indicator  is shown next to the primary column that determines the sort order of a list.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to sort.
2. Complete the following tasks to sort the rows in the list:
 - To change whether the rows in a list are sorted by ascending or descending order, click the name of the primary column in the heading row. Click the name of the column again to reverse the sort order.
 - To change the column that determines the sort order of resources, click the name of that column. For example, click the Name column to sort the list of resources alphabetically, or click the **Status** column to sort the list according to the status of resources.

Tip: When you click a column other than the primary sort column, that column becomes the new primary sort column. However, the previous primary column maintains its sort order (ascending or descending).

For example, if a list of resources is ordered by the Name column in ascending order and you click the Status column, the rows are sorted according to the statuses of resources. At the same time, the resources for each type of status are still sorted in ascending order by the Name column. There is a maximum of three columns that can be sorted upon at the same time.

3. Optional: To reset the order of the rows for a column to the default setting, right-click the heading row in the list and select **Reset Table Preferences**.

Showing, hiding, and reordering columns in lists

Each list in the GUI has a set of columns that contain information about resources, jobs, or alerts. You can change the order of columns, show columns that are hidden, and hide columns that are shown in these lists.

1. In the GUI, go to a list of resources, jobs, or alerts that you want to customize.
2. Complete the following tasks to customize the columns in the list:
 - To show or hide the columns in a list, right-click any of the column headings and select or clear the check box next to a column name.
 - To change the order of columns, click a column heading and drag it to a new position in the list.

Tip: The changes that you make to column settings are saved from session to session.

3. Optional: To reset the order and display of columns to the default setting, right-click any of the column headings and select **Restore Default View**.

Managing tasks

Tasks are used to provision storage and optimize resources in your storage environment. Use the Tasks page to manage all the tasks that are used by IBM Spectrum Control to provision storage and optimize resources. Use resource list and resource details pages to manage the tasks for specific resources and resource types.

Tasks are created when you take any of the following actions in the GUI:

- You can provision storage to assign storage to servers or hypervisors.
- You can analyze tiering to move volumes to higher or lower tiers.
- You can transform storage to move or convert volumes.
- You can balance pools to distribute the workload of volumes across pools on the same tier.

Use the following actions to manage tasks:

- View the status of tasks in a central location.
- Identify all the tasks that are related to a specific resource or a type of resource.
- Identify all the tasks that are encountering problems when the tasks are running.
- Administer the tasks. The actions that are available depend on the type of task that you select. For example, you can run or rename a tiering-analysis task, or export provisioning task details to a text file.
- Change the number of task runs that are displayed.

Tip: A *task run* is a single invocation of a task.


Viewing tasks

Use the Tasks page to view all the tasks that are used by IBM Spectrum Control to provision storage and optimize resources. Use resource list and resource details pages to view the tasks for specific resources and resource types.

Viewing all tasks

Use the Tasks page as a central location to view the tasks that IBM Spectrum Control is using to provision storage and optimize resources. You can view all the tasks for resources that are monitored by IBM Spectrum Control.

1. In the menu bar in the web-based GUI, go to **Home > Tasks**. The information about tasks is organized into columns. These columns include information about the status of tasks, the date and time of the last run for tasks, and the schedule for tasks.
2. Optional: View the status icons at the top of the page to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

Tip: To view descriptions of the icons and columns on the Tasks page, click the **Help** icon  in the upper-right corner of the page.

Viewing pending tasks

Some tasks might be pending because they can be run or scheduled only by an administrator. If tasks are pending, the number of pending tasks is displayed in the title bar of the GUI.

You must be assigned to the Administrator role to view the number of pending tasks.

Pending tasks are tasks that were not run and are not scheduled to run. Pending tasks might be created by a user who is not authorized to run or schedule the task. For example, users who are assigned to the Monitor or External Application role can have permission to create provisioning tasks by using a service class. However, if the service class specifies that administrator approval is required, these users can only create the provisioning task. The provisioning task is pending until an administrator runs or schedules it.

To view pending tasks, complete the following steps:

1. In the title bar of the GUI, look for the pending task count. If no number appears in the title bar, there are no pending tasks.
2. To open a Tasks page that shows only the pending tasks, click the number of pending tasks in the title bar.

Viewing tasks for a type of resource

Use resource list pages to view only the tasks for a specific type of top-level resource. Each resource list page has a **Tasks** tab that lists only the tasks for that type of resource. For example, you can use the list page for storage systems to view only the tasks that are associated with storage systems.

1. In the menu bar, select a type of top-level resource. For example, to view the tasks that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor.

Exception: Provisioning tasks that provision network-attached storage (NAS) file shares to servers or hypervisors are not shown on the resource list pages for servers or hypervisors. You must use the Tasks page to view these tasks.


2. Click the **Tasks** tab to view a list of all the tasks for the monitored resources of that type.

The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.

Tip: The status icon on the **Tasks** tab represents the most critical status of the tasks that are associated with the monitored resources. For example, if one of the tasks for storage systems failed, the following icon is shown on the tab:



3. Optional: View the status icons on the **Tasks** page for a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.


Tip: To view descriptions of the icons and columns on the **Tasks** page, click the **Help** icon .

Viewing tasks for a specific resource


Use resource detail pages to view only the tasks for a specific monitored resource. Each resource detail page has a **Tasks** link that lists only the tasks for the associated storage system, server, hypervisor, switch, or fabric.

1. In the menu bar, select a type of top-level resource. For example, if you want to view the tasks for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Right-click a resource in the list and select **View Details**. The resource detail page is displayed. In the **General** section, the number next to **Tasks** represents the number of tasks that are associated with the resource. For example, **Tasks (3)** indicates that you can view three tasks.

The icon next to the **Tasks** link represents the most critical status of the tasks that are associated with the resource. For example, if a task failed, the following icon is shown:




3. Click **Tasks** in the **General** section to view the tasks for the resource. The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.
4. Optional: View the status icons on the Tasks pane to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

Tip: To view descriptions of the icons and columns on the Tasks pane, click the **Help** icon .

Viewing task details

Use task detail pages to view detailed information about the tasks that IBM Spectrum Control uses to optimize resources and provision storage. You can use this information to troubleshoot the reason why a task failed, analyze details of a task before you implement task recommendations, or schedule future task runs.

On all the task detail pages, you can view information such as the task status, the duration and completion date of task runs, and the actions that you can take to manage the task. For tasks such as tiering analysis, balance analysis, and transform plan, you can view information about the recommendations that are generated by the most recent or earlier task runs. For provisioning tasks, you can view the status of the actions that the provisioning task completes for each of the resources that are configured in the task.





1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor, and click the **Tasks** tab.
2. Locate the task that you want to view in more detail.
3. Right-click the task row and select **View Details**. The task details page opens in a separate window. For example, the Tiering Analysis page opens if you select a tiering-analysis task row.
4. Optional: To view task logs, click **Open Logs**. You can use the information in the log file to troubleshoot any errors that might occur when a task is run.
5. Optional: For provisioning tasks, you can export information about the provisioning task to a text file for offline viewing. To export information, click the  icon at the top of the Provisioning page.

Viewing task logs

Task logs include detailed information about the status, actions, and progress of a task. Status information includes informational, warning, and error messages that are related to each action in a task that is taken during processing. You can use this information to troubleshoot any errors that might occur when a task is run.

The following examples show the type of information that you can view in the task logs:

- View the number of recommendations that are generated by tiering-analysis, balance-analysis, or transform-plan task runs.
- View the options that you have when you transform storage, analyze tiering, or balancing pools. For example, when you create a tiering-analysis task and the options that you choose are recorded in the task log.
- View error messages that are related to task processing.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor, and click the **Tasks** tab.
2. Locate the task that has a log file that you want to view.
3. Right-click the task row and select **Open Logs**. The Logs page shows log entries for the most recent task run.
4. Optional: On the Logs page, to view the log for a previous task run, select a task run from the **Select a log** list.
5. Optional: To view only the actions in a task run that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only actions that have the following statuses:
 -  Only error entries
 -  Only warning entries
 -   Error and warning entries
6. Optional: To view an explanation of the message that is associated with an action, click the link in the **ID** column.

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a task. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the task run.
- The status for each action in a task.
- The date and time when an action was completed. The date, time, and time zone of the action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Setting the number of task runs that are displayed

You can change the number of task runs that are shown in the GUI by configuring the settings on the History Retention page. By default, the information for the last five runs of a task is shown.

1. In the menu bar, go to **Settings > History Retention**.
2. Click **Edit**.
3. In the **Job logs** field, enter the number of runs that you want to retain for display.
4. Click **Save** to save your changes.

After you update the value for job logs, the list of task runs is automatically updated within a few minutes.

Managing tasks for provisioning

Provisioning tasks are used to create storage volumes and assign the volumes to servers or hypervisors. Provisioning tasks are also used to create network-attached storage (NAS) file shares. Use the Tasks page in the GUI to manage all the tasks that are used by IBM Spectrum Control to provision storage. Use resource list and resource details pages to manage the provisioning tasks for specific resources and resource types.

Use the GUI to complete the following provisioning actions:

- Create provisioning tasks
- Run provisioning tasks immediately
- Schedule provisioning tasks to run at a specified date and time
- Export information about provisioning tasks to text files
- Delete provisioning tasks

Related concepts:

“Provisioning storage” on page 57

IBM Spectrum Control guides you through the steps for provisioning storage volumes or network-attached storage (NAS) shares. You can provision storage volumes or NAS file shares to one or more servers, one or more hypervisors, or one cluster.

Creating provisioning tasks

You can create provisioning tasks with the GUI and assign storage volumes or NAS file shares to servers or hypervisors.

You must have Administrator privileges to create provision tasks.

Related tasks:

“Provisioning storage with the IBM Spectrum Control GUI” on page 311

You can provision volumes or shares to servers, hypervisors, and clusters with the GUI.

“Collecting information about shares on storage systems” on page 147

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Running provisioning tasks

Use the Tasks page to run provisioning tasks. You can run provisioning tasks immediately or you can schedule tasks to run at a date and time that you specify. When you run a provisioning task, the configuration changes that are contained in the task are implemented. For example, storage volumes are created and assigned to servers.

To run provisioning tasks, you must be the owner of the task or have Administrator privileges.

Provisioning tasks can run only once. You can run the task when you complete the steps for storage provisioning or you can save the provisioning task and run the task later.

The **Run Now** and **Schedule** actions are available for a provisioning task when the task is not running, and was not previously run.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the provisioning tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the provisioning task that you want to run.
3. Depending on when you want the provisioning task to run, choose one of the following options:
 - To run the provisioning task immediately, right-click the task row and select **Run Now**. The task status changes to Running and the provisioning task runs immediately.
 - To schedule the provisioning task to run at a specific date and time, complete the following steps:
 - a. Right-click the task row and select **Schedule**.
 - b. From the Schedule page, select a future date and time from the **Provision** lists, and click **Save**.

You can also run or schedule the provisioning task from the Provisioning detail page. On the Provisioning detail page, click **Execute** to run the task immediately or click **Schedule** to schedule the task to run at a specific date and time.


The provisioning task runs immediately or on the scheduled date and time. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View detailed information about the progress of the task on the Provisioning detail page. You can view the task duration and the status of the actions that the provisioning task completes for each of the resources that are configured in the task.
- View the task logs from the Tasks page or from the Provisioning detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Exporting provisioning task information

Use the Provisioning detail page to export details of a provisioning task to a text file. The task details that are shown on the Provisioning page are saved to the generated text file.

If you do not have authority to provision storage, you can export the provisioning task details to a file and email the file to a user with Administrator privileges. The administrator can run or schedule the provisioning task if they approve the task details.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the provisioning tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the provisioning task that you want to export.
3. Right-click the task row and select **View Details**.
4. From the Provisioning page, to start the export, click the  icon in the upper-right corner of the page.
5. Save the generated text file. Depending on how your web browser is configured, you can specify the file name and the location where the file is saved.

The provisioning task information is exported to a text file. The generated file includes all the provisioning task details that are listed on the Provisioning page. Columns that are hidden on the Provisioning page are included in the generated file. The information in the file is organized according to a default order and is not affected by the column order that is defined on the Provisioning page.

Managing tasks for tiering storage, balancing pools, and transforming storage

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

To schedule, run, edit, or delete tiering storage, balancing pools, and transforming-storage tasks, you must have Administrator privileges.

To complete any of the optimization tasks, such as balancing pools, you must complete the following steps:

1. Create the analysis task with the GUI. For example, to balance pools create a balance-analysis task.
2. Run the analysis task to generate the optimization recommendations. For example, run the balance-analysis task to generate the recommendations to move volumes.
3. Edit analyze tiering and balancing pool tasks by modifying the thresholds.
4. Implement the recommendations to optimize storage. For example, implement the recommendations to balance pools by running the analysis-execution task.

Use the Tasks, resource list, and resource detail pages to take the following actions for optimization tasks:

- Run the optimization analysis immediately to generate optimization recommendations.
- Schedule the optimization analysis to run at a specified date and time.
- Implement optimization recommendations immediately.
- Schedule the implementation of optimization recommendations for a specified date and time.
- Pause, cancel, and resume the implementation of optimization recommendations.
- Delete optimization tasks.
- View task logs.

Related tasks:

“Modifying the criteria for analyzing tiering” on page 324

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

“Modifying the criteria for balancing pools” on page 326

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

Creating analysis tasks

To create tasks for optimizing resources in your storage environment, use IBM Spectrum Control GUI.

You must have Administrator privileges to create tasks for optimizing resources in your storage environment.

Storage optimization tasks are created when you take any of the following actions in the GUI:

Tier storage

You can analyze tiering to move volumes to higher or lower tiers; a tiering-analysis task is created.

Balance pools

You can balance pools to distribute the workload of volumes across pools on the same tier; a balance-analysis task is created.

Transform storage

You can transform storage to move or convert volumes; a transform-plan task is created.

Related tasks:

“Tiering volumes by I/O density and I/O rate” on page 323

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

“Balancing pools” on page 325

You balance pools to distribute the workload of volumes across pools on the same tier and on the same storage virtualizer. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier and on the same storage virtualizer.

“Transforming and migrating volumes” on page 331

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Running analysis tasks

To run tiering analysis, balance analysis, and transform-plan tasks, use the Tasks page. When you run the analysis for any of these optimization tasks, recommendations are generated to optimize the storage resources that are configured in the task. For example, when you run a tiering-analysis task, recommendations are generated to tier volumes based on the criteria that you set in tiering policies.

To run tiering analysis, balance analysis, and transform-plan tasks, you must have Administrator privileges.

You can run tasks immediately or schedule tasks to run at a future date and time.

Tip: The **Run Now** action is available when the task is enabled and the task is not running.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tiering-analysis tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to run.
3. Right-click the task row and select **Run Now**. The task status changes to Running.

4. Optional: To schedule future task runs for tiering analysis or balance-analysis tasks, right-click the task row and select **Schedule Analysis**. On the Schedule Analysis page, select the time that the task is run and how often the task is run.

Restriction: You cannot schedule future task runs for transform-plan tasks.

5. Optional: To enable or disable a tiering analysis or balance-analysis task, select **Enabled** or **Disabled** on the Schedule Analysis page. When the task is enabled, the task runs according to the defined schedule. If the task is disabled, you can modify the task schedule but the task cannot run at the scheduled time.

When the task run completes, you can view the recommendations that are generated on the task detail page. For example, you can view the recommendations that are generated for a tiering-analysis task on the Tiering Analysis page. To check the progress of the task run, you can take the following actions:

- View the task status on the Tasks page.
- View information about the progress of the task on the task detail page. You can view the duration of the task run, and the date and time that the task run completed.
- View the task logs from the Tasks page or from the task detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Related tasks:

“Modifying the criteria for analyzing tiering” on page 324

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Implementing recommendations to optimize storage

Use the Tasks page and task detail pages to implement optimization recommendations to tier storage, balance pools, and transform storage. You can implement recommendations immediately or you can schedule the implementation of the recommendations for a future date and time.

Related tasks:

“Modifying the criteria for analyzing tiering” on page 324

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Implementing recommendations immediately:

To implement optimization recommendations to tier storage, balance pools, and transform storage, use the task detail pages. For example, use the Balance Analysis page to implement the recommendations to balance pools.

To implement storage optimization recommendations, you must have Administrator privileges.

IBM Spectrum Control uses analysis-execution tasks to implement optimization recommendations. When you implement recommendations for a tiering analysis, balance analysis, or transform-plan task, the following process takes place:

- An analysis-execution task is automatically created so that the recommendations can be implemented.

- The analysis-execution task runs immediately and the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

You can view analysis-execution tasks on the Tasks page. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

The **Execute** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
 - An analysis-execution task was not already run for the latest set of recommendations.
 - An analysis-execution task is not running or paused.
1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the analysis-execution tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
 2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to implement.
 3. Right-click the task row and select **View Details**. For example, the Balance Analysis page opens if you select a balance-analysis task.
 4. Click **Execute**. An analysis-execution task is created and immediately implemented.

The analysis-execution task runs immediately. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View the status of individual task recommendations on the Analysis Execution page.
- View the task logs from the Tasks page or from the Analysis Execution page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Scheduling the implementation of recommendations:

To schedule the implementation of optimization recommendations to tier storage, balance pools, and transform storage, use the Tasks page.

To schedule the implementation of storage optimization recommendations, you must have Administrator privileges.

The **Schedule Execution** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task is not running or paused.
- An analysis-execution task was not already run for the latest set of analysis recommendations. For each run of the tiering analysis, balance analysis, or transform-plan task, you can complete only one run of the analysis-execution task.

When you define the schedule, an analysis-execution task is created for the implementation of the optimization recommendations.

Tip: The analysis-execution task implements the recommendations that were current when the analysis-execution task was defined. For example, you define a tiering-analysis task that runs daily. You review the recommendations every day and, on a Tuesday, you schedule an analysis-execution task to run on the following Friday. When the analysis-execution task runs on Friday, the recommendations that were current on Tuesday are implemented, even if more recent recommendations are available.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tiering analysis tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to schedule.
3. Right-click the task row and select **Schedule Execution**.
4. On the Schedule Execution page, select the date and time that the recommendations are implemented and click **Save**.

An analysis-execution task is created that you can monitor on the Tasks page. The name of the analysis-execution task is based on the name of the related tiering analysis, balance analysis, or transform-plan task. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

When the analysis-execution task runs on the scheduled date, the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

Pausing, resuming, and canceling the implementation of recommendations

IBM Spectrum Control uses analysis-execution tasks to implement optimization recommendations to tier storage, balance pools, and transform storage. To pause, resume, and cancel analysis-execution tasks, use the Tasks page.

To pause, resume, or cancel analysis-execution tasks, you must have Administrator privileges.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the analysis-execution tasks that are related to servers, click **Servers > Servers**, and click the **Tasks** tab.
2. Locate the analysis-execution task that you want to pause, resume, or cancel.
3. Right-click the task row and select **Pause**, **Resume**, or **Cancel**.

The **Pause** action is only available when the analysis-execution task is running.

The **Resume** action is only available when the analysis-execution task is paused.

The **Cancel** action is available when the task is running or paused.

When you pause, resume, or cancel analysis-execution tasks, the following processing occurs:

Pause The status of the analysis-execution task changes to Paused and the recommendations that are being implemented are suspended.

Resume

The status of the analysis-execution task changes to Running and the recommendations are resumed.

Cancel

The status of the analysis-execution task changes to Canceled and the recommendations that are being implemented are stopped. Any recommendations that are already implemented for this analysis-execution task are not rolled back. For example, a recommendation is generated to move a volume from pool_a to pool_b. If this recommendation is already implemented before the task is canceled, the volume remains in pool_b. Use the **Analysis Execution** page to determine which recommendations were implemented for the task before the task was canceled.

You cannot restart an analysis-execution task that is canceled. If you want to implement the same set of recommendations again, you must complete the following steps:

1. Delete the analysis-execution task that you canceled.
2. Use the **Schedule Execution** action for the tiering analysis, balance analysis, or transform-plan task to create an analysis-execution task.

Restriction: The implementation of recommendations for optimizing the volumes is canceled except for the recommendations to optimize the volumes with mirrored volume copies. The implementation of recommendations for optimizing volumes with mirrored volume copies cannot be canceled.

Renaming tasks

You can change the name of a task by using the task detail page. For example, use the Provisioning page to change the name of a provisioning task. A task name is automatically generated when the task is created. If the original name is ambiguous or is not clear enough, you can change the task name to a unique task name that you provide.

You must have the following privileges to rename the different task types:

Provisioning

To rename provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types

To rename all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the task that you want to rename.
3. Right-click the task row and select **View Details**. The task detail page opens in a separate window. For example, the Provisioning page opens if you select a provisioning task.
4. Type the new name in the task name field. The naming conventions for task names are:
 - The maximum length is 64 characters.
 - Uppercase and lowercase letters are allowed.
 - Spaces and numbers are allowed.
 - The following special characters are allowed:

! # % & * + - / = ? ^ _ { } ()

5. Click **Rename**.

The task is renamed. You can view the renamed task on the Tasks page.

Deleting tasks

Use the Tasks page to delete tasks, such as provisioning, tiering analysis, balance analysis, transform plan, and analysis-execution tasks.

You must have the following privileges to delete the different task types:

Provisioning

To delete provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types

To delete all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

The following conditions apply for the different task types:

Analysis-execution

You can delete analysis-execution tasks that were canceled, already run, or scheduled to run in the future. You cannot delete tasks that are running or paused.

Other task types

You can delete all other tasks if the tasks were already run or if the tasks are scheduled to run in the future. For example, you can delete provisioning, tiering analysis, balance analysis, and transform-plan tasks that were already run. You cannot delete tasks that are running.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the task that you want to delete.
3. Right-click the task row and select **Delete**. Alternatively, you can delete the task by clicking **Delete** on the task detail page. The Delete page opens in a separate window.
4. If you are deleting a provisioning task that was not yet run, enter the reason that you are deleting the task.
5. Click **Delete**. To exit the window without deleting the task, click **Cancel**.

Planning copy data resources

To plan for your block storage requirements, you need to be able to see how you currently use block storage in your storage environment.

You can use the Copy Data view the following aspects of your block storage environment:

- You can see how much of your current allocated capacity is taken up by both the primary and backup data.
- You can see the ratio of allocated capacity that is taken up by primary data and plan for backup data according to your security policies.

- You can analyze storage device relationships to identify how redundancy is affecting available capacity.
- You can view how each type of storage you use is deployed in your environment.
- You can identify the volumes that are not backed up.

Monitoring copy data resources

You can monitor the use and type of block storage resources in Copy Data.

You can get a complete breakdown of how each type of storage is used across the block storage environment.

Remote Relationships

Detailed information on remote relationships includes the type of relationship that exists between the source and target storage systems.



Synchronous relationship

Host writes are delayed until the source receives confirmation that the data is written to the target volume.



Asynchronous relationship

Host writes can continue while data is being written to the target volume.

Where the type of volume is known, one of the following icons appears next to the volume in question.



Standard volume

A volume that is not thin-provisioned, compressed, encrypted, or encrypt-able.



Thin provisioned volume

The volume is only allocated the required amount of server space as required.



Compressed VDisk volume

Data is compressed as it is written to the VDisk.

Consistency Group

Details the status and type of the consistency group, and the number of relationships that exist in the group.

FlashCopy

Details the information on the last FlashCopy operation between each source and target storage system.

VDisk Mirrors

Details the information on the volumes, pools, tiers, and synchronization status of the VDisk mirrors.

- Identify target and source mirror disks at a glance.

- Easily see the disks that are compressed, thin, or normal.

HyperSwap

Details the information on the status of the master and auxiliary volumes, pools, and tiers of the HyperSwap relationships.

- Quickly identify the master and auxiliary volumes and pools.
- Drill down to see the amount of allocated disk for each volume.

Unprotected Volumes

Details the information on the volumes that are not backed up.

Chapter 7. Reporting

To help management plan capacity purchases and aware of the cost and the amount of the storage that is used by storage consumers, create chargeback and consumer reports.

Chargeback and storage consumer reports

In chargeback reports, you create reports that show the capacity and the cost of the block and file storage that is used by the types of storage consumers in your storage environment such as:

- Applications
- Departments
- Hypervisors
- Physical servers

In storage consumer reports, you create reports that show the capacity and cost of the block storage that is used by a single storage consumer, such as the capacity and the cost of the storage that is used by a single application.

To make the owners of the storage consumers aware of the capacity and cost of the storage they use, you schedule and send chargeback and storage consumer reports by email.

Show cost and capacity of tiered storage: If you want to show the block capacity of tiered storage and differentiate between the cost of storing data on tier 1 or on lower tiers of storage, you must assign your block storage pools to tiers. Click **Storage > Pools**. Right-click one or more pools, and then click **Tier** and the tier level from the menu.

To view the capacity and charge for the storage that is used by applications and departments, you must create models of the applications and departments in your business organization and map the storage resources to the applications and map the applications that the departments use to the departments. If you organize an application or department in hierarchies, the chargeback and consumer reports show the capacity and cost of the storage resources that the application and its child applications use and the capacity and cost of the storage resources that the department and its child departments use.

Storage capacity

You can choose the unit of measurement that is used to show the amount of storage that is consumed by your storage resources. In consumer and chargeback reports, you can show the storage that is consumed, for example, in PB, PIB, TB, TiB, GB, or GiB.

Storage costs

The cost of the storage that is used by storage consumers is included in all reports even if the storage costs are set to zero. If you want to send reports that show the capacity of the storage, but not the cost of the storage that is used, you can exclude the storage costs from the report.

Before you decide how you want to set the costs for the block storage that is used by the storage consumers, you must choose whether you want to charge for the storage that is allocated to the storage consumer or the storage that is assigned to the storage consumer.

By default, block storage is calculated by the amount of storage that is allocated to the volumes that are being used for storage by the consumer. Whether you choose to calculate by allocated or assigned capacity, it doesn't affect the calculation of capacity for standard volumes because standard volumes are fully allocated when they are created. However, the calculation of the capacity of thin provisioned and compressed volumes is affected depending on whether you choose to calculate capacity by allocated or assigned capacity.

For example, an application has three volumes, two fully allocated volumes with 1 TiB each and one thin-provisioned volume with an assigned capacity of 3 TiB. The thin-provisioned volume is allocated capacity in increments of 1 TiB. If you generate a report that calculates capacity by allocated capacity, and the thin-provisioned volume is currently allocated 1 TiB, the total amount of capacity for the application is 1 TiB + 1 TiB + 1 TiB = 3 TiB. However, if you choose to calculate by assigned space, the total capacity for the application is 1 TiB + 1 TiB + 3 TiB = 5 TiB. The cost is also affected because the cost is the total capacity of the block storage, which is multiplied by the unit cost or the unit cost per tier.

Restriction: If you choose to calculate block storage by allocated space, the block storage for FlashSystem A9000 and FlashSystem A9000R storage systems is calculated based on the virtual allocation percentage of the volumes. This restriction occurs because used space information for volumes is not available from FlashSystem A9000 and FlashSystem A9000R.

Default or custom storage costs

When you create or update a report, you can choose one of the following options:

Use default cost

Set or use default costs per unit and type of storage.

If you choose to use the default costs option, the changes that you make apply to all of the reports that use the default option to calculate the cost of storage. When you exclude costs from reports that use default costs, you exclude costs from all reports that use default costs to calculate the cost of storage.

Set custom cost

Use custom costs per unit and type of storage.

If you choose to use the custom costs option, the costs that you set are used to calculate the cost of storage only for that report. If you exclude costs from the custom report, the costs are only excluded from the custom report that you create or edit.

The reports that were created in IBM Spectrum Control Version 5.2.11 use the unit costs that were specified when the reports were configured.

Tip: If you want to change the configuration of a report, such as how storage costs are calculated, click **Reports > Email Reports**, select the report, and then click **Edit**.

Deleted and unavailable reports

If you delete a storage consumer, such as an application, the entry for that storage consumer is removed from all chargeback reports for applications. If you created a consumer report for a storage consumer that was deleted, the report is automatically deleted.

If you organize your applications or departments in hierarchies, the chargeback and consumer reports show the total capacity and costs for the parent application or department and its child applications and departments. However, if you create a consumer report for a specific application or department and then designate the application or department as a child application of another application, the consumer report is deleted.

For example, you create a consumer report for an application called German Online Sales. You then reorganize your sales applications in a hierarchy and the application for German Online Sales becomes a child application of European Sales. To include the capacity and cost of the storage that is used for German Online sales, you must create a consumer report for the application for European Sales.

Creating chargeback reports

Create chargeback reports that show how much capacity is used and the cost of the capacity that is used by applications, departments, hypervisors, and physical servers.

Chargeback reports help managers to realize the cost of the storage that they use, plan storage usage more efficiently, and reduce storage costs. For example, a manager gets a chargeback report that shows that the bulk of the storage that is used by a non-critical application is on tier-1. To reduce storage costs and use storage more efficiently, the manager can request the storage administrator to use tier-2 or lower tiers of storage for the application's data.

1. From the **Reports** menu, click **Email Reports**.
2. In the **Chargeback Reports** section, click **Create New**.
3. Type the unique name of the report, and enter the delivery and the scheduling details.
4. Choose the type of storage consumer from the list. If you choose hypervisor or physical server, the capacity and cost of the hypervisors and physical servers that belong to clusters are listed under the name of the cluster.
5. Choose how you want to calculate the capacity and cost of thin-provisioned storage.
6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

7. Choose whether you want to use default or customized costs for block and file storage. To change the default costs for block and file storage, click the lock (🔒) icon.
8. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	<ol style="list-style-type: none"> 1. Click the lock (🔒) icon. 2. In the Include Storage Cost in Report section, click No. <p>Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.</p>
Use Custom Cost	In the Include Storage Cost in Report section, click No .

9. Choose the types of storage that you want to include in the report.
10. Set the unit costs for the storage that is used. To get a complete picture of the capacity and cost of the block storage, include the cost of the data that is copied.
11. Click **Save**. If you set the frequency of the report to **Now**, you click **Save and Send**.

The report is saved and depending on the scheduling option is sent to the recipients or scheduled to be sent later. If you want to change the configuration of the report, you can click **Email Reports**, select the report, and then click **Edit**.

Chargeback reports

Configure the chargeback reports to include block and file storage and to calculate the cost of the block and file storage that is used by the storage consumers.

You configure, schedule, and send chargeback reports so that the owners and managers of storage resources:

- Know the cost of maintaining the block data and file data that they use
- Know the cost of maintaining block data on each tier of storage
- Know the cost of maintaining block copy data

Include total block and file capacity in the report

In the **Configure Report** section, you choose the types of storage that you want to show in the report. For example, if you click the **File Storage** check box, the names of the resources and the total file capacity that the resources use are shown in the report.

For block storage, you can show the total storage capacity and, if your storage is tiered, the total capacity by tier. You can also break down the total block capacity into total primary capacity and total copy capacity.

Tip: The recipient of the report, such as the application owner, sees that the capacity for copy data is high for tier-1 storage. To reduce costs and use the block storage more efficiently, the application owner can ask the storage administrator to place the volumes that are used for copy data on lower and less expensive storage tiers. The storage administrator can also analyze the tiering of the pools to make sure that the application's performance is not affected when the copy data volumes are moved to lower storage tiers.

Include block and file costs in the report

In the **Set Storage Cost** section, you set the cost per TiB of storage that the resources use. The total cost of the storage that is used is shown in the Cost column in the report.

For block storage, you set the costs as follows:

- If your storage environment is tiered, you set the price per TiB for each tier of storage
- If your storage environment is not tiered, you set the price per TiB for all of your block storage
- If your storage environment contains tiered and non-tiered storage, you can set the price per TiB for the tiered storage, but you can't set a price for the non-tiered storage. Although you can't set the cost for the non-tiered storage, the capacity of the non-tiered storage that is used is calculated.

If you assign your block storage pools to tiers, you can differentiate between the cost of expensive tier-1 storage and the cost of the less expensive storage that is used on the lower tiers of storage.

For file storage, you set a price for all of the file storage that is used by your resources.

When you configure chargeback reports, you can send the report immediately or schedule the report to be sent every day, every week, or every month.

Tip: If your storage environment changes, you should check whether you need to update your scheduled reports. For example, if you add a storage tier, you can edit the report and set the price for the new tier.

Creating consumer reports

Create consumer reports that show how much block capacity is used and the cost of the block capacity that is used by an application, department, hypervisor, or physical server.

Storage consumer reports help the managers and owners of applications, departments, hypervisors, and physical servers realize the cost of the block storage that they use, plan block storage usage more efficiently, and reduce block storage costs.

1. From the **Reports** menu, click **Email Reports**.
2. In the **Consumer Report** section, click **Create New**.
3. Type the unique name of the report, and enter the delivery and the scheduling details.
4. From the list, click the type of storage consumer and type or click the name of the storage consumer. To generate a report about the hypervisors or physical servers that belong to a cluster, you choose **Hypervisor** or **Physical Server**, and then choose or type the name of the cluster.
5. Choose whether you want to calculate block capacity by allocated or assigned space.
6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

7. Choose whether you want to use default or customized costs for block storage.
8. Choose the types of storage that you want to include in the report. To change the default costs for block storage, click the lock (🔒) icon.
9. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	<ol style="list-style-type: none"> 1. Click the lock (🔒) icon. 2. In the Include Storage Cost in Report section, click No. <p>Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.</p>
Use Custom Cost	In the Include Storage Cost in Report section, click No .

10. Set the unit costs for the block storage that is used. To get a complete picture of the capacity and cost of the block storage, include the cost of the data that is copied.
11. Click **Save**. If you set the frequency of the report to **Now**, you click **Save and Send**.

The report is saved and depending on the scheduling option is sent to the recipients or scheduled to be sent later. If you want to change the configuration of the report, you can click **Email Reports**, select the report, and then click **Edit**.

Consumer reports

Configure consumer reports to calculate the cost and capacity of the block storage that is used by a storage consumer such as an application, a department, a hypervisor, or a physical server.

You configure, schedule, and send consumer reports so that the owner or manager of the storage consumer knows the cost of maintaining the block storage that the storage consumer uses.

In the **Configure Report** section, you include the total block storage capacity and, if your storage is tiered, the total capacity by tier in your report. You can also include the total primary capacity and the total copy capacity of the storage that is used, and you can include the total capacity:

- For VDisk mirrors
- For FlashCopy volumes
- For remote mirrors such as Global Mirror, Global Copy, and Metro Mirror

In the **Set Storage Cost** section, you set the cost per TiB of the block storage that the storage consumer uses.

You set the costs as follows:

- If your storage environment is tiered, you set the price per TiB for each tier of storage

- If your storage environment is not tiered, you set the price per TiB for all of your block storage
- If your storage environment contains tiered and non-tiered storage, you can set the price per TiB for the tiered storage, but you can't set a price for the non-tiered storage. Although you can't set the cost for the non-tiered storage, the capacity of the non-tiered storage that is used is calculated.

Creating summary reports of the storage capacity

Configure reports that make managers aware of the capacity of the storage that is used by their applications, departments, hypervisors, or physical servers. The cost of the storage is not shown in the reports.

You can create reports that show the total capacity of the storage that is used by storage consumers and exclude the storage costs from the reports.

1. Click **Reports > Email Reports**.
2. Choose one of the following options:

Option	Steps
Create chargeback reports	In the Chargeback section, click Create New .
Create consumer reports	In the Consumer section, click Create New .

3. Enter the report details and schedule the report.
4. Choose one of the following options:

Option	Steps
Chargeback reports	Choose the type of storage consumer.
Consumer reports	Choose the type of storage consumer and the storage consumer.

5. To calculate the capacity that is assigned to the thin-provisioned volumes, click **Assigned space**. By default, the amount of space that is allocated to the thin-provisioned volumes is used to calculate the cost.
6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

7. In the **Configure Report** section, choose the capacity values that you want to include in the report.
8. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	<ol style="list-style-type: none"> 1. Click the lock (🔒) icon. 2. In the Include Storage Cost in Report section, click No. <p>Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.</p>

Storage cost option	Steps
Use Custom Cost	In the Include Storage Cost in Report section, click No .

- Click **Save**. If you set the frequency of the report to **Now**, you click **Save and Send**.

The reports include the capacity values and exclude the costs of the capacity that is used.

Configuring the email server

To send reports, you must configure the email server.

The email server that is used to send alert notifications is also used to send reports. If you didn't configure the email server for alert notifications, you must configure the email server before you create reports.

IBM Spectrum Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server. For example, if your email server requires authentication and TLS to send emails, then IBM Spectrum Control uses the authentication credentials that you provide when you set up the email server, and uses TLS to establish the connection with your email server. Because the email server that you set up is automatically trusted, you don't have any additional management tasks such as importing the server certificate into the keystores that are used by IBM Spectrum Control.

Tip: To change the configuration of the mail server, you click **Settings > Alert Notifications**.

- Type the host name, or IPv4, or IPv6 address of the email server.
- Type the port number that is used by the email server to send alert notifications and reports.
- Provide the user's credentials for authenticating with your email server. If your email server requires authentication to establish the connection, you must enter the user name and password.
- Optional: Type the email address, such as the email address of the administrator, that you want to use for receiving replies to alert notifications and reports. If you don't enter a value, and you entered an email address in the **User Name** field, then replies are sent to that address. Otherwise, replies are sent to no-reply@<Spectrum_Control_server_host_name>.
- Test the connection and then click **Save**.

You can create reports and send the reports by email.

Capacity metrics for reports

To gain insights into storage usage, review the capacity metrics that are collected, analyzed, and shown in reports.

Capacity metrics for chargeback and consumer reports

Learn more about how the capacity and costs for block and file storage is calculated in chargeback and consumer reports.

Chargeback and storage consumer reports

The same block capacity metrics are shown in chargeback and consumer reports. In chargeback reports, the block capacity and costs are shown for all applications, departments, hypervisors, and physical servers. In consumer reports, the block capacity and costs are shown for a specific application, department, hypervisor, or physical server. The capacity and cost of file storage is not included in consumer reports.

If you create reports about physical servers, servers that are virtual machines are excluded from the report.

Tip: To find out which servers are virtual machines, click **Servers > Servers**, right-click any column heading and click **Virtual Machine**.

Virtual machine capacity

If you assign the storage resources on a virtual machine to an application or department, the capacity of those storage resources is included in the calculation of capacity and cost for the application or department.

For example, three virtual machines are allocated capacity from a data store, which has a total capacity of 3 TiB. The capacity for the data store comprises:

- One volume on tier 1 with a capacity of 2 TiB
- One volume on tier 2 with a capacity of 1 TiB

The ratio of tier-1 storage to tier-2 storage is 2:1 or 66.66% to 33.33%.

The total capacity of the data store is split across 3 by 1,024 GiB disks and one of the disks is allocated to `Virtual_Machine_1`. The allocated space on `Virtual_Machine_1` is allocated to a database application. So, the total amount of space that is allocated to the database application is 1,024 GiB, which comprises 66.66% of tier-1 storage and 33.33% of tier-2 storage. When the report is generated for the database application, the block capacity for tier-1 storage is $1024 \text{ GiB} \times 66.66\%$, which is 682.67 GiB or 0.67 TiB. The block capacity for the tier-2 storage is $1024 \text{ GiB} \times 33.33\%$, which is 341.33 GiB or 0.33 TiB.

If you calculate the capacity and cost of storage by allocated space, the capacity and cost of the storage that is used by the application or department is shown in the chargeback and consumer reports.

Restriction: You might get unexpected results in chargeback and consumer reports if the volumes that are used to provide the capacity on the disks for the virtual machine are thin-provisioned and shared by two or more applications.

Hypervisor and physical server clusters

In chargeback and storage consumer reports, the capacity and cost of hypervisors and physical servers that are members of clusters are included in the capacity and cost of the cluster that they belong to. For example, `hypervisor_1` and `hypervisor_2` belong to `cluster_1`. To create the chargeback report about the hypervisors cluster, you choose **Hypervisors** as the consumer, and the capacity and the cost of the two hypervisors are included in the report entry for `cluster_1`. To create a report for `cluster_1` in a storage consumer report, you choose **Hypervisor** as the type of consumer, and `cluster_1` as the name of the consumer.

Block storage

The following information is provided about how block capacity is calculated:

Block or Used Capacity

In chargeback reports, the total block capacity for all of the storage resources is called Block capacity, whereas in consumer reports the total block capacity that the storage consumer uses is called Used Capacity.

The total capacity of these volumes is included in the calculation of block capacity:

- The total volume capacity of the volumes without copies that are allocated to the consumer
- The total volume capacity of the volumes with copies that are allocated to the consumer such as:
 - The total capacity of VDisk mirror volumes and their copies
 - The total capacity of FlashCopy volumes and their target volumes
 - The total capacity of Global Mirror, Global Copy, and Metro Mirror volumes and their target volumes

For example, an application is assigned all of the block storage volumes on a physical server. In this scenario, the block capacity value includes the capacity of all of the block volumes that are assigned to the physical server. If the assigned volume on the physical server has a VDisk mirror copy or the volume is a FlashCopy source volume, the capacity of the VDisk mirror copy or the capacity of the FlashCopy target volume is also included in the calculation.

Primary capacity

The total capacity of the following types of volumes is included in the calculation of primary capacity:

- The total volume capacity of the volumes without copies that are allocated to the consumer
- The total volume capacity of primary VDisk mirror copies that are allocated to the consumer
- The total volume capacity of the source volumes in FlashCopy relationships that are allocated to the consumer
- The total volume capacity of the Global Mirror, Global Copy, and Metro Mirror source volumes that are allocated to the consumer

VDisk Mirrors

The total capacity of the VDisk mirror copies for the volumes that are allocated to the consumer.

FlashCopy

The total capacity of the FlashCopy target volumes for the source volumes that are allocated to the consumer.

Remote Mirrors

The total capacity of the Global Mirror, Global Copy, and Metro Mirror target volumes for the source volumes that are allocated to the consumer.

Primary capacity by tier

The primary capacity that is allocated to the consumer for the specified tier.

Copy capacity by tier

The copy capacity that is allocated to the consumer for the specified tier, which is calculated as follows:

- The total capacity of the non-primary VDisk mirror copies
- The total capacity of the target volumes for volumes in FlashCopy relationships
- The total capacity of target volumes for Global Mirror, Global Copy, and Metro Mirror

In tiered environments, the source volume might be on tier 1 and the target volume on tier 3. In this case, the capacity of the source volume is included in the calculation for the total primary capacity for Tier 1 and the capacity of the target volume is included in the calculation for the total copy capacity for tier 3.




Figure 3. Calculating copy capacity

Restriction: For volumes that have multiple volume copies, only the capacity of the first copy and the capacity of the second copy of the volume are included in the calculation of copy capacity.

File storage

File capacity is the total amount of file storage space that is allocated to the consumer.

For example, if a server mounts a fileset export from a file storage system and a quota is defined for the fileset, the file capacity is the hard limit capacity that is defined for the quota.

If a quota is not defined for the fileset, the file capacity is the capacity of the file system that contains the fileset. If multiple fileset exports from the same file system are mounted and the hard limits that are defined for the quota exceed the total file system capacity, then the file capacity is the total capacity of the file system.

If a hypervisor data store is a mounted NFS export, file capacity is the capacity of the file system on which the export is based.

Creating resource reports

Notify your colleagues about the current state and potential issues with your storage environment by exporting the capacity and performance information that is shown in the GUI.

You can use the information that is shown about your storage resources and their related resources to create reports in HTML or PDF file formats that you can send

to your colleagues. Alternatively, you can export the information about your resources to a CSV file, which you can include in a report that you want to create about your storage resources.

You can go to the resource or internal resource page and create reports that contain the following types of information:

Table 71. Types of information in resource reports

Location	Types of information
Resource or internal resource page	Status, configuration, and capacity data
Performance tab for resource or internal resource	Key performance metrics
Capacity tab for resource or internal resource	Key capacity metrics

Restriction: You might not be able to view capacity and performance metrics for all of the resources.

For example, you monitor the pools for a SAN Volume Controller and you see that some of the pools are running out of capacity. To notify your colleagues, you export the information about the pools to an HTML file and send it to your colleagues. You want to include information about the input/output performance of the volumes for your storage systems in a report. To create the report, you export the information that is shown in the performance chart to a CSV file.

If you want to extract configuration, capacity, or performance information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to create your reports.

Related tasks:

“Exporting information to a file” on page 336

You can save information about resources, jobs, or alerts to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of resources, jobs, or alerts.

Related information:

 [REST API](#)



Learn how to use the REST API for IBM Spectrum Control to access information about resources and to generate custom configuration and performance reports.

Creating resource reports about capacity shortages in pools

To include information about the capacity of your resources in reports, use the reporting feature in IBM Spectrum Control.

You check the capacity of pools on your SAN Volume Controller storage system and see that the capacity of some of the pools is nearly depleted. So, you decide to notify your colleagues by including information about the pools in a report.

1. Click **Storage > Block Storage Systems**.
2. Double-click the storage system and click **Pools** in the navigation pane.
3. On the Pools page, right-click one of the column headings and click **Zero Capacity**.

Tip: If you want to export information only about the pools that are depleted, click , select **Zero Capacity**, and type Depleted in the filter  field.

4. Click **Actions** and choose one of the options.

You can get information about the capacity of the pools, whether the pools are depleted, and when the pools will become depleted.

Creating resource reports about the input/output performance of volumes

To include information about the performance of your resources in reports, use the reporting feature in IBM Spectrum Control.

You want to include information about the input/output performance of the volumes in a report.

1. Click **Storage > Volumes**.
2. Select all of the volumes in the list and click **View Performance**. By default, the top 10 volumes are ordered by total input/out rate and the reporting period is set to 24 hours.
3. Click **Actions > Export > Export as CSV**.

Information about the performance of the volumes with the highest input/output rates are exported to a CSV file, which you can download and include in a report.

Appendix A. Reference

View reference information that is related to IBM Spectrum Control. Topics include information about alerts, fabrics, commands, configuration and log files, performance metrics, protocols, standards, and accessibility features.

Alerts

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to configuration and performance conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Alerts are generated when particular conditions are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource

For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource. Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem V840, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. In general, the following types of conditions can trigger alerts:

- An attribute or configuration of a resource changed
- The capacity of a resource fell outside a specified range
- The performance of a resource fell outside a specified range
- The storage infrastructure was changed, such as a new or removed resource
- Data is not being collected for a resource

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over used or under used. IBM Spectrum Control provides many metrics for measuring performance and determining violations of the thresholds that you specify.

Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, IBM Tivoli Netcool/OMNIbus events, login notifications, entries in Windows event log or UNIX syslog, and emails.

Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

Acknowledging alerts: Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, IBM Tivoli Netcool/OMNIbus events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable IBM Spectrum Control to listen to SNMP traps. The default port number is 162, and the default community is public.

Related concepts:

“Alerting and event processing” on page 148

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Related tasks:

“Acknowledging alerts” on page 193

Some alerts in IBM Spectrum Control are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

Triggering conditions for alerts

Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

Related concepts:

“Alerting and event processing” on page 148

IBM Spectrum Control alerting functions can examine the attributes, capacity, and performance of resources and alert you when events or violations are detected. Determining when and how you are alerted to conditions and events within your storage environment is important to helping you maintain and administer storage resources. Many conditions and events can trigger alerts.

Related tasks:

“Adding resources” on page 112

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Triggering conditions for storage system alerts

You can set up IBM Spectrum Control so that it examines the attributes, capacity, and performance of a storage system and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following resources:

- Storage Systems (performance)
- “Storage Systems (general changes)” on page 373
- “Internal resources (common conditions)” on page 376
- “Clusters” on page 376
- “Disks” on page 377
- “Disk Groups” on page 377
- “Filesets” on page 377
- “File Systems” on page 378
- “File System Pools” on page 378
- “Host Connections” on page 378
- “I/O Groups” on page 378
- “Managed Disks” on page 378
- “Modules” on page 379
- “Network Shared Disks” on page 379
- “Nodes” on page 379
- “Pools” on page 380
- “Ports” on page 383
- “Quotas” on page 383
- “RAID Arrays” on page 383
- “Shares” on page 384
- “Volumes” on page 384

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a storage system. To view a complete list of attributes upon which you can alert, go to the Block Storage Systems page, File Storage Systems page, or Object Storage Systems page, right-click the storage system, and select **Edit Alert Definitions**.

Attributes that are automatically configured for alerts are marked with an asterisk (*).

Tips:

- The type of storage system determines which attributes and performance conditions are available for alerts. For example, triggering conditions for shares are available only for storage systems that are configured for file storage, such as Storwize V7000 Unified.

- For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Storage Systems (performance)

Define alerts that notify you when the performance of a storage system falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of volumes, disks, ports, and nodes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your storage infrastructure.

Examples:

- You can define an alert to be notified when the average number of I/O operations per second for read and write operations on a storage system's volumes is greater than or equal to a specified threshold. Use this alert to be notified when the workload of a volume is high and you might need to balance that load across other volumes to improve overall performance.
- You can define an alert to be notified when the percentage of the average response time that can be attributed to delays from host systems is greater than or equal to a specified threshold. Use this alert to be notified of slow hosts and fabrics that might not be working efficiently.
- You can also define an alert that notifies you when a metric is less than a specified threshold, such as if you want to identify volumes that might be under used.

Tips:

- The type of storage system determines the metrics that can be alerted upon. For a list of the metrics that are available for each type of storage system, see “Performance metrics” on page 403.
- A performance monitor must collect data about a resource before IBM Spectrum Control can determine whether a threshold is violated and an alert is generated for a performance condition.

Best practice: When you set thresholds for performance conditions, try to determine the best value so you can derive the maximum benefit without generating too many false alerts. Because suitable thresholds are highly dependent on the type of workload that is being run, hardware configuration, the number of physical disks, exact model numbers, and other factors, there are no easy or standard default rules.

A recommended approach is to monitor the performance of resources for a number of weeks and by using this historical data, determine reasonable threshold values for each performance condition. After that is done, you can fine-tune the condition settings to minimize the number of false alerts.

Storage Systems (general changes)

Table 72. Triggering attributes and conditions for general changes on storage systems

General Attributes	Triggering Conditions for Attributes
Firmware	<p>The firmware version of the microcode on a storage system. For the DS-series of storage systems, this value represents the SEA version of the firmware.</p> <p>To view information about the code bundles for the firmware versions of the DS-series, go to the IBM support site at http://www.ibm.com/support/entry/portal/Overview and search for code bundle information. An internet connection is required to access the support site.</p>
Last Successful Probe Last Successful Monitor	<p>A specified amount of time has passed since a probe or performance monitor was able to collect data about a storage system. You can use this alert to be notified when up-to-date configuration, status, or performance data is not being collected about a storage system and its existing data might be stale. This situation might occur if the resource, network, or IBM Spectrum Control server is unavailable.</p>
Performance Monitor Status*	<p>One of the following statuses is detected for a performance monitor:</p> <p>Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if the resource was rebooted during data collection, no valid performance data was provided by the resource, or a communication error occurred with the resource or its associated agent.</p> <p>Error A performance did not complete when it attempted to collect performance data about the resource. This status might occur if the resource cannot be reached during data collection, or if no configuration data is available for the resource.</p> <p>For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.</p>
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>

Table 72. Triggering attributes and conditions for general changes on storage systems (continued)

General Attributes	Triggering Conditions for Attributes
Status	<p>One of the following statuses is detected for a storage system:</p> <p>Not Normal An error or warning status was detected for a storage system.</p> <p>Warning A warning status was detected for a storage system. This status might occur if a storage system comes online or if its version changes.</p> <p>Error An error status was detected for a storage system. This status might occur if the cooling fans in a storage system are stopped and the internal temperature is too high or if a storage system goes offline.</p>

Table 73. Triggering attributes and conditions for capacity changes on storage systems.

Capacity Attributes	Triggering Conditions for Attributes
Assigned Volume Space	The total volume space in the storage system that is mapped or assigned to host systems, including child pool capacity.
Available Pool Space	<p>The total amount of the space in the pools that is not allocated to the volumes in the pools. To calculate available space, the following formula is used: (pool capacity - allocated space)</p> <p>For XIV systems, pool capacity is the physical capacity of the pools and does not include the virtual capacity of the pools. For other storage systems, pool capacity might also include the overhead space for the pools that are not formatted.</p>
Free Capacity	The amount of space in the pools on the storage system that is available for creating volumes.
Physical Allocation	The percentage of physical space in pools that is allocated to volumes, including child pools. The value is always less than or equal to 100% because you cannot allocate more physical space to the volumes than is available in the pools. This value is determined by the formula, $Allocated\ Space \div Capacity \times 100$. For example, if the space that is reserved for volumes is 50 GiB for a volume size of 200 GiB, physical allocation is 25%.

Table 73. Triggering attributes and conditions for capacity changes on storage systems (continued).

Capacity Attributes	Triggering Conditions for Attributes
Pool Shortfall	<p>The difference between the amount of virtual space that is committed to the volumes in the pools and the actual physical space that is available in the pools, which is expressed as a percentage. As the virtual space is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.</p> <p>This value is determined by the formula, <i>Unallocatable Volume Space</i> ÷ <i>Unallocated Volume Space</i> × 100.</p> <p>A pool shortfall occurs when you commit more space to the volumes in the pools than is physically available to the pools. If the physical space available to the pools is less than the committed virtual space, then the pools do not have enough space to fulfill the commitment to the virtual space.</p> <p>For example, the physical capacity of the pools is 70 GiB, but 150 GiB of virtual space was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% [(80 GiB ÷ 100 GiB × 100)].</p>
Pool Space	The amount of space in pools that are on the storage system.
Raw Disk Capacity	The total unformatted disk capacity of a storage system. When this value is calculated, IBM Spectrum Control does not include the capacity of storage system disks that become missing after data collection.
Remaining Unallocated Volume Space	The amount of the total volume capacity in the storage pool that is not allocated.
Reserved Pool Space	<p>The amount of unallocated space in the pool that is reserved for provisioning and optimization tasks.</p> <p>Pool space is reserved when a provisioning or optimization task is created, and allocated when the task is run.</p>
Shortfall	<p>The difference between the amount of virtual space that is committed to the volumes in the pools and the actual physical space that is available in the pools. As the virtual space is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.</p> <p>This value is determined by the formula, <i>Unallocatable Space</i> ÷ <i>Committed but Unallocated Space</i></p> <p>For example, the physical capacity of the pools is 70 GiB, but 150 GiB of virtual space was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB).</p>
Snapshot Space	The amount of space that is used by all of the snapshots of the file systems that are associated with the GPFS cluster.

Table 73. Triggering attributes and conditions for capacity changes on storage systems (continued).

Capacity Attributes	Triggering Conditions for Attributes
Total File System Capacity	The total amount of storage space on all of the file systems on the storage system or filer.
Unassigned Volume Space	The total volume space in the storage system that is not mapped or assigned to host systems.
Used Capacity	The amount of space in the pools on the storage system that is used by volumes.
Used Pool Space	The amount of space in the pool that is allocated to and used by volumes.
Used Volume Space	The amount of space in the storage system that is allocated to and used by volumes.
Volume Capacity	The total storage space on all the volumes in pools. For thin-provisioned and compressed volumes, this value includes virtual space. For volumes with parent pools, this value includes child pool capacity.

Internal resources (common conditions)

There are a number of alert conditions that are common to many of the internal resources in a storage system. These common conditions represent key changes in your storage infrastructure. For example, you can specify conditions that generate alerts when specific internal resources are added to or deleted from a storage system, or when current data isn't being collected about resources.

Table 74. Triggering attributes and conditions that are common to internal resources

General Attributes	Triggering Conditions for Attributes
New <i>resource</i>	A resource is detected for the first time. Use this alert to be notified when new physical and logical resources are added to a storage system.
Removed <i>physical resource</i> Deleted <i>logical resource</i>	<p>A previously monitored resource can no longer be found. Historical data about the resource is retained, but no current data is being collected. Use these alerts to be notified if a physical or logical resource is removed, deleted, or becomes unavailable.</p> <p>Physical resources include disks, RAID arrays, I/O groups, ports, nodes, host connections, clusters, and file systems. Logical resources include volumes, pools, filesets, and shares.</p>

Clusters

Table 75. Triggering attributes and conditions for general changes on clusters

General Attributes	Triggering Conditions for Attributes
New Cluster	A cluster is detected for the first time.
Removed NAS Cluster	A previously monitored NAS cluster can no longer be found. Historical data about the cluster is retained, but no current data is being collected. Use this alert to be notified if a cluster is removed or becomes unavailable.

Disks

Table 76. Triggering attributes and conditions for general changes on storage system disks

General Attributes	Triggering Conditions for Attributes
New Disk	A disk is detected for the first time.
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.

Table 77. Triggering attributes and conditions for capacity changes on disks.

Capacity Attributes	Triggering Conditions for Attributes
Available Disk Space	The amount of storage space that is available (not allocated) on the disk.
Disk Capacity	The total amount of storage space that is on the disk.

Disk Groups

Table 78. Triggering attributes and conditions for general changes on storage system disk groups

General Attributes	Triggering Conditions for Attributes
New Disk Group	A disk group is detected for the first time.
Removed Disk Group	A previously monitored disk group can no longer be found. Historical data about the disk group is retained, but no current data is being collected. Use this alert to be notified if a disk group is removed or becomes unavailable.
Status	<p>One of the following statuses is detected for a disk group:</p> <p>Not Normal An error or warning condition is detected on a disk group.</p> <p>Warning A warning condition is detected on a disk group. This condition might occur if any of the disks in the disk group goes offline.</p> <p>Error An error condition is detected on a disk group. This condition might occur under the following conditions:</p> <ul style="list-style-type: none">• The percentage of remaining, unallocated volume space in a disk group that is not available to be allocated is too high.• The available space in a disk group that is not reserved for volumes is too low.• The disk group virtual allocation exceeds the warning or critical threshold boundary value.

Filesets

Table 79. Triggering attributes and conditions for general changes on filesets

General Attributes	Triggering Conditions for Attributes
Deleted NAS Fileset	A previously monitored NAS fileset can no longer be found. Historical data about the fileset is retained, but no current data is being collected. Use this alert to be notified if a fileset is deleted or becomes unavailable.
New Fileset	A fileset is detected for the first time.
State	A fileset is linked to a file system, unlinked from a file system, or deleting.

File Systems

Table 80. Triggering attributes and conditions for general changes on file systems

General Attributes	Triggering Conditions for Attributes
New File System	A file system is detected for the first time.
Removed NAS File System	A previously monitored NAS file system can no longer be found. Historical data about the file system is retained, but no current data is being collected. Use this alert to be notified if a file system is deleted or becomes unavailable.

File System Pools

Table 81. Triggering attributes and conditions for general changes on file system pools

General Attributes	Triggering Conditions for Attributes
Deleted Pool	A previously monitored NAS pool can no longer be found. Historical data about the pool is retained, but no current data is being collected. Use this alert to be notified if a pool is deleted or becomes unavailable.
New Pool	A NAS pool is detected for the first time.

Host Connections

Table 82. Triggering attributes and conditions for general changes on host connections

General Attributes	Triggering Conditions for Attributes
New Host Connection	A host connection is detected for the first time.
Removed Host Connection	A previously monitored host connection can no longer be found. Historical data about the host connection is retained, but no current data is being collected. Use this alert to be notified if a host connection is removed or becomes unavailable.

I/O Groups

Table 83. Triggering attributes and conditions for general changes on I/O groups

General Attributes	Triggering Conditions for Attributes
New I/O Group	A new I/O group was detected for the first time.
Removed I/O Group	A previously monitored I/O group can no longer be found. Historical data about the I/O group is retained, but no current data is being collected. Use this alert to be notified if an I/O group is removed or becomes unavailable.

Managed Disks

Table 84. Triggering attributes and conditions for general changes on managed disks

General Attributes	Triggering Conditions for Attributes
Managed Disk Status	One of the following statuses is detected for a managed disk: Not Normal An error or warning status was detected on a managed disk. Warning A warning status was detected on a managed disk. Error An error status was detected on a managed disk.
New Managed Disk	A new managed disk was detected for the first time.

Table 84. Triggering attributes and conditions for general changes on managed disks (continued)

General Attributes	Triggering Conditions for Attributes
Removed Managed Disk	A previously monitored managed disk can no longer be found. Historical data about the managed disk is retained, but no current data is being collected. Use this alert to be notified if a managed disk is removed or becomes unavailable.

Table 85. Triggering attributes and conditions for capacity changes on managed disks.

Capacity Attributes	Triggering Conditions for Attributes
Available Space	The amount of unused storage space on the managed disk.
Total Space	The total amount of storage space that is on the managed disk on the storage system. This attribute is only available for Storwize V7000 storage systems that are configured as back-end storage.

Modules

Table 86. Triggering attributes and conditions for general changes on modules

General Attributes	Triggering Conditions for Attributes
Deleted Module	A previously monitored module can no longer be found. Use this alert to be notified if a module is removed or becomes unavailable.

Network Shared Disks

Table 87. Triggering attributes and conditions for general changes on NSDs

General Attributes	Triggering Conditions for Attributes
New NSD	A new NSD was detected for the first time.
Removed NSD	A previously monitored NSD can no longer be found. Historical data about the NSD is retained, but no current data is being collected. Use this alert to be notified if an NSD is removed or becomes unavailable.
Status	One of the following statuses is detected for an NSD: Not Normal An error or warning status was detected on an NSD. Warning A warning status was detected on an NSD. This status might occur if a storage system comes online or if its version changes. Error An error status was detected on an NSD.

Nodes

Table 88. Triggering attributes and conditions for general changes on nodes

General Attributes	Triggering Conditions for Attributes
New Node	A new node was detected for the first time.
Removed Node (Block storage)	A previously monitored node can no longer be found. Historical data about the node is retained, but no current data is being collected. Use this alert to be notified if a node is removed or becomes unavailable.
Removed NAS Node (File storage)	A previously monitored NAS node can no longer be found. Historical data about the node is retained, but no current data is being collected. Use this alert to be notified if a node is removed or becomes unavailable.

Table 88. Triggering attributes and conditions for general changes on nodes (continued)

General Attributes	Triggering Conditions for Attributes
Cloud Gateway Status (File storage)	<p>One of the following statuses is detected for a node. This attribute is available only for IBM Spectrum Scale.</p> <p>Not Running The gateway service is not running because, for example, the service is stopped, no cloud account is configured, or the connection to the cloud provider failed.</p> <p>Not Installed The node is not a cloud gateway.</p> <p>No File System The node is a cloud gateway and the gateway service is running, but the node is not yet assigned to a file system.</p> <p>Stopped The node is a cloud gateway for a file system, but the gateway service is stopped.</p> <p>No Cloud Account The node is a cloud gateway for a file system and the gateway service is running, but no cloud account is configured.</p> <p>Disconnected The node is a cloud gateway, the gateway service is running, and a cloud account is configured, but the connection to the cloud provider failed.</p>

Pools

Table 89. Triggering attributes and conditions for general changes on pools

General Attributes	Triggering Conditions for Attributes
Deleted Pool	A previously monitored pool can no longer be found. Historical data about the pool is retained, but no current data is being collected. Use this alert to be notified if a pool is deleted or becomes unavailable.
New Storage Pool	A new pool was detected for the first time.
Status	<p>One of the following statuses is detected for a pool:</p> <p>Not Normal An error or warning condition is detected on a pool.</p> <p>Warning A warning condition is detected on a pool. This condition might occur if a pool goes offline.</p> <p>Error An error condition is detected on a pool. This condition might occur under the following conditions:</p> <ul style="list-style-type: none"> • The percentage of remaining, unallocated volume space in a pool that is not available to be allocated is too high. • The available space in a pool that is not reserved for volumes is too low. • The pool virtual allocation exceeds the warning or critical threshold boundary value.

Table 90. Triggering attributes and conditions for capacity changes on pools.

Capacity Attributes	Triggering Conditions for Attributes
Allocated Space	The amount of space that is reserved for all the volumes in a pool.
Assigned Volume Space	<p>The amount of space on all the volumes in a pool that are mapped or assigned to host systems.</p> <p>For a thin-provisioning pool, this value includes the virtual capacity of thin-provisioned volumes, which might exceed the total space in the pool. For Hitachi VSP non-thin provisioning pool space, this value is the sum of assigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.</p>
Available Repository Space	<p>The amount of available, unallocated storage space on all extents in the repository of a pool for Track Space-Efficient (TSE) thin-provisioning. This attribute applies only to the DS8000 storage systems.</p> <p>You can use this alert to be notified about space-efficient volumes.</p>
Available Virtual Capacity	The amount of virtual capacity in a thin-provisioned pool that is not used by volumes.
Physical Allocation	<p>The percentage of physical capacity in the pool that is allocated to the regular volumes, the thin-provisioned volumes, and the volumes in child pools. This value is always less than or equal to 100% because you cannot allocate more physical space than is available in a pool.</p> <p>This value is determined by the formula, $\text{Allocated Space} \div \text{Capacity} \times 100$. For example, if the space that is reserved for volumes is 50 GiB for a volume size of 200 GiB, physical allocation is 25%.</p>
Pool Shortfall	<p>The difference between the amount of virtual space that is committed to the volumes in the pools and the actual physical space that is available in the pools, which is expressed as a percentage. As the virtual space is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.</p> <p>This value is determined by the formula, $\text{Unallocatable Volume Space} \div \text{Unallocated Volume Space} \times 100$.</p> <p>A pool shortfall occurs when you commit more space to the volumes in the pools than is physically available to the pools. If the physical space available to the pools is less than the committed virtual space, then the pools do not have enough space to fulfill the commitment to the virtual space.</p> <p>For example, the physical capacity of the pools is 70 GiB, but 150 GiB of virtual space was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% [(80 GiB ÷ 100 GiB) × 100].</p>
Repository Space	The amount of space on all extents in the repository of a pool. This attribute applies only to the DS8000 storage systems.

Table 90. Triggering attributes and conditions for capacity changes on pools (continued).

Capacity Attributes	Triggering Conditions for Attributes
Reserved Pool Space	<p>The amount of unallocated space in a pool that is reserved for provisioning and optimization tasks.</p> <p>Pool space is reserved when a provisioning or optimization task is created, and allocated when the task is run.</p>
Shortfall	<p>The difference between the amount of virtual space that is committed to the volumes in the pools and the actual physical space that is available in the pools. As the virtual space is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.</p> <p>This value is determined by the formula, <i>Unallocatable Space ÷ Committed but Unallocated Space</i></p> <p>For example, the physical capacity of the pools is 70 GiB, but 150 GiB of virtual space was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB).</p>
Total Reserved Capacity	<p>The total amount of space on the pool that is reserved for provisioning and optimization tasks. Pool space is reserved when a provisioning or optimization task is created, and allocated when the task is run.</p>
Unused Space	<p>The amount of space that is allocated to the volumes in a pool and is not yet used. This value is determined by the formula, <i>Allocated Space – Used Space</i>.</p>
Used Repository Space	<p>The amount of used space on all extents in the repository of a pool. This attribute applies only to the DS8000 storage systems.</p> <p>You can use this alert to be notified about space-efficient volumes.</p>
Used Volume Space	<p>The amount of space on the storage system that is used by volumes.</p>
User Reserved Capacity	<p>The amount of space in the pools on the storage system that is reserved for user-defined purposes.</p>
Virtual Allocation	<p>The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool.</p> <p>You can use this alert to be notified about space-efficient volumes.</p>
Virtual Capacity Limit	<p>The maximum amount of virtual storage space available to allocate to volumes in the storage pools that are associated with the storage system.</p> <p>You can use this alert to be notified about space-efficient volumes.</p>
Virtual Volume Space	<p>The total amount of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool.</p> <p>You can use this alert to be notified about space-efficient volumes.</p>

Ports

Table 91. Triggering attributes and conditions for general changes on ports

General Attributes	Triggering Conditions for Attributes
New Port	A new port was detected for the first time.
Removed Port	A previously monitored port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is deleted or becomes unavailable.
Status	One of the following statuses is detected for a port: Not Normal An error or warning condition is detected on a port. Warning A warning condition is detected on a port. Error An error condition is detected on a port.

Quotas

Table 92. Triggering attributes and conditions for general changes on quotas

General Attributes	Triggering Conditions for Attributes
Deleted Quota	A previously monitored quota can no longer be found. Historical data about the quota is retained, but no current data is being collected. Use this alert to be notified if a quota is deleted or becomes unavailable.
New Quota	A new quota was detected for the first time.

RAID Arrays

Table 93. Triggering attributes and conditions for general changes on RAID arrays

General Attributes	Triggering Conditions for Attributes
New RAID Array	A new RAID array was detected for the first time.
Removed RAID Array	A previously monitored RAID array can no longer be found. Historical data about the RAID array is retained, but no current data is being collected. Use this alert to be notified if a RAID array is removed or becomes unavailable.
Status	One of the following statuses is detected for a RAID array: Not Normal An error or warning condition is detected on a RAID array. Warning A warning condition is detected on a RAID array. Error An error condition is detected on a RAID array.

Table 94. Triggering attributes and conditions for capacity changes on RAID arrays.

General Attributes	Triggering Conditions for Attributes
RAIDArrayFreeSpace	The amount of unused storage space in the RAID array.
Total Space	The total amount of storage space in the RAID array.

Shares

Table 95. Triggering attributes and conditions for general changes on shares

General Attributes	Triggering Conditions for Attributes
New Share	A new share was detected for the first time.
Deleted Share	A previously monitored share can no longer be found. Historical data about the share is retained, but no current data is being collected. Use this alert to be notified if a share is deleted or becomes unavailable.

Volumes

Table 96. Triggering attributes and conditions for general changes on volumes

General Attributes	Triggering Conditions for Attributes
Deleted Volume	A previously monitored volume can no longer be found. Historical data about the volume is retained, but no current data is being collected. Use this alert to be notified if a volume is deleted or becomes unavailable.
New Volume	A new volume was detected for the first time.
Status	One of the following statuses is detected for a volume: Not Normal An error or warning condition is detected on a RAID array. Warning A warning condition is detected on a RAID array. Error An error condition is detected on a RAID array.

Table 97. Triggering attributes and conditions for capacity changes on volumes.

Capacity Attributes	Triggering Conditions for Attributes
Allocated Space	The total amount of storage space on a pool that is physically allocated to a volume.
Real Capacity	The total amount of storage space that is physically allocated to a volume. For thin-provisioned volumes, this value is less than the virtual capacity of the volume. In an XIV and IBM Spectrum Accelerate, this value represents the physical ("hard") capacity of the volume, not the virtual ("soft") capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.
Remaining Managed Space	The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.
Space	The amount of space in a pool that is allocated to a volume.
Unallocatable Space	The amount of space by which the capacity of a volume exceeds the physical capacity of the associated pool. In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool by creating volumes with more virtual capacity than can be physically allocated in the pool. This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Table 97. Triggering attributes and conditions for capacity changes on volumes (continued).

Capacity Attributes	Triggering Conditions for Attributes
Unallocated Space	<p>The amount of space in a pool that is not reserved for a volume. This value is determined by the formula: <i>Capacity – Allocated Space</i>.</p> <p>The value for Unallocated Space is available only for thin provisioned volumes.</p>
Uncompressed Used Capacity	<p>The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.</p>
Unused Space	<p>The amount of space that is allocated to a volume and is not yet used. This value is determined by the formula, <i>Allocated Space – Used Space</i>.</p> <p>The value for Unused Space is available only for SAN Volume Controller and Storwize family storage systems that are configured with block storage.</p>
Used Space	<p>The amount of allocated space that is used by a volume. For SAN Volume Controller and Storwize family storage systems that are configured with block storage, you can pre-allocate thin-provisioned volume space when the volumes are created. In these cases, the Used Space might be different from the Allocated Space. For compressed volumes on SAN Volume Controller and Storwize family storage systems that are configured with block storage, the Used Space reflects the size of compressed data that is written to disk. As the data changes, the Used Space might at times be less than the Allocated Space. For volumes that are not thin provisioned in SAN Volume Controller, Storwize family storage systems that are configured with block storage, and other storage systems, the values for Used Space and Allocated Space are equal. This value is accurate as of the most recent time that data was collected about a volume. Data collection is run on a set schedule and the used space on a volume can change rapidly.</p>

Triggering conditions for hypervisor alerts

You can set up IBM Spectrum Control so that it examines the attributes and capacity of a hypervisor and notifies you when changes are detected.

Alerts can notify you of general changes and capacity changes on the following resources:

- “Hypervisors” on page 386
- “Controllers” on page 386
- “Data Stores” on page 387
- “Disks” on page 387
- “Paths” on page 388
- “Virtual Machines” on page 388
- “VMDKs” on page 389

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a hypervisor. To view a complete list of hypervisor attributes upon which you can alert, go to **Servers > Hypervisors**, right-click the hypervisor, and select **Edit Alert Definitions**.

Attributes that are automatically configured for alerts are marked with an asterisk (*).

Tip: For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Hypervisors

Table 98. Triggering attributes and conditions for general hypervisor changes.

General Attributes	Descriptions of Triggering Conditions
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a hypervisor. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a hypervisor and its existing data might be stale. Data collection might be interrupted or not occur if the resource, network, or IBM Spectrum Control server is unavailable.
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Status	<p>One of the following statuses is detected on a hypervisor:</p> <p>Not Normal An error or warning status was detected on the hypervisor or its internal resources.</p> <p>Warning A warning status was detected on the hypervisor or its internal resources.</p> <p>Error (default) An error status was detected on the hypervisor or its internal resources. For example, an error status might occur when a hypervisor goes offline, or a server disk is disconnected or partially disconnected if it has multiple paths and one of the paths is disconnected.</p> <p>Unreachable One or more of the monitored resources for a hypervisor are not responding. This status might be caused by a problem in the network.</p>

Controllers

Table 99. Triggering attributes and conditions for general changes on controllers.

Controller Attributes	Descriptions of Triggering Conditions
New Disk Controller	A disk controller is detected for the first time. Use this alert to be notified of hardware additions on servers.

Table 99. Triggering attributes and conditions for general changes on controllers (continued).

Controller Attributes	Descriptions of Triggering Conditions
Removed Disk Controller	A previously monitored disk controller can no longer be found. Historical data about the controller is retained, but no current data is being collected. Use this alert to be notified if a controller is removed or becomes unavailable.

Data Stores

Table 100. Triggering attributes and conditions for general changes on data stores.

Data Store Attributes	Descriptions of Triggering Conditions
Deleted Logical Volume	A previously monitored logical volume on a data store can no longer be found. Historical data about the volume is retained, but no current data is being collected. Use this alert to be notified if a volume is deleted or becomes unavailable.
New Logical Volume	A logical volume on the data store is detected for the first time.
VMDKs	You can use a number of operators to determine when you are notified of the number of virtual machine disk files (VMDKs) on the data store. For example, you can select to be notified when the number of VMDKs changes, or when the number is or is not a specific number, or when the number lies outside a specific range.

Table 101. Triggering attributes and conditions for capacity changes on data stores.

Capacity Attributes	Descriptions of Triggering Conditions
Available Data Store Space	The amount of unused storage space on the hypervisor data store.
Available File System Space	The total amount of unused storage space on a file system.
Data Store Capacity	The total amount of storage space that is assigned to a data store.
Total File System Capacity	The total amount of storage space on a file system.
Used File System Space	The amount of used storage space on the hypervisor file system.
Used Space	The percentage of storage space that is used on the hypervisor data store.

Disks

Table 102. Triggering attributes and conditions for general changes on disks.

General Attributes	Descriptions of Triggering Conditions
Firmware	The version of the Licensed Internal Code on the disk changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Multipathing Policy	The multipathing policy that is in effect for a disk. For example, you can be notified when the policy changes, or when the policy is Round Robin, Load Balancing, Failover Only, or other another policy.
New Disk	A disk is detected for the first time. Use this alert to be notified of hardware changes on servers or hypervisors.
Paths	The number of access paths that are associated with the disk falls outside a specified range, or is equal to or not equal to a specified value.

Table 102. Triggering attributes and conditions for general changes on disks (continued).

General Attributes	Descriptions of Triggering Conditions
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.
Status*	<p>One of the following statuses is detected on a disk:</p> <p>Not Normal An error or warning status was detected on the disk.</p> <p>Warning A warning status was detected on the disk.</p> <p>Error (default) An error status was detected on the disk.</p>

Table 103. Triggering attributes and conditions for capacity changes on disks.

Capacity Attributes	Descriptions of Triggering Conditions
Available Disk Space	The amount of unused storage space on a hypervisor disk.
Capacity	The total amount of storage space assigned to a hypervisor disk.
Used Space	The amount of used storage space on a hypervisor disk .

Paths

Table 104. Triggering attributes and conditions for general changes on paths.

Path Attributes	Descriptions of Triggering Conditions
Deleted Path	A previously monitored access path for a server disk can no longer be found. This change might or might not affect the availability of the disk because there might be more than one path available.
New Path	An access path for a disk is detected for the first time.
Status*	<p>One of the following statuses is detected on a path:</p> <p>Not Normal An error or warning status was detected on the path.</p> <p>Warning A warning status was detected on the path.</p> <p>Error (default) An error status was detected on the path.</p>

Virtual Machines

Table 105. Triggering attributes and conditions for general changes on virtual machines.

Virtual Machine Attributes	Descriptions of Triggering Conditions
Deleted Virtual Machine	A previously monitored access path for a virtual machine can no longer be found. Historical data about the virtual machine is retained, but no current data is being collected. Use this alert to be notified if a virtual machine is deleted or becomes unavailable.
New Virtual Machine	A virtual machine is detected for the first time.

Table 105. Triggering attributes and conditions for general changes on virtual machines (continued).

Virtual Machine Attributes	Descriptions of Triggering Conditions
Status	<p>One of the following statuses is detected on a virtual machine:</p> <p>Not Normal An error or warning status was detected on the virtual machine.</p> <p>Warning A warning status was detected on the virtual machine.</p> <p>Error (default) An error status was detected on the virtual machine.</p>
VMDKs	You can use a number of operators to determine when you are notified of the number of VMDKs on the virtual machine. For example, you can select to be notified when the number of VMDKs changes, or when the number is or is not a specific number, or when the number lies outside a specific range.

Table 106. Triggering attributes and conditions for capacity changes on virtual machines.

Capacity Attributes	Descriptions of Triggering Conditions
Capacity	The total amount of storage space assigned to a virtual machine hosted by the hypervisor.

VMDKs

Table 107. Triggering attributes and conditions for general changes on virtual machine disks (VMDKs).

VMDK Attributes	Descriptions of Triggering Conditions
Deleted Virtual Disk	A previously monitored virtual disk can no longer be found. Historical data about the virtual disk is retained, but no current data is being collected. Use this alert to be notified if a virtual disk is deleted or becomes unavailable.
New Virtual Disk	A virtual disk is detected for the first time.

Table 108. Triggering attributes and conditions for capacity changes on VMDKs.

Capacity Attributes	Descriptions of Triggering Conditions
Size	The size of the VMDK on the hypervisor data store.
Used Space	The percentage of storage space that is used on the VMDK on the hypervisor data store. The used space is available if a Storage Resource agent is deployed.

Triggering conditions for switch alerts

You can set up IBM Spectrum Control so that it examines the attributes and performance of a switch and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following resources:

- Switches (performance)
- “Switches (general changes)” on page 391
- “Ports” on page 392

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key

properties of a switch. To view a complete list of attributes upon which you can alert, go to **Network > Switches**, right-click a switch, and select **Edit Alert Definitions**.

Attributes that are automatically configured for alerts are marked with an asterisk (*).

Switches (performance)

Define alerts that notify you when the performance of a switch falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of switch ports, including I/O, data, and error rates, and frame transfer sizes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your network infrastructure.

For example, you can define an alert to be notified when the port congestion index for a port is greater than or equal to a specified threshold. Port congestion represents the estimated degree to which frame transmission was delayed due to a lack of buffer credits. Use this alert to help identify port conditions that might slow the performance of the resources to which those ports are connected.

You can also be notified when a metric is less than a specified threshold, such as when you want to identify ports that might be under used.

For a complete list of switch metrics that can be alerted upon, see “Performance metrics for switches” on page 464.

Tips for performance conditions:

- A performance monitor must collect data about a resource before IBM Spectrum Control can determine whether a threshold is violated and an alert is generated for a performance condition.
- When you define a performance alert for the inter-switch connections category of a switch, the performance of ISL trunk connections is measured and alerted on. To alert on the performance of standard ISL connections, define performance alerts for the ports category of a switch.

Best practice: When you set thresholds for performance conditions, try to determine the best value so you can derive the maximum benefit without generating too many false alerts. Because suitable thresholds are highly dependent on the type of workload that is being run, hardware configuration, the number of physical disks, exact model numbers, and other factors, there are no easy or standard default rules.

A recommended approach is to monitor the performance of resources for a number of weeks and by using this historical data, determine reasonable threshold values for each performance condition. After that is done, you can fine-tune the condition settings to minimize the number of false alerts.

Switches (general changes)

Table 109. Triggering attributes and conditions for general changes on switches

General Attributes	Triggering Conditions for Attributes
Firmware	The version of the Licensed Internal Code on the switch changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Last Successful Probe Last Successful Monitor	A specified amount of time has passed since a probe or performance monitor was able to collect data about a switch. You can use this alert to be notified when up-to-date configuration, status, or performance data is not being collected about a switch and its existing data might be stale. This situation might occur if the resource, network, or IBM Spectrum Control server is unavailable.
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Performance Monitor Status*	<p>One of the following statuses is detected for a performance monitor:</p> <p>Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if the resource was rebooted during data collection, no valid performance data was provided by the resource, or a communication error occurred with the resource or its associated agent.</p> <p>Error A performance did not complete when it attempted to collect performance data about the resource. This status might occur if the resource cannot be reached during data collection, or if no configuration data is available for the resource.</p> <p>For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.</p>

Ports

Table 110. Triggering attributes and conditions for general changes on ports

General Attributes	Triggering Conditions for Attributes
Removed Port	A previously monitored port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is removed or becomes unavailable.
Speed	The negotiated speed of a port, which is measured in GiB per second. You can generate an alert when the port speed changes, or when it is greater than, less than, equal to, or is not equal to a specific value.
State	A port is online, enabled but offline, or disabled.
Status*	One of the following statuses is detected for a port: Not Normal An error or warning status is detected on a port. Warning A warning status is detected on a port. This status might occur if the switch is stopped, starting, or in service (being maintained, cleaned, or administered). Error An error status is detected on a port.

Triggering conditions for fabric alerts

You can set up IBM Spectrum Control so that it examines the attributes of a fabric and notifies you when changes are detected.

Alerts can notify you of general changes on the following types of fabric resources:

- “Fabrics”
- “Switches” on page 393

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a fabric. To view a complete list of attributes upon which you can alert, go to **Network > Fabrics**, right-click the fabric, and select **Edit Alert Definitions**.

Attributes that are automatically configured for alerts are marked with an asterisk (*).

Fabrics

Table 111. Triggering attributes and conditions for general changes on fabrics

Fabric Attributes	Triggering Conditions for Attributes
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a fabric. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a fabric and its existing data might be stale. This situation might occur if the fabric, network, or IBM Spectrum Control server is unavailable.

Table 111. Triggering attributes and conditions for general changes on fabrics (continued)

Fabric Attributes	Triggering Conditions for Attributes
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Status*	<p>One of the following statuses is detected on a fabric:</p> <p>Not Normal An error or warning status was detected on the fabric or its internal resources.</p> <p>Warning A warning status was detected on the fabric or its internal resources.</p> <p>Error (default) An error status was detected on the fabric or its internal resources.</p> <p>Unreachable One or more of the monitored resources for a fabric are not responding.</p>

Switches

Table 112. Triggering attributes and conditions for changes on switches

Fabric Attributes	Triggering Conditions for Attributes
Last Successful Monitor	<p>A specified amount of time has passed since the last successful performance monitor for switches in the fabric. You can use this alert to be notified when performance data is not being collected for a fabric's switches. This situation might occur if a switch, network, or IBM Spectrum Control server is unavailable during data collection.</p>

Table 112. Triggering attributes and conditions for changes on switches (continued)

Fabric Attributes	Triggering Conditions for Attributes
Performance Monitor Status*	<p>One of the following statuses is detected for performance monitors that are collecting data about switches in the fabric:</p> <p>Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a switch.</p> <p>Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if a switch was restarted during data collection, no valid performance data was provided by the switch, or a communication error occurred.</p> <p>Error A performance did not complete when it attempted to collect performance data about a switch. This status might occur if a switch cannot be reached during data collection, or if no configuration data is available for a switch.</p> <p>For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a switch, click Data Collection, and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.</p>

Triggering conditions for server alerts

You can set up IBM Spectrum Control so that it examines the attributes and capacity of a server and notifies you when changes are detected.

Alerts can notify you of general changes and capacity changes on the following resources:

- “Servers” on page 395
- “Controllers” on page 396
- “Disks” on page 397
- “Disk Groups” on page 398
- “File Systems and Logical Volumes” on page 398
- “Paths” on page 399
- “Shares” on page 400
- “Triggering conditions for the IBM Spectrum Control server” on page 400

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a server. To view a complete list of server attributes upon which you can alert, go to **Servers > Servers**, right-click the server, and select **Edit Alert Definitions**.

Attributes that are automatically configured for alerts are marked with an asterisk (*).

Tip: For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Servers

Table 113. Triggering attributes and conditions for general changes on servers.

General Attributes	Descriptions of Triggering Conditions
Agent State*	<p>A Storage Resource agent is in one of the following states:</p> <p>Not Normal An error or warning state was detected on a Storage Resource agent.</p> <p>Warning A warning state was detected on a Storage Resource agent. For example, a warning state might occur when an agent must be upgraded to the same version level as the IBM Spectrum Control server to which it is communicating.</p> <p>Error (default) An error state was detected on a Storage Resource agent. For example, an error state might occur when an agent was not able to be upgraded.</p>
Last Successful Probe	<p>A specified amount of time has passed since a probe was able to collect data about a server. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a server and its existing data might be stale. Data collection might be interrupted or not occur if the resource, network, or IBM Spectrum Control server are unavailable.</p>
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>

Table 113. Triggering attributes and conditions for general changes on servers (continued).

General Attributes	Descriptions of Triggering Conditions
Status*	<p>One of the following statuses is detected on a server:</p> <p>Not Normal An error or warning status was detected on the server or its internal resources.</p> <p>Warning A warning status was detected on the server or its internal resources. For example, a warning status might occur when an HBA or HBA to a server node is newly discovered, is missing, or is rediscovered.</p> <p>Error (default) An error status was detected on the server or its internal resources. For example, an error status might occur when a server goes offline, or a server disk is disconnected or partially disconnected if it has multiple paths and one of the paths is disconnected.</p> <p>Unreachable One or more of the monitored resources for a server are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.</p>

Table 114. Triggering attributes and conditions for capacity changes on servers.

Capacity Attributes	Descriptions of Triggering Conditions
Available Disk Space	The amount of unused disk space on the local and SAN-attached storage for the server. SAN-attached storage is assigned to the server from storage systems.
Available File System Space	<p>The amount of unused storage space in the file systems on the server.</p> <p>Unused file system space does not include space that is reserved for the operating system. For example, the unused space for tmpfs on UNIX operating systems is not included in this value.</p>
Assigned SAN Space	<p>The amount of disk space that is assigned to the server from storage systems.</p> <p>The disk space from storage systems is only available when SAN-attached storage is assigned to the server.</p>
File System Capacity	The amount of file system space on the server.
File System Space from Storage Systems	<p>The amount of file system space that is assigned to the server from storage systems.</p> <p>The file system space from storage systems is only available when SAN-attached storage is assigned to the server.</p>
Total Disk Space	The total amount of disk space for all the local and SAN-attached storage on the server. SAN-attached storage is assigned to the server from storage systems.
Used Space	The amount of used disk space on the local and SAN-attached storage for the server. SAN-attached storage is assigned to the server from storage systems.

Controllers

Table 115. Triggering attributes and conditions for changes on disk controllers.

Controller Attributes	Descriptions of Triggering Conditions
<ul style="list-style-type: none"> • Driver Version • Firmware • ROM Version 	<p>The version of the driver, firmware, or read-only memory (ROM) on a disk controller changes. You can use a number of operators to determine when you are notified of a version change, such as when the driver, firmware, or ROM is, or is not, a specific version, or when the version number contains a specific value.</p> <p>Use this alert for HBAs only.</p>
Last Data Collection	A specified amount of time since data was collected about a controller. Use this alert to be notified if data is not being collected about a controller or if the existing data is becoming too stale.
New Disk Controller	A disk controller is detected for the first time. Use this alert to be notified of hardware additions on servers.
Removed Disk Controller	A previously monitored disk controller can no longer be found. Historical data about the controller is retained, but no current data is being collected. Use this alert to be notified if a controller is removed or becomes unavailable.
Status*	<p>One of the following statuses is detected on a disk controller:</p> <p>Not Normal An error or warning status was detected on the controller.</p> <p>Warning A warning status was detected on the controller.</p> <p>Error (default) An error status was detected on the controller.</p>

Disks

Table 116. Triggering attributes and conditions for general changes on disks.

General Attributes	Descriptions of Triggering Conditions
Firmware	The version of the Licensed Internal Code on the disk changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Multipathing Policy	The multipathing policy that is in effect for a disk. For example, you can be notified when the policy changes, or when the policy is Round Robin, Load Balancing, Failover Only, or other another policy.
New Disk	A disk is detected for the first time. Use this alert to be notified of hardware changes on servers or hypervisors.
Paths	The number of access paths that are associated with the disk falls outside a specified range, or is equal to or not equal to a specified value.
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.
Status*	<p>One of the following statuses is detected on a disk:</p> <p>Not Normal An error or warning status was detected on the disk.</p> <p>Warning A warning status was detected on the disk.</p> <p>Error (default) An error status was detected on the disk.</p>

Table 117. Triggering attributes and conditions for capacity changes on disks.

Capacity Attributes	Descriptions of Triggering Conditions
Available Disk Space	The amount of unused storage space on a disk that is attached to the server.
Capacity	The total amount of storage space on a disk that is attached to the server.
Used Space	The amount of used storage space on a disk that is attached to the server.

Disk Groups

Table 118. Triggering attributes and conditions for general changes on disk groups (volume groups).

General Attributes	Descriptions of Triggering Conditions
Deleted Volume Group	A previously monitored volume group can no longer be found. Historical data about the volume group is retained, but no current data is being collected. Use this alert to be notified if a volume group is deleted or becomes unavailable.
New Volume Group	A volume group is detected for the first time.
Status*	One of the following statuses is detected on a disk group: Not Normal An error or warning status was detected on the disk group. Warning A warning status was detected on the disk group. Error (default) An error status was detected on the disk group.

Table 119. Triggering attributes and conditions for capacity changes on disk groups (volume groups).

Capacity Attributes	Descriptions of Triggering Conditions
Available Space	The amount of unused storage space on a server disk group.
Used Space	The amount of used storage space on a server disk group.
Volume Group Capacity	The total amount of storage space on a server volume group. This value is inclusive of all storage space and applies to all space values related to volume groups.

File Systems and Logical Volumes

Table 120. Triggering attributes and conditions for general changes on file systems and logical volumes.

General Attributes	Descriptions of Triggering Conditions
Available Inodes	The number of unused inodes on file systems on the operating system changes, falls outside a specified range, or is equal to or not equal to a specified value.
Deleted File System	A previously monitored file system is deleted or unmounted from a server. Historical data about the file system is retained, but no current data is being collected. This attribute applies to file systems on the following resources: <ul style="list-style-type: none"> Storage systems that are configured for file storage, including Storwize V7000 Unified and IBM Spectrum Scale Servers that are managed by Storage Resource agents

Table 120. Triggering attributes and conditions for general changes on file systems and logical volumes (continued).

General Attributes	Descriptions of Triggering Conditions
Deleted Logical Volume	A previously monitored logical volume can no longer be found. Historical data about the logical volume is retained, but no current data is being collected. Use this alert to be notified if a logical volume is removed or becomes unavailable.
New File System	A file system was detected for the first time. This alert applies to file systems on the following resources: <ul style="list-style-type: none"> Storage systems that are configured for file storage, including Storwize V7000 Unified and IBM Spectrum Scale Servers that are managed by Storage Resource agents
New Logical Volume	A logical volume is detected for the first time.
Used Inodes	The number of used inodes on file systems on the operating system changes. You can use a number of operators to determine when you are notified, such when the number of used inodes falls outside a specified range, or is equal to or not equal to a specified value.

Table 121. Triggering attributes and conditions for capacity changes on file systems and logical volumes.

Capacity Attributes	Descriptions of Triggering Conditions
Available File System Space	The amount of unused storage space on a file system on the server disk.
File System Capacity	The total amount of storage space on a file system on the server disk.
Logical Volume Capacity	The total amount of storage space on a logical volume on the server disk.
Used File System Space	The amount of used storage space on a file system on the server disk.
Used Space	The percentage of used storage space on a file system or logical volume on the server disk.

Paths

Table 122. Triggering attributes and conditions for general changes on paths.

Path Attributes	Descriptions of Triggering Conditions
Deleted Path	A previously monitored access path for a server disk can no longer be found. This change might or might not affect the availability of the disk because there might be more than one path available.
New Path	An access path for a disk is detected for the first time.
Status*	One of the following statuses is detected on a path: <p>Not Normal An error or warning status was detected on the path.</p> <p>Warning A warning status was detected on the path.</p> <p>Error (default) An error status was detected on the path.</p>

Shares

Table 123. Triggering attributes and conditions for general changes on shares.

Share Attributes	Descriptions of Triggering Conditions
Deleted Share	A previously monitored share can no longer be found. Historical data about the share is retained, but no current data is being collected. Use this alert to be notified if a share is removed or becomes unavailable.
New Share	A new share was detected for the first time.

Triggering conditions for the IBM Spectrum Control server

The server on which IBM Spectrum Control is installed is automatically monitored for conditions that might cause an interruption in product functions. When these conditions are detected, alerts are triggered and shown on the **Home > Alerts** page. You do not need to manually define alerts for these product-related conditions; they are automatically enabled.

Table 124. Triggering conditions for the IBM Spectrum Control server.

Triggering Condition	Explanation	Related Error Message
Database unavailable	The product database is not available. This database is the repository for information that is collected about the monitored resources in your environment.	ALR4112E, ALR4113E
High memory usage*	A high amount of memory is being used by a server process and might cause stability problems.	ALR4103W
Database alarm*	The system database or the database manager that hosts the product's database repository is reporting an alarm.	ALR4104W
High workload	The workload queue for the Device server is high and might cause performance issues.	ALR4105W
High number of external events	The server is receiving a high number of external events, such as CIM indications or SNMP traps. The high number of events might cause performance issues.	ALR4106W

Alert suppression actions

When you select triggering conditions for alerts, you can specify the conditions under which alerts are triggered and choose whether you want to suppress repeating alerts. Alerts can be suppressed to avoid generating too many alert log entries or too many actions when the triggering condition occurs often. You can view suppressed alerts in the constraint violation reports.

You can define the following options that enable you to specify conditions that suppress alerts.

Do not suppress alerts

Receive alert notifications whenever an alert violation is detected.

Only alert once until problem clears

Receive one notification for an alert violation, even if the condition is violated multiple times. This option is the default.

Only generate alerts every

Receive one notification when an alert is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.

Use this option when a single occurrence of an alert violation might be insignificant, but repeated occurrences can represent a potential problem.

Do not alert until condition has been violated for more than


Do not receive notifications for an alert unless the violation has not cleared for longer than the specified time.

Use this option to avoid email messages and other disruptive notifications if the same violation occurs repeatedly during the specified time period. In these cases, an alert is generated only for the first occurrence of the violation. Alerts for other occurrences during the time period are suppressed. When the suppression period ends, the cycle starts again.

Triggering actions and notifications for alerts

Specify how you are notified when alert conditions are detected on resources and any actions to take as a result of those alerts. These settings can be defined globally for a resource, or for specific alert definitions that apply to a resource.

Table 125. Locations where you can specify alert actions and notification settings.

Task	Location in GUI
Specify the notification settings for all the alert definitions that apply to a resource	Resource page and resource details page For example, to specify the global notification settings for a block storage system, go to Storage > Block Storage Systems , right-click a storage system, and select Edit Alert Notification Settings . These notification settings are applied globally to all the alert definitions that are specified for the selected storage system. You can override the global settings by specifying different settings for each alert definition.
Specify alert actions and notification settings for a specific alert definition	Resource page and resource details page For example, to specify the notification settings for an alert definition that applies to a block storage system, go to Storage > Block Storage Systems , right-click a storage system, and select Edit Alert Definitions . Click the envelope icon  next to the alert definition that you want to edit. The notification settings that you define override any global settings for the resource.

You can specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. [Learn more.](#)

Email Send emails to specific email addresses when an alert is detected on a monitored resource.

Netcool or OMNIbus EIF Probe Server

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP trap

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

How scripts are run

After an alert is triggered, scripts can be run from the server where IBM Spectrum Control is installed or from any server where a Storage Resource agent is deployed. For some alert conditions, scripts can be specified to run on the server where the change was detected.

You must include the script files that you want to run in the /scripts directory on the server where IBM Spectrum Control is installed. The default path of /scripts is determined by the operating system of the server:

- Windows: C:\Program Files\IBM\TPC\Data\scripts
- UNIX or Linux: /opt/IBM/TPC/Data/scripts or /usr/IBM/TPC/Data/scripts

Tip: When you run a script against a NAS filer, the script is run from the Storage Resource agent that is assigned to the file system where the condition was detected.

How script names are resolved

When you run scripts in response to alerts, the names of those scripts are resolved by using the following criteria:

- The operating system of the server where the product is installed.
- The script name that you specify on the **Script name** field when you define or edit an alert definition.

The operating system of the server or Storage Resource agent where you choose to run a script determines how that script that is run:

UNIX, Linux operating systems

A Storage Resource agent that is deployed on a UNIX or Linux operating system does not run scripts that have an extension. If the specified script name contains an extension, the agent ignores the extension and searches for a file of the same name (without an extension). For example, if you enter backup.vbs in the **Script name** field, a UNIX or Linux agent searches for a file with the name backup. If the script exists, the first line of the script is read and the appropriate interpreter is used to run the script. If a file cannot be found, no action is taken.

Windows operating systems

A Storage Resource agent that is deployed on a Windows operating system runs scripts that have an extension. The extension of a script file determines which interpreter is used to run the script. If you specify a script name that does not contain an extension, the agent searches for the file name in the following order (based on its extension): .com, .exe, .bat, .cmd, .vbs, .vbe, .js, .jse. If a file cannot be found, no action is taken.

Tip: You can run scripts with the same name across multiple operating systems. To run a script on both a UNIX or Linux and Windows operating systems, you must have two versions of that file in the /scripts directory. For example, if you want to run a provisioning script against Storage Resource agents on UNIX or Linux and Windows servers, you must have two versions of that file (provision.bat and provision) in /scripts.

Performance data

Use the IBM Spectrum Control to view the performance data of devices in your environment.

Performance metrics

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can view performance information or specify alerts for when the performance of a resource falls outside a specified range.

You can view performance metrics at the following locations:

Resource pages

To access a view of performance metrics for a resource, go to the following pages:

- Storage systems page, Storage system details page
- Volumes page
- Pools page
- Switches page, Switch details page
- Server details page > Related Resources
- Hypervisor details page > Related Resources

Cognos BI reports

To view detailed performance reports for multiple resources in the Cognos BI reporting tool, hover the mouse pointer over **Cognos** in the menu bar and select an option.

Related concepts:

“Monitoring the performance of resources” on page 217

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

Performance metrics for DS6000 and DS8000

Monitor the performance metrics that are collected for DS8000 and DS6000 storage systems.

Overview

Performance metrics are available for the following resources:

- “Volume metrics” on page 404
- “Disk metrics” on page 408
- “Port metrics” on page 410

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- I/O rate metrics
- Cache hit percentage metrics
- Average transfer size and HPF I/O metrics

Tip: Unless otherwise noted, you can view the volume metrics in Table 126, Table 127 on page 405, Table 128 on page 406, and Table 129 on page 408 for the following resources:

- Host connections
- Nodes
- Pools
- RAID arrays
- Storage systems
- Volumes

Table 126. Key volume metrics.

Metric	Description
Cache Holding Time ¹	The average number of seconds that I/O data for a storage system node is held in the cache. A short cache-holding time indicates adverse performance.
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Pool Activity Score ²	The activity level of pools, which is set to the following value: $[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$
Response Time (Read)	The average number of milliseconds to complete a read operation.

Table 126. Key volume metrics (continued).

Metric	Description
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ³	The average percentage of time that the volume is busy.
Write-Cache Delay Percentage	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from storage systems that are V7.3 or later with earlier versions.
Notes: <ol style="list-style-type: none"> 1. This metric is only available when you view the performance of nodes or storage systems. 2. This metric is only available when you view the performance of pools. 3. This metric is only available when you view the performance of volumes. 	

Table 127. I/O Rates.

Metric	Description
Average Transfer Rate (Cache-to-Disk) ¹	The average number of sectors or tracks per second that are transferred from the cache to the disks.
Average Transfer Rate (Disk-to-Cache) ¹	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
High Performance FICON® (Read) ²	The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.
High Performance FICON (Write) ²	The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.
High Performance FICON (Total) ²	The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.
Normal I/O Rate (Read)	The average number of nonsequential read operations per second.

Table 127. I/O Rates (continued).

Metric	Description
Normal I/O Rate (Write)	The average number of nonsequential write operations per second.
Normal I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
PPRC Transfer Rate	The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.
Record Mode Read I/O Rate ¹	The average number of I/O operations per second for record-mode read operations. For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.
Sequential I/O Rate (Read)	The average number of sequential read operations per second.
Sequential I/O Rate (Write)	The average number of sequential write operations per second.
Sequential I/O Rate (Total)	The average number of sequential I/O operations per second. This value includes both read and write operations.
Write-Cache Delay I/O Rate	The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.
Notes: 1. This metric is only available when you view the performance of volumes, RAID arrays, pools, nodes, and host connections. 2. This metric applies only to DS8000 storage systems.	

Table 128. Cache hit percentages.

Metric	Description
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 128. Cache hit percentages (continued).

Metric	Description
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Record Mode Read Cache Hit Percentage ¹	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Note: 1. This metric is only available when you view the performance of volumes, RAID arrays, pools, nodes, and host connections.	

Table 129. Average transfer size, HPF I/O, and miscellaneous metrics.

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
HPF I/O Percentage	The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.
Sequential I/O Percentage ¹	Percentage of all I/O operations that were completed for a RAID array over a specified time interval that were sequential operations.
Note: 1. This metric is only available when you view the performance of RAID arrays.	

Disk metrics

Disk performance metrics are divided into the following categories:

- Key disk metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in Table 130 and Table 131 on page 409 for the following resources:

- Nodes
- Pools
- RAID arrays
- Ranks
- Storage systems

Table 130. Key disk metrics.

Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Table 130. Key disk metrics (continued).

Metric	Description
Disk Utilization Percentage ¹	The average percentage of time that the disks that are associated with an array are busy. No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed for this property. This limitation applies only to DS8000 storage systems. If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
Note: 1. This metric is only available when you view the performance of RAID arrays.	

Table 131. Transfer size metrics.

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Average Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics

Port performance metrics are divided into the following categories:

- Key port metrics
- Port I/O rate metrics
- Port data rate metrics
- Port response time metrics
- Error rate metrics
- Remote mirror metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the port metrics in Table 132, Table 136 on page 412, Table 137 on page 415, and Table 138 on page 415 for ports and storage systems. The port metrics in Table 133 on page 411, Table 134 on page 411, Table 135 on page 412 are only available for ports.

Table 132. Key port metrics.

Metric	Description
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Port Utilization (Receive) ¹	The average percentage of time that the port is busy receiving data.
Port Utilization (Send) ¹	The average percentage of time that the port is busy sending data.
Port Utilization (Overall) ¹	The average percentage of time that the port is busy sending or receiving data.

Table 132. Key port metrics (continued).

Metric	Description
Bandwidth (Receive) ¹	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Send) ¹	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall) ¹	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Note: 1. This metric is only available when you view the performance of ports.	

Table 133. Port I/O rate metrics.

Metric	Description
FICON I/O Rate (Send)	The average number of send operations per second for Fibre Channel connection (FICON) usage.
FICON I/O Rate (Receive)	The average number of receive operations per second for FICON usage.
FICON I/O Rate (Total)	The average number of send and receive operations per second for FICON usage.
FCP I/O Rate (Send)	The average number of send operations per second for Fibre Channel Protocol (FCP) usage.
FCP I/O Rate (Receive)	The average number of receive operations per second for FCP usage.
FCP I/O Rate (Total)	The average number of send operations and receive operations per second for FCP usage.

Table 134. Port data rates.

Metric	Description
FICON Data Rate (Send)	The average number of MiB per second that is sent for FICON usage.
FICON Data Rate (Receive)	The average number of MiB per second that is received for FICON usage.
FICON Data Rate (Total)	The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.
FCP Data Rate (Send)	The average number of MiB per second that are sent for FCP usage.
FCP Data Rate (Receive)	The average number of MiB per second that are received for FCP usage.

Table 134. Port data rates (continued).

Metric	Description
FCP Data Rate (Total)	The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

Table 135. Port response times.

Metric	Description
FICON Response Time (Send)	The average number of milliseconds to complete a send operation for FICON usage.
FICON Response Time (Receive)	The average number of milliseconds to complete a receive operation for FICON usage.
FICON Response Time (Overall)	The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.
FCP Response Time (Send)	The average number of milliseconds to complete a send operation for FCP usage.
FCP Response Time (Receive)	The average number of milliseconds to complete a receive operation for FCP usage.
FCP Response Time (Overall)	The average number of milliseconds to complete a send or receive operation for FCP usage.

Table 136. Error rate metrics.

Tip: You can view the following metrics for ports and storage systems.

Metric	Description
Frame Errors (CRC Errors)	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.
Frame Errors (Error Frame)	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
Frame Errors (Invalid Relative Offset Rate)	The average number of times per second that frames are received with an invalid relative offset in the frame header.
Link Errors (Invalid Link Transmission)	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Link Errors (Invalid Transmission Words)	The average number of bit errors per second that are detected.

Table 136. Error rate metrics (continued).

Tip: You can view the following metrics for ports and storage systems.

Metric	Description
Link Errors (Link Failures)	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Link Errors (Primitive Sequence Protocol Error Rate)	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Link Errors (Sequence Timeouts)	The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.
Link Errors (Signal Loss)	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Link Errors (Sync Loss)	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Overload Errors (Extreme I/O Concurrency Percentage) ¹	The average percentage of I/O operations for which the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for ports on DS8000 cannot exceed 2000.
Overload Errors (Extreme I/O Concurrency Rate) ¹	The average number of times per second that the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for a port on DS8000 cannot exceed 2000.
Overload Errors (I/O Busy Percentage) ¹	The average percentage of I/O operations or exchanges for which the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of I/O operations exceeds an internal DS8000 threshold.

Table 136. Error rate metrics (continued).

Tip: You can view the following metrics for ports and storage systems.

Metric	Description
Overload Errors (I/O Busy Rate) ¹	The average number of times per second that the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of I/O operations or exchanges exceeds an internal DS8000 threshold.
Overload Errors (I/O Overrun Percentage) ¹	The average percentage of I/O operations that the port on DS8000 had to discard because the number of concurrent I/O operations or exchanges for the port exceeded 2000.
Overload Errors (I/O Overrun Rate) ¹	The average number of times per second that the port on DS8000 had to discard commands because the number of concurrent I/O operations or exchanges for the port exceeded 2000.
Port Protocol Errors (Credit Recovery Link Resets)	The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Port Protocol Errors (Duplicate Frames)	The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.
Port Protocol Errors (Link Reset Received)	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Port Protocol Errors (Link Reset Transmitted)	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Port Protocol Errors (Out of Order ACK)	The average number of times per second that an out-of-order acknowledge (ACK) frame is detected. An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.
Port Protocol Errors (Out of Order Data)	The average number of times per second that an out-of-order frame is detected.
Zero Receive Buffer Credit Percentage ¹	The amount of time, as a percentage, that the port on DS8000 had depleted its receive buffer credits. That is, the percentage of time that the port on DS8000 had no credit to provide to the sending port.
Zero Send Buffer Credit Percentage ¹	The amount of time, as a percentage, that the port on DS8000 had depleted its send buffer credits. That is, the percentage of time that the receiving port had no credit to provide to the port on DS8000.

Table 136. Error rate metrics (continued).

Tip: You can view the following metrics for ports and storage systems.

Metric	Description
Note:	
1. This metric is only available when you view the performance of ports.	

Table 137. Remote mirror metrics.

Metric	Description
PPRC Data Rate (Receive)	The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.
PPRC Data Rate (Send)	The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.
PPRC Data Rate (Total)	The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.
PPRC I/O Rate (Receive)	The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.
PPRC I/O Rate (Send)	The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.
PPRC I/O Rate (Total)	The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Receive)	The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Send)	The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Overall)	The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Table 138. Transfer size metrics.

Metric	Description
Average Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Average Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for storage systems that run IBM Spectrum Virtualize

Monitor the performance metrics that are collected for SAN Volume Controller, Storwize, FlashSystem V840, and FlashSystem V9000 storage systems.

Overview

Definitions are provided for the performance metrics that are collected for the following storage systems:

- FlashSystem V840
- FlashSystem V9000
- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V5000
- Storwize V7000
- Storwize V7000 Unified

The following terms are used in the performance metrics for these storage systems:

Stage To write data from a disk to the cache. The data is not prefetched data.

Destage

To write data from the cache to a disk.

Prestage

To write prefetched data from a disk to the cache.

The performance metrics are divided into the following categories:

- “Volume performance metrics”
- “Disk performance metrics” on page 425
- “Port performance metrics” on page 427
- “Node performance metrics” on page 433

Volume performance metrics

Volume performance metrics are divided into the following categories:

- Key metrics for volumes
- I/O rate metrics for volumes
- Cache hit percentage metrics for volumes
- Response time metrics for volumes
- Remote mirror metrics for volumes
- Volume cache metrics for volumes
- Volume copy cache metrics for volumes
- Miscellaneous metrics for volumes
- Legacy cache metrics for volumes

Tip: Unless otherwise noted, you can view the volume metrics in Table 139 on page 417, Table 140 on page 418, Table 141 on page 418, Table 142 on page 419, Table 143 on page 419, Table 144 on page 420, Table 145 on page 421, Table 146 on page 423, and Table 147 on page 424 for the following resources:

- Host connections

- I/O groups
- Nodes
- Pools
- Storage systems
- Volumes

Table 139. Key metrics for volumes.

Metric	Definition
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall Host Attributed Response Time Percentage	The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently. For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
Pool Activity Score ¹	The activity level of pools, which is set to the following value: $[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ²	The average percentage of time that the volume is busy.

Table 139. Key metrics for volumes (continued).

Metric	Definition
Write Cache Delay	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from storage systems that are V7.3 or later with earlier versions.
Notes: <ol style="list-style-type: none"> 1. This metric is also available when you view the performance of pools. 2. This metric is only available when you view the performance of volumes. 	

Table 140. I/O rate metrics for volumes.

Metric	Definition
Transfer Rate (Cache-to-Disk)*	The average number of sectors or tracks per second that are transferred from the cache to the disks.
Transfer Rate (Disk-to-Cache)*	The average number of sectors or tracks per second that are transferred per second from the disks to the cache.
Write-Cache Delay I/O Rate*	The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from storage systems that are V7.3 or later with earlier versions.
Note: *This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from storage systems that are V7.3 or later with earlier versions.	

Table 141. Cache hit percentages metrics for volumes.

Metric	Definition
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Table 141. Cache hit percentages metrics for volumes (continued).

Metric	Definition
Overall I/O Cache Hits (Write)	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 142. Response time metrics for volumes.

Metric	Definition
Peak Response Time (Read)	The worst response time measured for a read operation in the sample interval.
Peak Response Time (Write)	The worst response time measured for a write operation in the sample interval.

Table 143. Remote mirror metrics for volumes.

Metric	Definition
Global Mirror (Overlapping Write I/O Rate)	The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
Global Mirror (Overlapping Write Percentage)	The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
Global Mirror (Secondary Write Lag)	The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations. You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.
Global Mirror (Write I/O Rate)	The average number of write operations per second that are issued to the Global Mirror secondary site.

Volume cache metrics are only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.

Table 144. Volume cache metrics for volumes.

Metric	Definition
Cache Hits (Dirty Writes)	The percentage of all cache write hits that occur on data that is marked as modified in the volume cache. This value represents how effectively write operations are coalesced before the data is written to disk.
Cache Hits (Read)	The percentage of read operations that find data in the volume cache.
Cache Hits (Total)	The percentage of cache hits for read and write operations that are handled in the volume cache.
Cache Hits (Write)	The percentage of cache hits for write operations that are handled in the volume cache.
I/O Rate (Destage)	The average number of cache-to-disk transfer operations per second that are processed in the volume cache.
I/O Rate (Stage)	The average number of disk-to-cache transfer operations per second that are processed in the volume cache.
I/O Rate (Read)	The average number of read operations per second that are processed in the volume cache. This value includes operations that are started by hosts or by remote replication sources.
I/O Rate (Total)	The average number of read and write operations per second that are processed by the volume cache. This value includes requests sent by hosts and by remote replication sources.
I/O Rate (Write)	The average number of write operations per second that are processed by the volume cache. This value includes operations that are started by hosts or by remote replication sources.
Response Time (Destage)	The average number of milliseconds that it took to complete each destage operation in the volume cache. That is, the time that it took to do write operations from the volume cache to the disk.
Response Time (Stage)	The average number of milliseconds that it took to complete each stage operation in the volume cache. That is, the time that it took to do read operations from the disk to the volume cache.
Transfer Rates (Cache-to-Disk)	The average number of sectors that are transferred per second from the volume cache to the disks.

Table 144. Volume cache metrics for volumes (continued).

Metric	Definition
Transfer Rates (Disk-to-Cache)	The average number of sectors that are transferred per second from the disks to the volume cache.
Write Delay Percentage (Flush-through)	The percentage of write operations that are written to disk in flush-through mode in the volume cache.
Write Delay Percentage (Total Delay)	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions in the volume cache. The value is a percentage of all operations.
Write Delay Percentage (Write-through)	The percentage of write operations that are written to disk in write-through mode in the volume cache.
Write Delay Rate (Flush-through)	The average number of tracks per second that are written to disk in flush-through mode in the volume cache.
Write Delay Rate (Total Delay)	The average number of I/O operations per second that are delayed. The delay might occur because of space constraints in the write cache, or because of other conditions in the volume cache. The value is an average of all operations.
Write Delay Rate (Write-through)	The average number of sectors per second that are written to disk in write-through mode in the volume cache.

Volume copy cache metrics are only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.

Table 145. Volume copy cache metrics for volumes.

Metric	Definition
Cache Hits (Dirty Writes)	The percentage of all cache write hits that occur on data that is marked as modified in the volume copy cache. This value represents how effectively write operations are coalesced before the data is written to disk.
Cache Hits (Read-ahead)*	The percentage of all read cache hits that occur on pre-staged data.
Cache Hits (Read)	The percentage of read operations that find data in the volume copy cache.
Cache Hits (Total)	The percentage of cache hits for read and write operations that are handled in the volume copy cache.
Cache Hits (Write)	The percentage of cache hits for write operations that are handled in the volume copy cache.

Table 145. Volume copy cache metrics for volumes (continued).

Metric	Definition
I/O Rate (Destage)	The average number of cache-to-disk transfer operations per second that are processed in the volume copy cache.
I/O Rate (Prestage)	The average number of prefetch disk-to-cache transfer operations per second that are processed in the volume copy cache.
I/O Rate (Stage)	The average number of disk-to-cache transfer operations per second that are not prefetch operations and are processed in the volume copy cache.
I/O Rate (Read)	The average number of read operations per second that are processed in the volume copy cache. This value includes read operations that are associated with FlashCopy services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
I/O Rate (Total)	The average number of read and write operations per second that are processed by the volume copy cache. This value includes operations that are associated with FlashCopy services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
I/O Rate (Write)	The average number of write operations per second that are processed by the volume copy cache. This value includes read operations that are associated with FlashCopy services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
Response Time (Destage)	The average number of milliseconds that it took to complete each destage operation in the volume copy cache. That is, the time that it took to do write operations from the volume copy cache to the disk.
Response Time (Prestage)	The average number of milliseconds that it took to complete each prestage operation in the volume copy cache. That is, the time that it took to prefetch data from the disk into the volume copy cache.
Response Time (Stage)	The average number of milliseconds that it took to complete each stage operation in the volume copy cache. That is, the time that it took to do read operations from the disk to the volume copy cache.
Transfer Rates (Cache-to-Disk)	The average number of sectors that are transferred per second from the volume copy cache to the disks.

Table 145. Volume copy cache metrics for volumes (continued).

Metric	Definition
Transfer Rates (Disk-to-Cache)	The average number of sectors that are transferred per second from the disks to the volume copy cache.
Write Delay Percentage (Flush-through)	The percentage of write operations that are written to disk in flush-through mode in the volume copy cache.
Write Delay Percentage (Total Delay)	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions in the volume copy cache. The value is a percentage of all operations.
Write Delay Percentage (Write-through)	The percentage of write operations that are written to disk in write-through mode in the volume copy cache.
Write Delay Rate (Flush-through)	The average number of sectors per second that are written to disk in flush-through mode in the volume copy cache.
Write Delay Rate (Total Delay)	The average number of I/O operations per second that are delayed. The delay might occur because of space constraints in the write cache, or because of other conditions in the volume copy cache. The value is an average of all operations.
Write Delay Rate (Write-through)	The average number of sectors per second that are written to disk in write-through mode in the volume copy cache.
Note: *This metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.4 or later.	

Tip: Unless otherwise noted, you can view the volume metrics in Table 146 for the following resources:

- Nodes
- I/O groups
- Host connections
- Storage systems

Table 146. Miscellaneous metrics for volumes.

Metric	Definition
Cache to Host Transfer Response Time ¹	The average number of milliseconds that is taken to transfer a track from the cache to the host, including any queuing time that occurs because of throttling.
Non-Preferred Node Usage Percentage ²	The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.

Table 146. Miscellaneous metrics for volumes (continued).

Metric	Definition
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
Unaligned Write I/O Rate ³	The average number of write operations per second that are not aligned on a boundary between 4 KB physical blocks. Unaligned write operations can cause a significant decrease in efficiency of write operations to some types of back-end disks. You can ignore this metric for iSeries servers.
Notes: <ol style="list-style-type: none"> 1. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later. 2. This metric is only available when you view the performance of volumes, I/O groups, and host connections. 3. This metric is only available when you view the performance of pools, volumes, and host connections. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.4 or later. 	

Legacy cache metrics are only available for SAN Volume Controller and Storwize block storage systems whose firmware version is earlier than V7.3.

Table 147. Legacy cache metrics for volumes.

Metric	Definition
Write-Cache Percentages (Flush-through)	The percentage of tracks that are written to disk in flush-through mode. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache Percentages (Write-through)	The percentage of tracks that are written to disk in write-through mode. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache Percentages (Overflow)	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.
Write-Cache I/O Rate (Flush-through)	The average number of tracks per second that are written to disk in flush-through mode. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3.

Table 147. Legacy cache metrics for volumes (continued).

Metric	Definition
Write-Cache I/O Rate (Overflow)	The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache I/O Rate (Write-through)	The average number of tracks per second that are written to disk in write-through mode. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Dirty Write Percentage of Cache Hits	The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to storage systems that are running a version of IBM Spectrum Virtualize earlier than V7.3.

Disk performance metrics

Disk performance metrics are divided into the following categories:

- Key metrics for disks
- Response time metrics for disks
- Miscellaneous metrics for disks

Unless otherwise noted, you can view disk metrics for the following resources:

- Managed disks
- Pools
- Nodes
- I/O Groups
- Storage systems

Table 148. Key metrics for disks

Metric	Definition
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
I/O Rate (Read) ¹	The average number of read operations per second that are issued to the back-end storage resources.

Table 148. Key metrics for disks (continued)

Metric	Definition
I/O Rate (Write) ²	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total) ³	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
Notes: <ol style="list-style-type: none"> 1. The performance metrics for I/O Rate (Read) are available for pools, nodes, I/O groups, and storage systems. 2. The performance metrics for I/O Rate (Write) are available for pools, nodes, I/O groups, and storage systems. 3. The performance metrics for I/O Rate (Total) are available for pools, nodes, I/O groups, and storage systems. 	

Table 149. Response time metrics for disks

Metric	Definition
Queue Time (Read)	The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.
Queue Time (Write)	The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.
Queue Time (Overall)	The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Queue Time (Read)	The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Queue Time (Write)	The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Response Time (Read)	The longest time for a back-end storage resource to respond to a read operation.
Peak Back-end Response Time (Write)	The longest time for a back-end storage resource to respond to a write operation by a node.

Table 150. Miscellaneous metrics for disks

Metric	Definition
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.
Cache Destage (In-Flight I/O)*	The actual number of concurrent operations that are issued from the cache to the managed disk. The number of in-flight operations cannot exceed the target, but also varies over time, depending on the type of I/O operations and on the performance of the managed disk.
Cache Destage (Target I/O)*	The maximum number of concurrent operations that can be issued from the cache to the managed disk. This target changes over time, depending on changes in cache utilization and on changes in the performance of the managed disk.
Note: *This metric is only available when you view the performance of managed disks. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.	

Port performance metrics

Port performance metrics are divided into the following categories:

- Key metrics for ports
- I/O rate metrics for ports
- Data rate metrics for ports
- Response time metrics for ports
- Error rate metrics for ports
- Miscellaneous metrics for ports

Unless otherwise noted, you can view port metrics for the following resources:

- Ports
- Nodes
- I/O Groups
- Storage systems

Table 151. Key metrics for ports

Metric	Definition
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.

Table 151. Key metrics for ports (continued)

Metric	Definition
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
Data Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 152. I/O rate metrics for ports

Metric	Definition
Port-to-Disk I/O Rate (Receive)	The average number of exchanges per second that are received from back-end storage resources.
Port-to-Disk I/O Rate (Send)	The average number of exchanges per second that are sent to back-end storage resources.
Port-to-Disk I/O Rate (Total)	The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.
Port-to-Host I/O Rate (Receive)	The average number of exchanges per second that are received from host computers.
Port-to-Host I/O Rate (Send)	The average number of exchanges per second that are sent to host computers.
Port-to-Host I/O Rate (Total)	The average number of exchanges per second that are transmitted between host computers and the component.
Port-to-Local Node I/O Rate (Receive)	The average number of exchanges per second that are received from other nodes that are in the local cluster.
Port-to-Local Node I/O Rate (Send)	The average number of exchanges per second that are sent to other nodes that are in the local cluster.
Port-to-Local Node I/O Rate (Total)	The average number of exchanges per second that are transmitted between the component and other nodes that are in the local cluster.
Port-to-Remote Node I/O Rate (Receive)	The average number of exchanges per second that are received from nodes that are in the remote cluster.

Table 152. I/O rate metrics for ports (continued)

Metric	Definition
Port-to-Remote Node I/O Rate (Send)	The average number of exchanges per second that are sent to nodes that are in the remote cluster.
Port-to-Remote Node I/O Rate (Total)	The average number of exchanges per second that are transmitted between the component and nodes that are in the remote cluster.

Table 153. Data rate metrics for ports

Metric	Definition
Port-to-Disk Data Rate (Receive)	The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.
Port-to-Disk Data Rate (Send)	The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.
Port-to-Disk Data Rate (Total)	The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.
Port-to-Host Data Rate (Receive)	The average rate at which data is received from host computers. The rate is measured in MiB per second.
Port-to-Host Data Rate (Send)	The average rate at which data is sent to host computers. The rate is measured in MiB per second.
Port-to-Host Data Rate (Total)	The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.
Port-to-Local Node Data Rate (Receive)	The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Local Node Data Rate (Send)	The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Local Node Data Rate (Total)	The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Receive)	The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Send)	The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.

Table 153. Data rate metrics for ports (continued)

Metric	Definition
Port-to-Remote Node Data Rate (Total)	The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Table 154. Response time metrics for ports

Metric	Definition
Port-to-Local Node Response Time (Receive)	The average number of milliseconds to complete a receive operation from another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Local Node Response Time (Send)	The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Local Node Response Time (Overall)	The average number of milliseconds to complete a send or receive operation with another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Overall)	The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Receive)	The average number of milliseconds to complete a receive operation from a node that is in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Send)	The average number of milliseconds to complete a send operation to a node that is in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Overall)	The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

Metrics availability restrictions: The response time metrics are available for nodes, I/O groups, and storage system.

Table 155. Error rate metrics for ports

Metric	Definition
CRC Error Rate	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.
Link Errors (Invalid Link Transmission Rate)	The average number of bit errors per second that are detected.
Link Errors (Invalid Transmission Word Rate)	The average number of bit errors per second that are detected.
Link Errors (Link Failures)	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Link Errors (Primitive Sequence Protocol Error Rate)	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Link Errors (Signal Loss)	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Link Errors (Sync Loss)	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Port Congestion Index*	The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Table 155. Error rate metrics for ports (continued)

Metric	Definition
Port Protocol Errors (Port Send Delay Time)	The average number of milliseconds of delay that occur on the port for each send operation. The reason for these delays might be a lack of buffer credits. You cannot view zero buffer credit performance metrics for 16 Gbps Fibre Channel ports on storage systems that run IBM Spectrum Virtualize. Use the Port Send Delay Time metric if the Zero Buffer Credit Timer metric is not available.
Port Protocol Errors (Port Send Delay I/O Percentage)	The percentage of send operations where a delay occurred, relative to the total number of send operations that were measured for the port. Use this metric with the Port Send Delay Time metric to distinguish a few long delays from many short delays.
Port Protocol Errors (Zero Buffer Credit Percentage)	The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.
Port Protocol Errors (Zero Buffer Credit Timer)	The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset. If this metric is not available, use the Port Send Delay Time metric instead.
Note: *The performance metric for Port Congestion Index is only available for ports.	

Table 156. Miscellaneous metrics for ports

Metric	Definition
Port-to-Local Node Queue Time (Receive)	The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local cluster.
Port-to-Local Node Queue Time (Send)	The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster.

Table 156. Miscellaneous metrics for ports (continued)

Metric	Definition
Port-to-Local Node Queue Time (Overall)	The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local cluster.
Port-to-Remote Node Queue Time (Receive)	The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote cluster.
Port-to-Remote Node Queue Time (Send)	The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote cluster.
Port-to-Remote Node Queue Time (Overall)	The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote cluster.

Node performance metrics

Unless otherwise noted, you can view node metrics for the following resources:

- Nodes
- I/O Groups
- Storage systems

Table 157. Metrics for nodes.

Metric	Definition
CPU Utilization (System CPU)	The average percentage of time that the processors on nodes are busy doing system I/O tasks.
CPU Utilization (Compression CPU)	The average percentage of time that the processors used for data compression I/O tasks are busy.
System CPU Utilization (Core 1 to Core 4)	The approximate percentage of time that the processor core was busy with system I/O tasks.
Compression CPU Utilization (Core 1 to Core 4)	The approximate percentage of time that the processor core was busy with data compression tasks.
Node Utilization by Node	The average of the bandwidth percentages of those ports in the node that are actively used for host and MDisk send and receive operations. The average is weighted by port speed and adjusted according to the technology limitations of the node hardware.

Performance metrics for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R

Monitor the performance metrics that are collected for XIV systems, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R.

For XIV systems, IBM FlashSystem A9000, and IBM FlashSystem A9000R, performance metrics are available for the following resources:

- “Volume metrics”
- “Port metrics” on page 438

For IBM Spectrum Accelerate, performance metrics are available for the following resource:

- “Volume metrics”

Restriction: XIV systems, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R do not track performance statistics for volumes that were never used. Because there are no performance statistics, performance metrics are not shown for these volumes and their related components.

Volume metrics

Volume performance metrics are organized into the following categories:

- Key storage metrics
- Cache hit percentage metrics
- Response time metrics
- Miscellaneous metrics

Unless otherwise noted, you can view volume metrics for the following resources:

- Volumes
- Pools
- Modules
- Host connections
- Storage systems

Table 158. Key volume metrics

Metric	Description
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Table 158. Key volume metrics (continued)

Metric	Description
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Important: For XIV systems, all the key storage metrics are available for version 10.2.2 or later.	

Table 159. Cache hit percentage metrics

Metric	Description
Data Cache Hits (Read) ²	The percentage of all read data that is read from the cache.
Data Cache Hits (Write) ²	The percentage of all write data that is written to cache slots that are marked as modified.
Data Cache Hits (Overall) ²	The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.
Overall I/O Cache Hits (Read) ¹	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) ¹	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Table 159. Cache hit percentage metrics (continued)

Metric	Description
Overall I/O Cache Hits (Total) ¹	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
SSD Read Cache Hits (I/O) ²	The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric. This metric is only available for XIV systems that use solid-state drives as drives for caching.
SSD Read Cache Hits (Data) ²	The percentage of all read data that was read from cache memory on a solid-state drive. This metric is only available for XIV systems that use solid-state drives as drives for caching.
Notes: 1. For XIV systems, this metric is available in version 10.2.2 or later. 2. For XIV systems, this metric is available in version 10.2.4 or later.	

Table 160. Response time metrics

Metric	Description
Cache Hit Response Time (Read)	The average number of milliseconds to complete a read-cache hit operation.
Cache Hit Response Time (Write)	The average number of milliseconds to complete a write-cache hit operation.
Cache Hit Response Time (Overall)	The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.
Cache Miss Response Time (Read)	The average number of milliseconds to complete a read-cache miss operation.
Cache Miss Response Time (Write)	The average number of milliseconds to complete a write-cache miss operation.
Cache Miss Response Time (Overall)	The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.
Response Time by Transfer Size (Small)	The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Table 160. Response time metrics (continued)

Metric	Description
Response Time by Transfer Size (Medium)	The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
Response Time by Transfer Size (Large)	The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
Response Time by Transfer Size (Very Large)	The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.
SSD Read Cache Hit Response Time	The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric. This metric is only available for XIV systems that use solid-state drives as drives for caching.
Note: For XIV systems, all the response time metrics are available for version 10.2.2 or later.	

Table 161. Miscellaneous metrics

Metric	Description
Average Transfer Size (Read) ¹	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write) ¹	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall) ¹	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
I/O Transfer Size (Small) ²	The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
I/O Transfer Size (Medium) ²	The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
I/O Transfer Size (Large) ²	The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
I/O Transfer Size (Very Large) ²	The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Table 161. Miscellaneous metrics (continued)

Metric	Description
Data Transfer Size (Small) ²	The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
Data Transfer Size (Medium) ²	The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
Data Transfer Size (Large) ²	The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
Data Transfer Size (Very Large) ²	The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.
Volume Utilization ^{1, 3}	The average percentage of time that the volume is busy.
Pool Activity Score ^{1, 3}	<p>The activity level of pools, which is set to the following value: $[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$</p> <p>The value of this metric is the same for all pools in the storage system.</p>
Notes: <ol style="list-style-type: none"> For XIV systems, this metric is available in version 10.2.2 or later. For XIV systems, this metric is available in version 10.2.4 or later. This metric is only available when you view the performance of pools. 	

Port metrics

The following performance metrics for ports are available for XIV systems that are version 10.2.4 or later:

- I/O rate metrics
- Data rate metrics
- Response time metrics
- Miscellaneous port metrics

Restriction: Port metrics are not available for IBM Spectrum Accelerate.

Table 162. I/O rate metrics

Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.

Table 162. I/O rate metrics (continued)

Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.

Table 163. Data rate metrics

Metric	Description
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 164. Response time metrics

Metric	Description
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.

Table 165. Miscellaneous port metrics

Metric	Description
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall)	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Performance metrics for ESS

Monitor the performance metrics that are collected for IBM TotalStorage Enterprise Storage Server® systems.

Performance metrics are available for the following resources:

- “Volume metrics” on page 440
- “Disk metrics” on page 443

- “Port metrics” on page 445

When you view metrics for ESS systems, you must take into account the following differences between IBM Spectrum Control reports and the native reports of those systems:

- IBM Spectrum Control reports display port performance metrics as send and receive metrics, such as Data Rate (Send) and Data Rate (Receive).
- Reports that are included with storage systems (for example, reports based on data that is collected by the DS CLI) display port performance metrics as read and write metrics, such as Bytread and Bytewrite.

Volume metrics

Volume performance metrics are organized into the following categories:

- Key volume metrics
- I/O rate metrics
- Cache hit percentage metrics
- Miscellaneous metrics

Unless otherwise noted, you can view volume metrics for the following resources:

- Host connections
- Nodes
- Pools
- RAID arrays
- Storage systems
- Volumes

Table 166. Key volume metrics.

Metric	Description
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Cache Holding Time*	The average number of seconds that I/O data for a storage system node is held in the cache. A short cache-holding time indicates adverse performance.
Write-Cache Delay Percentage	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Table 166. Key volume metrics (continued).

Metric	Description
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Note: * This metric is only available when you view the performance of nodes and storage systems.	

Table 167. I/O rate metrics.

Metric	Description
Average Transfer Rate (Cache-to-Disk)	The average number of sectors or tracks per second that are transferred from the cache to the disks.
Average Transfer Rate (Disk-to-Cache)	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
Normal I/O Rate (Read)	The average number of nonsequential read operations per second.
Normal I/O Rate (Write)	The average number of nonsequential write operations per second.
Normal I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
PPRC Transfer Rate	The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.
Record Mode Read I/O Rate	The average number of I/O operations per second for record-mode read operations. For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.
Sequential I/O Rate (Read)	The average number of sequential read operations per second.
Sequential I/O Rate (Write)	The average number of sequential write operations per second.
Sequential I/O Rate (Total)	The average number of sequential I/O operations per second. This value includes both read and write operations.

Table 167. I/O rate metrics (continued).

Metric	Description
Write-Cache Delay I/O Rate	The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.
Write-Cache Delay Percentage	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

Table 168. Cache hit percentage metrics.

Metric	Description
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Table 168. Cache hit percentage metrics (continued).

Metric	Description
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Record Mode Read Cache Hit Percentage	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 169. Miscellaneous metrics.

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
Volume Utilization*	The average percentage of time that the volume is busy.
Note: * This metric is only available when you view the performance of volumes.	

Disk metrics

Disk performance metrics are organized into the following categories:

- Key disk metrics

- Transfer size metrics
- Utilization metrics

Unless otherwise noted, you can view disk metrics for the following resources:

- Nodes
- Pools
- RAID arrays
- Rank
- Storage systems

Table 170. Key disk metrics.

Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Table 171. Transfer size metrics.

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.

Table 171. Transfer size metrics (continued).

Metric	Description
Average Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Table 172. Utilization metrics.

Metric	Internal resources	Description
Disk Utilization	RAID array	The average percentage of time that the disks that are associated with an array are busy. No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed for this property. This limitation applies only to DS8000 storage systems. Tip: IBM Spectrum Control does not calculate a value for this column if there are multiple ranks in the extent pool where thin-provisioned volumes are allocated. This column displays value of N/A for the reports in which it is displayed. However, if there is only a single rank in the extent pool, IBM Spectrum Control does calculate the value for this column regardless of the thin-provisioned volumes.
Sequential I/O Percentage	RAID array	Percentage of all I/O operations that were completed for a RAID array over a specified time interval that were sequential operations.

Port metrics

Port performance metrics are organized into the following categories:

- Key port metrics
- I/O rate metrics
- Data rate metrics
- Response time metrics
- Remote mirror metrics
- Transfer size metrics

Table 173. Key port metrics.

Metric	Resources	Description
Bandwidth (Send)	Port	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Receive)	Port	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall)	Port	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Data Rate (Send)	Port, storage system	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	Port, storage system	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	Port, storage system	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Send)	Port, storage system	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	Port, storage system	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	Port, storage system	The average number of send operations and receive operations per second.
Port Utilization (Send)	Port	The average percentage of time that the port is busy sending data.
Port Utilization (Receive)	Port	The average percentage of time that the port is busy receiving data.

Table 173. Key port metrics (continued).

Metric	Resources	Description
Port Utilization (Overall)	Port	The average percentage of time that the port is busy sending or receiving data.
Response Time (Send)	Port, storage system	The average number of milliseconds to complete a send operation.
Response Time (Receive)	Port, storage system	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	Port, storage system	The average number of milliseconds to complete a send or receive operation.

Table 174. I/O rate metrics.

Metric	Resources	Description
FICON I/O Rate (Send)	Port	The average number of send operations per second for Fibre Channel connection (FICON) usage.
FICON I/O Rate (Receive)	Port	The average number of receive operations per second for FICON usage.
FICON I/O Rate (Total)	Port	The average number of send and receive operations per second for FICON usage.
FCP I/O Rate (Send)	Port	The average number of send operations per second for Fibre Channel Protocol (FCP) usage.
FCP I/O Rate (Receive)	Port	The average number of receive operations per second for FCP usage.
FCP I/O Rate (Total)	Port	The average number of send operations and receive operations per second for FCP usage.

Table 175. Data rate metrics.

Metric	Resources	Description
FICON Data Rate (Send)	Port	The average number of MiB per second that is sent for FICON usage.
FICON Data Rate (Receive)	Port	The average number of MiB per second that is received for FICON usage.
FICON I/O Rate (Total)	Port	The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.

Table 175. Data rate metrics (continued).

Metric	Resources	Description
FCP Data Rate (Send)	Port	The average number of MiB per second that are sent for FCP usage.
FCP Data Rate (Receive)	Port	The average number of MiB per second that are received for FCP usage.
FCP Data Rate (Total)	Port	The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

Table 176. Response time metrics.

Metric	Resources	Description
FICON Response Time (Send)	Port	The average number of milliseconds to complete a send operation for FICON usage.
FICON Response Time (Receive)	Port	The average number of milliseconds to complete a receive operation for FICON usage.
FICON Response Time (Overall)	Port	The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.
FCP Response Time (Send)	Port	The average number of milliseconds to complete a send operation for FICON usage.
FCP Response Time (Receive)	Port	The average number of milliseconds to complete a receive operation for FICON usage.
FCP Response Time (Overall)	Port	The average number of send operations and receive operations per second for FCP usage.

Table 177. Remote mirror metrics.

Metric	Resources	Description
PPRC Data Rate (Send)	Port, storage system	The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.

Table 177. Remote mirror metrics (continued).

Metric	Resources	Description
PPRC Data Rate (Receive)	Port, storage system	The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.
PPRC Data Rate (Total)	Port, storage system	The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.
PPRC I/O Rate (Send)	Port, storage system	The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.
PPRC I/O Rate (Receive)	Port, storage system	The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.
PPRC I/O Rate (Total)	Port, storage system	The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Send)	Port, storage system	The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Receive)	Port, storage system	The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Overall)	Port, storage system	The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Table 178. Transfer size metrics.

Metric	Resources	Description
Average Transfer Size (Send)	Port, storage system	The average number of KiB that are transferred per send operation.
Average Transfer Size (Receive)	Port, storage system	The average number of KiB that are transferred per receive operation.

Table 178. Transfer size metrics (continued).

Metric	Resources	Description
Average Transfer Size (Overall)	Port, storage system	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for IBM Spectrum Scale

Monitor the performance metrics that are collected for IBM Spectrum Scale storage systems.

Overview

The performance metrics are divided into the following categories:

- “Node performance metrics”
- “File system performance metrics” on page 451

Node performance metrics

You can view the following metrics for each GPFS cluster node.

Table 179. Metrics for nodes.

Metric	Definition
CPU Utilization (User)	The average percentage of time that the processors on a node are busy doing user tasks.
CPU Utilization (System)	The average percentage of time that the processors on a node are busy doing system tasks.
CPU Utilization (Total)	The average percentage of time that the processors on a node are busy doing user tasks and system tasks.
Memory Used (Cache and Buffer)	The average percentage of memory on a node that is used for cache and buffer memory.
Memory Used	The average percentage of memory that is used on a node. This value does not include the memory that is used for cache and buffer memory.
Memory Used (Total)	The average percentage of memory that is used on a node. This value includes the memory that is used for cache and buffer memory.
I/O Rate (Read)	The average number of read operations per second.
I/O Rate (Write)	The average number of write operations per second.
I/O Rate (Total)	The average number of read operations and write operations per second.

File system performance metrics

You can view file system metrics for the following resources:

- File systems
- IBM Spectrum Scale storage systems

Tip: For a storage system, the metrics contain summary values for all the file systems on the storage system.

Table 180. Metrics for file systems.

Metric	Definition
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Maximum Data Rate (Read)	The maximum number of MiB per second that are transferred for read operations.
Maximum Data Rate (Write)	The maximum number of MiB per second that are transferred for write operations.
Maximum Data Rate (Total)	The maximum number of MiB per second that are transferred for read operations and write operations.
I/O Rate (Read)	The average number of read operations per second.
I/O Rate (Write)	The average number of write operations per second.
I/O Rate (Total)	The average number of read operations and write operations per second.
Maximum I/O Rate (Read)	The maximum number of read operations per second.
Maximum I/O Rate (Write)	The maximum number of write operations per second.
Maximum I/O Rate (Total)	The maximum number of read operations and write operations per second.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation from the file system.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation from the file system.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read operation or a write operation from the file system.

Performance metrics for IBM FlashSystem 840 and IBM FlashSystem 900

To review trends in performance for IBM FlashSystem 840 and IBM FlashSystem 900 storage systems, you add performance metrics to performance charts. Use the charts to monitor the performance of the storage systems.

Overview

The performance metrics are described in the following sections:

- “Volume metrics”
- “Disk metrics” on page 453
- “Port metrics” on page 453

Restrictions:

- IBM FlashSystem 840 and IBM FlashSystem 900 storage systems do not track performance for individual volumes. The volume metrics that are available are provided to allow comparisons across different storage systems, but represent performance that is measured for ports rather than for volumes.
- IBM Spectrum Control supports only Fibre Channel adapters and ports that are installed on storage systems. For IBM FlashSystem 840 and IBM FlashSystem 900 storage systems that have other types of ports, performance data is only available for the drives.
- Response time metrics are not available for all IBM FlashSystem 840 and IBM FlashSystem 900 storage systems, depending on the firmware release of the storage system. For details of the firmware that is supported by IBM Spectrum Control, check the storage information in the [IBM Spectrum Control interoperability matrix at `http://www.ibm.com/support/docview.wss?uid=swg21386446`](http://www.ibm.com/support/docview.wss?uid=swg21386446).

Volume metrics

IBM FlashSystem 840 and IBM FlashSystem 900 storage systems do not track the performance of individual volumes. Therefore, the performance metrics that you can view represent the sum of the measured performance of all the ports in the storage system.

Table 181. Metrics for volumes

Metric	Description
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Table 181. Metrics for volumes (continued)

Metric	Description
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Total)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Disk metrics

You can view the following performance metrics for drives or flash modules on IBM FlashSystem 840 and IBM FlashSystem 900 storage systems.

Table 182. Metrics for disks

Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Flash Health Percentage	The overall flash module health. The health percentage is calculated based on the number of unusable blocks on the flash module.

Port metrics

You can view the following port performance metrics for Fibre Channel ports on IBM FlashSystem 840 and IBM FlashSystem 900 storage systems. No performance data is available for InfiniBand, iSCSI, or FCoE ports.

Table 183. Metrics for ports

Metric	Description
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Table 183. Metrics for ports (continued)

Metric	Description
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Total)	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for EMC VMAX and VNX storage systems

Monitor the performance metrics that are collected for EMC VMAX and VNX storage systems.

Overview

The performance metrics are described in the following sections:

- “VMAX storage systems” on page 455
- “VNX storage systems” on page 459

Restriction: For EMC, performance metrics are only available for block storage systems.

VMAX storage systems

The following performance metrics are available for VMAX resources:

- “Volume metrics”
- “Disk metrics” on page 457
- “Port metrics” on page 459

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- Cache hit metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in Table 184, Table 185 on page 456, and Table 186 on page 457 for the following resources:

- Volumes
- Directors
- Storage systems

Table 184. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read) ¹	The average number of milliseconds to complete a read operation.
Response Time (Write) ¹	The average number of milliseconds to complete a write operation.

Table 184. Key volume metrics (continued)

Metric	Description
Response Time (Overall) ¹	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ²	The average percentage of time that the volume is busy.
Normal I/O Rate (Read)	The average number of nonsequential read operations per second.
Normal I/O Rate (Write)	The average number of nonsequential write operations per second.
Normal I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
Sequential I/O Rate (Read)	The average number of sequential read operations per second.
Sequential I/O Rate (Write)	The average number of sequential write operations per second.
Sequential I/O Rate (Total)	The average number of sequential I/O operations per second. This value includes both read and write operations.
Note: 1. This metric is only available when you view the performance of volumes or storage systems. 2. This metric is only available when you view the performance of volumes.	

Table 185. Cache hit metrics

Metric	Description
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 185. Cache hit metrics (continued)

Metric	Description
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 186. Transfer size metrics

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Disk metrics

Disk performance metrics for VMAX block storage systems are divided into the following categories:

- Key disk metrics
- Transfer size metrics

Table 187. Key disk metrics

Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
Disk Utilization Percentage	The approximate utilization percentage of a particular array over a particular time interval, i.e. the average percent of time that the disk associated with the array was busy.

Table 188. Transfer size metrics

Metric	Description
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics

You can view the metrics below for ports and storage systems.

Table 189. Port metrics

Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

VNX storage systems

The following performance metrics are available for VNX resources:

- Volume metrics
- Disk metrics
- Port metrics

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in Table 190 and Table 191 on page 460 for the following resources:

- Volumes
- Nodes
- Storage systems

Table 190. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.

Table 190. Key volume metrics (continued)

Metric	Description
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.

Table 191. Transfer size metrics

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Disk metrics

Disk performance metrics for VNX storage systems are divided into the following categories:

- Key disk metrics
- Transfer size metrics

Table 192. Key disk metrics

Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Table 193. Transfer size metrics

Metric	Description
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics

You can view the metrics below for ports and storage systems.

Table 194. Port metrics

Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for other storage systems

Monitor the performance metrics that are collected for SMI-S BSP systems such as DS5000, DS4000, and DS3000, EMC, Hitachi, NetApp and other storage systems.

Overview

Performance metrics are available for the following resources:

- “Volume metrics”
- “Port metrics” on page 463

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- Cache hit percentage metrics
- Average transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in Table 195 on page 462, Table 196 on page 462, and Table 197 on page 463 for the following resources:

- Volumes
- Nodes
- Storage systems

Table 195. Key volume metrics.

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ¹	The average percentage of time that the volume is busy.
Note: 1. This metric is only available when you view the performance of volumes.	

Table 196. Cache hit percentages.

Metric	Description
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Table 196. Cache hit percentages (continued).

Metric	Description
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 197. Average transfer size metrics.

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Port metrics

Port performance metrics are divided into the following categories:

- Key port metrics
- Transfer size metrics

Tip: You can view the port metrics in Table 198 and Table 199 on page 464 for ports and storage systems.

Table 198. Key port metrics.

Metric	Description
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.

Table 198. Key port metrics (continued).

Metric	Description
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 199. Transfer size metrics.

Metric	Description
Average Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Average Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for switches

Monitor the performance metrics that are collected for physical switches, switch ports, and inter-switch connections.

Performance metrics for switches are divided into the following categories:

- Key port metrics
- I/O rate metrics
- Peak data rate metrics
- Frame error rate metrics
- Port protocol error rate metrics
- Link error rate metrics
- Miscellaneous metrics

Table 200. Key port metrics.

Metric	Description
Bandwidth Percentage (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth Percentage (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth Percentage (Overall)	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Table 200. Key port metrics (continued).

Metric	Description
Data Rate (Send)	The average rate at which data is sent by the port. A send operation is a read operation that is processed, or a write operation that is initiated by the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. A receive operation is a write operation that is processed, or a read operation that is initiated by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Total Port Error Rate	The average number of times per second that an error was detected on the port. This rate is a summation of all the other error rates for the port.

Table 201. I/O rates.

Metric	Description
Port Frame Rate (Send)	The average number of frames per second that are sent through the port.
Port Frame Rate (Receive)	The average number of frames per second that are received by the port.
Port Frame Rate (Total)	The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Table 202. Peak data rates.

Metric	Description
Peak Data Rate (Send) ¹	The highest rate at which data is sent by the port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.
Peak Data Rate (Receive) ¹	The highest rate at which data is received by the port. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.
Notes: 1. This metric is only available for ports on Brocade switches. You must use a CIM agent to collect data for this metric. In the Storage Management Initiative Specification (SMI-S), peak data rates for a port are reported in the CIM_FcPortRateStatistics class. For SMI-S Version 1.1, this is an optional class, meaning that it might not be supported by a given vendor. This can result in IBM Spectrum Control displaying zeros for the peak rates, even when there is non-zero traffic flowing through the port. When a vendor does support this class, the peak data rate represents the peak value of the associated metric counter (as reported by the CIM agent used to access the device) over a recent, vendor-defined measurement window (for example, 20 milliseconds) of the IBM Spectrum Control measurement window (for example, 15 minutes).	

Table 203. Frame error rates.

Metric	Description
Bad EOF CRC Error Rate ¹	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
CRC Error Rate	The percentage of nonsequential write operations that are handled in the cache.
Discarded Class 3 Frame Rate	The average number of class 3 frames per second that are discarded.
Error Frame Rate ¹	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
F-BSY Frame Rate ²	The average number of F-BSY frames per second that are generated. An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.
F-RJT Frame Rate ²	The average number of F-RJT frames per second that are generated. An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.
Long Frame Rate	The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.
Short Frame Rate ²	The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.
Notes: 1. This metric is only available for ports on Brocade switches. 2. You must use an SNMP agent to collect data for this metric.	

Table 204. Port protocol error rates.

Metric	Description
Class 3 Receive Timeout Frame Rate ¹	The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Class 3 Send Timeout Frame Rate ¹	The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Credit Recovery Link Reset Rate	The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Discarded Frame Rate ¹	The average number of frames per second that are discarded because host buffers are unavailable for the port.
Link Reset Received Rate	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Link Reset Transmitted Rate	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Port Congestion Index	The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Table 204. Port protocol error rates (continued).

Metric	Description
Zero Buffer Credit Percentage	The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.
Zero Buffer Credit Rate	The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Notes: <ol style="list-style-type: none"> 1. This metric is available only for ports on Brocade switches. 2. You must use an SNMP agent to collect data for this metric. 	

Table 205. Link error rates.

Metric	Description
Encoding Disparity	The average number of disparity errors per second that are received.
Invalid Link Transmission Rate	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Invalid Word Transmission Rate	The average number of bit errors per second that are detected.
Link Failure Rate	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Loss of Signal Rate	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

Table 205. Link error rates (continued).

Metric	Description
Loss of Sync Rate	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Primitive Sequence Protocol Error Rate	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Table 206. Miscellaneous port metrics.

Metric	Description
Link Quality Percentage	The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port.
Port Frame Size (Overall)	The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.
Port Frame Size (Receive)	The average size of a frame, in KiB, that is received by the port.
Port Frame Size (Send)	The average size of a frame, in KiB, that is sent through the port.

Capacity data

Use the IBM Spectrum Control GUI to view the capacity data that is collected about the storage systems in your environment.

Capacity metrics for block storage systems

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

Alphabetical lists of the capacity and space usage metrics that you can add to charts are provided in the following sections:

- “Storage system capacity metrics” on page 470
- “Pool capacity metrics” on page 473
- “Volume capacity metrics” on page 480

Tip: IBM Spectrum Control displays capacity values in base 2 (GiB), while the XIV management GUI displays capacity values in base 10 (GB, TB). Even though

different units of measurement are used, the storage values are equivalent. For more information about units of measurement, see “Units of measurement for storage data” on page 69.

Storage system capacity metrics

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for storage systems:

Allocated Space (GiB)

The amount of space that is allocated to the regular and thin-provisioned volumes in the pools. If the pool is a parent pool, the amount of space that is allocated to the volumes in the child pools is also calculated.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the **Total Volume Capacity (GiB)** column. If a pool doesn't have thin-provisioned volumes, the value for allocated space is the same as the value for total volume capacity.

Allocated space is the same as used space on all storage systems with the following exceptions:

- SAN Volume Controller storage systems that are thin provisioned
- Storwize family storage systems that are configured with block storage and that are thin provisioned

Available for: All storage systems.

Assigned Volume Space (GiB)

The total volume space in the storage system that is mapped or assigned to host systems, including child pool capacity.

Available for: All storage systems.

Available Pool Space (GiB)

The total amount of the space in the pools that is not allocated to the volumes in the pools. To calculate available space, the following formula is used:

$(\text{pool capacity} - \text{allocated space})$

For XIV systems, pool capacity is the physical capacity of the pools and does not include the virtual capacity of the pools. For other storage systems, pool capacity might also include the overhead space for the pools that are not formatted.

Available for: All storage systems.

Compression Savings (%)

The estimated percentage of the capacity of the resource that is saved by using data compression.

The following formula is used to determine the compression savings percentage:

$(\text{original size} - \text{compressed size}) \div (\text{original size}) \times 100$

For example, the compression savings percentage for a storage system is 75%. Therefore, compression reduced the size of the data that is written by 75%. The compressed size is 25% of the original size, which is the equivalent of a compression factor of 4 to 1.

Inline compression is a software feature that is supported by FlashSystem A9000, FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and storage systems that run IBM Spectrum Virtualize.

Data Deduplication Savings (%)

The estimated percentage of the capacity of the resource that is saved by data deduplication.

The following formula is used to determine the deduplication savings percentage:

$$(\text{original size} - \text{deduplicated size}) \div (\text{original size}) \times 100$$

For example, the data deduplication savings percentage for a storage system is 75%. Therefore, data deduplication reduced the size of the data that is written by 75%. The deduplicated size is 25% of the original size, which is the equivalent of a data deduplication factor of 4 to 1.

Available for: FlashSystem A9000 and FlashSystem A9000R only.

Data Reduction (%)

The estimated percentage of the capacity of the resource that is saved by using data deduplication and data compression.


The following formula is used to determine the data reduction savings percentage:

$$(\text{original size} - \text{size after data reduction}) \div (\text{original size}) \times 100$$

For example, the data reduction savings percentage for a storage system is 75%. Therefore, data reduction reduced the size of the data that is written by 75%. The reduced size is 25% of the original size, which is the equivalent of a data reduction factor of 4 to 1.

Available for: FlashSystem A9000, FlashSystem A9000R, XIV storage systems, and storage systems that run IBM Spectrum Virtualize. For storage systems that do not support data deduplication, the data reduction savings percentage is equivalent to the compression savings percentage.

Physical Allocation (%)



The percentage of physical capacity in the pools that is allocated to the regular volumes, the thin-provisioned volumes, and the volumes in child pools. Check the value for physical allocation to see:

- Whether the physical capacity of the pools is fully allocated. That is, the value for physical allocation is 100%.
- Whether you have sufficient capacity to provision new volumes with storage
- Whether you have sufficient capacity to allocate to the compressed and thin-provisioned volumes in the pools

Available for: All storage systems.

Pool Capacity (GiB)

The total amount of storage space in the pools, which might include overhead space if the disks for the pools aren't formatted. For XIV systems and IBM Spectrum Accelerate, capacity represents the physical or 'hard' capacity of the pool, not the virtual 'soft' capacity. Pools that are allocated from other pools are not included in the total pool space.

Available for: All storage systems.

Pool Shortfall (%)

The percentage of space that is over committed to the pools with thin-provisioned volumes. For example, you commit 100 GiB of space to a thin-provisioned volume in a pool with a capacity of 50 GiB. As the space is allocated to the thin-provisioned volume in increments of 10 GiB, the space available for allocation decreases and the shortfall in capacity becomes more acute.

To calculate the shortfall, the following formula is used:

$$[(\text{unallocatable space} \div \text{committed but unallocated space}) \times 100]$$

A pool shortfall occurs when you commit more space to the volumes in the pools than is physically available to the pools. If the physical space available to the pools is less than the committed virtual space, then the pools do not have enough space to fulfill the commitment to the virtual space.

For example, the physical capacity of the pools is 70 GiB, but 150 GiB of virtual space was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, 100 GiB is committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% $[(80 \text{ GiB} \div 100 \text{ GiB}) \times 100]$.

Available for: DS8000, FlashSystem storage systems, EMC VMAX, VNX, and VNXe storage systems, Hitachi Virtual Storage Platform, SAN Volume Controller, Storwize family storage systems that are configured with block storage, XIV systems, and IBM Spectrum Accelerate storage systems.

Total Volume Capacity (GiB)

The total amount of storage space that can be made available to the regular and thin-provisioned volumes in the pools. If the pool is a parent pool, it also includes the storage space that can be made available to the volumes in the child pools.

Available for: All storage systems.

Unallocatable Volume Space (GiB)

The amount of space that cannot be allocated to volumes because the physical capacity of the pools cannot meet the demands for virtual space. The following formula is used to calculate this value:

$$[\text{Total Volume Capacity} - \text{Pool Capacity}]$$

Available for: All storage systems.

Unallocated Volume Space (GiB)

The total amount of remaining space that can be allocated to the volumes in the pools. The following formula is used to calculate this value:

[Total Volume Capacity - Allocated Space]

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pools. For Hitachi VSP non-thin provisioned pool space, the unallocated volume space is always zero.

Available for: All storage systems.

Unassigned Volume Space (GiB)

The total amount of space in the volumes that are not assigned to hosts.

Available for: All storage systems.

Virtual Allocation (%)

The percentage of the physical capacity that is committed to the virtual capacity of the volumes in the pools. If the value exceeds 100%, the physical capacity doesn't meet the demands for virtual capacity.

To calculate virtual allocation, the following formula is used:

$$[(\text{total volume capacity} \div \text{pool capacity}) \times 100]$$

For example, if the virtual allocation percentage is 200% for a storage pool with a physical capacity of 15 GiB, then the virtual capacity that is committed to the volumes in the pools is 30 GiB. Twice as much space is committed to the pools than is physically available to the pools. If the virtual allocation percentage is 100% and the physical capacity is 15 GiB, then the virtual capacity that is committed to the pools is 15 GiB. The total physical capacity that is available to the pools is allocated to the volumes in the pools.

An allocation percentage that is higher than 100% is considered to be aggressive because insufficient physical capacity is available to the pools to satisfy the allocation of the committed space to the compressed and thin-provisioned volumes in the pools. In such cases, you can check the **Pool Shortfall (%)** value to determine how critical the shortage of space is for the storage system pools.

Available for: All storage systems.

Pool capacity metrics

If sufficient data is collected about the pools in your data center, you can view charts that compare the capacity of the pools with the space that is allocated to the pools and the space that is still available in the pools. In the **Zero Capacity** column on the Pools page, you can see the date, based on the space usage trends for the pool, when the pool runs out of available space.

Tip: To order the pools in the table by the amount of space available to the pools, click **Filter by column**, and then click **Zero Capacity**.

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for pools:

Allocated Space (GiB)

The amount of space that is allocated to the regular and thin-provisioned volumes in the pool. If the pool is a parent pool, the amount of space that is allocated to the volumes in the child pools is also included.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Total Volume Capacity column. If a pool does not contain thin-provisioned volumes, this value is the same as Total Volume Capacity.

Allocated space is the same as used space on all storage systems, except for storage systems that run IBM Spectrum Virtualize. These storage systems might have more allocated space than used space if the storage administrator preallocated some space for thin-provisioned volumes when the volumes were created.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Assigned Volume Space (GiB)

The total amount of space in the volumes that is assigned to hosts. For a thin-provisioning pool, this value includes the virtual capacity of thin-provisioned volumes, which might exceed the total space in the pool. For Hitachi VSP non-thin provisioning pool space, this value is the sum of assigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Available for: All storage systems.

Available Pool Space (GiB)

The amount of space that is available to create new volumes in the pool. If the pool is a parent pool, the amount of space that is allocated to the volumes in the child pools is also included.

Available for: All storage systems. For FlashSystem A9000 and FlashSystem A9000R, this value represents virtual space rather than physical space.

Available Repository Space (GiB)

The available, unallocated storage space in the repository for Track Space-Efficient (TSE) thin-provisioning.

Available for: DS8000 thin-provisioned pools.

Available Soft Space (GiB)

The amount of virtual storage space that is available to allocate to volumes in a storage pool.

Available for: XIV systems, and IBM Spectrum Accelerate storage systems.

Capacity (GiB)

The total amount of storage space in the pool, which might include overhead space if the disks for the pool aren't formatted. For XIV systems and IBM Spectrum Accelerate, capacity represents the physical or 'hard' capacity of the pool, not the virtual 'soft' capacity.

Available for: All storage systems.

Compression Savings (%)

The estimated percentage of the capacity of the resource that is saved by using data compression. This value is relative to all volumes in the pool, not just relative to the compressed volumes in the pool.

Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, XIV storage systems with firmware version 11.6 or later, and storage systems that run IBM Spectrum Virtualize.

Data Reduction (%)

The estimated percentage of the capacity of the resource that is saved by using data deduplication and data compression.

The following formula is used to determine the data reduction savings percentage:

$$(\text{original size} - \text{size after data reduction}) \div (\text{original size}) \times 100$$

For example, the data reduction savings percentage for a storage system is 75%. Therefore, data reduction reduced the size of the data that is written by 75%. The reduced size is 25% of the original size, which is the equivalent of a data reduction factor of 4 to 1.

Available for: FlashSystem A9000, FlashSystem A9000R, XIV storage systems, and storage systems that run IBM Spectrum Virtualize. For storage systems that do not support data deduplication, the data reduction savings percentage is equivalent to the compression savings percentage.

Enterprise HDD Available Space (GiB)

The amount of storage space that is available on the Enterprise hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Enterprise HDD Capacity (GiB)

The total amount of storage space on the Enterprise hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Nearline HDD Available Space (GiB)

The amount of storage space that is available on the Nearline hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

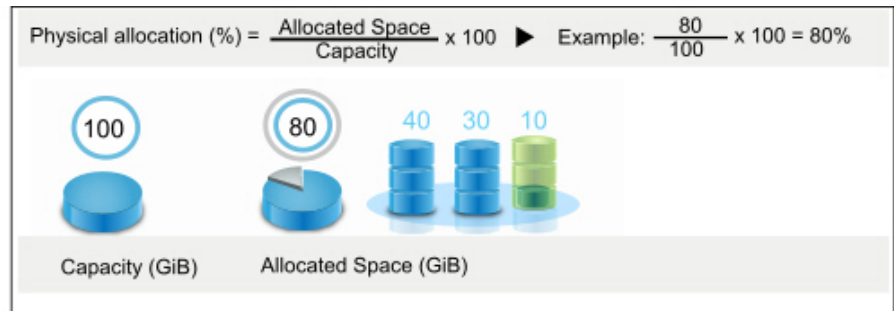
Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Nearline HDD Capacity (GiB)

The total amount of storage space on the Nearline hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Physical Allocation (%)



The percentage of physical capacity in the pool that is allocated to the regular volumes, the thin-provisioned volumes, and the volumes in child pools. This value is always less than or equal to 100% because you cannot allocate more physical space than is available in a pool. Check the value for physical allocation to see:

- Whether the physical capacity of the pool is fully allocated. That is, the value for physical allocation is 100%.
- Whether you have sufficient capacity to provision new volumes with storage
- Whether you have sufficient capacity to allocate to the compressed and thin-provisioned volumes in the pool

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Repository Capacity (GiB)

The total storage capacity of the repository for Track Space-Efficient (TSE) thin-provisioning.

Available for: DS8000 thin-provisioned pools.

Shortfall (%)

The difference between the remaining unallocated volume space and the available capacity of the associated pool, expressed as a percentage of the remaining unallocated volume space. The shortfall represents the relative risk of running out of space for overallocated thin-provisioned volumes. If the pool has sufficient available capacity to satisfy the remaining unallocated volume space, no shortfall exists. As the remaining unallocated volume space grows, or as the available pool capacity decreases, the shortfall increases and the risk of running out of space becomes higher. If the available capacity of the pool is exhausted, the shortfall is 100% and any volumes that are not yet fully allocated have run out of space.

If the pool isn't thin-provisioned, the shortfall percentage equals zero. If shortfall percentage isn't calculated for the storage system, the field is left blank.

The following formula is used to calculate this value:

$$[(\text{Unallocatable Space} \div \text{Committed but Unallocated Space}) \times 100]$$

You can use this percentage to determine when the amount of over-committed space in a pool is at a critically high level. Specifically, if the physical space in a pool is less than the committed virtual space, then the pool does not have enough space to fulfill the commitment to virtual space. This value represents the percentage of the committed virtual space

that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

Example: The remaining physical capacity of a pool is 70 GiB, but 150 GiB of virtual space was committed to thin-provisioned volumes. If the volumes are using 50 GiB, then 100 GiB is still committed to the volumes (150 GiB – 50 GiB) with a shortfall of 30 GiB (70 GiB remaining pool space – 100 GiB remaining commitment of volume space to the volumes).

Because the volumes are overcommitted by 30 GiB based on the available space in the pool, the shortfall is 30% when the following calculation is used:

$$\frac{[(100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space}) \div 100 \text{ GiB unallocated volume space}] \times 100}{}$$

Available for: DS8000, FlashSystem storage systems, Hitachi Virtual Storage Platform, SAN Volume Controller, and Storwize family storage systems that are configured with block storage. For FlashSystem A9000 and FlashSystem A9000R, this value is not available.

Soft Space (GiB)

The amount of virtual storage space that is configured for the pool.

Available for: XIV systems and IBM Spectrum Accelerate storage systems.

SSD Available Space (GiB)

The amount of storage space that is available on the solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller.

SSD Capacity (GiB)

The total amount of storage space on the solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Available for: DS8000, FlashSystem storage systems, and SAN Volume Controller.

Tier 0 Flash Available Space (GiB)

The amount of storage space that is available on the Tier 0 flash solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: Storwize family storage systems that are configured with block storage.

Tier 0 Flash Capacity (GiB)

The total amount of storage space on the Tier 0 flash solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: Storwize family storage systems that are configured with block storage.

Tier 1 Flash Available Space (GiB)

The amount of storage space that is available on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for retiring the volume extents in the pool.

Available for: Storwize family storage systems that are configured with block storage.

Tier 1 Flash Capacity (GiB)

The total amount of storage space on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for retiring the volume extents in the pool.

Available for: Storwize family storage systems that are configured with block storage.

Total Volume Capacity (GiB)

The total amount of storage space that can be made available to the regular and thin-provisioned volumes in the pool. If the pool is a parent pool, it also includes the storage space that can be made available to the volumes in the child pools. For Hitachi VSP non-thin provisioned pool space, this value is the sum of the capacity of regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Available for: All storage systems.

Unallocatable Volume Space (GiB)

The amount of space that cannot be allocated to volumes because the physical capacity of the pool cannot meet the demands for virtual space. The following formula is used to calculate this value:

[Total Volume Capacity - Pool Capacity]

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool by creating volumes with more virtual capacity than can be physically allocated in the pool. This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool. For Hitachi VSP non-thin provisioned pool space, this value is always zero.

Available for: All storage systems, except FlashSystem A9000, FlashSystem A9000R, XIV, and IBM Spectrum Accelerate.

Unallocated Volume Space (GiB)

The total amount of remaining space that can be allocated to the existing volumes in the pools. The following formula is used to calculate this value:

[Total Volume Capacity - Allocated Space]

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool. For Hitachi VSP non-thin provisioned pool space, the unallocated volume space is always zero.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Unassigned Volume Space (GiB)

The total amount of space in the volumes that are not assigned to hosts. For a thin-provisioning pool, this value includes the virtual capacity of thin-provisioned volumes, which might exceed the total space in the pool. For Hitachi VSP non-thin provisioning pool space, this value is the sum of unassigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Available for: All storage systems.

Unused Space (GiB)

The amount of allocated space that is not used for storage by the regular and thin-provisioned volumes in the pool. If the pool is a parent pool, the amount of allocated space that is not used for the volumes in the child pools is also calculated. The following formula is used to calculate this value:

[Allocated Space - Used Space]

Available for: Storage systems that run IBM Spectrum Virtualize.

Used Space (GiB)

The amount of allocated space that is used by the regular and thin-provisioned volumes in the pool. If the pool is a parent pool, the amount of allocated space that is used for the volumes in the child pools is also included.

For storage systems that run IBM Spectrum Virtualize and are configured with block storage, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Used Space might be different from the Allocated Space for pools that contain thin-provisioned volumes.

The Used Space reflects the size of compressed data that is written to disk for compressed volumes in pools on FlashSystem A9000, FlashSystem A9000, XIV systems, and storage systems that run IBM Spectrum Virtualize and are configured with block storage. As the data changes, the Used Space value can fluctuate as the compression and deduplication factors fluctuate.

For pools with volumes that are not thin provisioned, the values for Used Space and Allocated Space are equal. This value is accurate as of the most recent time that data was collected about a pool. Data collection is run on a set schedule and the used space on a volume can change rapidly. Therefore, the value in this column might not reflect the current state of a pool.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Virtual Allocation (%)

The percentage of the physical capacity that is committed to the virtual capacity of the volumes in the pool. If the value exceeds 100%, the physical capacity doesn't meet the demands for virtual capacity. The following formula is used to calculate this value:

$$[(\text{Total Volume Capacity} \div \text{Pool Capacity}) \times 100]$$

This value is available for all pools.

For Hitachi VSP non-thin provisioned pool space, the following formula is used to calculate this value:

$$[(\text{Allocated Space} \div \text{Pool Capacity}) \times 100]$$

Example: If the allocation percentage is 200% for a total storage pool size of 15 GiB, then the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive because insufficient physical capacity is available in the pool to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Shortfall (%) to estimate how critical the shortage of space is for a pool.

You can hover the mouse pointer over the percentage bar to view values for the total volume capacity and capacity. Total volume capacity total is the storage space on all of the volumes in the pool. For thin-provisioned volumes, this value includes virtual space. Capacity is the total amount of storage in the pool.

Available for: All storage systems.

Volume capacity metrics

You use the capacity chart to detect capacity shortages for the following types of volumes:

- Space-efficient volumes such as compressed volumes and thin-provisioned volumes
- Regular volumes that use Easy Tier to re-tier volume extents

You can review the allocation of space to space-efficient volumes to detect capacity shortfalls. You can also review the space usage of volumes that use Easy Tier to distribute volume extents across Enterprise HDD, Nearline HDD, and SSDS drives.

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the chart for volumes:

Allocated Space (GiB)

The amount of space that is allocated to the compressed, thin-provisioned, or the Easy Tier volume. Typically, the space that is allocated to the compressed or thin-provisioned volume is less than the capacity of the volume. For Easy Tier volumes, allocated space is the space that is allocated to the volume's extents on the Enterprise HDD, Nearline HDD, or SSD drives.

Available for: All storage systems.

Capacity (GiB)

The capacity of the compressed or the thin-provisioned volume, which comprises the sum of the allocated and the unallocated space. If the disks for the pool aren't formatted, the capacity of the volume might include the

overhead space. For thin-provisioned volumes in XIV systems pools or IBM Spectrum Accelerate pools, capacity is the physical or 'hard' capacity of the volume.

Available for: All storage systems.

Enterprise HDD Capacity (GiB)

The total amount of storage space on the Enterprise hard disk drive that the Easy Tier volume uses for re-tiering the volume extents.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Nearline HDD Capacity (GiB)

The total amount of storage space on the Nearline hard disk drive that the Easy Tier volume uses for re-tiering the volume extents.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

SSD Capacity (GiB)

The total amount of storage space on the solid-state drive that the Easy Tier volume uses for re-tiering the volume extents.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Available for: DS8000, FlashSystem storage systems, SAN Volume Controller, and Storwize family storage systems that are configured with block storage.

Available for: DS8000, FlashSystem storage systems, and SAN Volume Controller.

Tier 0 Flash Capacity (GiB)

The total amount of storage space on the Tier 0 flash solid-state drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: Storwize family storage systems that are configured with block storage.

Tier 1 Flash Capacity (GiB)

The total amount of storage space on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: Storwize family storage systems that are configured with block storage.

Used Space (GiB)

The amount of allocated space that is used by the compressed, thin-provisioned, or Easy Tier volume.

For compressed and thin-provisioned volumes, used space might not be the same as allocated space on the following types of storage systems:

- SAN Volume Controller

- Storwize family block storage systems
- FlashSystem storage systems

For thin-provisioned volumes, used space might not be the same as allocated space because you can preallocate space to the thin-provisioned volumes when the volumes are created. For compressed volumes, used space might not be the same as allocated space because more space is used to read the data than is used to write the data to the disk. For regular volumes that use Easy Tier on the storage systems that are listed, used space is the same as allocated space.

Available for: SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem V840, and FlashSystem V9000.

Capacity metrics for file storage systems

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

Alphabetical lists of the capacity and space usage metrics that you can add to charts are provided in the following sections:

- “Capacity metrics for filesets”
- “Capacity metrics for file systems ”
- “Capacity metrics for file system pools” on page 483

Capacity metrics for filesets

If sufficient data is collected about the filesets in your file storage systems, you can view charts that compare the used space of the filesets. The following metric is displayed in the capacity chart for filesets:

Used Space (GiB)

The amount of storage space that is used by the fileset. Used space is not provided for filesets that are cache targets.

Available for: All storage systems.

Capacity metrics for file systems

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for file systems:

Available Inodes

The number of inodes that are available in a file system.

Available Space (GiB)

The amount of storage space that is available (not allocated) on a file system.

Capacity (%)

The percentage of the total storage space on the file system that is used by files and directories.

Available for: All storage systems.

Maximum Inodes

The total number of inodes that the file system can contain. This value consists of the available inodes and the used inodes for the file system.

Available for: All storage systems.

Physical Capacity (GiB)

The raw capacity of the partition where a file system resides.

Total Capacity (GiB)

The amount of storage space on the file system.

Available for: All storage systems.

Used Inodes

The number of used inodes on a file system. An inode is the internal structure that describes the individual files or directories in the file system metadata. An inode contains the node, type, owner, and location of a file or directory.

Available for: All storage systems.

Used Inodes (%)

The percentage of inodes that are already used on the file system. Each time that you create a file or directory on the file system, an inode is allocated to the file or directory.

Available for: All storage systems.

Internal Used Space (GiB)

The amount of storage space that is unavailable (allocated) on a file system. For IBM Spectrum Scale, this value does not include the space that is used by migrated data on external pools.

Capacity metrics for file system pools

If sufficient data is collected about the file system pools in your data center, you can view charts that compare the capacity of the pools with the space that is allocated to the pools and the space that is still available in the pools.

The internal resource for IBM Spectrum Scale storage systems is called **Pools**. For all other file storage systems, the internal resource is called **File System Pools**.

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for file system pools:

Available Space (GiB)

The amount of unallocated storage space in the pool. Available space usually comprises the space that can be used for storage. However, if the pool is not formatted, the amount of overhead space might also be included in the calculation.

Available for: All storage systems.

Capacity (%)

The percentage of space in the file system pool that is being used.

Available for: All storage systems.

Total Capacity (GiB)

The total amount of storage space in the storage pool. If the pool is not formatted, the total amount of storage space might also include overhead space.

Available for: All storage systems.

Used Space (GiB)

The amount of space in the file system pool that is being used.

Available for: All storage systems.

Capacity metrics for object storage systems

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

Capacity metrics for containers

If sufficient data is collected about the containers in your object storage systems, you can view charts that compare the used space of the containers. You can add the following metrics to the capacity chart for containers:

Available Objects

The number of objects that you can add to the container on the object storage system.

Available for: IBM Spectrum Scale storage systems.

Available Space (GiB)

The amount of file system space that can be used to store object data for the container.

Available for: IBM Spectrum Scale storage systems.

Objects

The number of objects in the container on the object storage system.

An object stores data content, such as files, videos, images, virtual machine snapshots, and other unstructured data. In IBM Spectrum Scale, objects are stored as files on the GPFS file system.

Available for: IBM Spectrum Scale storage systems.

Objects Quota

The number of objects that can be stored in the container. This column contains a value only if an objects quota is defined for the container on the OpenStack Swift object storage system.

Available for: IBM Spectrum Scale storage systems.

Objects Quota (%)

The percentage of the objects quota for the container that was used when the storage system was last probed. The quota limits the number of objects that can be stored in the container. This column contains a value only if an objects quota is defined for the container on the OpenStack Swift object storage system.

If the number of objects that are used exceeds the objects quota, the percentage exceeds 100%. The number can exceed the quota if the quota is configured after the object files are stored on the file system.

Available for: IBM Spectrum Scale storage systems.

Space Quota (%)

The percentage of the space quota for the container that was used when the storage system was last probed. The quota limits the amount of file system space that can be used by the container. This column contains a value only if a space quota is defined for the container on the OpenStack Swift object storage system.

If the used space value for the container exceeds the space quota value, the percentage exceeds 100%. The used space can exceed the space quota if the quota is set after object files are uploaded to the file system or if an object file without size information is uploaded to the file system.

Available for: IBM Spectrum Scale storage systems.

Space Quota (GiB)

The amount of file system space that can be used to store object data for the container. This column contains a value only if a space quota is defined for the container on the OpenStack Swift object storage system.

Available for: IBM Spectrum Scale storage systems.

Used Space (GiB)

The amount of file system space that is used by the objects in the container. If a probe is run immediately after objects are added to a container, the used space value might not reflect the updates until after the next scheduled probe is run.

Available for: IBM Spectrum Scale storage systems.

Capacity metrics for tiers

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

Based on historical space usage, you can plan the capacity growth for tiers.

You can view values for the following metrics in the tier charts:

Capacity (GiB)

The total amount of storage space in the tier, which comprises the total storage space of the pools in the tier.

Allocated Space (GiB)

The amount of storage space that is used by pools in the tier.

Cinder driver for IBM Spectrum Control

The Cinder driver for IBM Spectrum Control enables cloud storage that is powered by OpenStack to use a IBM Spectrum Control installation for block storage provisioning.

IBM Spectrum Control provides block storage provisioning capabilities that the storage administrator can employ to define the properties and characteristics of storage volumes within a particular service class. For example, a block storage service class can define RAID levels, tiers of storage, and various other storage characteristics. For more on service classes, see “Configuring IBM Spectrum Control for provisioning” on page 294.

The Cinder driver enables an OpenStack cloud to use these defined service classes, which are available in your Cinder environment as Cinder volume types. For each block storage service class, a Cinder volume type is created. Your cloud users can create and use volumes of this type.

To download the IBM Spectrum Control Cinder driver and installation instructions go to https://www.ibm.com/developerworks/servicemanagement/sm/spectrum_control/index.html and click the Downloads tab.

Planning to deploy the Cinder driver

Before using the Cinder driver for IBM Spectrum Control you need to create volume types and service classes for the service to use. You also need to decide the connection protocol your storage will be using (FC or iSCSI).

Volume types within OpenStack

New volumes are created within OpenStack by specifying the type of volume to be created. The volume types are high-level descriptions that identify the class of storage a volume can provide. For instance, a cloud user might want to create a new volume with Gold class storage. The user can employ the OpenStack Horizon GUI or Cinder CLI to provision a new volume. The user can then specify that the type of volume they want to create is of type Gold. The OpenStack Cinder scheduler then allocates from those volume drivers that can provision a volume of type Gold.

The Cinder driver works with this information through the medium of service classes. A service class is used to identify classes of storage within the storage infrastructure that can be configured for such services as thin provisioning, virtualized volumes, and varying RAID levels. Service classes can be configured to use storage across multiple devices and device types that still meet the constraints that are defined within the service class. The Cinder driver can discover the service classes that are configured within IBM Spectrum Control and create corresponding volume types within the OpenStack environment. The user can then provision a volume from the service classes that are already defined within IBM Spectrum Control.

Creating a service class and volume type

To provide a new class or type of volume, the storage administrator creates a service class within IBM Spectrum Control. The definition includes the characteristics of the service class and the storage pools that are available to the service class. When the service class is created, the Cinder driver detects the new service class and automatically creates a corresponding Volume Type entry within the OpenStack infrastructure.

Specifying connection protocols for storage (FC and iSCSI)

When the Cinder driver is requested to attach a volume to a Nova node, the Cinder driver submits required access information to the IBM Spectrum Control service. The service then assigns the volume to all available host ports associated with the Nova node.

When the IBM Spectrum Control service assigns the volume to the ports, the service returns all available access information back to the Cinder driver. The Cinder driver chooses the optimal access to give to the Nova node to initialize the connection. By default, the Cinder driver attempts to choose the Fibre Channel method of access. The Cinder driver determines whether the IBM Spectrum Control server can enable Fibre Channel access and if the Nova node can use Fibre Channel access. If both are available, Fibre Channel connection information is returned. If Fibre Channel access is not available, the Cinder driver checks whether the IBM Spectrum Control service is able to assign iSCSI access to the volume. If iSCSI access is available to the Nova node, the Cinder driver returns the iSCSI access information to the Nova node.

An OpenStack cloud administrator can change which access protocol the Cinder driver attempts to assign first. The setting is the `tpc_preferred_protocol` parameter in the `cinder.conf` driver configuration section. By changing the preferred protocol to iSCSI, the Cinder driver first attempts to provide iSCSI access. If unsuccessful, the Cinder driver falls back to Fibre Channel access. If the preferred protocol cannot be met but an alternative protocol is available, then the alternative protocol is chosen.

Enabling an OpenStack cloud for IBM Spectrum Control

Enabling your OpenStack cloud to leverage your IBM Spectrum Control installation is a two step process. Configuration scripts for enabling IBM Spectrum Control functionality and policies for the Cinder driver are available in the `tpc-scripts` directory.

Download this zip file and extract its contents to a location on the IBM Spectrum Control server. Open a command prompt to the `tpc-scripts` directory and run the setup script for your platform.

Note: You must restart the IBM Spectrum Control Device, Data, and Web servers to complete the setup process.

A natively installable package containing the Cinder driver is available in the `cinder-packages` directory. Copy the appropriate `.rpm` or `.deb` file to your Cinder node and install it using the operating system's native packaging tool.

For example, if the Cinder node is running Red Hat, copy the `openstack-tpcdriver-1.0.0-1.x86_64.rhel.rpm` file and install the file using the command `rpm -ivh openstack-tpcdriver-1.0.0-1.x86_64.rhel.rpm`.

Upon installation of the `.rpm` or `.deb` file, you can run the `tpcdriver-config.sh` utility to help determine appropriate Cinder settings for your installation.

Copy the appropriate settings into your `cinder.conf` file and restart the Cinder volume service to enable the Cinder driver.

Once the Cinder driver is installed and configured, you can verify the installation and configuration by navigating to the Volumes page of the Horizon Admin panel to check that the service classes configured within IBM Spectrum Control are now available as volume types.

Setting up the Cinder driver environment

To use the Cinder driver, you need to configure role-to-group mappings for an authorized Cinder user, add storage system mappings, and define capacity pools and service classes.

Configuring volume creation and attachment

After you install the IBM Spectrum Control, but before you configure the Cinder node, use the IBM Spectrum Control GUI to set up volume creation and attachment.

1. Log in to the IBM Spectrum Control GUI.
2. Configure role-to-group mappings for an authorized OpenStack Cinder user in a selected group to connect IBM Spectrum Control.
 - a. In the menu bar, go to **Settings > User Management**.
 - b. On the **User Management** window, click **Add Group**.
 - c. On the **Add Group** window, enter the name of the group to which the user belongs, and click **Search**.
 - d. Select the group name and assign the **External Application** role to this group.
 - e. Click **OK**.
3. Add storage subsystems to manage to IBM Spectrum Control. For more information, see “Adding storage systems” on page 115.
 - a. In the menu bar, go to **Storage > Storage Systems**.
 - b. Click **Add Storage System**.
 - c. Click the **Storage System** icon.
 - d. On the **Discover** page, enter the information, and click **Next**.
 - e. On the **Configure** page, enter the information, and click **Configure**.
4. Define capacity pools for the storage subsystems to be used by the service classes. For more information, see “Capacity pools” on page 56.
 - a. In the menu bar, go to **Storage > Pools**.
 - b. Select one or more pools of a storage system and select **Action > Add to Capacity Pool**.
 - c. Click **Create Capacity Pool**.
 - d. In the **Create Capacity Pool** page, specify a name for the capacity pool.
 - e. Click **OK**.
5. Define the service classes to be used by the Cinder driver. For more information, see “Service classes” on page 52.
 - a. In the menu bar, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Service Classes**.
 - b. Click **Create Service Class**.
 - c. In the **Select Type** page, click on the **Block** icon to create a block storage service class.
 - d. In the **Define Properties** page, specify a name for the service class.
 - e. Specify the requirement for provisioning block-storage of this service class, and click **Next**.

- f. In the **Capacity Pool** page, select **Only allow provisioning from selected capacity pools**, select the capacity pools, and click **Next**.
 - g. In the **Authorize Users** page, disable **Approval required**. Select **Allow administrators and specified users to provision**.
 - h. Select the user name
 - i. Click **Add User**, and then click **Finish**.
6. Install the IBM Spectrum Control OpenStack patch and the Cinder driver. See “Installing the Cinder driver” on page 491.
 7. Use the OpenStack Horizon GUI to verify that the service classes are listed as volume types on Cinder nodes.
 - a. Log in to the OpenStack Horizon GUI as a user with admin roles.
 - b. Select **Admin tab > System Panel > Volumes**.
 - c. Verify that the service classes created on IBM Spectrum Control are listed under Volume Types.
 - OR
 - a. Log in to the OpenStack Cinder node.
 - b. Use the CLI cinder command to list volume type: `# cinder type-list`.
 - c. Verify that the service classes you created on IBM Spectrum Control are listed.
 - d. Ensure that the correct settings for IBM Spectrum Control are added to the Cinder configuration file `/etc/cinder/cinder.conf`. Check that the Cinder volume service was restarted.
 8. iSCSI volume attachment

To use iSCSI connections, the storage system must be configured for iSCSI access. You must assign IP addresses to specific Ethernet ports on the storage subsystem itself. For each storage subsystem, consult the product documentation on how to configure Ethernet ports for iSCSI access.

After the storage system is configured for iSCSI access, the IBM Spectrum Control server populates information about the iSCSI access with the next probe of the storage subsystem. You can manually trigger this action from the IBM Spectrum Control GUI by running the Start Probe command under the **Data Collection** menu for storage subsystems.

When you assign a volume to a server with iSCSI host ports, the IBM Spectrum Control generates and assigns a pseudo-random 16 character CHAP secret to any host definitions it creates on the storage subsystem. If the host port already belongs to a host definition on the storage system, the server does not carry out any CHAP secret modification.

Note: In IBM Spectrum Control version 5.2.1, iSCSI ports appear as Fibre Channel ports within the IBM Spectrum Control GUI. The WWPN of the iSCSI port is the iSCSI name of the port itself.

In IBM Spectrum Control 5.2.2, iSCSI ports appear as Fibre Channel ports within the IBM Spectrum Control GUI. The WWPN of the iSCSI port is the iSCSI name of the port itself.

Requirements for installing the Cinder driver

To offer cloud services for provisioning block files, you must ensure that you have the correct versions of IBM Spectrum Control and OpenStack.

Requirements for IBM Spectrum Control

You need to install IBM Spectrum Control with IBM Spectrum Control Advanced Edition license. You also need:

- A local operating system, LDAP, or Windows Domain group to add to IBM Spectrum Control
- A user added as a member of this group for use by the Cinder driver to connect to IBM Spectrum Control
- A supported version of DB2 client or DB2 server
- Service classes available to the OpenStack Cinder node(s) as volume types
- Storage systems available for OpenStack volume creation and attachment support

Storage platform support

The following storage platforms are supported:

- DS8000
- IBM SAN Volume Controller 6.1 or later
- Storwize V7000 6.1 or later
- XIV 11.0 or later

Requirements for OpenStack

You will need:

- OpenStack (Havana release)
- The OpenStack Cinder node must run on one of the following operating systems:
 - Red Hat Enterprise Linux 6.4 or higher
 - Ubuntu 12.04
- The python-cinder (included in python-cinderclient package) and python-rdfliib library packages installed on the Cinder node:
 - Red Hat Enterprise Linux 6.4 or higher.

The Red Hat python-rdfliib library is available in the Extra Packages for Enterprise Linux (EPEL) repository. Visit the [Fedora wiki](#) for information on installing the EPEL repository on your system.

```
# yum install python-cinderclient
# yum install python-rdfliib
```

- Ubuntu 12.04

```
# sudo apt-get install python-cinderclient
# sudo apt-get install python-pip
# sudo pip install rdfliib==3.2.1
```

Note: The python-rdfliib debian package available in the Ubuntu precise archives contains rdfliib 2.4.2. The IBM Spectrum Control Cinder driver requires version 3.2.1 or higher of the rdfliib library. Use the pip command above in order to install the correct version of the rdfliib library.

Cinder driver package contents

To enable an OpenStack cloud to leverage a IBM Spectrum Control installation, run the installation script and use a configuration utility to determine appropriate Cinder settings for your installation.

The tpc-scripts directory contains the platform-specific IBM Spectrum Control Server OpenStack scripts to enable or disable the IBM Spectrum Control functionality and policies for the Cinder driver.

- AIX Scripts: setup_aix.sh and unsetup_aix.sh
- Linux Scripts: setup_linux.sh and unsetup_linux.sh
- Windows Scripts: setup_windows.bat and unsetup_windows.bat

The setup script configures a IBM Spectrum Control server with the files required to enable the supported OpenStack functionality on the IBM Spectrum Control server. For more information about installing the Cinder driver, see “Requirements for installing the Cinder driver” on page 489 and “Installing the Cinder driver.”

The unsetup script reconfigures a a IBM Spectrum Control server with the files required to disable supported OpenStack functionality on the a IBM Spectrum Control server. For more information about uninstalling the Cinder driver, see “Uninstalling the Cinder driver” on page 493.

When IBM Spectrum Control Cinder driver is installed, you are ready to enable OpenStack and IBM Spectrum Control to work together. The cinder-packages directory contains the platform specific Cinder driver installer packages . Both the IBM Spectrum Control server and the OpenStack Cinder node need to be patched.

- Red Hat Enterprise Linux 6.4 or higher: openstack-tpcdriver-1.0.1-1.x86_64.rhel.rpm
 - Ubuntu 12.04: openstack-tpcdriver_1.0.1-1_all.deb
1. Log on to IBM Spectrum Control server as root for Linux/AIX or as Administrator for Windows.
 2. Change directory to tpc-scripts.
 3. Run the setup script, providing the path to the IBM Spectrum Control installation location. For example:

Option	Description
AIX	# ./setup_aix.sh /opt/IBM/TPC
Linux	# ./setup_linux.sh /opt/IBM/TPC
Windows	> setup_windows.bat "C:\Program Files\IBM\TPC"

4. The installer package installs the OpenStack tpc.py driver in the proper location. It also installs the tpcdriver-config.sh script which generates the additional Cinder driver-related entries needed for the Cinder configuration file (/etc/cinder/cinder.conf).

Installing the Cinder driver

To install the Cinder driver, run the installation script and use a configuration utility to determine appropriate Cinder settings for your installation.

The IBM Spectrum Control server and all OpenStack Cinder nodes need to be patched to enable IBM Spectrum Control OpenStack functions and policies for the Cinder driver. Review “Requirements for installing the Cinder driver” on page 489 before you proceed with the installation.

1. Run the IBM Spectrum Control OpenStack patch installation script for AIX, Linux, or Windows. On IBM Spectrum Control:
 - a. Log on as root for Linux/AIX or **Administrator** for Windows.
 - b. Change directory to tpc-scripts.

- c. Run the setup script, specifying the path to the IBM Spectrum Control server installation location:
 - AIX# ./setup_aix.sh /opt/IBM/TPC
 - Linux# ./setup_linux.sh /opt/IBM/TPC
 - Windows: setup_windows.bat "C:\Program Files\IBM\TPC"
- d. When prompted if you want to proceed with the update, enter Y or y.

Note: You must restart the IBM Spectrum Control Device, Data, and Web servers to complete the setup process.

2. The IBM Spectrum Control OpenStack patch installation script attempts to determine the DB2 settings. If the script cannot find a setting, you are prompted to enter the value. Descriptions of the settings are in the following table:

Option	Description
Hostname of the DB2 server	The host name of the IBM Spectrum Control database server.
Port of the DB2 server	The port that the IBM Spectrum Control database server is running on.
DB2 profile user (AIX/Linux only)	The user name of the DB2 user on the server this installation script is being run on. This setting is only relevant for AIX and Linux.
DB2 user	The user name of the DB2 user with permissions to access the IBM Spectrum Control database.
DB2 password	Password for the DB2 user.
Database name	The name of the IBM Spectrum Control database.
Schema name	The schema name of the IBM Spectrum Control database.

3. Run the Cinder driver installation script for Red Hat and Ubuntu installations. On the OpenStack Cinder nodes:
 - a. Ensure that the current user has administrative privileges - **root** for Red Hat and Ubuntu.
 - b. Change directory to cinder-packages.
 - c. Install the package for your OpenStack distribution:
 - Red Hat Enterprise Linux 6.4: # rpm -ivh openstack-tpcdriver-1.0.0-1.x86_64.rhel.rpm
 - Ubuntu 12.04: # dpkg -i openstack-tpcdriver_1.0.0-1.x86_64.deb

Note: If you installed any packages manually (for example, with pip or from source), you can force the installation by overriding dependency checks. On Red Hat systems, the --nodeps flag is passed to the rpm command. On Ubuntu systems, the --force-depends flag is passed to the dpkg command.

4. Run the tpcdriver-config.sh configuration script. You are prompted to provide IBM Spectrum Control server information as the script runs: # /usr/local/bin/tpcdriver-config.sh Refer to "Configuring the Cinder driver" on page 493 for more details on the Cinder configuration settings specific to IBM Spectrum Control.
5. Copy the Cinder driver settings that are generated by the tpcdriver-config.sh script to the [DEFAULT] section of the Cinder configuration file, /etc/cinder/cinder.conf.
6. Restart the Cinder volume service.

Uninstalling the Cinder driver

To uninstall the Cinder driver, run uninstallation scripts for AIX, Linux, and Windows installations.

Use IBM Spectrum Control OpenStack patch uninstallation scripts for AIX, Linux, and Windows to uninstall an instance of the Cinder driver.

1. Log on to IBM Spectrum Control with administrative privileges - **root** for Linux/AIX or **Administrator** for Windows.
2. Change directory to tpc-scripts.
3. Run the unsetup script, specifying the path to the IBM Spectrum Control server installation location:
 - AIX

```
# ./unsetup_aix.sh /opt/IBM/TPC
```
 - Linux

```
# ./unsetup_linux.sh /opt/IBM/TPC
```
 - Windows

```
unsetup_windows.bat "C:\Program Files\IBM\TPC"
```
4. When prompted if you want to proceed with the update, enter **Y** or **y**.
5. You must restart the IBM Spectrum Control Device, Data, and Web servers to complete the unsetup process.
6. On the OpenStack Cinder node(s), log on as root for Red Hat and Ubuntu.
7. Configure the Cinder volume service to no longer use the IBM Spectrum Control Cinder volume driver by editing the Cinder configuration file `/etc/cinder/cinder.conf` and removing all references to the IBM Spectrum Control Cinder volume driver. Parameters for the driver are prefixed with the string `tpc_`.
8. Restart the Cinder volume service.
9. After the Cinder volume services are no longer configured to use the IBM Spectrum Control driver, the native package can be removed from the Cinder volume node. For example:
 - Red Hat

```
# rpm -e openstack-tpcdriver
```
 - Ubuntu

```
# dpkg -r openstack-tpcdriver
```

Configuring the Cinder driver

To make IBM Spectrum Control volumes and service classes available to Cinder services, you must configure the Cinder driver.

On IBM Spectrum Control, you must plan user and service class details:

- Identify which IBM Spectrum Control user the Cinder driver can use to log on. The user's group must be added to **User Management** as **External Application**.
- Identify the preferred service class. This class is used when no volume type is specified in the Cinder configuration file. All block service classes available to the user are available to the Cinder driver as volume types.

The `tpcdriver-config.sh` script installed on the Cinder node as part of the rpm or deb file helps to configure the Cinder volume service to use the Cinder driver. Configure the Cinder driver by editing the Cinder configuration file `/etc/cinder/cinder.conf`.

Edit the following `volume_` settings in the Cinder configuration file:

- `volume_driver`: The name of the Cinder driver to use for volume creation. By default, name is set as `cinder.volume.drivers.lvm.LVMISCSIDriver` in the Cinder configuration file. To use the Cinder driver, change the setting to `cinder.volume.drivers.tpc.TPCDriver`.
- `volume_name_template`: The template that is used by IBM Spectrum Control to name OpenStack volumes. By default, template is set as `volume-%s` in the Cinder configuration file. You can change this value to one more easily identifiable as an OpenStack volume (for example, `openstackvol-%s`) when you are viewing or searching volumes. The template is a global setting that affects volume names for all Cinder drivers.

You must add the following `tpc_` settings to the Cinder configuration file:

- `tpc_ip`: The IP address of the IBM Spectrum Control server.
- `tpc_login`: The user name that the Cinder driver will log on with.
- `tpc_password`: The password that the Cinder driver will log on with.
- `tpc_service_class`: The default IBM Spectrum Control service class that the Cinder driver uses if no volume type is specified.
- `tpc_preferred_protocol`: The preferred protocol if more than one is available. The preferred protocol can be set to `iSCSI` or `FC`; these values are case-sensitive. This flag is optional - if not specified, the default is `FC`.

After you edit the Cinder configuration file, restart Cinder services with the command: `service service-name restart`. For Ubuntu, the *service-name* parameter is `cinder-volume`. For RedHat and SLES, the *service-name* parameter is `openstack-cinder-volume`.


Use the Horizon GUI to create, attach, and delete volumes.


Learn more about the Cinder driver

You can learn more about the Cinder driver from Service Management Connect resources such as developer wikis, updates, and Beta packages.

IBM Spectrum Control Service Management Connect contains the OpenStack package and links to

- IBM Spectrum Control developerWorks wiki
- IBM Spectrum Control forums
- Updates to this package
- IBM Spectrum Control Beta material

To learn more, go to  https://www.ibm.com/developerworks/servicemanagement/sm/spectrum_control/index.html and click the Downloads tab.

The OpenStack website contains information on installing, configuring, and managing an OpenStack cloud. To learn more, go to  <http://www.openstack.org>.

Limitations and known issues for the Cinder driver

The Cinder driver is subject to the following limitations and known issues.

Decommissioning a service class and volume type

To remove a volume type from OpenStack, you must first remove the service class that is associated with the volume type in IBM Spectrum Control. After the service class is removed, the OpenStack Administrator must then remove the corresponding volume type in OpenStack.

For example, you can remove a service class (Gold) from IBM Spectrum Control, and then remove the type of volume that is associated with the Gold service class. If a user requests the type of volume that is associated with the Gold service class, the request is rejected.

If you remove the volume type from OpenStack but do not remove the associated service class, the volume type is re-created the next time the Cinder driver checks for available service classes.

iSCSI support with XIV Storage System

- For an XIV system with multiple Internet Small Computer System Interface (iSCSI) ports, only one port is populated by the Cinder driver.
- For the XIV system, IBM Spectrum Control is unable to use Challenge-Handshake Authentication Protocol (CHAP) secrets for iSCSI host definitions. CHAP information is not available through the CLI for the XiV system. iSCSI attachments that are made through the Cinder driver are assigned without CHAP secrets defined.

Mixed iSCSI and Fibre Channel Volume Assignments

IBM Spectrum Control attempts to assign all ports that belong to a Nova node. When the unmanaged server representing the Nova node contains both Fibre and iSCSI ports, the volume is available to the Fibre Channel and iSCSI ports on the Nova node. In some cases, when you detach a volume from a host, the ports for the access method in use by the Nova node are removed but the access for the alternative connectivity method remains. This access can cause errors to delete the volume from the storage system. In this case, the volume is removed from the Cinder database but it requires manual steps to remove the storage volume from the storage system.

Unable to Attach Storage Volumes

In some circumstances, attempts to attach a storage volume to a compute instance fails even though the volume is assigned to the Nova node's host connection information. To resolve this problem, verify that the Nova node can connect to the storage system through the Ethernet LAN or Fibre SAN. If an XIV Storage System is added (regardless of whether the volume is on an XIV Storage System), restart the IBM Spectrum Control Device server. Use the startup scripts in the scripts directory where IBM Spectrum Control is installed.

Opening IBM Spectrum Control on Windows operating systems

You can open IBM Spectrum Control CLIs and GUIs and administer IBM Spectrum Control on Windows operating systems.

Opening IBM Spectrum Control GUIs and CLIs

To manage and monitor storage resources, open IBM Spectrum Control GUIs and CLIs.

You can open the following GUIs and CLIs:

- “Opening IBM Spectrum Control GUI”
- “Opening DB2 Command Window”
- “Opening IBM Data Studio Administration client on Windows operating systems”
- “Opening IBM Tivoli Monitoring Services” on page 497

Opening IBM Spectrum Control GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none">1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.2. On the Start page, right-click, and then click the All apps taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start > All Programs .

2. Click **IBM Spectrum Control > IBM Spectrum Control**.

Opening DB2 Command Window

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none">1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.2. On the Start page, right-click, and then click the All apps taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start > All Programs .

2. Click **IBM DB2 > Command Line Tools > Command Window**.

Opening IBM Data Studio Administration client on Windows operating systems

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none">1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.2. On the Start page, right-click, and then click the All apps taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start > All Programs .

2. Click **IBM Data Studio > Data Studio Administration Client**.

Opening IBM Tivoli Monitoring Services

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none">1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.2. On the Start page, right-click, and then click the All apps taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start > All Programs .

2. Click **IBM Tivoli Monitoring > IBM Tivoli Monitoring Services**.

Accessing administration tools

To manage and maintain IBM Spectrum Control, access the Windows system administration tools.

To complete tasks in IBM Spectrum Control, you must open the following administration and maintenance facilities:

- “Accessing the Control Panel”
- “Accessing Administrative Tools”
- “Accessing Windows Services” on page 498
- “Accessing Computer Management” on page 498
- “Accessing Programs and Program Features” on page 498

Accessing the Control Panel

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start .

2. Click **Control Panel**

Accessing Administrative Tools

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows 7, Windows 2008	Click Start .
Windows XP, Windows Vista	Click Start > Control Panel .

2. Click **Administrative Tools**

Accessing Windows Services

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows 7, Windows 2008	Click Start .
Windows XP, Windows Vista	Click Start > Control Panel .

2. Click **Administrative Tools > Services**

Accessing Computer Management

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start .

2. Click **Control Panel > Computer Management**

Accessing Programs and Program Features

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start .

2. Click **Control Panel > Programs and Program Features**.

Accessing Window Run

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows 7, Windows 2008, Windows Vista, Windows XP	Click Start .

2. Click **Run**.

Appendix B. Accessibility features for IBM Spectrum Control

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Spectrum Control:

- Keyboard-only operation in the GUI.
- A Knowledge Center that includes the following accessibility features:
 - The Knowledge Center is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
 - All documentation for IBM Spectrum Control is available in Adobe Portable Document Format (PDF) by using the Adobe Acrobat Reader. To access PDFs, go to http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.14/com.ibm.spectrum.sc.doc/fqz0_r_printable_doc.html.
 - All images in the Knowledge Center are provided with alternative text, so that visually impaired users can understand the contents of the images.
- Interfaces that are commonly used by screen readers.

The setting for the automatic-refresh button in the screen reader is toggled to the ON position by default. If you want the screen reader to read the previous text, complete the following steps:

1. Navigate to the Accessibility Settings Navigation region by using the arrow keys. The region is located after the IBM Spectrum Control application title.
2. Click **Enter** to toggle the automatic-refresh button to the OFF position. (An alert sounds to make you aware that the turn-off automatic-refresh toggle button was pressed.)
3. To move backward to the previously read text so that the screen reader can read it again, use the arrow keys. You can move backwards and forwards through the page.
4. When you are ready to move on, click **Enter** to toggle the automatic-refresh button to the ON position and to refresh the page. (An alert sounds to make you aware that the turn-on, automatic-refresh toggle button was pressed.)

Tip: Alternatively, let the toggle setting persist, and refresh as needed by pressing the F5 key.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

IBM and accessibility

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at  <http://www.ibm.com/able/>.

Notices

This information was developed for products and services offered in the U.S.A. This material may 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
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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 states 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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 the names and addresses used by an actual business enterprise 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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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>.

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

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

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

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

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

Glossary

This glossary provides terms and definitions for the IBM Spectrum Control software and products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website ([opens in new window](#)).

A

active management server

A management server from which the storage environment can be monitored and managed. The active management server replicates its database to the standby server.

B

business continuity

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

C

CIM See Common Information Model.

CIM agent

The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

CKD See count key data.

cluster

1. In IBM System Storage DS8000, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit,

any operational cluster can take over the processing of a failing cluster.

2. A collection of complete systems that work together to provide a single, unified computing capability.
3. In Storwize® V7000, a pair of nodes that provides a single configuration and service interface.
4. A loosely coupled collection of independent systems (or nodes) organized into a network for the purpose of sharing resources and communicating with each other.

Common Information Model (CIM)

An implementation-neutral, object-oriented schema for describing network management or systems management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

community name

The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

consistency group

A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

copy set

The set of source volumes or target volumes involved in a FlashCopy operation.

count key data (CKD)

A data recording format that uses self-defining record formats in which each record on a volume is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

D

data collection

The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include Domain Name System (DNS) probes, web page analyzers, or database analyzers. See also discovery.

data source

A storage resource or agent that provides data about a storage environment.

discovery

The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also data collection.

E

ECKD See extended count key data.

enterprise repository

A component of the data server that records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by the data server using Java Database Connectivity (JDBC).

event An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

extended count key data (ECKD)

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

F

fileset See file set.

file set

1. A subset of a file system that provides granularity of functions such as snapshots or quotas within the file system.
2. A hierarchical grouping of files managed as a unit for balancing workload across a cluster.

FlashCopy

1. Pertaining to a point-in-time copy where a virtual copy of a volume is created. The target volume maintains the contents of the volume at the point in time when the copy was established. Any subsequent write operations to the source volume are not reflected on the target volume.
2. An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a volume.

freeze An operation in which a storage system blocks I/O from the host system to the affected volumes on the primary site. A freeze operation stops mirroring between the primary and secondary volumes to ensure data consistency at the secondary site. See also thaw.

G

Global Copy

A non-synchronous long-distance copy option for data migration and backup. See also Global Mirror, Metro Mirror, remote mirror and copy.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

Global Mirror

An optional capability of the remote mirror and copy feature that provides a two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also Global Copy, Metro Mirror, remote mirror and copy.

grain size

The unit size for allocating space on thin-provisioned volumes, such as 32, 64, 128, and 256 kibibyte. The grain size is defined when a volume is created.

GUID See Globally Unique Identifier.

H

Hardware Management Console (HMC)

In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.

heat map

A color-coded data chart in which colors are used to differentiate values in a data set.

HMC See Hardware Management Console.

host volume

A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing.

hypervisor

Software or a physical device that enables multiple instances of operating systems to run simultaneously on the same hardware.

J

journal volume

A volume that holds a consistent copy of data until a new consistent copy is formed. The journal volume restores the last consistent point during a recovery.

L

launch-in-context

An operation in which a user starts a secondary application from a primary application to perform a specific task. Using the parameters, navigation instructions, and user credentials that are supplied by the primary application, the secondary application opens to the specific place in which to complete the task.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

LUN See logical unit number.

M

Management Information Base (MIB)

In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

management server

A system that provides a central point of control for managing data replication.

management server relationship

A connection between two replication servers, where one server acts as the active server and replicates the data that is necessary for the standby server to take control of the replication environment.

Metro Global Mirror

A three-site, high availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

Metro Mirror

A function of the remote mirror and copy feature that constantly updates a secondary copy of a volume to match changes made to a source volume. See also Global Copy, Global Mirror, remote mirror and copy.

MIB See Management Information Base.

N

namespace

The scope within which a Common Information Model (CIM) schema applies.

native interface

An interface that is specific to a system or subsystem.

P

ping

1. The command that sends an Internet Control Message Protocol (ICMP) echo-request packet to a gateway, router, or host with the expectation of receiving a reply.
2. A job that tracks the availability of assets and that is performed by an

agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network.

pool A grouping of storage space that consists of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

practice volume
A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

primary site
A physical or virtual site that is made up of hardware, network, and storage resources. Typically, production operations run at the primary site. Data can be replicated to a secondary site for disaster recovery and failover operations. See also secondary site.

primordial pool
Storage capacity that is unallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

probe A data collection job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units.

R

recovery point objective
The maximum amount of data loss that can be tolerated during a service interruption.

remote mirror and copy
A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also Global Copy, Global Mirror, Metro Mirror.

resource
In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, and storage systems.

role A job function that identifies the tasks that a user can perform and the resources

to which a user has access. A user can be assigned one or more roles.

role pair

The association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

S

scan A data collection job that monitors storage usage and file statistics on the resources in an environment.

secondary site
A physical or virtual site that is made up of the hardware, network, and storage resources that support the recovery needs of the primary site. When a failure occurs at the primary site, operations can continue at the secondary site. See also primary site.

session
A collection of source and target volumes that are managed to create consistent copies of data. The type of data replication that is associated with the session determines the actions that can be conducted for the volumes.

site awareness
The association of a location with each storage system in a session. Site awareness ensures that a volume can only be selected for a session if it matches the location of the site. Site awareness helps to prevent both reversing a hardware relationship and selecting volumes at the wrong location.

SMI-S See Storage Management Initiative Specification.

standby management server
A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

Storage Management Initiative Specification (SMI-S)

A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable

interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

storage pod

A logical entity of components of a system consisting of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

storage resource group

A named collection of logically related resources that are monitored by IBM Spectrum Control. Monitored resources can include fabrics, switches, computers, storage systems, and other storage resource groups.

T

target volume

A volume that receives data from a host volume or another intermediate volume.

thaw An operation in which a storage system releases the block of the I/O from the host system to the affected volumes on the primary site. A thaw operation can occur after a freeze operation ends and consistency is formed on the secondary site. See also freeze.

thin provisioning

A mechanism that provides the ability to define logical volume sizes that are larger than the physical capacity installed on the system.

track space-efficient volume (TSE volume)

A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume to the target volume.

TSE volume

See track space-efficient volume.

V

virtualization

The substitution of virtual resources for actual resources, where the virtual resources have the same functions and

external interfaces as their counterparts, but differ in attributes, such as size, performance, and cost. Virtualization is commonly applied to physical hardware resources by combining multiple physical resources into shared pools from which users receive virtual resources.

virtual storage area network (VSAN)

A fabric within the storage area network (SAN).

VSAN See virtual storage area network.

Z

zone A logical grouping of switches, switch ports, and their attached devices in a fabric.

zone alias

A name that is given to a collection of one or more zone members to be managed together.

zone set

A group of zones that function together on a fabric.

Index

A

- about this document x
- access path host replacement 300
- accessibility features
 - for IBM Spectrum Control 499
- acknowledging alerts 193
- active data 251
- activity deviation percentage
 - formula 327
- activity levels
 - formula 327
- activity limits
 - formula 327
- adding 267, 273, 274
 - NetApp devices 116
- adding applications 274
- adding hypervisors
 - ESX and ESXi 122
 - vCenter Server system 122
- adding resources 112, 267
 - fabrics 120
 - hypervisors 122, 123
 - rollup servers 132
 - servers with agents 130
 - servers without agents 126
 - storage systems 115
 - switches 120
 - vCenter Server system 124
- adding subdepartments 273
- administering general groups 283
- agents
 - CIM 75
- alert conditions and violations 64, 148, 369
- alert for resource
 - suppressing notifications 166
- alert notification settings
 - modifying 189, 190
- alerting 64, 148, 179, 369
 - alert log 64, 148, 369
 - defining alerts for attribute and capacity changes 151, 167
 - defining alerts for performance metrics 170
 - defining general group alerts for attribute and capacity changes 176
- notifications 64, 148, 369
- alerts 64, 148, 369
 - acknowledging 193
 - actions 400, 401
 - configuring for performance monitors 143
 - configuring for probes 137
 - notifications 401
 - removing 193
 - suppression actions 400
 - triggering conditions 370
 - viewing violations of capacity threshold 254
 - viewing violations of performance thresholds 235

- alias 73
- analysis period
 - performance data, change 332
- analyze tiering
 - analysis period, change 332
 - conditions, setting 87
 - hypervisors 323
 - servers 323
 - storage pools 323
 - storage virtualizers 323
 - thresholds, modify 324
 - volume placement 87
 - volumes 323
 - volumes, collocate 90
- application
 - scenario 95
- application context 269
- applications
 - adding applications to departments for monitoring 274
 - application hierarchy 262
 - creating 262
 - creating applications for monitoring 263
 - defining alerts for attribute and capacity changes 167
 - defining alerts for performance metrics 170
 - monitoring 261
 - removing 269
 - removing applications for monitoring 269
 - resources 269
 - scenarios
 - viewing capacity usage trends 97
 - viewing information 261
 - viewing status 261
- applications and departments
 - scenarios
 - monitoring capacity usage 91
- architecture
 - description 37
- audience vii

B

- balance pools 86
 - criteria 327
 - modify 326
 - volumes, collocate 90
- Balance Pools 326
- best practice performance guidelines for storage systems that run IBM Spectrum Virtualize
 - viewing 233
- block storage
 - capacity charts 469
 - capacity metrics 469
- block storage provisioning
 - on vSphere Web Client extension 316
- blogs 15

C

- capacity
 - adding to a capacity pool 301
 - alerts 254
 - block storage systems 238
 - capacity view 242
 - containers 238
 - controls for the capacity view 243
 - file storage systems 238
 - file system pools 238
 - file systems 238
 - filesets 238
 - metrics 243, 469
 - monitoring 238
 - non-tiered storage 321
 - object storage systems 238
 - pools 238
 - removing from a capacity pool 303
 - tiered storage 321
 - tiers 238
 - viewing threshold violations 254
 - volumes 238
- capacity charts
 - block pools 469
 - block storage systems 469
 - block volumes 469
 - containers 484
 - object storage systems 484
 - tiers 485
- capacity license model
 - category of drives 30, 33
 - uncategorized drives 30, 33
- capacity metrics
 - capacity charts 469
 - containers 484
 - object storage systems 484
 - pools 469
 - storage systems 469
 - tiers 485
 - volumes 469
- capacity pool 52, 294
 - adding resources to 301
 - creating 301
 - deleting 304
 - modifying 303
 - removing resources from 303
 - viewing capacity of 303
 - viewing details of 303
 - viewing resources of 303
- capacity view
 - accessing 242
 - applications 245
 - bookmarks 251
 - containers 245
 - controls 243
 - file system pools 245
 - file systems 245
 - filesets 245
 - metrics 243, 245
 - pools 245
 - storage systems 245

- capacity view (*continued*)
 - tiers 245
 - volumes 245
- CIM agents 75
- Cinder 486
- Cinder driver 486, 494
- Cisco MDS 9000 74
- cloud computing 486
- cloud configuration 52, 294
- cloud object storage
 - , vaults, accessibility risk 214
 - , vaults, inaccessible 213
- cloud storage 251, 253
- collecting data
 - GPFS snapshots 147
- collection
 - overview 47
- comments, sending x
- components 37
- computer alerts
 - triggering conditions 394
- concept diagram
 - VASA provider 63
- condition of resources
 - how the condition of resources is determined 204
- condition of resources in the web-based GUI
 - viewing the condition of top-level resources 208
- conditions and violations 64, 148, 369
- configuring 52, 295
 - alerts for performance monitors 143
 - alerts for probes 137
 - license compliance 33
 - license compliance after an upgrade 33
 - license model 33
 - license model after an upgrade 33
 - NetApp devices 116
- containers
 - capacity charts 484
 - capacity metrics 484
- controls for capacity view
 - export capacity information 243
 - hide and show resources 243
 - hide and show the legend 243
 - metrics 243
 - remove resources 243
 - select metrics 243
 - specify a time range 243
 - specify granularity 243
 - view chart 243
 - view chart in a new window 243
 - view table 243
- controls for performance view
 - export performance information 225
 - hide and show resources 225
 - hide and show the legend 225
 - remove resources 225
 - select metrics 225
 - specify a time range 225
 - specify granularity 225
 - view chart 225
 - view chart in a new window 225
 - view table 225
- creating 263, 267, 272, 295

- creating applications 263
- creating departments 272
- creating subcomponents 267
- custom alert for resource
 - editing 158
- custom alerts for applications
 - defining 173
 - editing 173
- custom alerts for general groups
 - defining 182
 - editing 182

D

- dashboard 197
- data collection 47
 - creating performance monitors in the web-based GUI 140
 - modifying performance monitors in the web-based GUI 142
 - performance monitors 140
 - starting and stopping performance monitors in the web-based GUI 144
 - verifying that a performance monitor is running in the web-based GUI 142
 - verifying that a probe is running in the web-based GUI 136
 - viewing all performance monitors in the web-based GUI 145
 - viewing performance monitor logs in the web-based GUI 146
 - viewing performance monitors for a resource in the web-based GUI 145
 - viewing performance monitors in the web-based GUI 145
- data paths
 - changing the layout 293
 - exporting a data path to a file 294
 - fabrics 287
 - hypervisors 287
 - navigating a data path 292
 - overview 287
 - printing a view 293
 - resources that you can view 288
 - saving the view settings 293
 - servers 287
 - storage systems 287
 - topology view 288
 - topology view for data paths
 - changing the layout 293
 - exporting a data path to a file 294
 - navigating a data path 292
 - printing a view 293
 - saving the view settings 293
 - viewing information about internal resources 292
 - viewing information about top-level resources 292
 - zooming in 292
 - zooming out 293
 - viewing information about internal resources 292
 - viewing information about top-level resources 292
 - zooming in 292
 - zooming out 293

- data store
 - creating in vSphere Web Client extension 316
- datastore
 - See* data store
- default
 - zone 73
- defining alerts for attribute and capacity changes 151, 167
- defining application alerts for performance metrics 170
- defining general group alerts for attribute and capacity changes 176
- defining general group alerts for performance metrics 179
- department
 - scenarios
 - comparing storage in each department 93
- department context 276
- departments
 - applications 276
 - creating departments for monitoring 272
 - monitoring 270
 - removing 276
 - removing departments for monitoring 275
 - viewing information 270
 - viewing status 270
- departments and applications
 - scenarios
 - different levels of business hierarchy 91
 - relational 97
- details 268, 275
- discover 147
- domain port zoning 73
- Domain_ID 73

E

- EB
 - See* EiB
- EiB 69
- enclosure license model 30, 33
- ESSp performance metrics 439
- ESX and ESXi
 - adding 122
- exported file systems 147
- exporting information in the GUI 336
- exporting information in the web-based GUI 278
- exportPerformanceData script 237
- external storage 251, 253

F

- fabric 72
- fabric alerts
 - triggering conditions 392
- Fabric Manager
 - zone 73
 - zone set 73
- fabrics
 - adding 120, 122

- fabrics (*continued*)
 - agent types for monitoring 122
 - exporting performance data 237
 - modifying properties 333
 - performance data, exporting 237
 - viewing data paths 287
 - viewing information in the GUI 194
 - viewing status in the web-based GUI 203
- FCID zoning 73
- file share
 - access path 300
- file storage provisioning
 - on vSphere Web Client extension 318
- filtering lists in the web-based GUI 338
- filters 264
 - assign resources 264
 - sorting rows 339

G

- general group hierarchy 280
- general groups
 - add resources 280
 - defining alerts for attribute and capacity changes 176
 - defining alerts for performance metrics 179
 - subgroups, add 281
- GiB 69
- glossary 505
- GPFS snapshots
 - collecting data 147
- groups
 - general groups 279
 - adding resources 282
 - administering actions 283
- GUI
 - data collection
 - performance monitors 140
 - exporting information about
 - applications and departments 278
 - how information is organized 194
 - internal, related, and object resources, list 107
 - managing tasks 340
 - monitoring fabrics 72
 - monitoring switches 72
 - resources that can be monitored 107
 - saving information about applications and departments 278
 - tasks
 - managing provisioning tasks 344
 - managing tiering storage, balancing pools, and transforming-storage tasks 346
 - viewing all performance monitors 145
 - viewing information about resources 194
 - viewing pending tasks 340
 - viewing tasks 340

H

- hardware zoning 73

- hierarchies
 - application 277
 - department 277
 - viewing information 277
- host definition
 - default 311
 - for fibre_channel ports 311
 - provisioning 311
- how the condition of resources is determined 204
- hypervisor alerts
 - triggering conditions 385
- hypervisors 323
 - adding 122, 123
 - modifying properties 333
 - optimized volumes, collocate 90
 - tiering, analyze 323
 - user roles 112
 - viewing data paths 287
 - viewing information in the GUI 194
 - viewing status in the web-based GUI 203

I

- I/O rate
 - volumes, tiering 87
- IBM Copy Services Manager 46
- IBM Spectrum Control
 - architecture 37
 - component overview 37
 - overview 37
 - VASA provider 62
- IBM Spectrum Control limitations and known issues 495
- IBM Spectrum Control REST API
 - using 77
- IBM Spectrum Scale
 - file systems, view performance metrics 450
 - nodes, view performance metrics 450
- IBM Spectrum Virtualize
 - best practice guidelines 234
 - host connections, view performance metrics 416
 - I/O groups, view performance metrics 416
 - managed disks, view performance metrics 416
 - nodes, view performance metrics 416
 - pools, view performance metrics 416
 - ports, view performance metrics 416
 - volumes, view performance metrics 416
- inactive data 251
- Installation
 - Cinder driver package contents 491
- interfaces 46
 - overview 38
 - Tivoli Storage Productivity Center for Replication 46

K

- KB
 - See* KiB

- KiB 69

L

- license
 - capacity limit 30
 - compliance
 - configuring 30, 33
 - managed capacity 30
 - model
 - capacity 30
 - enclosure 30
 - modelconfiguring
 - capacity 33
 - enclosure 33
- license restrictions 23
- licenses 22
- Linux 486
- lists
 - hiding columns 339
 - reorganizing columns 339
 - showing columns 339
- local server 66

M

- managed capacity 30
- master server 66
- max deviation percentage
 - default range 327
- MB
 - See* MiB
- metricsESS 439
- metricsperformance reports 403
- metricsXIV 434
- MiB 69
- modifying 297
- monitor
 - capacity and space usage 95
- Monitoring Agent
 - displaying reports 76

N

- N_Port address identifier 73
- N_Port_Name 73
- Node_Name 73
- non standard zones
 - fabric assist 74
 - protocol 74
 - quick loop 74
- notifications for alerts 401

O

- open source 486
- Opening the management GUI for a resource 335
- OpenStack 486
 - IBM Spectrum Control 488
 - Installing IBM Spectrum Control 491
 - Uninstalling IBM Spectrum Control 493
 - Cinder driver 493
- overview 52

P

- PB
 - See PiB
- performance 403
 - , creating support package 100
 - , exporting data 100
 - , storage systems, comparing 102
 - , troubleshooting 103
 - best practice guidelines 234
 - best practice guidelines for storage systems that run IBM Spectrum Virtualize 233
 - controls for the performance
 - view 225
 - hot spots 85
 - metrics
 - disk 403
 - ESS 439
 - IBM Spectrum Accelerate 434
 - port 403
 - volume 403
 - XIV 434
 - monitoring 217
 - performance view 224, 233
 - viewing threshold violations 235
- performance alert for applications
 - editing 170
- performance alert for general groups
 - editing 179
- performance alert for resource
 - editing 155
- performance data
 - analysis period, change 332
 - exporting 237
 - exporting performance data 237
- performance metrics
 - IBM Spectrum Scale 450
 - IBM Spectrum Virtualize 416
- performance view
 - accessing 224, 233
 - bookmarks 233
 - controls 225
 - favorites 233
 - saving URLs 233
 - viewing the performance of internal resources 229
 - viewing the performance of related resources 229
- PiB 69
- pools
 - balance 86
 - balance, analyze 326
 - capacity view 245
 - hot spots, identify 86
 - modifying properties 333
 - tiering, analyze 323
- port name zoning 73
- port zoning 73
- PortId zoning 73
- probes
 - automated scheduling 135
 - configuring alerts 137
- problems
 - report xi
- provisioning 147, 311
 - custom requirements for 299
 - overview 57, 294

provisioning (*continued*)

- scenarios
 - constrained set of storage 98
 - using vSphere Web Client extension 98
- shares 314, 315
- volumes 312, 313

R

- reader feedback, sending x
- reclamation
 - block storage systems 328
 - excluding 331
 - recommendations 329, 330
 - storage systems 330
 - viewing volumes 329
 - volumes 328, 331
- Reclamation
 - storage 61
 - volumes 61
- Red Hat 486
- removing 269, 275
- removing alerts
 - alerts
 - removing 193
- removing applications 269
- removing departments 275
- removing resources 132
- reporting
 - dashboard 197
- reporting problems xi
- resource alert definition
 - viewing 185
- resources
 - adding 112
 - adding agentless servers for monitoring 126
 - adding fabrics for monitoring 120
 - adding hypervisors for monitoring 122, 123, 124
 - adding rollup servers 132
 - adding storage systems for monitoring 115
 - adding switches for monitoring 120
 - adding to a capacity pool 301
 - manually adding resources to applications 267
 - modifying properties 333
 - removing 132
 - removing from a capacity pool 303
- restrictions
 - based on license 23
 - based on role 23
- role restrictions 23
- roles 70
- rollup mode
 - icon 66
 - master server 66
 - subordinate server 66
- rollup server connections 66
- rollup servers
 - adding 132
 - user roles 112

S

- SAN Volume Controller
 - publications vii
- script
 - performance data 237
- sending comments x
- server alerts
 - triggering conditions 394
- servers
 - adding
 - importing servers from a list 128
 - importing servers with agents from a file 130
 - adding by deploying an agent 130
 - adding without agents 126
 - modifying properties 333
 - optimized volumes, collocate 90
 - tiering, analyze 323
 - viewing data paths 287
 - viewing information in the GUI 194
 - viewing status in the web-based GUI 203
- service x
- service class 52, 294, 295, 297
 - deleting 298
 - satisfying custom requirements of a 299
 - viewing properties of 297
 - viewing storage constraints of 297
 - viewing users who can provision by using 297
 - viewing volumes or shares associated with 297
- service management connect x
- shares 147
- shortfall 253
- SMC x
- social media
 - blogs 15
 - videos 15
- software zoning 73
- sorting rows in lists 339
- space usage
 - block storage systems 238
 - containers 238
 - file storage systems 238
 - file system pools 238
 - file systems 238
 - filesets 238
 - monitoring 238
 - object storage systems 238
 - pools 238
 - tiers 238
 - volumes 238
- starting
 - IBM Spectrum Control web-based GUI 39
- status of resources
 - viewing the status of resources 209
- status of resources in the web-based GUI
 - acknowledging the status of a resource 210
 - acknowledging the status of a top-level resource 210
 - acknowledging the status of an internal resource 211

- status of resources in the web-based GUI *(continued)*
 - viewing the overall status of resources 208
- storage pools
 - tier level, setting 322
 - tiering, analyze 323
- storage provisioning
 - vSphere Web Client extension 63
- storage reporting
 - vSphere Web Client extension 63
- storage system alerts
 - triggering conditions 371
- storage systems
 - adding 115
 - exporting performance data 237
 - IBM Spectrum Scale, adding
 - openstack access, configuring 119
 - performance data collection, configuring 118
 - probe, configuring 118
 - sudo access, configuring 117
 - modifying properties 333
 - performance data, exporting 237
 - performance, comparing 102
 - user roles 112
 - viewing data paths 287
 - viewing information in the GUI 194
 - viewing status in the web-based GUI 203
- storage virtualizers 323
 - tiering, analyze 323
 - volumes, analyze placement 85
 - volumes, optimize 85
- subcomponents
 - creating subcomponents for existing applications 267
- subdepartments
 - creating 273
- subordinate servers
 - adding 66
 - removing 66
- support x
- Support
 - contact xi
- suppression actions 400
- suppressions for alerts 400
- switch alerts
 - triggering conditions 389
- switches
 - adding 120, 122
 - agent types for monitoring 122
 - modifying properties 333
 - viewing information in the GUI 194
 - viewing status in the web-based GUI 203

T

- tasks
 - analysis execution 348, 349
 - creating provisioning tasks in the web-based GUI 344
 - creating tiering storage, balancing pools, and transforming-storage tasks in the GUI 347
 - deleting in the web-based GUI 352

- tasks *(continued)*
 - exporting provisioning tasks in the web-based GUI 345
 - implementing optimization recommendations in the web-based GUI 348
 - managing in the GUI 340
 - managing provisioning tasks in the GUI 344
 - managing tiering storage, balancing pools, and transforming-storage tasks in the GUI 346
 - pausing, resuming, and canceling analysis-execution tasks in the web-based GUI 350
 - running provisioning tasks in the web-based GUI 344
 - running tiering analysis, transform-plan, and balance-analysis tasks in the web-based GUI 347
 - scheduling the implementation of optimization recommendations in the web-based GUI 349
 - setting the number of task runs displayed 343
 - viewing all tasks in the web-based GUI 340
 - viewing details in the web-based GUI 342
 - viewing in the GUI 340
 - viewing pending tasks 340
 - viewing task logs in the web-based GUI 342
 - viewing tasks for a resource in the web-based GUI 341
 - viewing tasks for a resource type in the web-based GUI 341

TB

- See* TiB
- temporary zone sets 73
- thresholds
 - viewing violations of performance 235
- TiB 69
- tier levels
 - storage pools, setting 322
- tiering
 - analysis period 332
 - I/O density 324
 - I/O rate 324
 - pools 321
 - pools, analyze 321
 - pools, assign tier levels 321
- tiering thresholds
 - I/O rate, modify 324
- Tivoli Storage Productivity Center
 - Cinder driver requirements 490
- Tivoli Storage Productivity Center for Replication 46
- topology view for data paths 288
- TPCD.cfg
 - saveNonRoot parameter, add 147
- trademarks 503
- translations
 - browser locale requirement ix
- triggering actions 401

- triggering conditions
 - alerts 370
 - computer alerts 394
 - fabric alerts 392
 - hypervisor alerts 385
 - server alerts 394
 - storage system alerts 371
 - switch alerts 389
- troubleshooting
 - block storage systems 238
 - capacity 238
 - containers 238
 - file storage systems 238
 - file system pools 238
 - file systems 238
 - filesystems 238
 - object storage systems 238
 - performance 217
 - pools 238
 - tiers 238
 - volumes 238

U

- Ubuntu 486
- units of measurement 69
- upgrading
 - configuring license compliance 33
- user roles
 - hypervisors 112
 - rollup servers 112
 - storage systems 112
 - VMware vSphere data sources 112
- users
 - assigning roles 70
- using 264
- REST API 77

V

- VASA provider
 - architecture 63
 - IBM Spectrum Control 62
- VASA provider for IBM Spectrum Control
 - architecture 63
- vCenter Server system
 - adding 122, 124
- videos 15
- viewing
 - scenarios
 - connectivity information 99
 - NPIV connections 99
 - viewing details of a department 275
 - viewing details of an application 268
- Viewing
 - the overall status of a storage environment 197
- viewing application details 268
- viewing department details 275
- viewing the status of resources 209
- VMware vSphere
 - provisioning 98
- volume collocation 90
- volumes
 - compressed, view capacity charts 238

- volumes (*continued*)
 - thin-provisioned, view capacity charts 238
 - tiering 87
 - tiering by I/O density 324
 - tiering by I/O rate 324
 - tiering, analyze 323
 - workload 326
- volumes, shares 311
- VSANs 74
- vSphere resources
 - managing 316, 318
- vSphere storage
 - provisioning 63
- vSphere storage reports
 - viewing 63
- vSphere Web Client extension
 - managing resources 316
 - provisioning block storage 316
 - provisioning file storage 318
 - provisioning storage 63
 - viewing storage reports 63

W

- web-based GUI
 - acknowledging the status of a resource 210
 - acknowledging the status of a top-level resource 210
 - acknowledging the status of an internal resource 211
 - alert severities 150
 - customizing lists 337
 - data collection
 - creating performance monitors 140
 - modifying performance monitors 142
 - starting performance monitors 144
 - stopping performance monitors 144
 - verifying that a performance monitor is running 142
 - verifying that a probe is running 136
 - exporting information about resources, jobs, and alerts 336
 - filtering 338
 - Monitoring the status of resources 203
 - Opening the management GUI for a resource 335
 - saving information about resources, jobs, and alerts 336
 - status of resources 203
 - tasks
 - analysis execution 348, 349
 - creating provisioning tasks 344
 - creating tiering storage, balancing pools, and transforming-storage tasks 347
 - deleting 352
 - exporting provisioning tasks 345
 - implementing optimization recommendations 348

- web-based GUI (*continued*)
 - tasks (*continued*)
 - pausing, resuming, and canceling analysis-execution tasks 350
 - running provisioning tasks 344
 - running tiering analysis, transform plan, and balance-analysis tasks 347
 - scheduling the implementation of optimization recommendations 349
 - viewing all tasks 340
 - viewing information about internal, object, and related resources 201
 - viewing information about top-level resources 200
 - viewing performance monitor logs 146
 - viewing performance monitors 145
 - viewing performance monitors for a resource 145
 - viewing task details 342
 - viewing task for a resource 341
 - viewing task for a resource type 341
 - viewing task logs 342
 - viewing the condition of top-level resources 208
 - viewing the overall status of resources 208
- Windows 496, 497
- workload
 - volumes 326
- WWN zoning 73

X

- XIVperformance metrics 434

Z

- zone 73
 - concepts 73
 - default 73
 - hardware 73
 - membership 73
 - non standard 74
 - port 73
 - software 73
 - supported switches 74
- zone set 73
 - temporary 73
- zones 72
- zoning
 - automatic 72
 - during provisioning 72



Product Number: 5725-F93, 5725-G33, 5725-Y23, 5725-Y24

Printed in USA

SC27-6588-06

