

IBM Hyper-Scale Manager for IBM Spectrum Accelerate Family

IBM XIV, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate

Lisa Martinez

Bert Dufrasne



Storage



International Technical Support Organization

**IBM Hyper-Scale Manager for IBM Spectrum Accelerate
Family**

April 2019

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

Fourth Edition (April 2019)

This edition applies to IBM FlashSystem A9000 and IBM FlashSystem A9000R with software Version 12.3 with IBM Hyper-Scale Manager V5.5.

© Copyright International Business Machines Corporation 2017, 2018. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	vii
Now you can become a published author, too!	viii
Comments welcome	viii
Stay connected to IBM Redbooks	ix
Chapter 1. Management tools	1
1.1 Management graphical user interface	2
1.1.1 Working with the Storage Management GUI	3
1.1.2 Systems & Domains Views	23
1.1.3 Statistics Views	24
1.1.4 Pools & Volumes Views	24
1.1.5 Hosts & Clusters Views	25
1.1.6 Remote Views	25
1.1.7 Access Views	25
1.1.8 Management Server	26
1.2 Installing and configuring Hyper-Scale Manager	26
1.2.1 Installing Hyper-Scale Manager	26
1.2.2 Starting the GUI	29
1.2.3 Using the installation wizard	30
1.2.4 Log in to IBM Hyper-Scale Manager after the initial setup	33
1.2.5 Adding a system after the initial setup	35
1.2.6 Removing a system or modifying a system IP address	36
Chapter 2. Storage management and configuration	37
2.1 Storage management interfaces	38
2.1.1 IBM Hyper-Scale Manager	38
2.2 Configuring pools	39
2.2.1 Creating pools	39
2.2.2 Managing storage pools with the GUI	44
2.2.3 Managing storage pools with the CLI	55
2.3 Volumes	57
2.3.1 Creating volumes	57
2.3.2 Managing volumes with the GUI	60
2.3.3 Managing volumes with the command-line interface	72
2.4 Snapshots	74
2.4.1 How to create a snapshot of a volume	75
2.4.2 Snapshot deletion priority	77
2.4.3 Managing snapshots with the GUI	78
2.4.4 Managing snapshots with the CLI	82
2.4.5 Snapshot Scheduler	83
2.5 Consistency groups	87
2.5.1 Creating consistency groups	88
2.5.2 Adding volumes to a consistency group	90
2.5.3 Managing consistency groups by using the GUI	91
2.5.4 Managing consistency groups with the CLI	96

2.5.5 Snapshot Groups	99
Chapter 3. Security	101
3.1 Physical access security	102
3.2 Native user authentication	102
3.2.1 Local credential repository	102
3.2.2 Managing user accounts	108
3.2.3 Security considerations for Hyper-Scale Manager	112
3.2.4 Managing user accounts by using the XCLI	114
Chapter 4. Multitenancy	119
4.1 Multitenancy	120
4.1.1 Domains	120
4.1.2 Domain attributes	121
4.1.3 Global administrators and domain administrators	121
4.1.4 Domain policy	121
4.2 Managing domains with the GUI	122
4.2.1 Creating domains	123
4.2.2 Customizing the Domain view	125
4.2.3 Adding a user to a domain	126
4.2.4 Domain properties	128
4.2.5 Managing domains	129
4.2.6 Managing within a domain	136
4.2.7 Removing a domain	141
4.3 Performance classes (quality of service)	143
4.3.1 Creating a performance class	144
4.3.2 Associating a domain with an existing performance class	145
Related publications	147
IBM Redbooks	147
Other publications and online resources	147
Help from IBM	149

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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

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

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

DS8000®	IBM Spectrum™	Redbooks (logo)  ®
HyperSwap®	IBM Spectrum Accelerate™	System Storage®
IBM®	MicroLatency®	XIV®
IBM FlashSystem®	Redbooks®	

The following terms are trademarks of other companies:

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

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

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

Preface

This IBM® Redbooks® publication describes storage management functions and their configuration and use with the IBM Hyper-Scale Manager management graphical user interface (GUI) for IBM XIV® Gen3, IBM FlashSystem® A9000 and A9000R, and IBM Spectrum™ Accelerate software.

The web-based GUI provides a revolutionary object-centered interface design that is aimed toward ease of use together with enhanced efficiency for storage administrators. The first chapter describes general features of the GUI and installation of the IBM Hyper-Scale Manager server. Subsequent chapters illustrate some typical GUI actions, among many other possibilities, to manage and configure the storage systems, to define security roles, and to set up multitenancy.

For most of the GUI-based actions that are illustrated in this book, the corresponding XIV Storage System command-line interface (XCLI) commands are also shown.

This edition applies to IBM Hyper-Scale Manager V5.4.

IBM Hyper-Scale Manager based GUI information regarding host attachment and replication is covered in *IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM XIV Storage System: Host Attachment and Interoperability*, SG24-8368 and *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401. See also *IBM HyperSwap and Multi-site HA/DR for IBM FlashSystem A9000 and A9000R*, REDP-5434.

Authors

This book was produced by a team of specialists from around the world.

Lisa Martinez has been working in the Washington Systems Center—Storage Team as a storage specialist since January 2012. Her focus has been with pre-sales support for IBM DS8000® and the IBM Spectrum Accelerate™ family of storage products (IBM FlashSystem A9000 and A9000R, XIV, and IBM Spectrum Accelerate). She is also the lead instructor for FlashSystem A9000 and A9000R, XIV, and IBM Spectrum Accelerate customer-based workshops. Prior experience includes roles as a storage architect in the Specialty Services Area in GTS and test architect in disk storage focusing on system level test for XIV for three years and copy services for DS8000. Lisa holds degrees in Computer Science from New Mexico Highlands University and Electrical Engineering from the University of New Mexico.

Bert Dufrasne is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage® disk products at the International Technical Support Organization (ITSO), San Jose Center. He has worked at IBM in various IT areas. He has authored many IBM Redbooks publications, and he has also developed and taught technical workshops. Before Bert joined the ITSO, he worked for IBM Global Services as an Application Architect. He holds a master's degree in electrical engineering.

Thanks to the authors of the previous edition:

Deepthi Anne
Roger Eriksson
Andrew Greenfield
Detlef Helmbrecht

Bruce Spell
Markus Oscheka
Ann Sargent
Stephen Solewin
Vadim Steckler
IBM

Special thanks to the following people for their contributions to this publication:

Randi Galron
John Hyams
Dan Ravid
Max Vidgerovich
Moshe Weiss
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author, all at the same time. Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run 2 - 6 weeks in length, and you can participate either in person or as a remote resident working from your home base.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

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

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

ibm.com/redbooks

- Send your comments in an email:

redbooks@us.ibm.com

- Mail your comments:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Management tools

This chapter provides an introduction to the Hyper-Scale Manager (HSM) management graphical user interface (GUI) for IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software.

The web-based GUI provides a revolutionary object-centered interface design that is aimed toward ease of use together with enhanced efficiency for storage administrators. The interface has the following features:

- ▶ Customizable dashboard
- ▶ Integrated real-time statistics
- ▶ A single-page application so that you can view all relevant information for every object at a glance

The following additional information is included in this chapter:

- ▶ Installing the Hyper-Scale Manager GUI
- ▶ Logging in to the GUI
- ▶ Adding a system to the inventory
- ▶ Removing a system or changing the IP address of a system in the inventory

Important: The material that is presented in this chapter was prepared with an early version of the Hyper-Scale Manager V5.5 user interface (UI). Some contents and layout might differ from the version that you are using.

1.1 Management graphical user interface

The Hyper-Scale Manager, which was introduced for use with FlashSystem A9000 and A9000R, provides a new web-based management GUI for all members of the IBM Spectrum Accelerate family, including the XIV Storage System Gen3, FlashSystem A9000 and A9000R, and any IBM Spectrum Accelerate software installations. By using the HSM GUI, you can manage and monitor one or more supported IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, or IBM Spectrum Accelerate systems in real time from a web browser.

The following functional highlights are included:

- ▶ Customizable dashboard
- ▶ Context-oriented UI in a single-page web-based application so that you can view all relevant information for every object at a glance
- ▶ Smart view of object relationships and dependencies in a visual map (for example, a view of all volumes that are mapped to a host)
- ▶ Instant object-centered management and monitoring
- ▶ Advanced filters for focusing on the required object
- ▶ Quick tracing of objects and fast navigation between objects
- ▶ One-click selection and operation options (map, unmap, delete, and so on)
- ▶ Health score of all systems in the inventory
- ▶ Integrated statistics information

The management GUI is accessed by directing your web browser at the Hyper-Scale Manager server (`http://<Hyper-ScaleManager server IP address>:8080` or for secure `https://<Hyper-ScaleManager server IP address>:8443`).

Figure 1-1 shows the Hyper-Scale Manager login window.

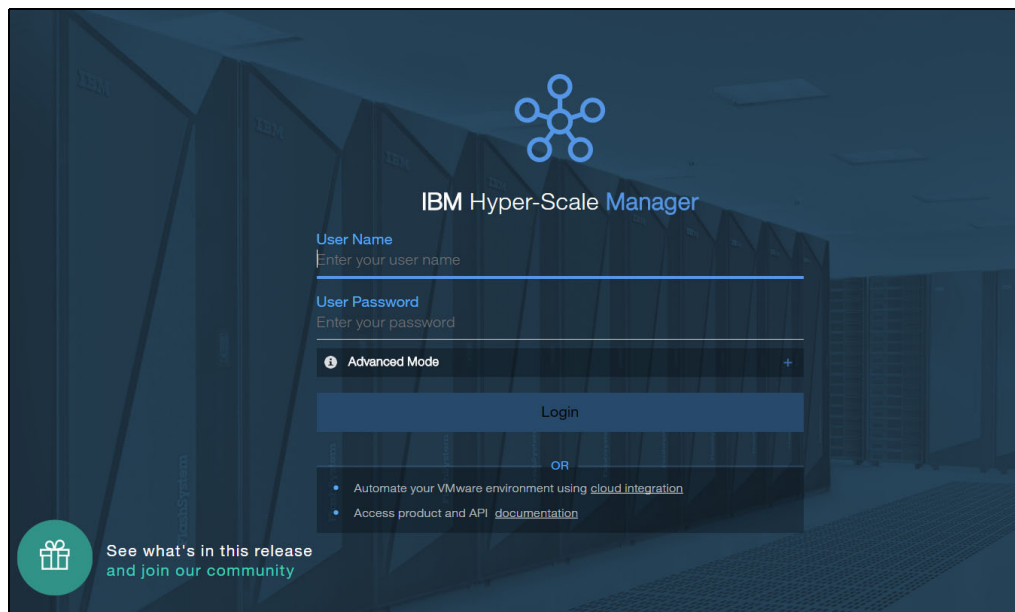


Figure 1-1 IBM Hyper-Scale Manager login window

IBM Hyper-Scale Manager is also the portal for access to IBM Hyper-Scale Manager Cloud Integration and the product and application programming interface (API) documentation.

Tip: If Hyper-Scale Manager is not installed and configured in your environment, see 1.2, “Installing and configuring Hyper-Scale Manager” on page 26 for summarized instructions.

The Advanced Mode option is not present on the logon pane if the Manager Server security option was set to “Don’t restrict access”. For more information, see 3.2.3, “Security considerations for Hyper-Scale Manager” on page 112.

1.1.1 Working with the Storage Management GUI

The Storage Management GUI is made up of static and dynamic elements. The static elements consist of the top navigation pane and the side menu. The remainder of the display is dynamic and controlled by the user.

When you initially enter the Storage Management GUI, it displays the Dashboard view by default. (See “Dashboard” on page 5.)

Figure 1-2 shows the *New Tabs* view, which is represented by the white rectangle on the left. You can use search and filter options to display storage objects, such as spools, volumes, hosts, and users, according to your search criteria.

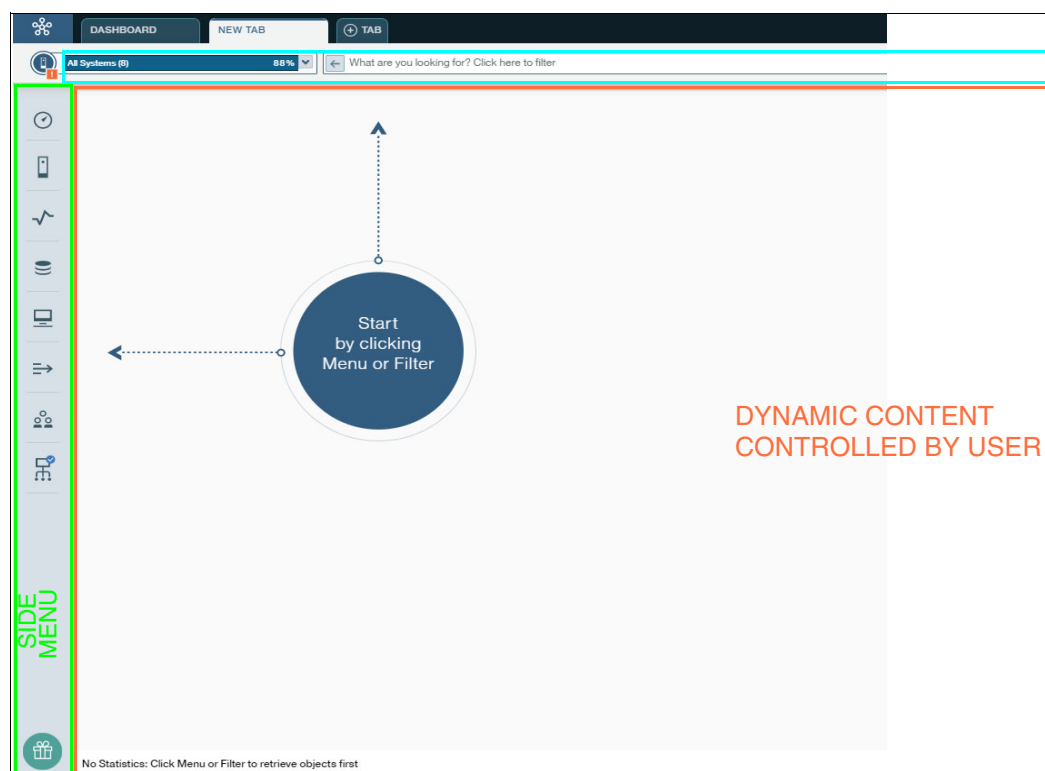






Figure 1-2 Static and dynamic elements of IBM Hyper-Scale Manager

GUI icons

Additional icons are used throughout the GUI. Running your mouse over the following icons reveals additional information. See Table 1-1.






Table 1-1 GUI icons and descriptions





Icon	Description
	Active alerts
	Additional information required
	Missing information
	Additional information provided

Top navigation pane

The top navigation pane of the GUI gives you quick access to system information, configuration changes, and system documentation. See Table 1-2.

Table 1-2 Top navigation pane icons, their names, and descriptions

Icon	Name	Description	Quick links for more information
	Dashboard	For the overall health of attached inventory	"Dashboard" on page 5
	New Tab	For customized searches of the attached inventory	"Tabs" on page 9
	Alerts	For displaying information about the highest severity issues of the attached inventory	"Alerts" on page 18
	Commands Log	Displays the history of XCLI commands that are used during the active session	
	New Item	For quickly adding new items, such as volumes, users, and hosts, to the attached inventory	"New item" on page 20

Icon	Name	Description	Quick links for more information
	Switch to dark	For quickly switching between the dark and light dashboard views of the GUI. With HSM V5.4, this icon is in the upper left corner of the dashboard.	"Dashboard" on page 5
	Settings	For changing the Hyper-Scale Manager settings	"Settings" on page 20
	Help	For additional GUI information and a link to the IBM Knowledge Center for product information	"Help" on page 21
	User	For logging out of the GUI	"User" on page 21

Dashboard

The Dashboard gives you an overall inventory status view of all the systems health, data reduction savings, provisioned capacity, and real-time statistics, as shown in Figure 1-3. You can filter to see just one type of hardware (such as FlashSystem A9000), as shown in Figure 1-4 on page 6.

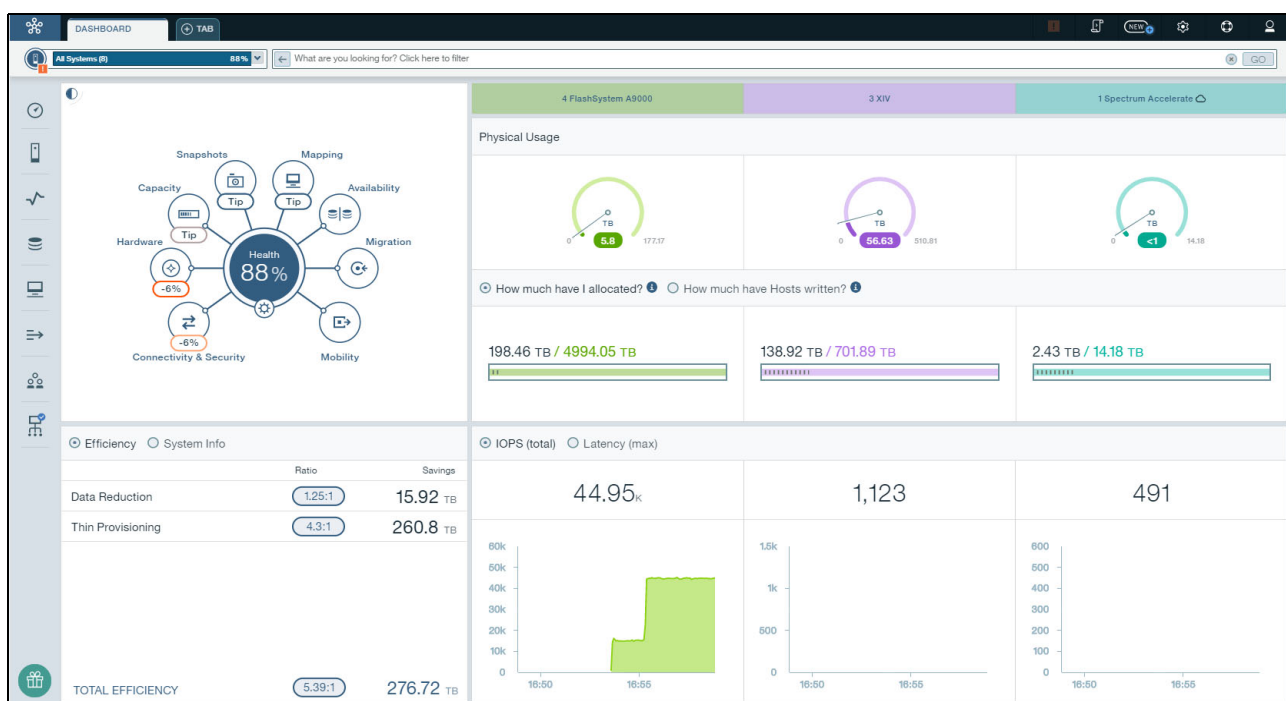


Figure 1-3 Dashboard for all systems that are configured in HSM

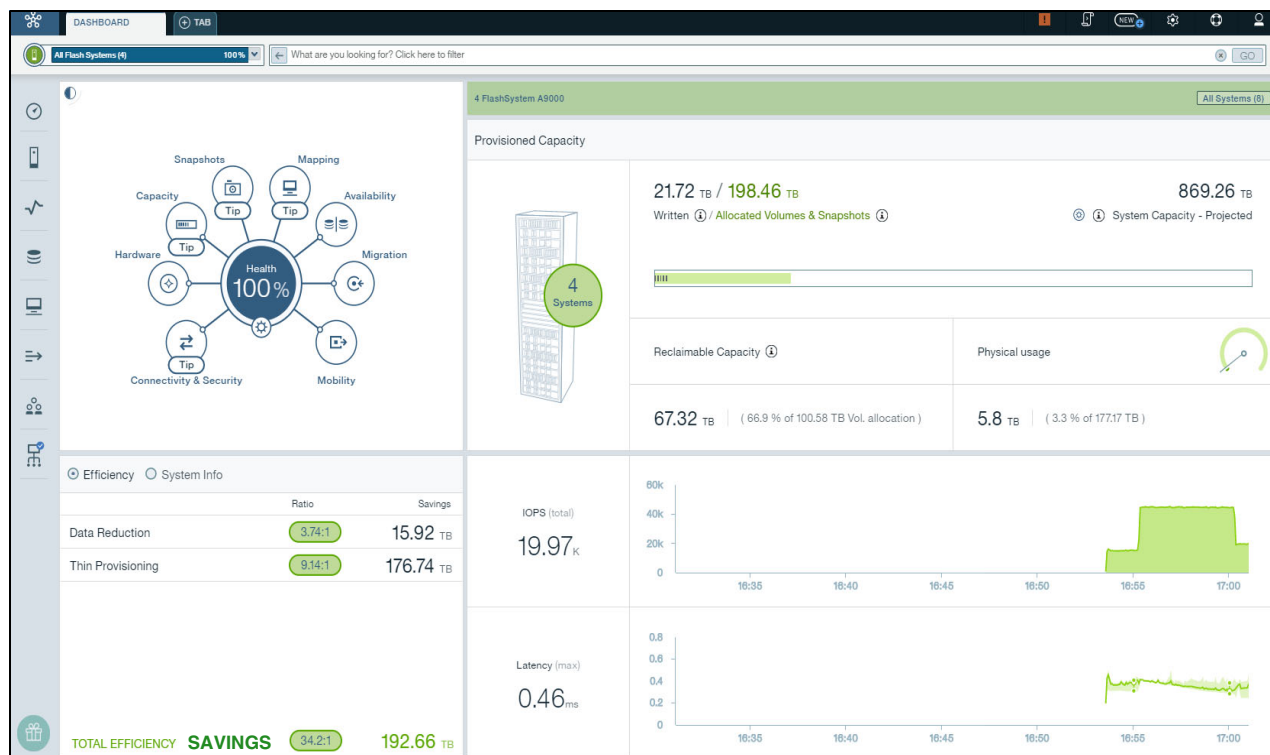



Figure 1-4 Dashboard for one type of system

Note: Click **Switch to Dark** icon  in the upper left corner of the dashboard to change the color scheme of the Dashboard display.

The following information is available in the Dashboard view.

Health widget (Hub)

The Health widget, which is shown in Figure 1-5, displays the number of systems in the inventory and a high-level view of the health of the managed inventory.

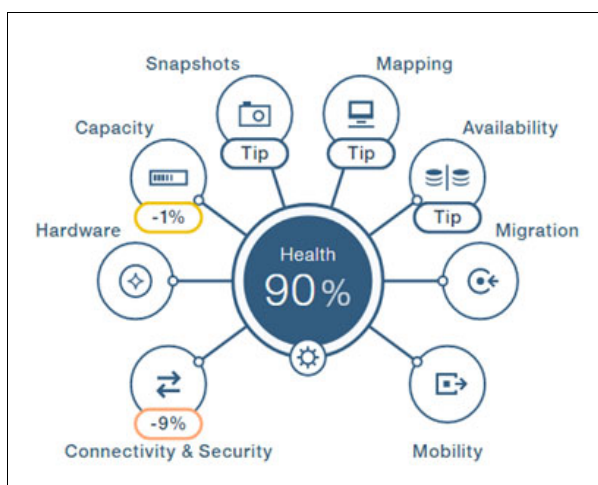


Figure 1-5 Dashboard Health widget

If the health of the system is less than 100%, the corresponding elements indicate a warning or alert with a small number in the circle. Clicking the element that indicates an issue displays additional information about the issue. Clicking one of the alerts opens the table view with the associated information and actions, as shown in Figure 1-6. You can also view available tips, as shown in Figure 1-7.

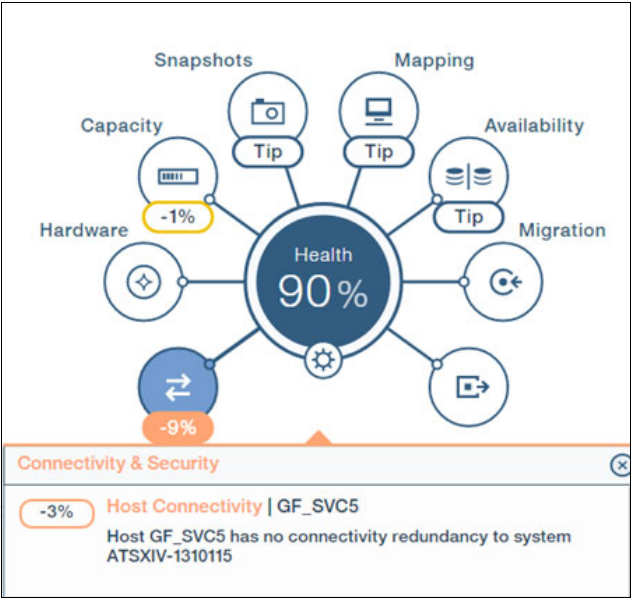


Figure 1-6 Dashboard Health widget that shows the system with issues

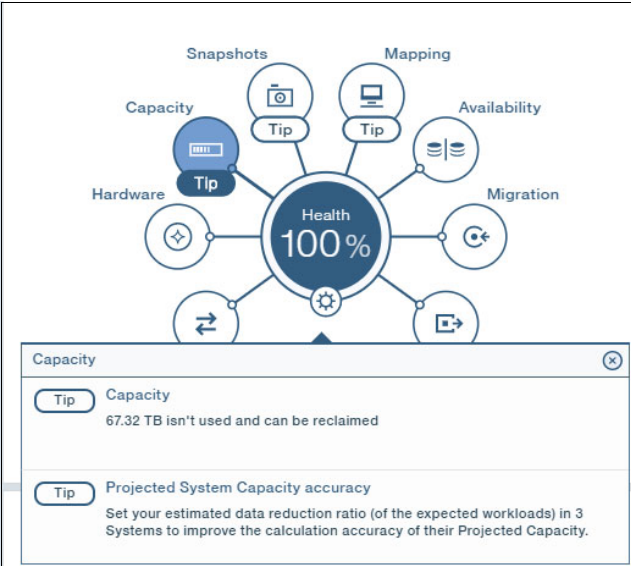


Figure 1-7 Display available tips

Total Efficiency

The Efficiency section of the Dashboard displays the total efficiency as a result of data reduction, which includes compression and deduplication, and thin provisioning of written data, as shown in Figure 1-8.

Efficiency System Info		
	Ratio	Savings
Data Reduction	1.25:1	15.94 TB
Thin Provisioning	4.3:1	260.8 TB
TOTAL EFFICIENCY		
	5.39:1	276.74 TB

Figure 1-8 Dashboard Data Reduction Saving section

Capacity

The Capacity section of the Dashboard displays current and projected information about various capacity aspects of all, or a subset, of the systems in the inventory. A subset of the systems in the inventory is shown in Figure 1-9.

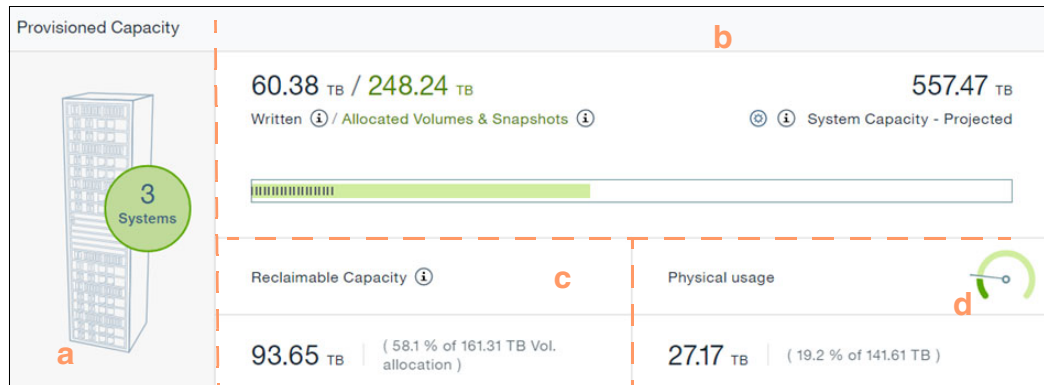


Figure 1-9 Dashboard Capacity section

The following information is contained in the Capacity section of the Dashboard, as indicated by the corresponding letters in Figure 1-9. This view applies when looking specifically at the same type of systems. The view for all systems in the inventory, if there are different types (including FlashSystem A9000, XIV, and IBM Spectrum Accelerate) look different, as shown in Figure 1-10 on page 9.

- ▶ Number of systems in the inventory (a)
- ▶ Written, allocated, and projected system capacity of the systems in the inventory (b)
- ▶ Reclaimable capacity of the systems in the inventory (c)
- ▶ Physical usage of the systems in the inventory (d)

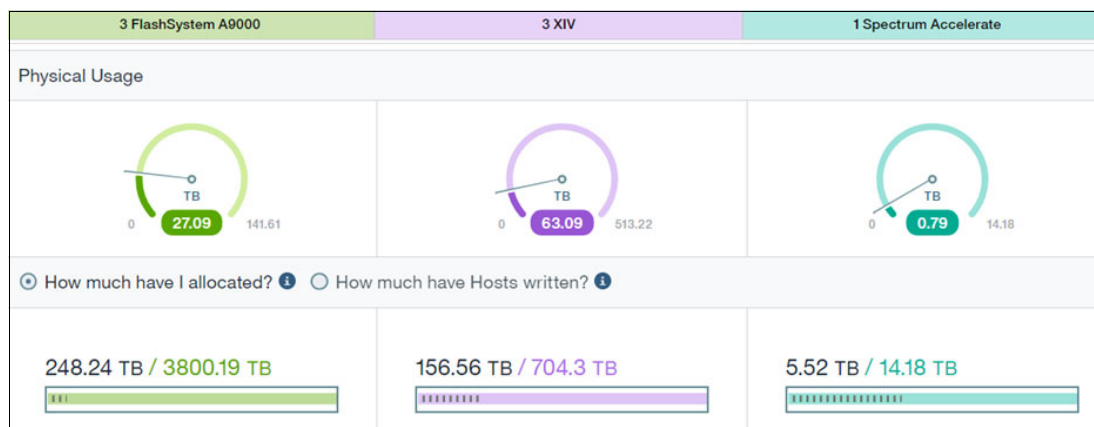


Figure 1-10 Physical usage view for all systems in the inventory

For more detailed information about data reduction savings, see the chapter about *Capacity management and planning in IBM FlashSystem A9000 and A9000R Architecture and Implementation (Version 12.3.1)*, SG24-8345.

Real-time statistics

The real-time statistics section of the Dashboard displays latency and input/output operations per second (IOPS) information for all systems in the inventory. See Figure 1-11. The graphs are grouped by color based on the type.

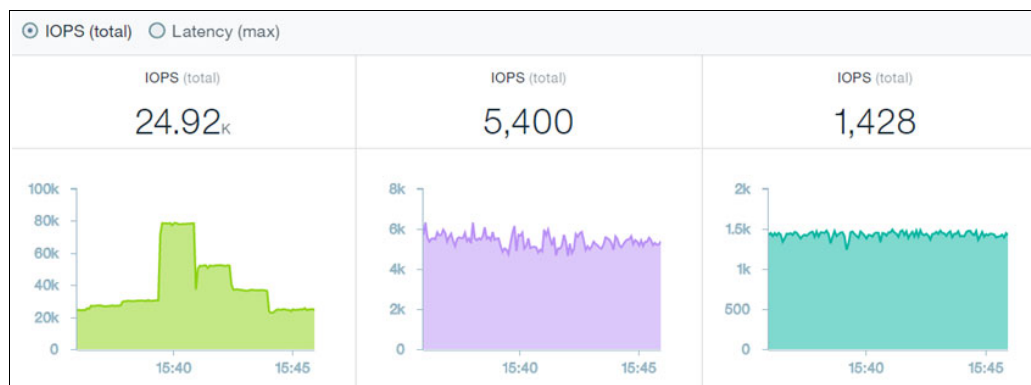


Figure 1-11 Dashboard real-time statistics section

Tabs

Information that is displayed in a tab is determined by choices that are made by the user. The following capabilities and information are available in most of the tabs:

- Smart filtering:
 - Search anything in your environment in seconds by any parameter or condition.
 - Use smart filtering for policy control to find non-valid objects by searching.
 - Additional wildcard filtering by using an asterisk (*) and OR.
- Contextual table views:
 - Can be configured with any column and mixture among mapping, mirroring, and snapshots.
 - No pre-made views with separate focus points.
 - The user builds a view with their own interests and tasks.

- ▶ Flexible layout:
 - Ability to shrink and grow the table view.
 - Maximize the inspector area.
- ▶ Integrated statistics:
 - Inside the GUI in real time. Helps in daily decision making.
 - Shows the statistics of a chosen, historic or current element, by comparing or drilling down.
- ▶ Hub relations:
 - Mini-dashboard for chosen objects.
 - Ability to troubleshoot and monitor objects quickly.
- ▶ Actions:
 - Always visible for any chosen object.
 - User can immediately change, modify, or add anything and view properties of any object, whether aggregated or single.
 - Fast modification with one click.

Getting started with tabs

Clicking the **TAB** icon in the top navigation pane causes a new tab to open and it displays a large blank window with a navigation menu that shows how to get started in a new tab (the “Start by clicking Menu or Filter” circle). After a menu item or filter is selected, the window in the left portion of the main window that will display information according to filters that are selected by the user. The right part of the window displays information that is associated with the filters that are selected by the user.

Start by clicking **Menu** or **Filter** to select one of the menu items from the side menu or use the filter to begin drilling down to the information that you want, as shown in Figure 1-12. You can maximize the left side of the window by clicking the “Expand Table” arrow. Additionally, in V5.4 you can drag the separator either left or right to adjust the amount of information that is seen in either the table or inspector view.



Figure 1-12 New tab view

Selecting filters

Clicking in the filter field in the upper-left portion of the table view displays a list of commonly used filters, as shown in Figure 1-13.

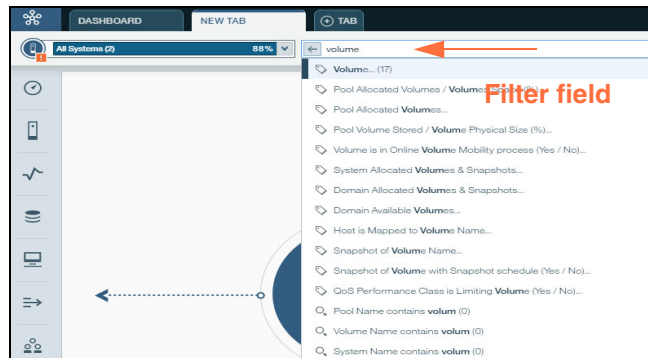


Figure 1-13 Filters list

Select one of the filters (for example, Volume) in the list to display additional searchable subfilters or to view all matching items for the selected filter.

You can use the subfilters to quickly drill down to the information that you want. In this example, the System can be selected from the drop-down list, then other filters can be applied with. See Figure 1-14.

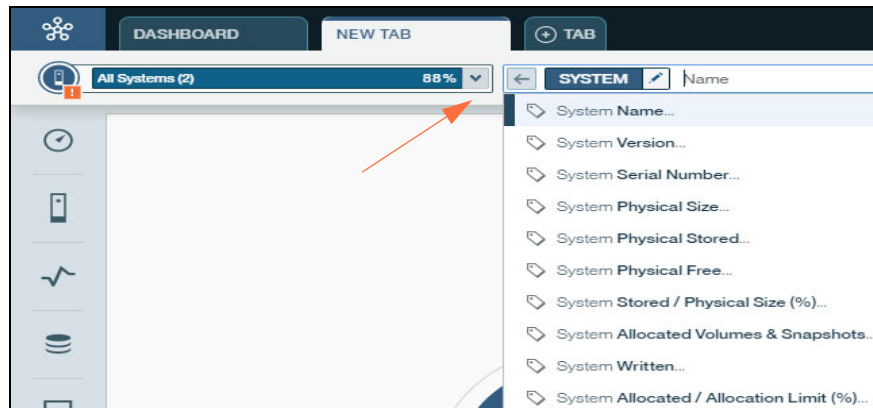


Figure 1-14 System list and filters

In addition, you can now anchor the view by selecting a system or system type from the system list. Everything that you select, either in the filter or from the side menu, is specific to what you anchored in the system list. An example is shown in Figure 1-15, where all the flash systems were selected from the list and then the systems view was selected from the side menu.

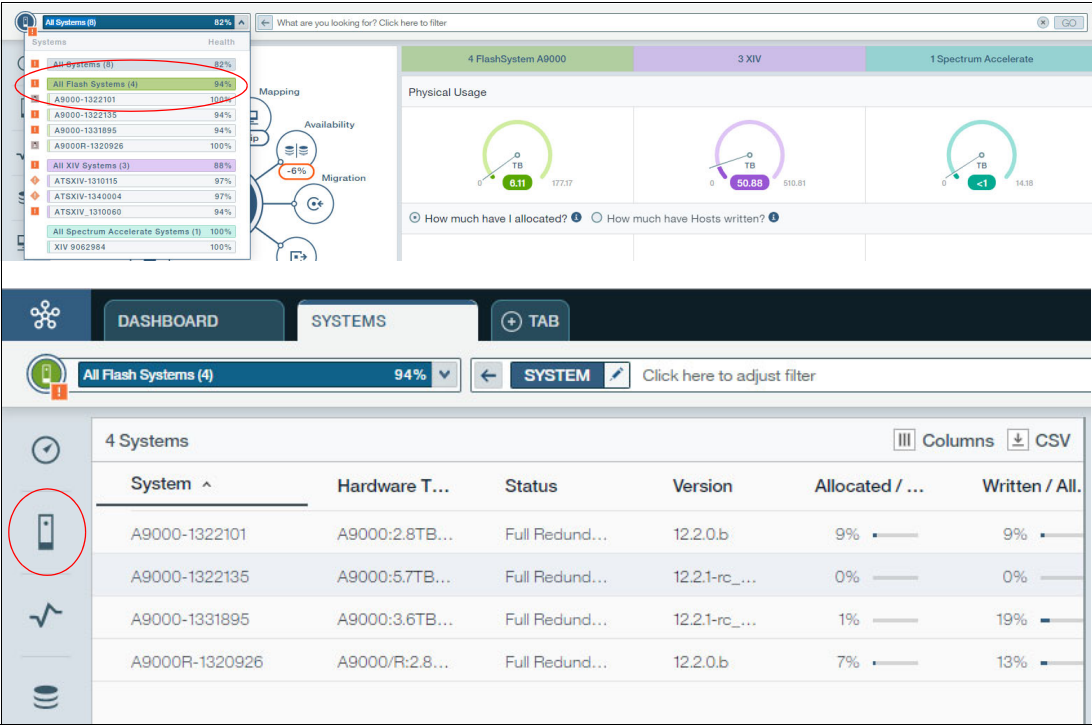


Figure 1-15 System anchor

Wildcards can be used in the search field by using the asterisk (*) to search for more than one option. You can see the options in the drop-down menu (see Figure 1-16) or by typing the wildcard option in the filter.




Figure 1-16 Filter search options

Figure 1-17 shows an example of typing in a search with a wildcard before selecting a filter. This type of search shows all the matching objects based on the entered search criteria. This example shows all objects that begin with the letter “t.”

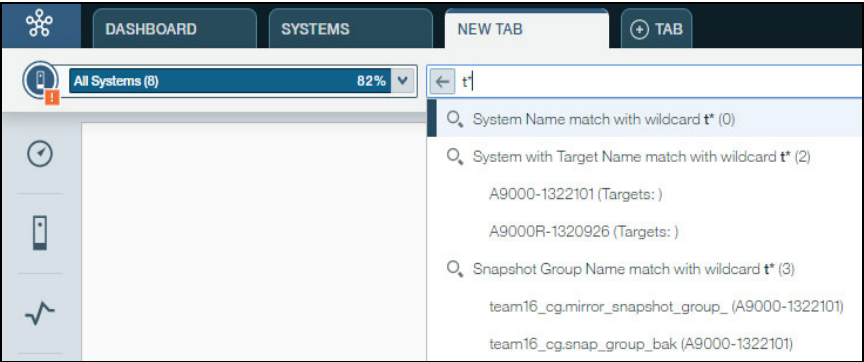


Figure 1-17 Search with a wildcard

Use the drop-down list to select a system in the inventory. With each additional filter selection, an additional field displays so that you can drill down to the information that you want. Remember to use the scroll bar on the right side of the list to see all of the available systems. See Figure 1-18.

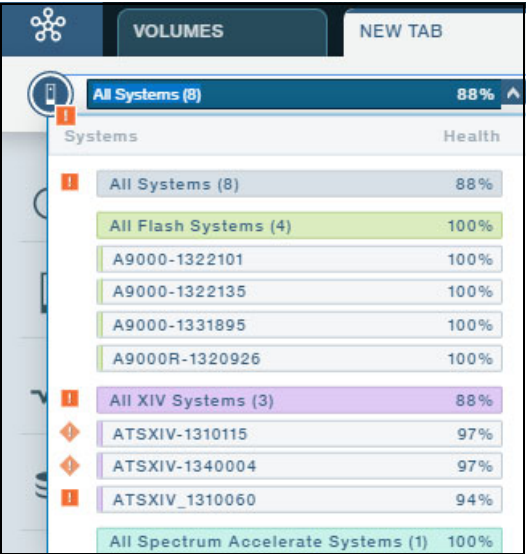


Figure 1-18 System selector drop-down menu

When all the filters are selected, click the Go icon to display the results of the search or press the Enter key. In versions before 5.4, the Search icon or a magnifying glass icon is used instead of the Go icon.

If available, the lower portion of the table view displays statistical information that is based on the selection that was made. The capability to access the historic or current statistics of a chosen element for comparisons or for drilling down helps you in daily decision making. This is shown in Figure 1-19. You can have two statistical views open at the same time, under the table view.

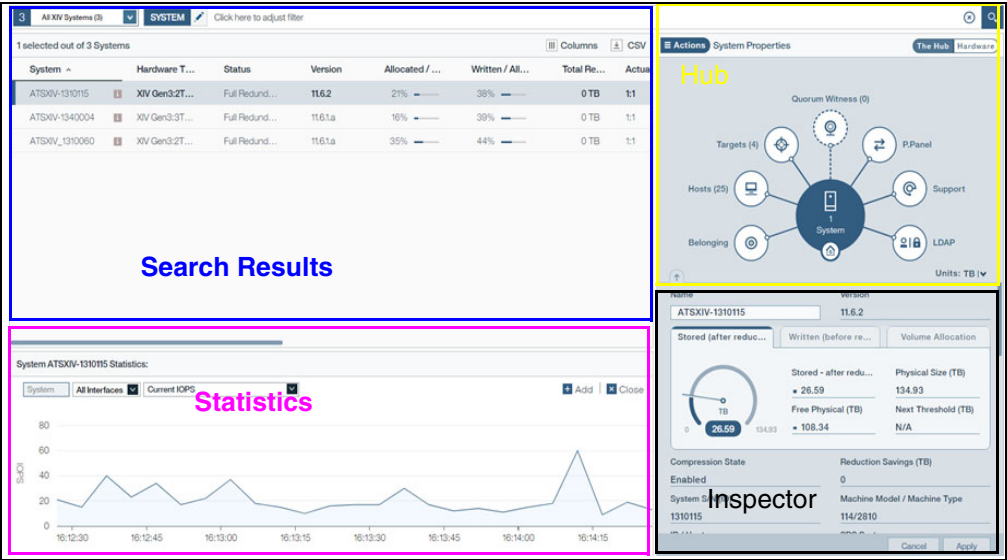


Figure 1-19 Table view sections

You can also compare the statistics of related elements by selecting several elements in the list that is displayed.

You can use the statistics view in Version 5.4 to separate host latency from what is seen internally in the FlashSystem A9000 and A9000R from the time that is spent from the host to the system. This feature does require Version 12.2.1 on the systems to use it. This feature is found under Historical Latency when a host is selected. An example is shown in Figure 1-20.

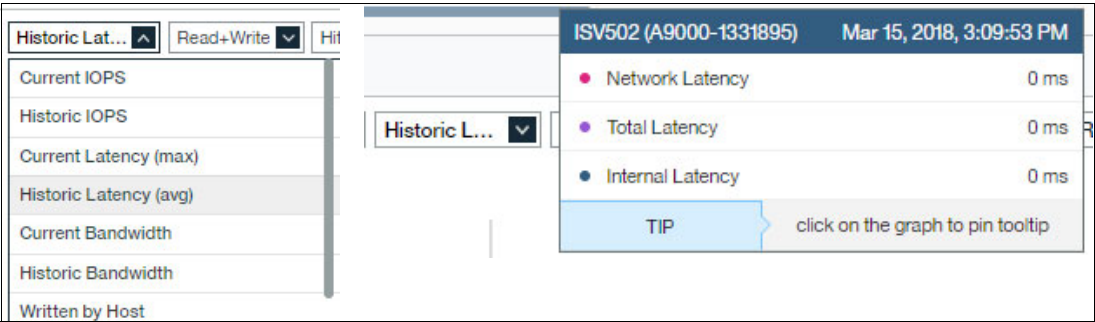


Figure 1-20 Finding network and system latency

With HSM Version 5.5 and system software Version 12.3, you can now, at a system level, separate internal latency from external latency. This view can be obtained by a hovering the mouse cursor over the latency graph in the dashboard view, as shown in Figure 1-21.

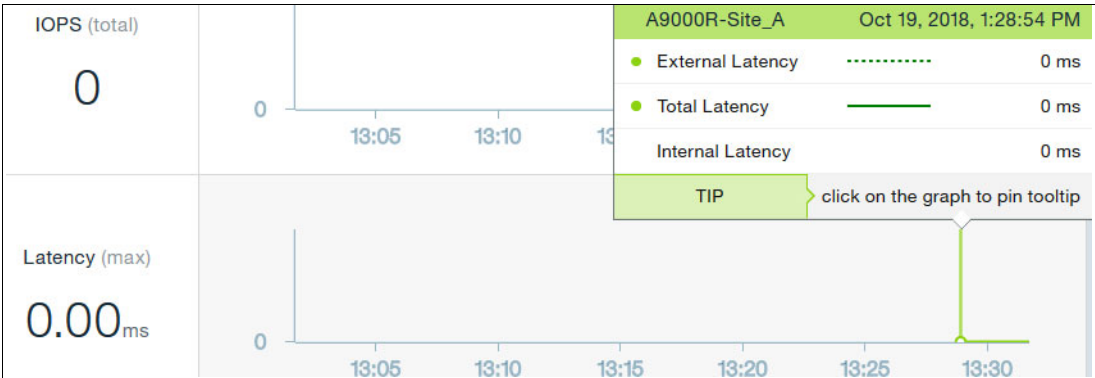


Figure 1-21 System level internal and external latency

Working with the Statistics view

Historical statistics have other filters for filtering to more specific information. The Range menu gives you the option to choose specific time frames, as shown in Figure 1-22.



Figure 1-22 Current and historic time range filter

Additional filters are available for historical statistics so that you can drill down to specific information. See Figure 1-23.

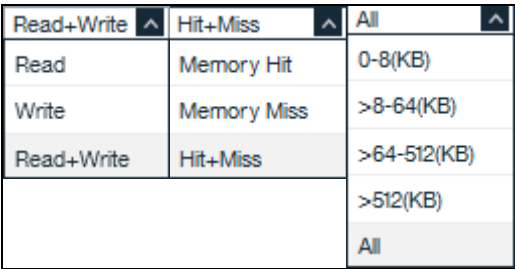


Figure 1-23 Historical statistics filters

The statistical views include current and historic IOPS; latency and bandwidth; and physical usage, physical usage + forecast, written/allocated, written/projected, written/projected + forecast, and projected/actual. New with Version 5.4, you can get a breakdown of the total latency, and separate network latency from storage system internal latency.

After you collect the current or historical data, a graphical representation of the data appears. Click the plus sign (+) icon in the upper-right corner to open a second view to allow comparative statistical analysis. Click the x icon in the upper-right corner to close the views. An example of this is shown in Figure 1-24.

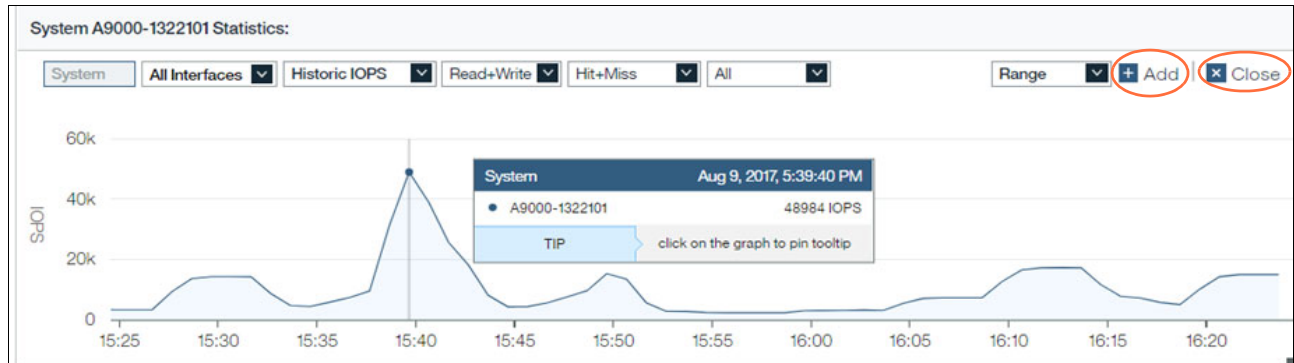


Figure 1-24 Cloned statistics section

Note: If the browser is zoomed out (Ctrl -) and maximized, you might be able to open a maximum of three different statistical graphs.

Working with the Hub view

The Hub section of the table view gives you quick access to related information about the selected item.

Each element of the Hub displays related information about the selected item below the Hub (the inspector), as shown in Figure 1-25. A dashed circle indicates that there are no associations with this category.

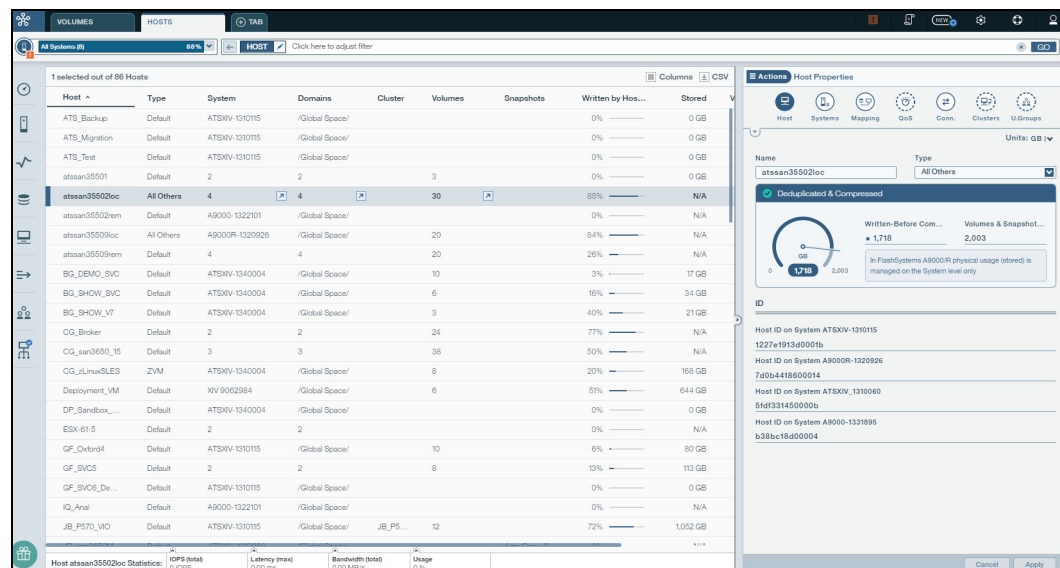


Figure 1-25 Hosts view

Beginning with Hyper-Scale Manager V5.3, when the HUB is minimized, the elements are displayed as icons on a single line, as shown in Figure 1-25.

Working with the System Hardware view

You can use the System Hardware view to click a component in the graphical representation of the hardware and drill down to specific information about that component. Hovering the cursor over any of the components in the hardware displays what that component is, as seen in Figure 1-26. Clicking the component opens the details about it.

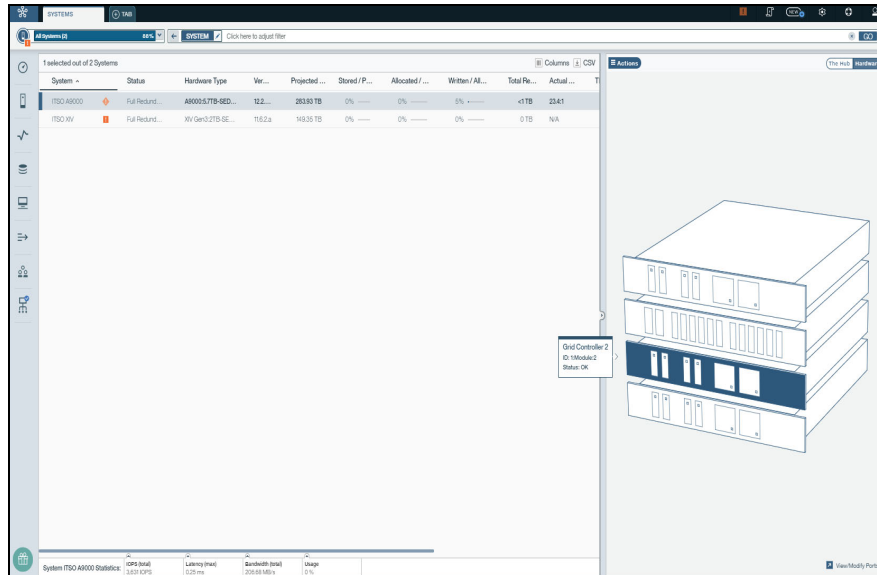


Figure 1-26 System Hardware view

Working with the Actions menu

Use the table view to perform certain actions against the selected system. The actions vary depending on the filtered selection. For an example of a list of actions for a system, see Figure 1-27. You can either select the Actions menu in the right pane or right-click an object in the table view to open the same menu.

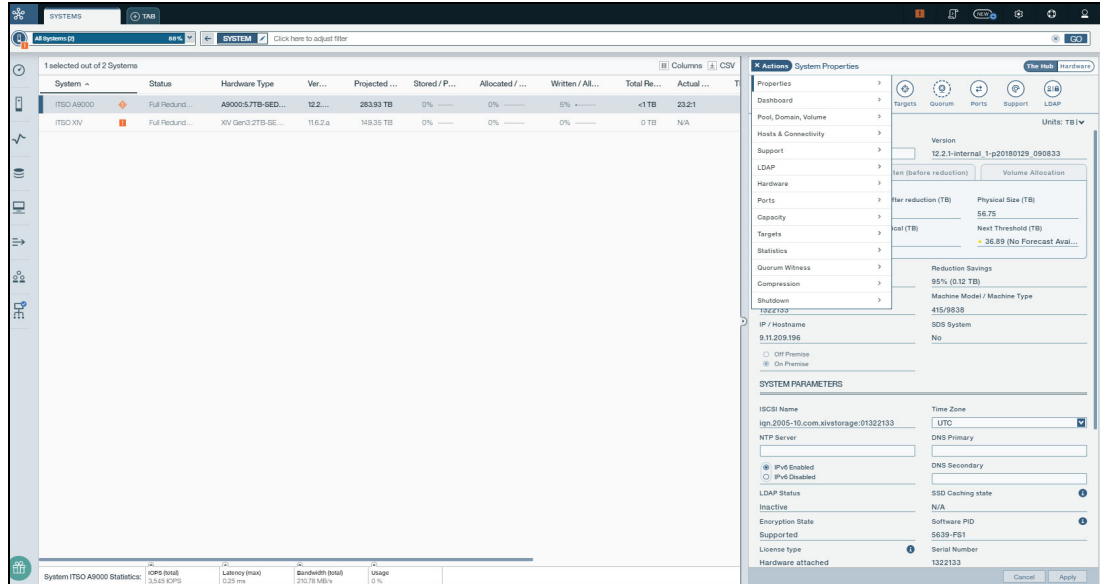


Figure 1-27 System Actions list

Tip: To close the Actions menu, click the Actions menu again or click anywhere outside the menu. When opening the Actions menu with a right-click, you can close it by clicking outside the menu or with the Esc key.

Alerts

The Alerts window displays information about all of the alerts for the systems in the inventory. If there are active alerts on any system that is managed by the Hyper-Scale Manager, an icon displays on the top menu bar, as shown in Figure 1-28. The icon shape might change based on the severity of the alerts. Clicking the icon opens a quick view of the active alerts.

All Alerts (2)			Date
Minor	Spectrum_Accelerate: is using default admin password		20/10/16 11:24
Warning	XIV_Gen3: Pool XIV_Notes is 81% full		20/10/16 11:24

Figure 1-28 Alerts window

With Hyper-Scale Manager V5.4, the active alerts can also be accessed by using the icon next to the system list and within the system list. An example is shown in Figure 1-29.

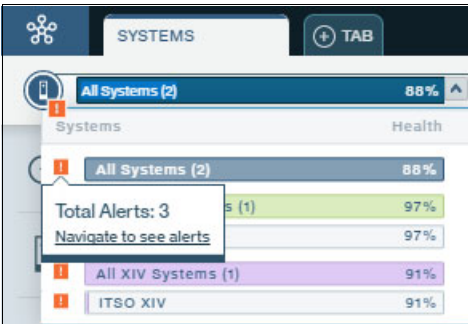


Figure 1-29 System alerts

Click an alert in the list to begin filtering to additional information about the alert. To close the All Alerts view, click the icon again or anywhere outside the view.

Commands log

The Hyper-Scale Manager issues `xcli` commands to the system when the administrator creates or modifies an object. Those commands are shown for an active session under the commands log icon, as shown in Figure 1-30.

#	Command	Description	System
2	vol_create vol="testvol_002" size_blocks="19531250...	Adding Volume testvol_002	A9000R-1320926
1	vol_create vol="testvol_001" size_blocks="19531250...	Adding Volume testvol_001	A9000R-1320926

Figure 1-30 Commands log

The commands that were run through the Hyper-Scale Manager can then be copied or exported in CSV format for later use.

Reminder: Commands log is generated for an active session. It is cleared with a logout from the Hyper-Scale Manager.

New item

Click the **New** icon to access a list of items that can be created. See Figure 1-31.



Figure 1-31 Create New window

Clicking any of the items in the list opens the table view. The table view displays a section for entering the required information to add the item. For more information see Chapter 2, “Storage management and configuration” on page 37.

Settings

This window is used by a storage administrator to access the Management Server configuration, the Cloud Integration server, and the Capacity Trend/Forecast Report. This menu is shown in Figure 1-32.




Figure 1-32 Settings window for changing preferences

The Management Server option is available for quick access to the Inventory view where systems can be added or deleted. This option is used to maintain Management Server-related access information. If access is restricted, the user must log in with advanced mode credentials to perform these actions.

The Cloud Integration option opens the Cloud Integration server in another browser window. This is a version of IBM Spectrum Connect that is installed with the Hyper-Scale Manager.

The Capacity Trend/Forecast Report option generates and downloads a compressed PDF file for all the systems in the inventory when it is selected from this menu.

Help

The Help menu contains videos and documentation links for assistance with the Hyper-Scale Manager, links to IBM Knowledge Center and access to the Hyper-Scale API documentation. Help is accessed through the top menu bar by using the Help icon, as shown in Figure 1-33. You also see a file on what is new in the version that is installed and tips on navigating and using the Hyper-Scale Manager and a direct link to download the XCLI. Beginning with V5.3, you also have an option to view details about the Hyper-Scale Manager.

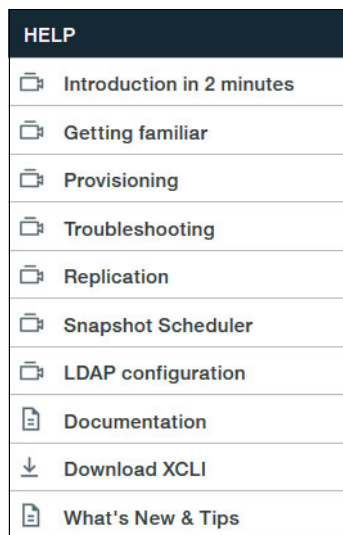


Figure 1-33 Help menu

User

Use the User icon to sign out of the GUI and return to the IBM Storage Manager login window. See Figure 1-34.




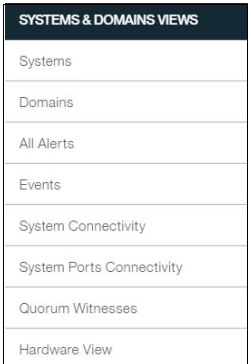



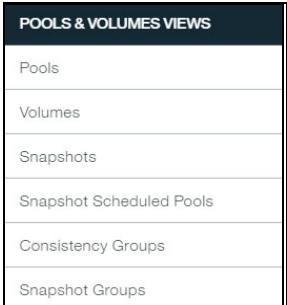







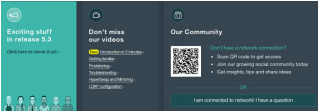
Figure 1-34 Sign Out window

Side menu

Use the Side menu for quick access to the following filters that are listed in Table 1-3.

Table 1-3 Side menu icons and menus

Icons	Menus	Quick links for more information
		This icon opens the main dashboard view
		1.1.2, “Systems & Domains Views” on page 23
		1.1.3, “Statistics Views” on page 24
		1.1.4, “Pools & Volumes Views” on page 24

Icons	Menus	Quick links for more information
	<div>HOSTS & CLUSTERS VIEWS</div> <div>Mapped Volumes</div> <div>Mapped Snapshots</div> <div>Hosts</div> <div>Clusters</div> <div>Host Connectivity</div> <div>Quality of Service</div>	1.1.5, “Hosts & Clusters Views” on page 25
	<div>REMOTE VIEWS</div> <div>Replication Overview</div> <div>Replication Details</div> <div>Hyper-Scale Mobility</div> <div>Migrating Volumes</div> <div>Mirror / Mobility Connectivity</div> <div>Migration Connectivity</div>	1.1.6, “Remote Views” on page 25
	<div>ACCESS VIEWS</div> <div>Users</div> <div>User Groups</div>	1.1.7, “Access Views” on page 25
	<div>MANAGEMENT SERVER</div> <div>Manager Configuration</div> <div>Add Systems to Monitor</div> <div>Security (recommended)</div>	1.1.8, “Management Server” on page 26
		Connect to the online community

1.1.2 Systems & Domains Views

Choosing any of the systems or domains views displays the associated data in the table view:

- ▶ **Systems:** Displays a list of the systems in the inventory.
- ▶ **Domains:** Displays a list of all domains in all systems in the inventory.
- ▶ **All Alerts:** Displays a list of all alerts from all systems in the inventory.
- ▶ **Events:** Displays a list of all events from all systems in the inventory.
- ▶ **System Connectivity:** Displays all target connections of all systems in the inventory.
- ▶ **System Ports Connectivity:** Displays connectivity at the port level.

- ▶ **Quorum Witnesses:** Displays details about the quorum witnesses that are used in IBM HyperSwap®.
- ▶ **Hardware View:** Displays a list of all systems in the inventory and capacity information.

Clicking the top portion of the table view reveals more filter options for drilling down to specific information.

1.1.3 Statistics Views

Choosing any of the statistics filters displays associated data and statistics in the table view:

- ▶ **System & Ports Performance Statistics:** Displays all systems in the inventory and associated interface performance statistics in a graphical format.
- ▶ **Volume Performance Statistics:** Displays all volumes in all systems in the inventory and associated volume performance statistics in a graphical format.
- ▶ **Host & Ports Performance Statistics:** Displays all hosts in all systems in the inventory and associated performance statistics in a graphical format.
- ▶ **Quality of service (QoS) Performance Statistics:** Displays QoS performance class information from all systems in the inventory and associated QoS statistics in a graphical format.
- ▶ **System Capacity History & Forecast:** Displays all systems in the inventory and associated interface capacity history and future forecast.
- ▶ **Pool Capacity History & Forecast:** Displays all pools in the inventory and associated interface capacity history and future forecast.
- ▶ **Host Capacity History & Forecast:** Displays all hosts in the inventory and associated interface capacity history and future forecast.

Clicking the column titles of the table view allows you to add and remove columns from the table view.

1.1.4 Pools & Volumes Views

Choosing any of the pools or volumes filters displays associated data in the table view:

- ▶ **Pools:** Displays all pools in all systems in the inventory.
- ▶ **Volumes:** Displays all volumes in all systems in the inventory and associated volume statistics in a graphical format.
- ▶ **Snapshots:** Displays all snapshots from all systems in the inventory and snapshot statistics in a graphical format.
- ▶ **Snapshot Scheduled Pools:** Displays all the pools that include a Snapshot schedule for all the systems in the inventory.
- ▶ **Consistency Groups:** Displays all consistency groups in all systems in the inventory.
- ▶ **Snapshot Groups:** Displays a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

To add and remove columns from the table view, click the table view column tiles.

1.1.5 Hosts & Clusters Views

Choosing any of the physical host or cluster filters displays the associated data in the table view:

- ▶ **Mapped Volumes:** Displays all volumes in all systems in the inventory and associated volume statistics in a graphical format.
- ▶ **Mapped Snapshots:** Displays all mapped snapshots in all systems in the inventory and associated volume statistics in a graphical format.
- ▶ **Hosts:** Displays all physical hosts and clusters in the inventory and associated physical host statistics in a graphical format.
- ▶ **Clusters:** Displays all clusters in the inventory and associated cluster statistics in a graphical format.
- ▶ **Host Connectivity:** Displays all physical hosts and their connection type in the inventory and physical host statistics in a graphical format.
- ▶ **Quality of service (QoS):** Displays all QoS information in the inventory.

To add and remove columns from the table view, click the table view column tiles.

1.1.6 Remote Views

Choosing any of the remote filters displays the associated data in the table view:

- ▶ **Replication Overview:** Displays all mirrored volumes and consistency groups in the inventory and their relation type.
- ▶ **Replication Details:** Displays details for mirrored volumes and consistency groups.
- ▶ **Hyper-Scale Mobility:** Displays all online Hyper-Scale Mobility volumes in the inventory and their associated statistics.
- ▶ **Migrating Volumes:** Displays the volumes that are currently in a migration process.
- ▶ **Mirroring / Mobility Connectivity:** Displays connectivity information for the selected system.
- ▶ **Migration Connectivity:** Displays target parameters for the selected system.

To add and remove columns from the table view, click the table view column tiles.

1.1.7 Access Views

Choosing the access filter displays associated data in the table view:

- ▶ **Users:** Displays all users with access to systems in the inventory. The amount of system access is determined by the user category that is assigned to each user by the storage administrator.
- ▶ **User Groups:** Displays a group of application administrators who share a set of snapshot creation permissions.

1.1.8 Management Server

To access the Management Server menu, the storage administrator must be logged in advanced mode if the server configuration is restricted. This menu includes the following options:

- ▶ **Manager Configuration:** Select this option to access the Inventory (Home) window where the storage administrator can enter the Hyper-Scale Manager Access Code, System Machine Account, and Notifications configuration to a specific system in the inventory.
- ▶ **Add System to Monitor:** Add a system to the inventory.
- ▶ **Security (recommended):** Select this option to modify restricted access, access code, and Auto Logout Time.

For more information about security with the Management Server, see 3.2.3, “Security considerations for Hyper-Scale Manager” on page 112.

1.2 Installing and configuring Hyper-Scale Manager

To start the web browser-based Storage Management GUI, the Hyper-Scale Manager server must be installed and configured.

1.2.1 Installing Hyper-Scale Manager

This procedure describes how to install the Hyper-Scale Manager application.

Before you begin

Download the installation file from Fix Central:

<http://www.ibm.com/support/fixcentral/>

The name of the installation file is IBM-Hyper-Scale-Manager_x.x.x-x.bin.

Important: Keep up with the latest version to support new functions in the FlashSystem A9000 and A9000R and the latest updates in the Hyper-Scale Manager.

Prerequisites

The following list shows the installation prerequisites:

- ▶ The installation takes place on a 64-bit Red Hat Enterprise Linux (RHEL) 6.x or later 64-bit operating system or CentOS 6.8 or later.
- ▶ The server requires the following components and memory:
 - At least two CPU cores and 4 GB of RAM
 - At least 1 GB of disk space in the /tmp directory
 - 76 GB free space under the installation path (The path is configurable, and it does not have to be the home directory.)

Important: To determine the right CPU and RAM specifications for your environment, refer to the release notes that are included with the Hyper-Scale Manager on Fix Central.

Determining that the requirements are met

To determine the amount of available disk space, enter the following command:

```
$ df -h %DIRECTORY%
```

If less than 1 GB of disk space and less than 76 GB of free disk space are available in the /tmp installation directory, follow these steps:

1. Delete unnecessary files from the /tmp/ installation directory to meet the disk space requirement.
2. Extend the file system that contains the /tmp/ installation directory. If necessary, contact the system administrator for information about extending file systems.

Important:

Note the following advice:

- ▶ The installation must be performed by the root user or a user with root privileges (sudo).
- ▶ For the best performance, install Hyper-Scale Manager on a dedicated, independent server.
- ▶ The Hyper-Scale Manager can be installed on a physical or virtual server.

Installing

Complete the following steps to install Hyper-Scale Manager:

1. Transfer (usually File Transfer Protocol (FTP) or Secure Copy Protocol (SCP) works) the downloaded installation package file into your RHEL or CentOS server by using the **root** user name.
2. Log in to your RHEL or CentOS server and change to the directory that contains the installation package file. Run the following command (Example 1-1) to verify that run permissions are set on the installation package file.

Example 1-1 Command for setting run permissions

```
$ chmod +x IBM-Hyper-Scale-Manager_5.5.bin
```

3. Run the command that is shown in Example 1-2 to start the installation.

Example 1-2 Command to start the installation

```
# ./IBM-Hyper-Scale-Manager_5.5.bin
```

4. Accept the installation location or select another location and press Enter. Enter y to approve the installation, as shown in Example 1-3, or n to cancel the installation.

Example 1-3 Approving the installation

```
IBM Hyper-Scale Manager is going to be installed under a designated user
'msms'.
Press 'Enter' to accept the installation location [/home/msms] or a different
path to override >
Press any key to continue the installation
IBM Hyper-Scale Manager will be installed under /home/msms/hyperscale.
Please approve to continue with the installation [y/n] >y
```

5. Accept the license agreement.

6. When the installation is complete, verify that all of the required ports are opened on the operating system (OS) firewall:
 - Outbound only: Simple Mail Transfer Protocol (SMTP): Port 25
 - XIV Gen3, FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software communication: Port 7778
 - Web GUI: Port 8443 (HTTPS) or 8080 (HTTP)

Upgrading

It is possible to upgrade Hyper-Scale Manager by using a similar process to the installation process:

1. Transfer (usually FTP or SCP works) the downloaded installation package file into your Red Hat Enterprise Linux (RHEL) management server by using the root user name.
 2. Log in to your management server with the root user and make the file executable with the command `chmod +x IBM-Hyper-Scale-Manager_5.x.bin`, as shown in Example 1-1 on page 27.
 3. Run the `./IBM-Hyper-Scale-Manager_5.x.bin` command that is shown in Example 1-2 on page 27 to start the installation. If an older version of the GUI exists on the server, a menu shows three options, as shown in Example 1-4:
- Upgrade
 - Uninstall
 - Exit

Example 1-4 Menu with upgrade option

```
-----
-----  IBM Hyper-Scale Manager 5.x  -----
-----

IBM Hyper-Scale Manager version 5.x is already installed on this server.
Please select one of the following:
1) Upgrade
2) Uninstall
3) Exit
Please enter your choice > 1

IBM Hyper-Scale Manager will be upgraded from version 5.x to 5.y.
Are you sure you want to upgrade now? [y/n]
```

Note 1: The upgrade option only works from 5.0.x to 5.1.x or higher. You cannot upgrade from earlier versions of Hyper-Scale Manager. What you can do is to transfer your already configured systems from the old Hyper-Scale Manager to the new installation.

However, be aware that the transfer does not merge with the existing configuration, it replaces it.

To do the systems configuration transfer from pre-V5.x Hyper-Scale Manager, complete the following steps:

1. Copy the `systems.xml` file from `<$MSMS_HOME>/server/data` to the same path in the 5.x server installation.
2. Restart 5.x HSM service with `service xivms stop` and `service xivms start`.

Note 2: The `systems.xml` file is editable in a text editing tool, so if you want to merge, it is possible to copy and paste two files together. But no automatic process is created.

1.2.2 Starting the GUI

To connect and log in to the GUI, complete the following steps:

1. Open one of the following internet browsers:
 - Google Chrome version 49 or later
 - Mozilla Firefox version 45 or later
2. Connect with one of the following URLs:
 - `http://<hostname>:8080`
 - `https://<hostname>:8443`
 - `http://<IP address>:8080`
 - `https://<IP address>:8443`
3. The Hyper-Scale Manager Login window opens. Enter your credentials and click **Login** to enter the Management GUI (Use the default user name `admin` and the default password `adminadmin` unless the defaults were previously changed).

1.2.3 Using the installation wizard

When you start IBM Storage Manager for the first time, an installation wizard welcome window opens. The wizard is used to apply the initial system configuration. Follow these steps:

1. Click **Let's Start**, as shown in Figure 1-35, to apply the initial system configuration.

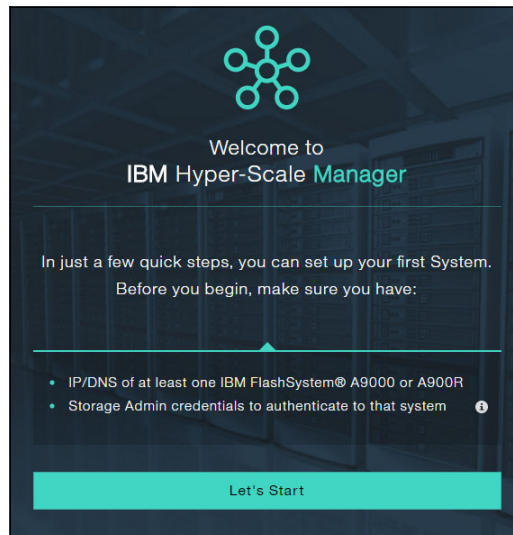


Figure 1-35 Installation wizard welcome window

Note: These screen captures were taken with a Hyper-Scale Manager version before V5.3. If installing V5.3 or later, the dialog boxes appear slightly different, but the main steps are the same.

2. Add your first system by entering at least one Internet Protocol (IP)/Domain Name System (DNS) address to add that system to your inventory. Click **Continue** to proceed to the next window. See Figure 1-36.




Figure 1-36 Installation wizard window for adding a system

3. Enter the default user credentials and click **Continue** to begin your setup. The default *admin* user name comes with the storage administrator role. See Figure 1-37.

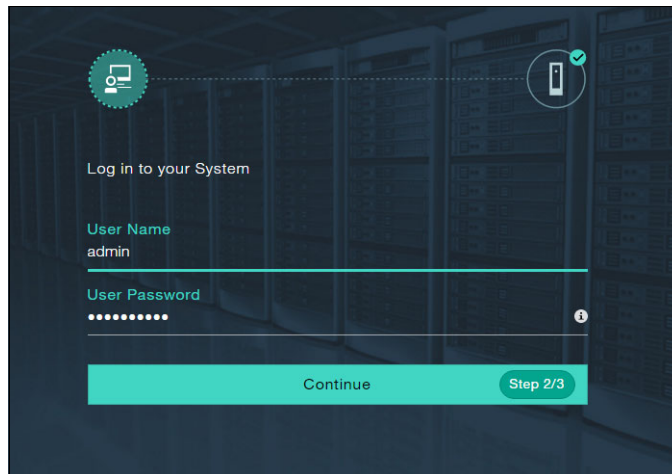
The image shows a login window for the installation wizard. At the top, there is a progress bar with three icons: a server rack, a smartphone, and a smartphone with a checkmark. Below the progress bar, the text "Log in to your System" is displayed. Underneath, there are two input fields: "User Name" with the value "admin" and "User Password" with a masked password ".....". A "Continue" button is located at the bottom right, and a "Step 2/3" indicator is next to it.

Figure 1-37 Installation wizard login window

Tip: The default user name and password with storage administrator access are shown:

- ▶ User name: admin
- ▶ Password: adminadmin

4. Any administrative action on IBM Hyper-Scale Manager that is performed from the GUI requires the Hyper-Scale Manager access code. See Figure 1-38.

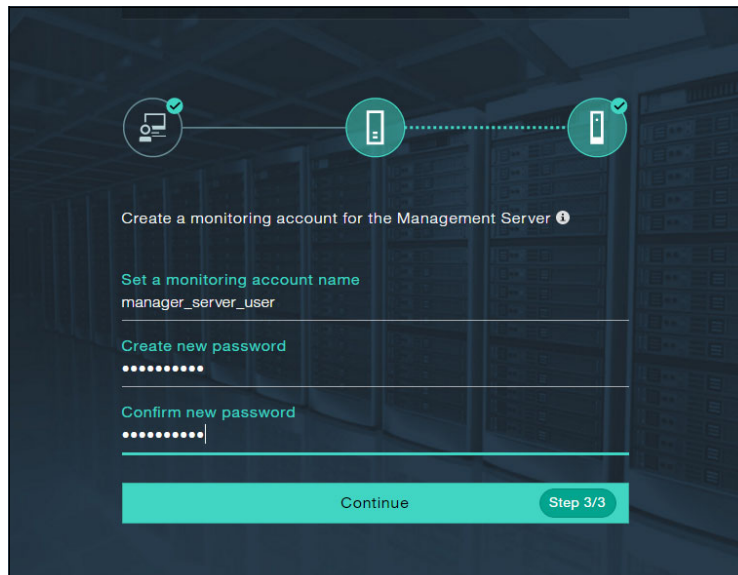
The image shows a window for creating a monitoring account for the Management Server. At the top, there is a progress bar with three icons: a server rack, a smartphone, and a smartphone with a checkmark. Below the progress bar, the text "Create a monitoring account for the Management Server" is displayed. Underneath, there are three input fields: "Set a monitoring account name" with the value "manager_server_user", "Create new password" with a masked password ".....", and "Confirm new password" with a masked password ".....". A "Continue" button is located at the bottom right, and a "Step 3/3" indicator is next to it.

Figure 1-38 Installation wizard window to define the Hyper-Scale Manager access code

Note: If multiple Hyper-Scale Manager servers are installed and manage the same storage systems, make sure to use the same user credentials for the monitoring account on all of your management servers.

This account information is used by Hyper-Scale Manager to connect to IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software systems:

- This user must be defined.
- This user does not change the configuration.
- This user's name default is `manager_server_user`, and it can be edited.
- This user can be defined in Lightweight Directory Access Protocol (LDAP). (Ensure that this user is added to all IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software systems storage admin groups in the LDAP). If you are using an LDAP domain, enter the name and password that exists in the LDAP server. If your LDAP server requires a domain, use UPN notation (for example, `name@domain.ibm.com`)
- This user must have a storage administrator role (similar to the admin user).
- This user must be defined with the same password on all IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software systems in the Hyper-Scale Manager inventory.

Define a System Machine Account and password to connect to and monitor your system inventory from Hyper-Scale Manager. Click **Continue** to complete the initial setup. See Figure 1-39.

Tip: The default System Machine Account name is `manager_server_user`. The storage administrator (admin) defines the password.

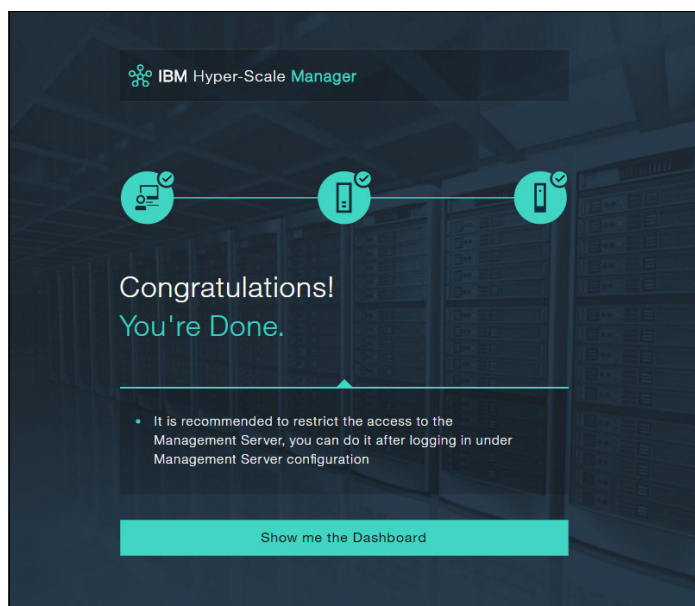


Figure 1-39 Installation complete wizard window

1.2.4 Log in to IBM Hyper-Scale Manager after the initial setup

Enter the URL for the IBM Hyper-Scale Manager server in one of the supported web browsers:

`http://<Hyper-Scale Manager server IP address>:8080`

Alternatively, enter the following URL for secure connectivity:

`https://<Hyper-Scale Manager server IP address>:8443`

The login window is displayed, as shown in Figure 1-40. Enter your assigned user name and password, then click **Login**.

The default user name is admin. The default password is adminadmin. The user with these credentials can log on in advanced mode and further configure Hyper-Scale Manager as explained in 1.2.5, “Adding a system after the initial setup” on page 35 and 1.2.6, “Removing a system or modifying a system IP address” on page 36. The administrator should also change the default password.

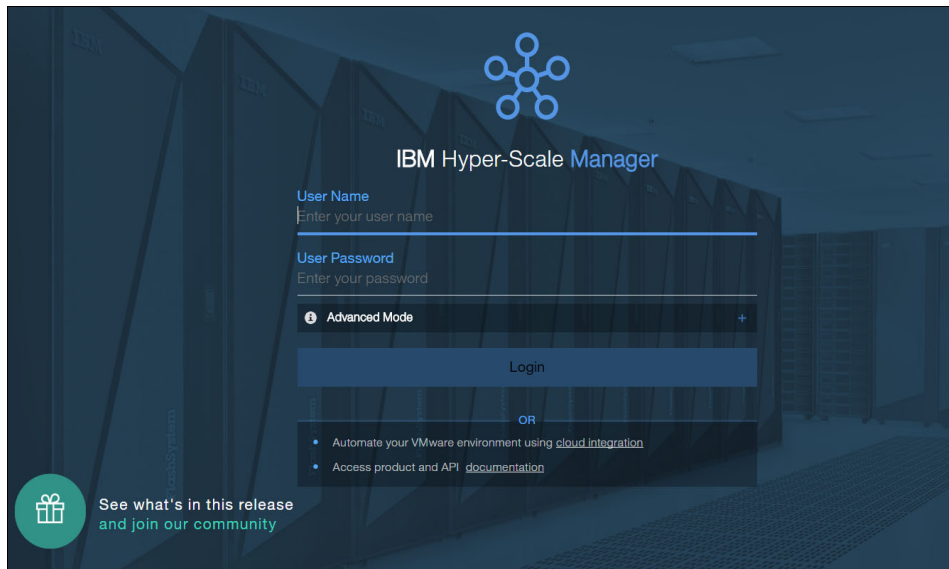
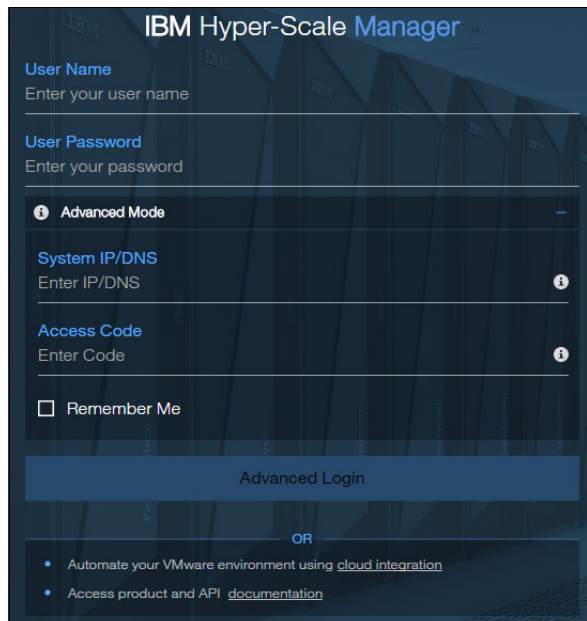


Figure 1-40 IBM Hyper-Scale Manager login window

It is possible to restrict advanced access. When advanced access is restricted, you must specifically log on in advanced mode to further configure the Hyper-Scale Manager settings. Figure 1-41 shows the Advanced Mode login window.



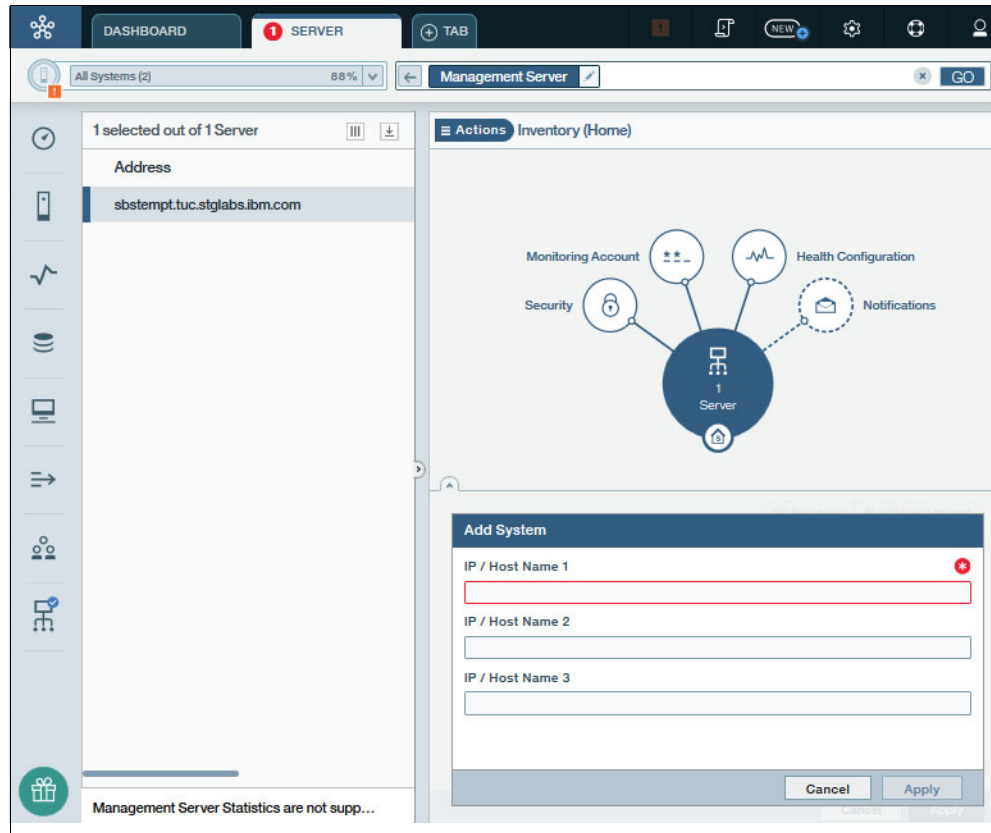
The screenshot displays the 'IBM Hyper-Scale Manager' login interface in 'Advanced Mode'. The background is a dark blue with a faint architectural pattern. The title 'IBM Hyper-Scale Manager' is at the top in white. Below it, the 'User Name' field is labeled 'Enter your user name'. The 'User Password' field is labeled 'Enter your password'. A section titled 'Advanced Mode' with an information icon contains three fields: 'System IP/DNS' (labeled 'Enter IP/DNS'), 'Access Code' (labeled 'Enter Code'), and a 'Remember Me' checkbox. Below these fields is a large blue button labeled 'Advanced Login'. Underneath the button is the word 'OR' in blue. At the bottom, there are two bullet points: 'Automate your VMware environment using [cloud integration](#)' and 'Access product and API [documentation](#)'.

Figure 1-41 Advanced mode login window

1.2.5 Adding a system after the initial setup

If you are logged in, you can add a system to your inventory after the initial setup by completing the following steps:

1. Log in to the Hyper-Scale Manager.
2. Use the **+New** button and click **System**. See Figure 1-42.



3. Enter an IP or DNS address for the system that you are adding.
4. Click **System Machine Account** on the Hub to add or change the system machine account information.
5. Click **Notifications** on the Hub to enable and add users to receive notifications from the system that is added.
6. Click **Apply** to save the information that was entered. Click **Cancel** to not save the information that was entered.

Important: If you have enabled the security restriction, you must log in to the Hyper-Scale Manager by using the access code to add or remove systems from the inventory.

1.2.6 Removing a system or modifying a system IP address

When the storage administrator is logged on in advanced mode, the storage administrator can remove a system from the inventory or change the IP address of a system in the inventory. To remove a system from your inventory or to change the IP address of a system in your inventory, complete the following steps:

1. Log in to the Hyper-Scale Manager.
2. In the menu on the left, select **Hyper-Scale Manager**, select **Management Server**, then click **Manager Configuration**.
3. Click the small blue menu icon in the upper-right corner of the section that contains the system to remove. See Figure 1-43.

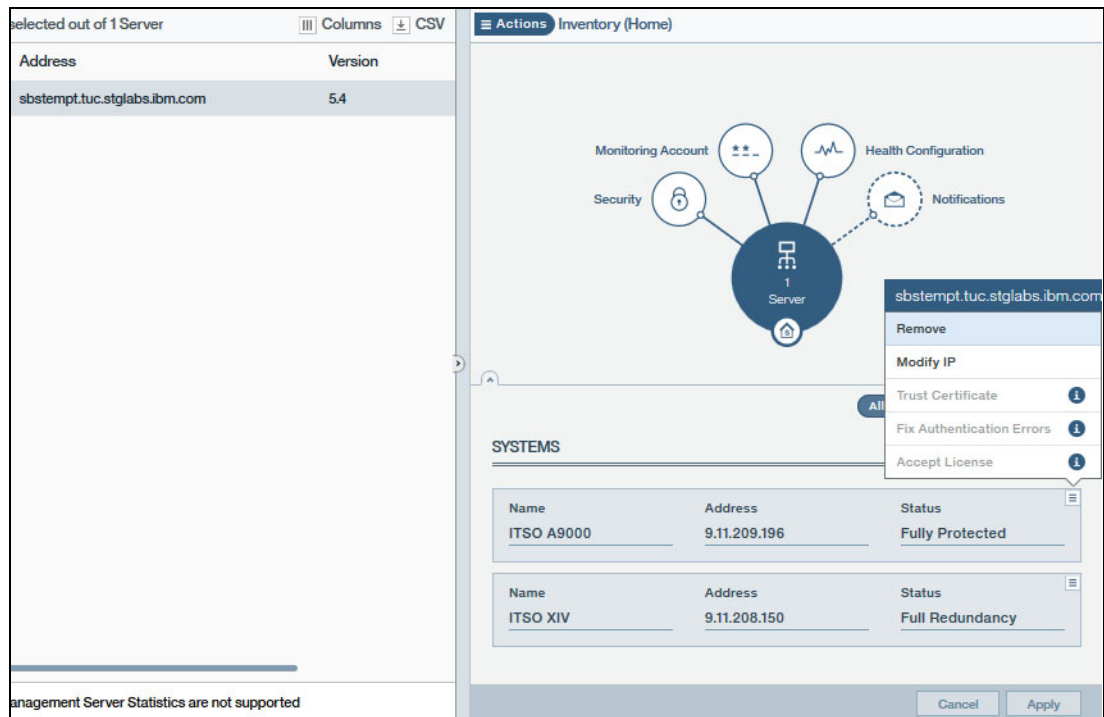


Figure 1-43 Removing a system from inventory or modifying the IP address of a system in inventory

4. Click **Remove** to delete the system from the inventory. Click **Modify IP** to change the IP address of a system.
5. Click **Apply** to save the changes. Click **Cancel** to leave the information unchanged.

If you plan to manage multiple systems, the user and password must be the same on all the systems in the inventory to manage them from a single login, including the default administrator user's password. If a user's password is not the same on all systems, the user cannot see and manage the system with different login credentials.



Storage management and configuration

This chapter covers the IBM Spectrum Accelerate family, which includes IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software, storage management, and configuration by using the Hyper-Scale Manager V5.5 graphical user interface (GUI) and command-line interface (CLI) management tools. The topics include the definition of systems in the GUI and the creation of a logical configuration by using the GUI and CLI. The following topics are included:

- ▶ Creating and managing pools
- ▶ Creating and managing volumes
- ▶ Creating and managing snapshots
- ▶ Creating and managing consistency groups and snapshot groups

Basic information about the topics is provided. This chapter also includes a section that guides you through typical scenarios that you might encounter when you manage IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software with the new GUI.

Important: See the following IBM Redbooks publications for information business continuity functions (mirroring, replication, and migration) and for host attachment topics:

- ▶ *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401
- ▶ *IBM HyperSwap and Multi-site HA/DR for IBM FlashSystem A9000 and A9000R*, REDP-5434
- ▶ *IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM XIV Storage System: Host Attachment and Interoperability*, SG24-8368
- ▶ *IBM FlashSystem A9000 and A9000R, IBM XIV, and IBM Spectrum Accelerate with IBM SAN Volume Controller Best Practices*, REDP-5408

2.1 Storage management interfaces

You can manage IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software deployments with the same GUI and CLI interfaces. The steps that are needed for IBM XIV Gen3, IBM FlashSystem A9000, and IBM Spectrum Accelerate are the same.

2.1.1 IBM Hyper-Scale Manager

IBM Hyper-Scale Manager is software that is installed on a server for managing and monitoring IBM storage systems. Hyper-Scale Manager reduces operational complexity and enhances capacity planning through integrated management for large and multiple-site storage system deployments. Storage administrators can use a web browser to log in to Hyper-Scale Manager and employ its advanced web-based GUI for managing and monitoring in real time, multiple IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate based storage systems.

Graphical user interface

The web-based GUI in Hyper-Scale Manager V5.5 provides an object-centered interface design that is aimed at ease of use with maximum productivity for storage administrators.

The GUI is installed as part of IBM Hyper-Scale Manager. The GUI can be accessed by using a web browser and pointing to the address or URL of Hyper-Scale Manager with the port that is used to open the connection (in our case, port 8080). After Hyper-Scale Manager is installed, you are ready to define the systems to manage and start their logical configuration.

Note: As with any view in the GUI, the tabular view can be customized to add or remove the columns that are shown. On any view, click the title as shown in Figure 2-1, and select the options that you want to see. Each view differs, depending on the subject. The customized view can also be saved for later use.

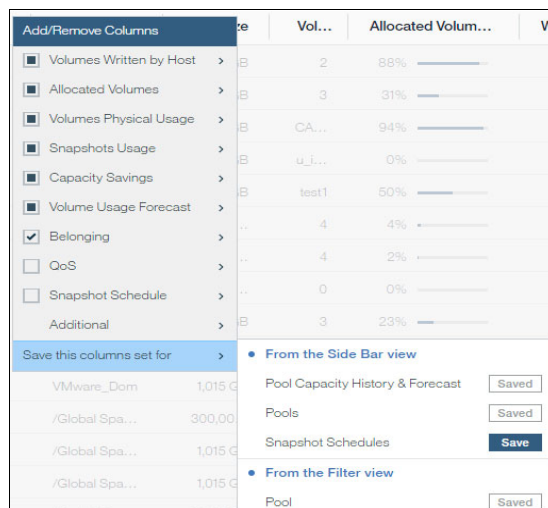


Figure 2-1 Customize the tabular view in the GUI

2.2 Configuring pools

The concept of storage pools is purely administrative. Although the hardware resources within IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software are virtualized in a global sense, the available capacity in the system can be administratively portioned into separate and independent storage pools. *Storage pools* function as a virtual way to effectively manage a related group of logical volumes that are provisioned in a similar manner, consistency groups, and their snapshots.

A pool that is configured in IBM FlashSystem A9000 and A9000R is capacity that is virtual and not calculated on the physical or hard capacity of the system due to the data reduction techniques unique to these systems. A pool is still configured to be a specific size. However, the amount of capacity that is used depends on the pattern matching, data deduplication, and compression savings that are realized based on the data that is written from the host system.

Tip: Starting with system software V12.3, you can define a unique storage pool that encompasses the overall system capacity.

The IBM XIV Gen3 and IBM Spectrum Accelerate software can also reserve physical hard capacity, in which case the pool is called *regular*.

The size of a storage pool can always be resized (to grow or shrink). When you are decreasing the size of a pool, the only limit is that the new pool size cannot be smaller than the total size of configured volumes in a pool or not less than the amount of data that is written to an overallocated pool.

Volumes that are not in a consistency group can be moved between storage pools without any limitations (assuming that adequate space exists in the new pool). Volumes that are part of a consistency group can be moved together as a group.

Snapshot space must be reserved during pool definition if any of the volumes in the pool will be duplicated with snapshots. This capacity is included as part of the usable capacity in the storage pool. A pool can be resized later to add or remove snapshot space, as needed.

Important: The space that is reserved for snapshots is not available for use by volumes.

2.2.1 Creating pools

We describe how to create storage pools by using the GUI and the CLI. Because this GUI is new, we spend extra time illustrating how this simple task is performed so that the process and views can be fully explained. The CLI is straightforward, but we included useful hints for the steps that are taken.

Note: Whether you create pools with the GUI or the CLI, the size that is created might differ from the size that is requested. For the IBM FlashSystem A9000 and A9000R, the created size might be larger based on a rounding-up factor to the nearest 103 GB boundary when the requested size is within 5% below that boundary. For XIV and IBM Spectrum Accelerate, the rounding up factor is based on a 17 GB boundary.

Creating pools with the GUI

The creation and resizing of storage pools are straightforward. The name of the storage pool must be unique within a system. Because Hyper-Scale Manager can have multiple systems in the inventory, the name uniqueness is still at a system level, not based on all systems that are managed by Hyper-Scale Manager.

The size of the pool must consider volumes to be added to the pool and the future activity within the storage pool (especially for snapshot propagation from creating too many snapshots). On the depletion of space in a pool, the system progressively deletes snapshots (based on a deletion priority) in the pool to clear space for more write requests.

The system enables the assignment of the entire effective capacity to user-created storage pools. The storage pool is initially empty, and it does not contain volumes. When you create a storage pool, the size cannot be zero capacity.

To create a storage pool, complete the following steps:

1. Click the **New** icon in the top menu bar, and select **Pool** as shown in Figure 2-2.



Figure 2-2 Create New menu

- The Create Pool wizard opens on the right side of the window, as shown in Figure 2-3. Note the red asterisk (*) icon to the right of certain field names. This asterisk indicates the required fields.

Figure 2-3 Creating a pool

- Choose the unit of measurement, GB or TB, from the drop-down list. GB is the default setting.
- The option to select Regular or Thin is available but is disabled if the system selection is FlashSystem A9000 or A9000R. If *Thin* is selected, an extra field is displayed to select the hard size (or physical size) and the pool size (or soft size).
- Enter the snapshot reserved size and pool size. The snapshot size can be zero if no snapshots are planned for the volumes in this pool. It can also be updated later if the snapshot requirements change.

Snapshot size: The snapshot size is a subset of the total pool size. It does not allocate more space.

Snapshot reserve: The system preemptively deletes snapshots if the snapshots fully use the allocated space of the storage pool. Ensure that adequate space is allocated for the snapshot reserve when you define a storage pool. A preferred practice is to allocate 10% - 20% of the pool size for snapshots.

- Enter the pool size that you want. The “i” indicates that additional information is available for this field. Hover the pointer over this i to display information about the pool size. For IBM FlashSystem A9000 and A9000R, it is rounded up to 103 GB units, as shown in Figure 2-4, and for IBM XIV Gen3 and IBM Spectrum Accelerate software, it is rounded up to 17 GB units.

Figure 2-4 Pool size additional information

Note: IBM FlashSystem A9000 and A9000R use the concept of allocation unit (AU) size for volumes. The AU size is set at 103 GB. The minimum volume size that can be created is 1 GB. However, volumes that are created with a specified size within 5% less than the AU size are rounded to the AU size. For example, if you create a volume and specify a size of 98 GB, a volume of 103 GB is created on the system.

In IBM XIV Gen3 and IBM Spectrum Accelerate, the minimum volume size is 17 GB, and all volumes are in increments of 17 GB.

7. Select the system in which you want the pool created. Click the arrow icon to show a drop-down list that displays the systems that are managed by Hyper-Scale Manager. An example of this list is shown in Figure 2-5. Select the system from the list to close the list of systems and insert your selection in the field.

The screenshot shows a user interface for selecting a system and domain. At the top, there is a dropdown menu labeled 'System' with 'FlashSystem A90001' selected. Below this is a table listing available systems with columns for 'System', 'Allocated', and 'Effective Size'. The table lists four systems: FlashSystem A90001, FlashSystem A90002, Spectrum Accelerate, and XIV Gen3. Below the table is another dropdown menu labeled '/Global Space/' with an arrow icon. Below this is a dashed box containing a link that says 'CREATE DOMAIN'. At the bottom, there is a table with two columns: 'Domain' and 'System'. The first row shows '/Global Space/' under 'Domain' and 'FlashSystem A90001' under 'System'.

System	Allocated	Effective Size
FlashSystem A90001	15.81TB	1,750TB (x5)
FlashSystem A90002	114.46TB	1,750TB (x5)
Spectrum Accelerate	2.58TB	19.16TB (x1)
XIV Gen3	198.54TB	650TB (x1.997)

Domain	System
/Global Space/	FlashSystem A90001

Figure 2-5 System and domain selection for pool

8. The domain selection is next, and it is also a required field. Click the arrow icon to expand the list. If the pool will not be in a domain, leave it default as /Global Space/.

If the pool will be associated with a domain, either select the domain from the list or choose **Create Domain**, which starts the wizard to create a domain. Refer to 4.2.1, “Creating domains” on page 123 for more details. Figure 2-5 shows an example.

The last section is only for IBM XIV Gen3 pool creation, where it is possible to choose compressed volumes or uncompressed volumes.

9. Two extra options are available depending on prior selections. If the system is XIV or IBM Spectrum Accelerate and **Thin** was selected, then you can select **Read-only** or **No IO** for the pool. Additionally, you can select **Compressed** or **Uncompressed** for XIV.

10. Click **Create** to create the pool with your specified settings, as shown in Figure 2-6. Any of these settings can be modified later, if needed.

Figure 2-6 Completed fields for a new pool

Creating pools with the command-line interface

All of the operations that are performed with a pool in the GUI can be performed with the CLI. To obtain a list of all of the pool-related commands in the CLI, enter the following command in the CLI session:

```
help category=storage-pool
```

The output of this command is shown in Figure 2-7.

Category	Name	Description
storage-pool	pool_list	Lists all Storage Pools or the specified one.
storage-pool	cg_move	Moves a Consistency Group, all its volumes and all their snapshots and Snapshot Sets from one Storage Pool or Grouped Pool to another.
storage-pool	pool_create	Creates a Storage Pool.
storage-pool	pool_resize	Resizes a Storage Pool.
storage-pool	pool_change_config	Changes Storage Pool configuration.
storage-pool	pool_rename	Renames a specified Storage Pool.
storage-pool	pool_delete	Deletes a Storage Pool.
storage-pool	pool_config_snapshots	Changes the Storage Pool snapshot settings.
storage-pool	vol_move	Moves a volume and all its snapshot from one Storage Pool to another.

Figure 2-7 Storage pool commands in the CLI

Alternatively, while you are in the CLI session, you can enter the **pool_** command followed by pressing the Tab key to see all of the commands that start with these characters. The output of this command is shown in Example 2-1.

Example 2-1 All of the commands that begin with "pool_"

```
ITSO_2_A9000R>>pool_  
pool_change_config      pool_config_snapshots  pool_create  
pool_delete             pool_list  
pool_rename             pool_resize
```

To create a pool, enter the **pool_create** command. To see all of the options that are associated with the creation of a pool, enter the command followed by pressing the Tab key to see them, as shown in Example 2-2. Of these parameters, only three are mandatory: Pool, size, and snapshot_size. You do not need to enter the name of the system, which was required with the GUI. You are not required to enter the name with the CLI because the CLI connects directly to a single system and not through Hyper-Scale Manager.

Example 2-2 Parameters for the pool_create command

```
ITS0_2_A9000R>>pool_create  
pool=          size=          snapshot_size=          perf_class=          domain=
```

The remaining fields are optional. The system administrator can use the **perf_class** option to associate a quality of service (QoS) performance class with the pool. The default is no performance class. To create the pool, enter the following command as shown in Example 2-3.

Example 2-3 Create a pool

```
ITS0_2_A9000R>>pool_create pool=itso_pool3 size=500 snapshot_size=20
```

In this example, we created a pool with a size of 500 GB and a snapshot reserve size of 20 GB. Other actions that can be performed on a pool, such as resizing a pool and deleting a pool, are covered in the next section.

2.2.2 Managing storage pools with the GUI

In this section, we describe how to manage storage pools by using the GUI management tool. Managing pools with the GUI is straightforward. The related tasks can be reached through the Pools & Volumes Views and the Hub view. You can also access the pool actions from other places, such as systems and domains and by right-clicking the pool. In this section, we use the Pools & Volumes views and then the Pools view.

Figure 2-8 shows an example of the Pools view. To manage a pool, select the specific pool and click the **Actions** menu (or right-click), which is shown in Figure 2-9.

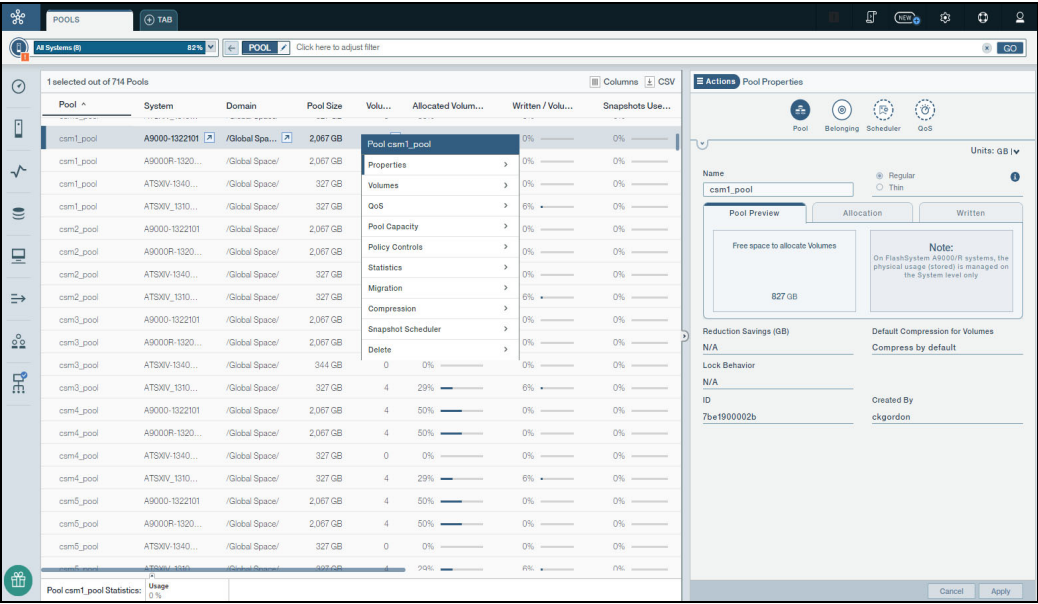


Figure 2-8 Pool view

Pool ITSO_p1	
Properties	>
Volumes	>
QoS	>
Pool Capacity	>
Policy Controls	>
Statistics	>
Migration	>
Compression	>
Snapshot Scheduler	>
Delete	>

Figure 2-9 Pool actions menu

The following actions are available:

- **Properties:** The administrator can view and modify pool properties.
- **Volumes:** The administrator can view and create volumes.
- **QoS:** The administrator can create and remove performance classes on a pool.
- **Pool Capacity:** The administrator can generate capacity reports.
- **Policy Controls:** The administrator can view and modify pool thresholds.
- **Statistics:** The administrator can view performance and capacity statistics.
- **Migration:** The administrator can migrate volumes from another system to this pool.
- **Compression:** The administrator can change between compressed pool and uncompressed pool for IBM XIV Gen3.'

- ▶ **Snapshot Scheduler:** The administrator can set up automatic scheduled snapshots.
- ▶ **Delete:** The administrator can delete a pool and all of its volumes.

Properties

First, we look at the properties for a pool. The available options under **Actions** → **Properties** are shown in Figure 2-10. The storage administrator can view the properties of a pool. The storage administrator can modify the pool's name, size, and domain association. For IBM XIV Gen3 and IBM Spectrum Accelerate software, the administrator can also change lock behavior.

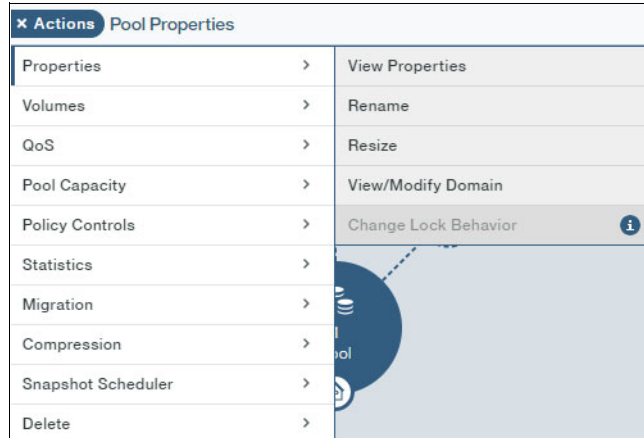


Figure 2-10 Options for pool properties

Selecting the View Properties option, two sections are available: the Hub view and the Pool Properties. Both sections are shown in Figure 2-11. The top section is the Hub view. You can use any elements of the Hub to drill down for more information and for pool-related options, such as Belonging and QoS.

The lower section lists the pool properties. From this view, you can modify the pool attributes. All editable fields are shown in a lighter shaded area (name, snapshot size, and pool size). Changes can be made and applied by modifying the fields and clicking **Apply** (not shown).

Figure 2-11 shows the pool properties.

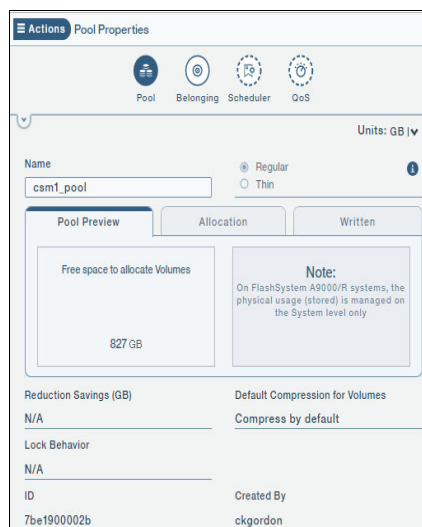


Figure 2-11 View the pool properties

Volumes

The storage administrator can use the Volumes action to view and create volumes within the pool. These options are shown in Figure 2-12.

x Actions Pool Properties	
Properties	> View Volumes
Volumes	> Create Volumes here

Figure 2-12 Volumes view for a pool

The storage administrator can optionally look at the pool's volume summary from this menu by selecting **View Volumes** and **Belonging**, which displays the Hub and Summary views as shown in Figure 2-13. The storage administrator also optionally can create snapshots of these volumes from this view by selecting **NEW** in the lower-right corner.

This action starts the Create Snapshot view, which is reviewed in 2.4.1, “How to create a snapshot of a volume” on page 75.

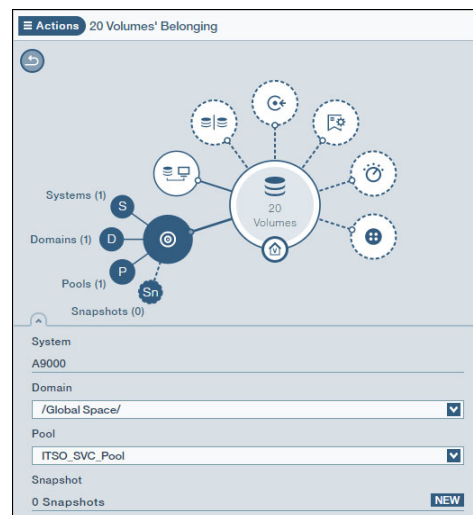


Figure 2-13 Volume summary and Belonging selected for a pool

Clicking any of the elements within the Hub shows more details about that element in the inspector area as shown in Figure 2-13.

Figure 2-14 illustrates what the Hub and Summary view looks like on the right side of the window when you list the volumes within a pool by using this method.

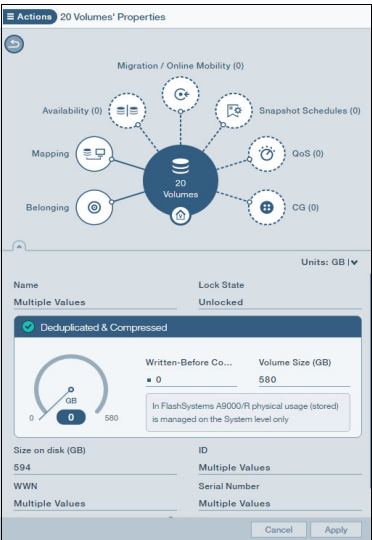


Figure 2-14 Hub view of the volumes in the same pool

Quality of service (QoS)

The storage administrator can easily set or cancel a performance class to limit the traffic to a specific pool with QoS under the Actions menu as shown in Figure 2-15.




Figure 2-15 Setting QoS

By selecting the option to set QoS, you can optionally use an existing performance class or create a new performance class. When you create a performance class, the Create QoS Performance Class window opens. You must enter the name of the class and select bandwidth (BW) or input/output operations per second (IOPS) as the limiting factor. You can also specify whether the performance class is shared among all the members that are assigned this QoS or if each member is limited individually.

Figure 2-16 shows an example of this view.




Figure 2-16 Create QoS Performance Class

Pool capacity

The storage administrator can select **Pool Capacity** to generate a capacity and trend forecasting report for the pool. See Figure 2-17.

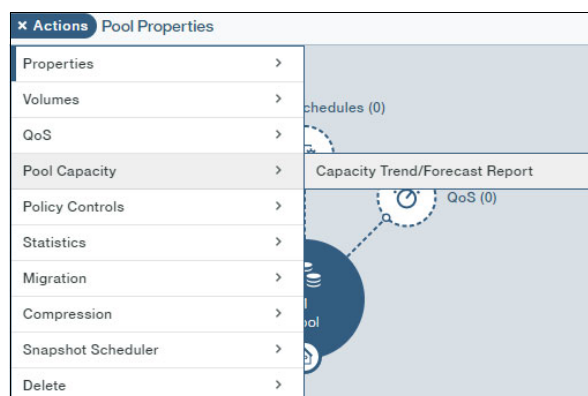


Figure 2-17 Pool capacity reports

The report is generated in PDF format and saved on the system from which the browser was started. The compressed file that contains the report is identified by the name and appended with the date and time that it was generated, as shown:

```
# 03/04/16 12:20 PM 69,097 capacity_planning_reports_2016-10-10_14-05-33.zip
```

This report requires at least 30 days of monitoring by the Hyper-Scale Manager (HSM) to generate the compressed file and includes information about system capacity, domain usage, and pool usage.

Policy controls

Policy controls are the capability to set or modify pool thresholds. You first select **Pools** in the Pools & Volumes Views and then select a pool in the list of pools. Under the Actions menu, or with a right-click, select **Policy controls** → **View/Modify Pool Thresholds** as shown in Figure 2-18.

Pool ITSO_SVC_Pool	580,081...	18	90%
Properties	5,008 GB	3	23%
Volumes			
QoS			
Pool Capacity			
Policy Controls	View/Modify Pool Thresholds		
Statistics			
Migration			
Compression			
Snapshot Scheduler			
Delete			

Figure 2-18 Pool thresholds

After View/Modify Pool Thresholds is selected, the Pool Thresholds view of the pool's properties is shown on the right side of the window. You can use the systemwide settings or customize the settings for this pool. When **Custom Threshold** is selected, as shown in Figure 2-19, the storage administrator can adjust the percentages for each value under volumes and snapshots, and decide whether a particular notification is needed based on severity. Changes to the thresholds are seen by all users of that system.

Pool Thresholds (%)

Derived From System

Custom Threshold

Volume Usage

Snapshot Usage

Informational

Warning

Minor

Major

Critical

☐

N/A

☐

N/A

☐

N/A

☐

N/A

☐

N/A

Informational

Warning

Minor

Major

Critical

☐

N/A

☒

80

☒

90

☒

95

☐

N/A

Figure 2-19 Customizing pool thresholds

Tip: The critical threshold is not set, by default. The critical threshold can easily be added if more notification is needed than the major threshold by selecting the **Critical** option and entering a percentage.

Important: Thresholds for volume usage cannot be defined for FlashSystem A9000 and A9000R running code V2.1.x or later as the volume's physical capacity is managed at the system level. The fields are uneditable for systems at code V12.1 or later.

Statistics

The Hyper-Scale Manager also provides you with the option to view the volumes and snapshots performance statistics and capacity statistics from the pool's Actions menu, as shown in Figure 2-20.

Actions

Pool Properties

Properties

Volumes

QoS

Pool Capacity

Policy Controls

Statistics

Migration

Compression

Delete

Volumes Performance Statistics

Show Capacity Statistics

Snapshots Performance Statistics

Figure 2-20 Volume statistics in a pool

When selecting the performance option, the Volume Statistics view for all of the volumes within the pool is displayed, as shown in the example in Figure 2-21.




Figure 2-21 Volume statistics within a pool

Note: Hovering with the cursor over any point on the line in the Statistics view opens a summary for that particular point, as shown in Figure 2-22.

8/12 Active Volume		Oct 11, 2016 10:39:37 AM
Total	8 Volume	424.51K IOPS
Max	A90001_Exchange_vol_8 (FlashSystem A90001)	57.66K IOPS
	A90001_Exchange_vol_6 (FlashSystem A90001)	57.11K IOPS
	A90001_Exchange_vol_7 (FlashSystem A90001)	53.47K IOPS
Min	A90001_Exchange_vol_2 (FlashSystem A90001)	52.66K IOPS
	A90001_Exchange_vol_3 (FlashSystem A90001)	49.01K IOPS
	A90001_Exchange_vol_1 (FlashSystem A90001)	48.47K IOPS
TIP	click on the graph to pin tooltip	

Figure 2-22 Summary of volume statistics within a pool

Additional information can be seen by selecting the capacity option or the snapshots option.

The storage admin can also view current and historical statistics for a single volume, as shown in Figure 2-23. This can be seen by selecting a single volume from the pool view and from the volumes view. You can find more information about the volume view in “Statistics” on page 80.




Figure 2-23 Statistics for a single volume

The storage administrator can view the pool usage by expanding the statistics view at the bottom of the page. An example of the expanded graph is shown in Figure 2-24. Notice the two available options to view usage and usage plus forecast for the selected pool's volumes. This example is showing a range of one month.




Figure 2-24 Pool usage

New with Version 5.4 is the ability to view internal systems and network latency (both current and historical) for a host and a single system (the system version must be at Version 12.2.1 or later). If a host is mapped to volumes on multiple systems, then simply selecting the system from the drop-down list in the Statistics view enables this view. An example is shown in Figure 2-25. With this function, administrators have more insight into latency issues.

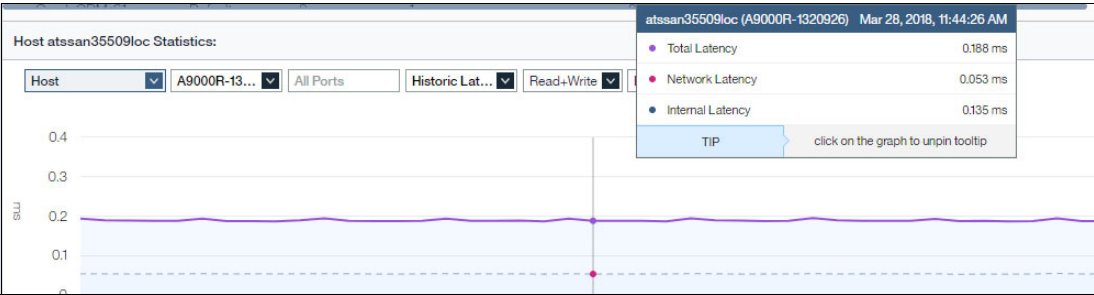


Figure 2-25 Internal and network latency statistics view

Migration

The Migration option within a pool is another path for the storage administrator to initiate data migration from another system to this system. The Migrate remote Volume to Pool menu option is shown in Figure 2-26.

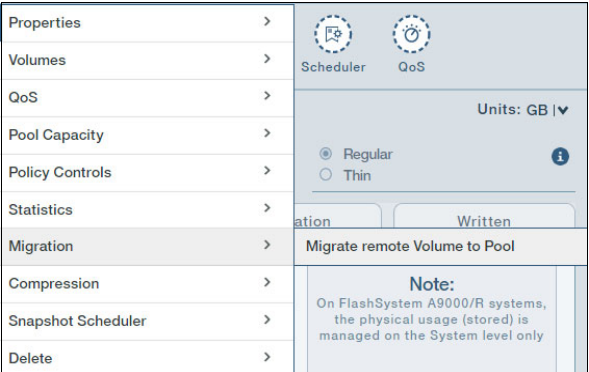


Figure 2-26 Migrate volumes to this pool

The selection of this option opens the data migration wizard, as shown in Figure 2-27.

Migrate remote Volume to Pool

Target *

Target LUN *

☐ Existing Volume Destination Volume *

☒ Create Volume

☐ Don't Update Source ☒ Active

☒ Update Source ☐ Inactive

Figure 2-27 Volume migration wizard

Compression

A pool's compressed state on IBM XIV Gen3 can be seen in the pool properties when a pool is selected from the pools view. An example of this is shown in Figure 2-28.

Name: CG_Pool ☐ Regular ☒ Thin

Pool Preview	Allocation (soft)	Stored (hard)
Free space to allocate Volumes (soft)	Free space to store Volume data (hard)	Reaching Warning threshold in (8/21/17)
241 GB	1,532 GB	256 GB

Reduction Savings (GB): 18 Default Compression for Volumes: Don't compress by default

Lock Behavior: Read-only

Figure 2-28 IBM XIV Gen3 pool compression state

Only the IBM XIV Gen3 can have a mixture of uncompressed and compressed pools. To change the pool from uncompressed to compressed on the XIV Gen3 select the Compression option in the Actions menu, as shown in Figure 2-29.

Pool Properties

Actions

- Properties >
- Volumes >
- QoS >
- Pool Capacity >
- Policy controls >
- Statistics >
- Migration >
- Compression >**
- Delete >

Change Default Compression

Figure 2-29 Compression setting

This setting can be changed only on IBM XIV Gen3. When set, it creates volumes with the given setting (compressed or uncompressed) as shown in Figure 2-30.



Figure 2-30 Selection Compressed or Uncompressed in IBM XIV Gen3

Reminder: All compressed pools in IBM XIV Gen3 must be thin-provisioned.

Snapshot Scheduler

The Snapshot Scheduler is a new option that is available with Hyper-Scale Manager V5.3 and later. It allows the system administrator to automate snapshots on a defined schedule. Doing this at the pool level automatically applies to existing and new volumes and consistency groups within the pool.

To work with the scheduler, select a pool in the pools view. Then, under the **Actions** menu, select **Snapshot Scheduler**. There are several options available, as shown in Figure 2-31, which are discussed in more detail in 2.4, “Snapshots” on page 74.

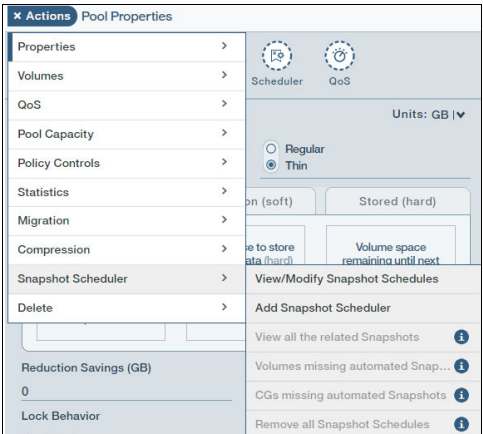


Figure 2-31 Snapshot scheduler

Deleting a pool

To delete a pool, select the pool in the Pools view. Then, under the Actions menu, select **Delete** → **Delete Pool(s)**, as shown in Figure 2-32. The capacity of the deleted pool is returned to the free space (either in the global space or in a domain, depending on what the pool is associated with).




Figure 2-32 Deleting a pool

Note: If a pool contains volumes or consistency groups, the system does not delete the pool, but the GUI prompts you with an option to go to the Volumes view to delete the volumes, as shown in Figure 2-33.




Figure 2-33 Volumes and consistency groups must be deleted or moved from a pool

Optionally, you can move the volumes or consistency group to another pool instead of deleting the objects so that you can delete the pool.

2.2.3 Managing storage pools with the CLI

All of the operations that are performed with the GUI on pools can also be performed by using the CLI. A list of pool-related commands is shown in Figure 2-34. To generate this list of commands during a CLI session, enter the following command:

```
help category=storage-pool
```

storage-pool	pool_list	Use the pool_list command to list all storage pools or the specified one.
storage-pool	cg_move	Use the cg_move command to move a consistency group, all its volumes, and all their snapshots and snapshot sets from one storage pool to another.
storage-pool	pool_create	Use the pool_create command to create a storage pool.
storage-pool	pool_resize	Use the pool_resize command to resize a storage pool.
storage-pool	pool_change_config	Use the pool_change_config command to change a storage pool configuration.
storage-pool	pool_rename	Use the pool_rename command to rename the specified storage pool.
storage-pool	pool_delete	Use the pool_delete command to delete a storage pool.
storage-pool	pool_config_snapshots	Use the pool_config_snapshots command to change storage pool snapshot settings.
storage-pool	vol_move	Use the vol_move command to move a volume and all its snapshot from one storage pool to another.

Figure 2-34 Storage pool-related CLI commands

Note: The commands that are shown in this section are queried during a CLI session.

Not all of the commands that relate to storage pools begin with `pool`. However, by using the CLI in session mode, you can press the Tab key to generate the commands that begin with “`pool_`”, as shown in Example 2-4.

Example 2-4 Generate a list of commands that begin with “pool_”

```
ITS0_2_A9000R>>pool_  
pool_change_config      pool_config_snapshots  pool_create  
pool_delete            pool_list  
pool_rename            pool_resize
```

You can use the Tab key to complete the command if it is unique (if you entered `pool_cr` and then pressed the Tab key). You also can issue all of the options, for example, the available options for modifying a pool, setting QoS, setting a performance class (`perf_class`), changing the lock behavior, and setting the pool thresholds, within a command. All of these options are shown in Example 2-5.

Example 2-5 Generate a list of options for a specific command

```
ITS0_2_A9000R>>pool_change_config  
pool=          lock_behavior=      perf_class=      code=  
severity=      threshold=  
restore_thresholds=  hysteresis=
```

Tip: The options for any CLI command can be listed by entering the command followed by pressing the Tab key.

To list the pools in a system or systems, issue the **`pool_list`** command. The sample output of this command is shown in Figure 2-35.

Name	Size (GB)	Snap Size (GB)	Total vols (GB)	Empty (GB)	Used by vols (GB)	Used by Snaps (GB)
itso_d1_p1	2067	413	1550	103	0	0
itso_d2_p1	2584	413	855	1240	0	0
itso_pool2	2067	0	250	1757	0	0
itso_pool1	6201	413	533	5168	0	0
test_pool	142022	7028	119503	15401	31836	0

Figure 2-35 List of pools

To create a pool, enter the following command:

```
pool_create pool=itso_p3 size=2050 snapshot_size=400 lock_behavior=read_only  
sparse=yes
```

The **`snapshot_size`** is a mandatory field, which must be a positive integer value (including 0). If **`sparse`** is included, the `lock_behavior` must be specified whether the behavior that you want is not the default of `read_only`. If the parameter is not included, the default is applied to the pool.

Even with the CLI command, the system automatically adjusts the pool size and snapshot size, as shown in Example 2-6.

Example 2-6 Automatically adjusted size

Name	Size (GB)	Snap Size (GB)
itso_pool3	2067	413

Run the following command to resize an existing pool:

```
pool_resize pool=itso_pool3 size=5000 snapshot_size=800
```

Run the following command to delete a pool:

```
pool_delete pool=itso_pool3
```

Approve or deny deletion by responding *y* or *n* when you are prompted. You can also use the *-y* parameter with the command to approve deletion without the prompt.

Tip: Use the *-y* parameter at the end of a command that requires confirmation so that the command is automatically approved. This parameter is useful for scripting.

2.3 Volumes

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software offer logical volumes as the basic data storage element for allocating usable storage space to attached hosts. This logical unit concept is known, and it is widely used by other storage systems and vendors. IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software facilitate the creation of volumes without planning volume layout or optimizing physical layer for performance.

Traditionally, logical volumes are defined within various Redundant Array of Independent Disks (RAID) arrays, where their segmentation and distribution are manually specified. The result is often a suboptimal distribution within and across storage modules that depends on the administrator's knowledge and expertise.

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software use the grid concept and distribute data evenly across flash or disk enclosures.

Within IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software, each volume is equally distributed across the flash or disk enclosures. This method dramatically simplifies storage provisioning, letting the system automatically lay out your volume in an optimal way.

2.3.1 Creating volumes

After you define the storage pools, the next logical step in system configuration is the definition of the volume. Volumes can be created with either the GUI or the CLI. Both of these methods are described in this section. However, with this version of the Hyper-Scale Manager, you can start by creating a volume and be prompted to create a pool.

Creating volumes with the GUI

To create a volume, click the **NEW** icon and select **Volume**, as shown in Figure 2-36.




Figure 2-36 Create volumes

This action opens the Create Volume view (Figure 2-37).

You must enter the required fields as indicated by the red asterisks (*). In this window, you can also use the drop-down list to select the units for the volume, the domain, the system, and the pool. If the Hyper-Scale Manager is managing multiple systems, including XIV and IBM Spectrum Accelerate, these are displayed in the System drop-down list. An example of the creating a new volume is shown in Figure 2-37.




Figure 2-37 Create Volume wizard

Tip: The information content is displayed when you move the cursor over the **i** in the circle next to the field name.

After you enter the information, select **Create** to complete the task. The new volumes are displayed, as shown in Figure 2-38.

VOLUME Click here to adjust filter						
144 Volumes CSV						
Volume ^	Volum...	Written by ...	System	Snaps...	Pool	Reduction Status
A90001_Exchange...	1,015 GB	100% <div></div>	FlashSystem...	0	A90001_Exchange	Deduplicated & Compres... ✓
A90001_Exchange...	1,015 GB	100% <div></div>	FlashSystem...	0	A90001_Exchange	Deduplicated & Compres... ✓

Figure 2-38 List of volumes

Alternatively, you can create volumes from the Pools view under **Actions** → **Volumes** → **Create Volumes here** as shown in “Volumes” on page 47. This action opens the same Create Volume view as shown in Figure 2-38 on page 58 except that the System, Domain, and Pool fields are completed.

Creating volumes with the CLI

All of the operations that are available in the GUI can also be performed through the CLI. When using the CLI, you are communicating directly to a system and do not need to specify which system in the command (as is required with the GUI). To obtain a list of all of the volume-related commands, enter the following command in a CLI session:

```
help category=volume
```

The output of this command is shown in Figure 2-39.

Category	Name	Description
volume	vol_list_extended	Use the vol_list_extended command to return the attributes of the volume which are not returned by vol_list.
volume	vol_list	Use the vol_list command to list all volumes or a specific one.
volume	reservation_list	Use the reservation_list command to list volume reservations.
volume	reservation_key_list	Use the reservation_key_list command to list reservation keys.
volume	reservation_clear	Use the reservation_clear command to clear reservations of a volume.
volume	vol_create	Use the vol_create command to create a new volume.
volume	vol_create_many	Use the vol_create_many command to create new volumes.
volume	vol_rename	Use the vol_rename command to rename a volume.
volume	vol_format	Use the vol_format command to formats a volume.
volume	vol_lock	Use the vol_lock command to lock a volume so that it is read-only.
volume	vol_unlock	Use the vol_unlock command to unlock a volume, so that it is no longer read-only and can be written to.
volume	vol_copy	Use the vol_copy command to copy a source volume onto a target volume.
volume	vol_delete	Use the vol_delete command to delete a volume.
volume	vol_resize	Use the vol_resize command to resize a volume.
volume	vol_set_external_id	Use the vol_set_external_id command to set a volume's external identifier.
volume	vol_clear_external_id	Use the vol_clear_external_id command to clear a volume's external identifier.
volume	vol_diff	Use the vol_diff command to display the allocated partitions' difference between two volumes.

Figure 2-39 CLI list of volume-related commands.

Alternatively, you can obtain a list of all of the CLI commands that begin with vol by entering vol in the session, followed by pressing the Tab key twice. This action displays the commands, as shown in Example 2-7.

Example 2-7 CLI commands that begin with vol

ITS0_2_A9000R>>vol_			
vol_allocation_bitmap	vol_bind_list	vol_clear_external_id	vol_copy
vol_create			
vol_create_many	vol_delete	vol_diff	vol_diff_copy
vol_format			
vol_list	vol_list_extended	vol_lock	vol_mapping_list
vol_move			
vol_rename	vol_resize	vol_set_external_id	vol_unlock

You can run a search on an object, such as volume, by entering the following command:

```
help search=volume
```

This command results in all of the commands that are related to volumes, including commands that are related to mapping, consistency groups, and snapshots.

To create a volume, run the following command:

```
vol_create vol=itso_vol7 pool=itso_pool1 size=40
```

The volume is logically formatted at creation time, which means that any read operation results in returning all zeros as a response. To format a volume, run the following command:

```
vol_format vol=itso_vol7
```

All of the data that is stored on the volume is lost and unrecoverable. If you want to bypass the warning message, type **-y** directly after the CLI command.

To list the existing volumes in a system and in a particular pool, run the following command:

```
vol_list pool=itso_pool2
```

For a list of all of the volumes in a system, run this command:

```
vol_list
```

The result of this command is similar to the output that is shown in Figure 2-40.

Name	Size (GB)	Consistency Group	Pool	Creator	Written (GB)	Est. Data Reduction Savings (%)	Est. Min Delete Size (GB)
itso_d1_v1	516		itso_d1_p1	admin	0	0	0
itso_d1_v2_001	516		itso_d1_p1	admin	0	0	0
itso_d1_v2_002	516		itso_d1_p1	admin	0	0	0
itso_d2_v_001	250		itso_d2_p1	admin	0	0	0
itso_d2_v_002	250		itso_d2_p1	admin	0	0	0
itso_d2_v_003	250		itso_d2_p1	admin	0	0	0
itso_global_test	103		itso_d2_p1	admin	0	0	0
itso_p3_vol_001	50		itso_pool2	admin	0	0	0
itso_p3_vol_002	50		itso_pool2	admin	0	0	0
itso_p3_vol_003	50		itso_pool2	admin	0	0	0
itso_p3_vol_004	50		itso_pool2	admin	0	0	0
itso_p3_vol_005	50		itso_pool2	admin	0	0	0
itso_vol1_001	20	itso_cg1	itso_pool1	admin	0	0	0
itso_vol1_002	20	itso_cg1	itso_pool1	admin	0	0	0

Figure 2-40 vol_list output

For more information, see the following resources:

- ▶ *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide*, SC27-8711
- ▶ *IBM FlashSystem A9000 Command-Line Interface (CLI) Reference Guide*, SC27-8559
- ▶ *IBM XIV Gen3 Command-Line Interface (CLI) Reference Guide*, GC27-3914
- ▶ *IBM Spectrum Accelerate Command-Line Interface (CLI) Reference Guide*, SC27-6697

2.3.2 Managing volumes with the GUI

To manage volumes from the Storage Management GUI, select **Volumes** from the Pools & Volumes Views. This option lists all of the volumes for the systems that are managed by the GUI in a single table. If you select any of the volumes in the table, you enable the Actions menu.

The available actions are listed:

- ▶ **Properties:** View and modify volume properties.
- ▶ **Snapshots:** Create, view, and restore snapshots.
- ▶ **Mirror:** Set up and modify volume mirroring.
- ▶ **HyperSwap:** define and manage HyperSwap volumes
- ▶ **Mapping:** View and modify volume mapping.
- ▶ **Consistency Group:** Move a volume to a consistency group or remove a volume from a consistency group.
- ▶ **QoS:** Enable or cancel a performance class with a volume.
- ▶ **Pool Capacity:** Create a capacity report.
- ▶ **Statistics:** View the performance statistics of a volume.
- ▶ **Migration:** Set up migration to the selected volume.
- ▶ **Online Volume Mobility:** Move the volume to another system of the same type.

- Volume Data: Format the volume and copy the volume.
- Compression: Compress or decompress a volume in IBM XIV Gen3.
- Delete: Delete the selected volume or volumes.

Other ways exist to do to volume actions, for example, from within a pool. However, in this section, we focus on the Volumes view and look at each main option under the Actions menu.

Properties

First, we look at the Properties action. By selecting this option, the storage administrator can view the properties of a volume, rename a volume, resize a volume, and change the lock state of a volume. The options are shown in Figure 2-41.

x Actions Volume Properties	
Properties >	View Properties
Snapshots >	Rename
Mirror >	Resize
HyperSwap >	Change Lock state
Mapping >	View/Move to Pool

Figure 2-41 Volume properties

When **View Properties** is selected, the data is displayed as shown in Figure 2-42.

The screenshot shows the 'Volume Properties' window. At the top, there's a 'Volume' hub with a central circle labeled '1' and several surrounding icons representing different properties: Availability (1), Mapping, Belonging, QoS (0), CG (1), Migration (0), and Online Mobility (0). Below the hub, there's a section titled 'Deduplicated & Compressed' with a green checkmark. It features a gauge showing '15' GB of usage out of a total '90' GB. Below this, there's a note: 'In FlashSystems A9000/R physical usage (stored) is managed on the System level only'. At the bottom, there's a table of properties:

Name	team06_HA1	Lock state	Unlocked
Size on disk (GB)	90	ID	56a33118303c12
WWN	6001738CCCE051BE000000000...	Serial Number	N/A
SSD Caching state	N/A	Created By	lisa

At the bottom right, there are 'Cancel' and 'Apply' buttons.

Figure 2-42 Volume Properties and Hub view

The volume properties are displayed on the lower half of the view in the inspector. The meter displays how much was written by the host before compression. IBM XIV Gen3 and IBM Spectrum Accelerate display additional information about compression and reduction savings. The equivalent FlashSystem A9000 and A9000R information is managed on a system level.

The upper portion shows the Hub with all of the associated options for the volume. Circles with a solid line indicate an active association, and they can be expanded for more information. Circles with a dashed line indicate no associations with the volume. A storage administrator can click one to create an association for that volume (such as a mapping, a consistency group, or migration).

Note: This view is the default view that is shown when you select a volume. The view changes slightly when multiple volumes are selected.

Snapshots

The next option under the Actions menu is snapshots. This option is shown in Figure 2-43. The options include creating snapshots, viewing snapshots, and restoring the volume data from a snapshot.

x Actions Volume Properties	
Properties	Create Snapshot(s)
Snapshots	View Snapshots
Mirror	Restore from Snapshot

Figure 2-43 Volume snapshot options

To create a snapshot of a volume in the GUI, select **Create Snapshot(s)**. After you select this option, the view that is shown in Figure 2-44 displays on the right side of the browser window. The system appends the name of the volume with `.snapshot_` or you can enter another name.

QoS (0)

Availability (1)

Mapping

Belonging

1 Volume

CG (1)

Migration (0)

Online Mobility (0)

Create Snapshot

Name

team06_HA1.snapshot_

Deletion Priority

1 - Last

1 - Last

2

3

4 - First

Cancel Apply

Figure 2-44 Create volume snapshot

You can then select the deletion priority for the snapshot from the drop-down list. The default deletion priority is 1 - Last. Deletion priority is described in more detail in 2.4.2, “Snapshot deletion priority” on page 77.

The snapshot is created after you click **Apply**. A message that indicates the success or failure of the command is displayed briefly at the bottom of the view.

Important: Snapshot deletion occurs automatically. Older snapshots are deleted without warning whenever a new snapshot is taken and space is needed for the new volume or snapshot data. Deletion occurs based on the deletion priority. Priority 0 - Golden Snapshot means that the snapshot is not deleted automatically.

In addition to a single snapshot of a volume, you can create snapshots of multiple volumes. This capability is not the same as creating a snapshot of a consistency group. Snapshots are created of all of the selected volumes, but might not be consistent.

An example is shown in Figure 2-45 with the selection of the volumes.

VOLUME Click here to adjust filter

2 selected out of 144 Volumes CSV

Volume ^	Volum...	Written by ...	System
XIV_PoolA_Worklo...	103 GB	100%	XIV Gen3
XIV_PoolA_Worklo...	103 GB	100%	XIV Gen3

Actions 2 Volumes' Properties

- Properties >
- Snapshots >
- Mirroring >
- Mapping >

Create Snapshot(s)

View Snapshots

Restore from Snapshot

Figure 2-45 Create multiple snapshots selection

After the volumes are selected, the Create Snapshot window opens. With multiple snapshots, you can select only the deletion priority. The name of the snapshots is generated automatically, and it can be changed later under the Snapshots view. An example of creating multiple snapshots at the same time is shown in Figure 2-46.

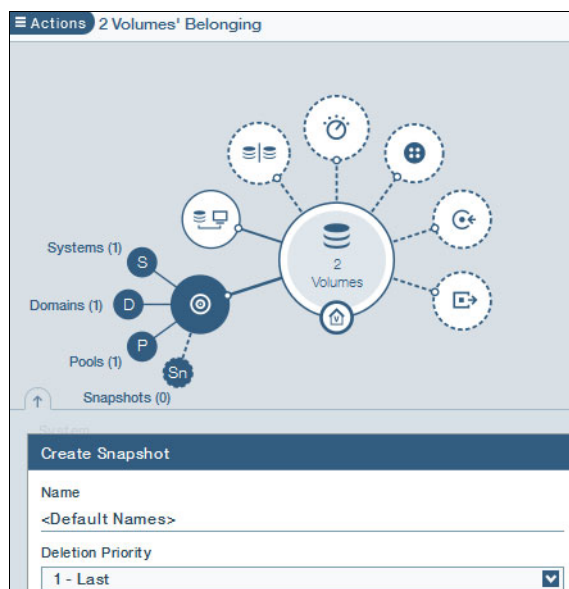


Figure 2-46 Create snapshots for multiple volumes at the same time

Volumes from different pools can also be selected at the same time to create snapshots. The “volume belonging” then shows the number of pools instead of the name of the pool.

Note: If the Inspector is maximized, you do not see the expanded view of the Belongings element in the Hub view.

For more information about snapshots, see 2.4, “Snapshots” on page 74.

Mirror

This option allows the storage administrator to create and manage mirrored volumes and is discussed in detail in *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.

HyperSwap

This option allows the storage administrator to create and manage a HyperSwap relationship and is discussed in detail in *IBM HyperSwap and Multi-site HA/DR for IBM FlashSystem A9000 and A9000R*, REDP-5434.

Mapping

Mapping volumes to a host is the next option under the Actions menu as shown in Figure 2-47. First, the storage administrator selects the volume or multiple volumes. Then, the storage administrator selects **Mapping** → **View/Modify Mapping** to display the Volume Mapping view on the right side of the window.

x Actions 2 Volumes' Properties	
Properties >	View/Modify Mapping
Snapshots >	Unmap All
Mirror >	
HyperSwap >	
Mapping >	

Figure 2-47 Mapping volumes to a host

If no mappings exist for the volumes, click the plus sign (+) to add a mapping to a host or cluster, as shown in Figure 2-48.



Figure 2-48 Create a volume mapping

This action displays the option to select a host or a cluster, the host name, and the logical unit number (LUN) for the volume. If multiple volumes are selected for mapping, the LUN option is set to **Auto**. An example of mapping a single volume to a host is shown in Figure 2-49.

Figure 2-49 Volume mapping host selection

Click **Apply** after the required selections are made and the volumes are ready to be used by the host system. For more information about host attachment and mapping, see *IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM XIV Storage System: Host Attachment and Interoperability*, SG24-8368.

Alternatively, as previously mentioned, a volume can be mapped from the Hub view by clicking the mapping circle, as shown in Figure 2-50, which opens the same mapping view starting in Figure 2-48 on page 65.

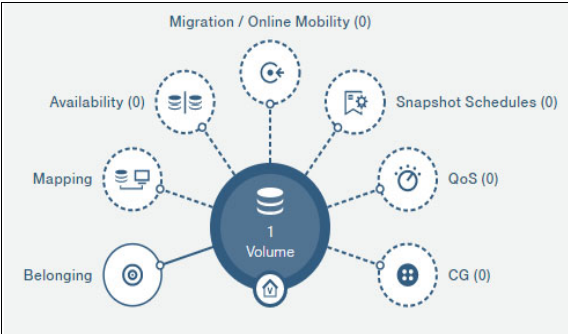


Figure 2-50 Volume mapping in the Hub view

Volumes can also be mapped from the Host view by selecting **Hosts** from the **Hosts & Clusters Views** in the left menu. Then, select the specific host to which you map volumes. This procedure was changed in HSM V5.4 and is shown in Figure 2-51. By hovering your cursor over a LUN ID that displays “Map here...”, you can create a volume and map it, map existing volumes, and map snapshots.

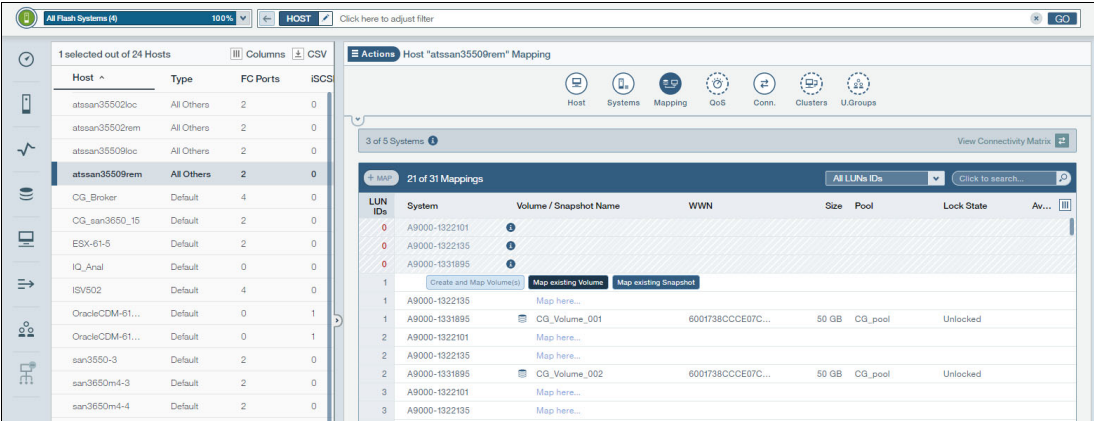


Figure 2-51 Volume mapping by host

You can also select a volume and perform actions on it. An example of the menu is shown in Figure 2-52.




Figure 2-52 Volume actions within mapping

Consistency groups

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software enable a higher level of volume management by grouping volumes and snapshots into sets that are called *consistency groups*. This grouping is useful for cluster-specific volumes.

Figure 2-53 shows the volume's Actions menu for consistency groups. It includes the following options:

- ▶ Move a volume to a consistency group
- ▶ Remove a volume from a consistency group

The **Remove from Group** option is available if the volume belongs to a consistency group.

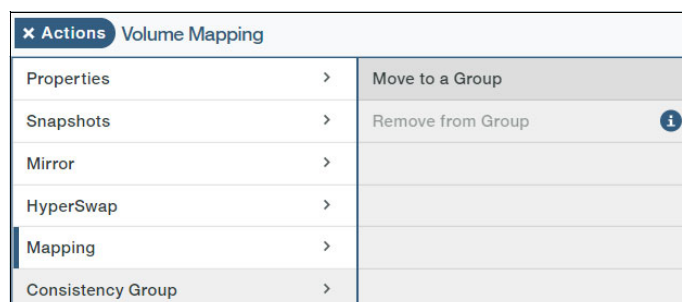


Figure 2-53 Consistency group options for a volume

Consistency groups are covered in detail in 2.5, “Consistency groups” on page 87.

Quality of service

The quality of service (QoS) feature is managed through the definition of performance classes and then the association of a volume with a performance class. To associate a volume with a performance class, select **QoS** on the volume's Actions menu as shown in Figure 2-54.

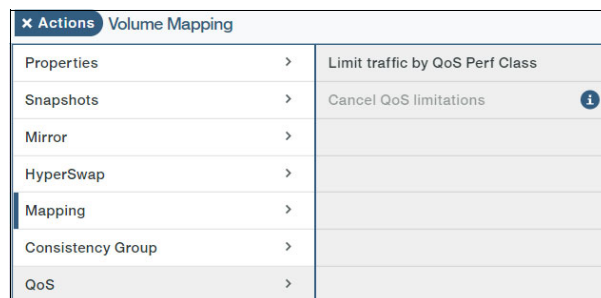


Figure 2-54 Associate a volume with a performance class

Then, the storage administrator optionally can assign an existing performance class to the volume or create a performance class to associate with the volume. To create a performance class, click **Create QoS performance class** as shown in Figure 2-55.




Figure 2-55 Create a performance class

After this option is selected, a new tab opens so that you can create the performance class as shown in Figure 2-56.

↑

Name

ITSO_QoS

System

A9000

BW Limit (MBps per interface) ⓘ

IOPS limit (per interface)

100

☒ Shared ⓘ
 ☐ Independent

Figure 2-56 Creating a performance class

In this example, we opted to limit the IOPS to 100. An option also exists to make the performance class shared or independent. For a *shared* performance class, all of the members that are assigned to this class share the 100 IOPS. If the *independent* option is selected, each member that is associated with this class is limited to 100 IOPS.

After the performance class is created, you can return to the Volume tab, as shown in Figure 2-57, to associate the new class to the volume. The performance class association can be canceled at any time. All changes that relate to QoS are immediate.

QoS Name

ITSO_QoS

CREATE QoS PERFORMANCE CLASS

QoS	ITSO_QoS
-----	----------

Figure 2-57 Volume tab that is pending QoS application

Important: If more than one object (volume, host, domain, or pool) is associated with a performance class, all of those objects either share the limitation that is defined in that performance (when *shared* is selected) or each object is limited by the performance class (when *independent* is selected).

As with most of the options in the Actions menu, you can get to the same view by selecting the QoS circle in the Hub as shown in Figure 2-58.




Figure 2-58 Select QoS option from the Hub view

For more information about QoS, see *IBM FlashSystem A9000 and A9000R Architecture and Implementation (Version 12.3.1)*, SG24-8345.

The Hyper-Scale Manager gives the storage administrator the option to view statistics from the volume's Actions menu as shown in Figure 2-59. With this option, the storage administrator can view the performance and capacity statistics of a volume.

✕ Actions Volume QoS Performance Class	
Properties	Show Performance Statistics
Snapshots	
Mirror	
HyperSwap	
Mapping	
Consistency Group	
QoS	
Pool Capacity	
Statistics	

Figure 2-59 Viewing volume statistics

Additionally, at the bottom of the tabular view for the volumes, the statistics for a selected volume or volumes are visible. An example is shown in Figure 2-60.

Volume A90001_Exchange_vol_1 Statistics:	IOPS (total) 56,587 IOPS ↑	Latency (max) 0.26 ms ↑	Bandwidth (total) 488.28 ... ↑
------------------------------------------	----------------------------	-------------------------	--------------------------------

Figure 2-60 Volume Statistics view

This view can be expanded to show up to two or three graphs (depending on the size of your browser window) at a time as shown in Figure 2-61 by clicking the up arrow in one of the blocks to expand the IOPS (total), latency (avg), or bandwidth (avg) in Figure 2-60. This can be either current data or historical data. Any of the graphs can be closed by clicking the x in the upper-right corner of the graph.




Figure 2-61 Volume statistics

Note: All statistics views can show current and historical data.

Migration

Under the Migration option of the Actions menu, the storage administrator can set up a migration from another system to this volume. This method is a secondary path to set up migration. Before you use this path, you must first establish a connection to the source system (where the data to migrate is located).

Other options that are available under this action include changing the activation state (such as activate or pause), testing the connection, and disconnecting the migration. These options are shown in Figure 2-62.

X Actions Volume Properties	
Properties >	Migrate to this Volume
Snapshots >	Change Activation State i
Mirror >	Test i
HyperSwap >	Disconnect i
Mapping >	
Consistency Group >	
QoS >	
Pool Capacity >	
Statistics >	
Migration >	

Figure 2-62 Migrate to a volume

For more information about migration, see *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.

Online volume mobility

You can move a volume to another system with Online Volume Mobility feature. This feature is only available for IBM FlashSystem A9000 and between XIV Gen 3 systems and IBM Spectrum Accelerate. The Online Volume Mobility feature is described in detail in *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.

X Actions Volume Properties	
Properties >	Mapping Availability Mobility Scheduler
Snapshots >	
Mirror >	
HyperSwap >	Lock State
Mapping >	Locked
Consistency Group >	
QoS >	
Pool Capacity >	Before Compression (GB) Volume
Statistics >	200
Migration >	systems A9000/R physical usage (stored) is man
Online Volume Mobility >	Move Volume to other System
Volume Data >	View Online Volume Mobility i
Compression >	Change activation state i
Snapshot Scheduler >	Start Proxy i
Delete >	End Proxy i
N/A	Cancel Online Volume Mobility i

Figure 2-63 Online Volume Mobility actions

For more information about Online Volume Mobility, see *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.

Volume data

You can use the Volume Data action to format or copy a volume. These menu options are shown in Figure 2-64.

* Actions Volume Properties	
Properties	> Format
Snapshots	> Copy This Volume
Mirror	>
HyperSwap	>
Mapping	>
Consistency Group	>
QoS	>
Pool Capacity	>
Statistics	>
Migration	>
Online Volume Mobility	>
Volume Data	>

Figure 2-64 Volume data actions

Formatting a volume destroys all of the data that existed on the volume. A formatted volume returns zeros as a response to any read command. The volume is formatted logically, and no data is written to the physical storage space that is allocated for that volume.

A *volume copy* differs from a snapshot because a volume copy does not create a point-in-time copy as a child device of the original volume. The volume copy command copies all of the data in a volume at the time that the command is issued into a separate and empty volume, which is independent of the source volume.

The advantage of a volume copy over a snapshot is that a volume copy is independent, it is not at risk of automatic deletion if pool space becomes constrained. Therefore, a volume copy target can also be in a different pool than the source volume.

To copy a volume, select the volume, then select **Actions** → **Volume Data** → **Copy This Volume**. This action opens the Copy Volume view. Complete the necessary steps by specifying the destination pool and the volume within that pool. If the volume does not exist, you optionally can create a volume, as shown in Figure 2-65.

Copy Volume

Destination Pool *

CS_A9000_Pool_001

Destination Volume *

CREATE VOLUME

Volume

CS_Volume_100GB_002

CS_Volume_100GB_003

Figure 2-65 Copy a volume

Compression

This option is available for IBM XIV Gen3. It is not available for other systems and volumes that are not in a thin pool. The options here include compressing a volume that meets the criteria, decompressing a compressed volume or canceling the action. An example of the options available for a compressed volume is shown in Figure 2-66.

* Actions Volume Properties	
Properties >	Compress i
Snapshots >	Decompress
Mirror >	Cancel i
HyperSwap >	
Mapping >	
Consistency Group >	
QoS >	
Pool Capacity >	
Statistics >	
Migration >	
Online Volume Mobility >	
Volume Data >	
Compression >	

Figure 2-66 Compression actions

Snapshot schedule

The snapshot scheduler is set up at a pool level. Under a volume view, this option allows the storage administrator to view an existing snapshot schedule. The view also indicates that no schedule exists and allows the user to go to the pool view to set up a new schedule.

Delete

The last item in the list under the Actions menu offers you the option to delete the selected volume or volumes. The GUI prompts for confirmation before the selected volumes are deleted.

Deleting a volume also deletes all of the associated snapshots, even snapshots that are part of a snapshot group. A volume can be deleted even if the volume is in the lock state. A volume that is mapped to a host or part of a consistency group cannot be deleted.

Important: After a volume or a snapshot is deleted, all of the data that was stored on the volume is lost, and the data cannot be restored.

2.3.3 Managing volumes with the command-line interface

All of the actions that can be performed on volumes in the GUI can also be run with the CLI. To get a list of the commands that relate to volumes, issue the following command in the CLI session:

```
help category=volume
```


This **help** command displays all of the commands that relate to volumes that are available in the CLI with a description of the command. Entering a command in a CLI session and then pressing the Tab key displays either the next character or a list of possible commands that start with the same letters. For example, if you entered `vol_` and then pressed the Tab key, the information that is shown in Example 2-8 is displayed.

Example 2-8 List of commands that start with vol_

```
ITS0_2_A9000R>>vol_
vol_allocation_bitmap  vol_bind_list          vol_clear_external_id
vol_copy               vol_create              vol_create_many
vol_delete             vol_diff                vol_diff_copy
vol_format             vol_list                vol_list_extended
vol_lock               vol_mapping_list        vol_move
vol_rename             vol_resize              vol_set_external_id
vol_unlock
```

The use of the Tab key in this way can help you use the correct syntax and options within a command, too. See Example 2-9.

Example 2-9 Parameter list for the command vol_list

```
ITS0_2_A9000R>>vol_list
vol=          pool=          cg=          show_proxy=  managed=     domain=
wwn=
```

The CLI gives you the ability to list all volumes without any of the options or list a subset of volumes, including the associated snapshots, by using an option, such as `pool`, as shown in Example 2-10.

Example 2-10 List all of the volumes in a pool

```
ITS0_2_A9000R>>vol_list pool=itso_pool1
```

Name	Size (GB)	Master Name	Consistency Group	Pool	Creator
itso_vol1_001	20		itso_cg1	itso_pool1	admin
itso_vol1_002	20		itso_cg1	itso_pool1	admin
itso_vol12_copy	20			itso_pool1	admin
itso_vol6	40			itso_pool1	admin
itso_vol7	40			itso_pool1	admin
itso_vol13	150			itso_pool1	admin
itso_vol14	150			itso_pool1	admin
itso_vol15	150			itso_pool1	admin
itso_vol1_001.snapshot_00003	20	itso_vol1_001		itso_pool1	admin
itso_vol1_002.snapshot_00002	20	itso_vol1_002		itso_pool1	admin

To format an existing volume to remove all data and reuse the volume, run the following command:

```
vol_format vol=itso_vol6
```

All of the data on the volume is lost and unrecoverable. To bypass the warning message, use **-y** after the **vol_format** command.

Resizing a volume by using the CLI is supported only for increasing the size of the volume. To increase the size of a volume, run the following command:

```
vol_resize vol=itso_vol_001 size=500
```

The system automatically increases the size of the volume over the requested size to the nearest 103 GB boundary if the volume is within 5% of that number. The system rounds up to only that number. In our example, the resulting size of the volume, by using the command that is shown in Example 2-11, is 516 GB.

Example 2-11 Increasing the size of a volume

```
ITS0_2_A9000R>>vol_resize vol=itso_vol1_001 size=500
Warning:  Are you sure you want to increase volume size? y/n: y
Command run successfully.
```

You can then list the volume to see its new size by issuing the following command:

```
vol_list vol=itso_vol1_001
```

Shrinking a volume is supported in the GUI (and used with caution). Shrinking a volume is not supported in the CLI. If the resize command is issued to shrink a volume, the CLI command fails with a `CAN_NOT_SHRINK_VOLUME` error message.

To move a volume from one pool to another pool to either clear space or consolidate similar volumes in a single pool, issue the **vol_move** command as shown:

```
vol_move vol=itso_vol7 pool=itso_pool2
```

This volume move is a simple task because the pools are only logical containers for volumes and capacity. The pointer to the volume is moved from one pool to another pool without affecting the location of the data.

Volumes can also be renamed without affecting the host I/O. To rename an existing volume, issue the following command:

```
vol_rename new_name="vol_new_name" vol="vol_old_name"
```

To delete an existing volume, issue the following command:

```
vol_delete vol="vol_to_delete"
```

2.4 Snapshots

A *snapshot* is a point-in-time copy of a volume's data.

The snapshot of IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software uses several innovative technologies to ensure minimal impact on system performance. Snapshots are space-efficient (redirect on write), use pointers, and contain only partitions with data that changed from the original volume. If a partition did not change, both the snapshot and the volume point to the same partition.

Because the source volume and the snapshot volume differ over time, snapshots require capacity. Space for snapshots must be set aside when you define a storage pool. The snapshot reserve space requires a minimum of 400 GB in IBM FlashSystem A9000 and A9000R, and 17 GB in IBM XIV Gen3 and IBM Spectrum Accelerate software.

An application can use many volumes on IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software. For example, a database application can span several volumes for application data and transaction logs. In this case, the snapshot for the volumes must occur at the same moment in time so that the data and logs are consistent. Creating a snapshot group of the consistency group ensures that all of the volumes that are assigned to the group are snapped at the same moment, which ensures data consistency throughout the group.

The creation and management of snapshots with IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software are straightforward. Snapshots are easy to perform. This section guides you through the lifecycle of a snapshot, providing examples of how to interact with the snapshots by using the GUI and the CLI.

2.4.1 How to create a snapshot of a volume

A snapshot of a volume or snapshots of multiple volumes can be created by using either the GUI or the CLI. To use the GUI, select the volume (or volumes) that you want. Select **Actions** → **Snapshots** → **Create Snapshot(s)**. You can see this path in Figure 2-67.




Figure 2-67 Create a volume snapshot

The view on the right side of the GUI displays the Inspector for this volume and the window to create the snapshot. If a single volume was selected, you can rename the volume. If multiple volumes were selected for snapshots, the name is generated by the system but you can change the name later.

The deletion priority can be selected now. An example of multiple snapshots with deletion priority selection is shown in Figure 2-68.

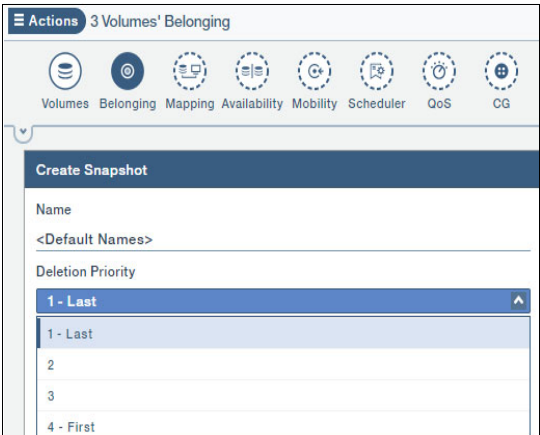


Figure 2-68 Create snapshots of multiple volumes

The new snapshots can be viewed by clicking the **Sn** element in the Hub, as shown in Figure 2-69.




Figure 2-69 Hub view that is expanded for snapshots

After you click the **Sn** circle, the GUI shows the table view with all of the snapshots and their properties and the Inspector view, as shown in Figure 2-70.

Name	Source Volume
Multiple Values	Multiple Values
Modification Status	Deletion Priority
Not Modified	4 - First
Size (GB)	Size On Disk (GB)
540	540
Created (System Time)	Created on Source
Multiple Values	Multiple Values
Lock state	ID
Locked	Multiple Values

Figure 2-70 Multiple snapshots Hub view

The deletion priority in this example is set to **4 - First**, but it can be changed by using the drop-down list. The deletion priority is the only field that can be changed when multiple items are selected in the table view (this is the default if there are multiple snapshots of a volume). When a single snapshot is selected, the name and size can also be modified, as shown in Figure 2-71.

Name	Source Volume
demo_vol_001.snapshot_	demo_vol_001
Modification Status	Deletion Priority
Not Modified	4 - First
Size (GB)	Size On Disk (GB)
100	103

Figure 2-71 Modifying the snapshot properties

The table view contains the specific details about each snapshot, as shown in Figure 2-72. The columns view can be fully shown by moving the slider at the bottom of the window or by expanding the table view. The column cell views can be increased to see the entire field (such as the name of the snapshot) or decreased to see less by hovering the cursor over the name of the column to see the dividers and then by moving the divider to the wanted width. Columns can be added or removed by right-clicking the column title, clicking the **Columns** icon to see the additional fields that are available, or right-clicking the column title area, as shown in Figure 2-72.

Snapshots				Columns	CSV	Action
	System	Pool	Source Volume	Add/Remove Columns		
...	A9000R-Site_A	ITSO_A	ITSO_QW_Volu...	<input type="checkbox"/> Properties	>	
...	A9000R-Site_A	ITSO_A	ITSO_QW_Volu...	<input type="checkbox"/> Belonging	>	
...	A9000R-Site_A	ITSO_A	ITSO_QW_Volu...	<input checked="" type="checkbox"/> Mapping	>	
1...	A9000R-Site_A	ITSO_A	ITSO_SS_Volum...	<input type="checkbox"/> QoS	>	

Figure 2-72 Adding columns to the snapshots table

From the view that is shown in Figure 2-72, you can see other details:

- ▶ The default name of the snapshot follows a specific format. The first part is the name of the volume followed by the word snapshot and then a number, or count, of snapshots for that volume. A snapshot is the same size as the volume.
- ▶ The locked property of the snapshot. By default, a snapshot is locked when it is created, which means that it is read-only.
- ▶ The modified property is displayed next to the locked field. The value changes to yes after a snapshot is unlocked.

2.4.2 Snapshot deletion priority

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software use an automatic snapshot deletion mechanism to protect against overutilizing the snapshot space in a pool. Snapshot space overutilization can occur as new data is written to a volume, which means that the snapshot now must contain the original data from the point in time that the snapshot was taken. It can also happen as snapshots are created.

Each snapshot has a deletion priority property that is set by the user at the time of creation (and it can be modified later). Five deletion priorities from priority 4 to priority 0 (which is also known as the *Golden Snapshot*). Priority 0 is used to protect snapshots that must *not* be subject to automatic deletion and is only available on XIV and IBM Spectrum Accelerate systems and snapshot groups on FlashSystem A9000 and A9000R. The system uses this priority to determine which snapshot to delete first.

When the system needs to delete a snapshot to make room for a new snapshot (or space in the pool), it deletes the oldest snapshot with deletion priority 4. The system deletes all of the snapshots with priority 4 before it deletes snapshots with priority 3. Then, the system deletes the snapshots with priority 2. The last snapshots that are deleted by the system are the deletion priority 1 snapshots. The default deletion priority in the GUI is priority 1.

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software provide alerts based on the percentage of the snapshot space that is used in a pool. As you receive higher-level alerts about snapshot usage, the snapshot space of the affected pool can be manually expanded to accommodate more snapshots or snapshots can be manually deleted before the automatic deletion occurs. For more information about viewing and modifying the default snapshot thresholds, see “Policy controls” on page 49.

2.4.3 Managing snapshots with the GUI

To manage snapshots from the Storage Management GUI, select **Pools & Volumes Views** → **Snapshots**. This action lists all of the snapshots on the systems in a single table view. Selecting any of the snapshots in the table enables the Actions menu.

The following actions are available:

- ▶ Properties: View and modify snapshot properties, including the lock state
- ▶ Mapping: Map snapshots, unmap snapshots, and modify snapshot mappings
- ▶ QoS: Set and cancel a performance class for a snapshot
- ▶ Statistics: View statistics for a snapshot
- ▶ Snapshot Data: Format, modify, duplicate, copy, and restore a snapshot
- ▶ Delete: Delete a snapshot

Properties

The Properties view of a snapshot also includes the ability to rename a snapshot, change the deletion priority, resize a snapshot, and change the lock state of a snapshot. These options are shown in Figure 2-73.

x Actions Snapshot Properties	
Properties >	View Properties
Mapping >	Rename
QoS >	Change Deletion Priority
Statistics >	Resize
Snapshot Data >	Change Lock state
Delete >	

Figure 2-73 Snapshot properties options

Although these options are easy to run, let’s look at unlocking a snapshot. By default, a snapshot is locked when it is created, and is a read-only snapshot. By unlocking a snapshot, you can modify the data in the snapshot for testing, development, or data mining without affecting the volume.

If the data needs to be maintained in the snapshot, you can create a duplicate snapshot for the data while the original snapshot is unlocked and modified. After the modifications and testing are complete, the snapshot can be deleted, keeping the duplicate snapshot that still points to the volume. You can also copy a snapshot to perform the same activities on the copy instead of a duplicate or the original snapshot.

To unlock a snapshot, select a snapshot from the Snapshots view, then select **Actions** → **Properties** → **Change Lock State** to change the state as shown in Figure 2-74.




Figure 2-74 Change the lock state of a snapshot

The state of the snapshot can be toggled between unlocked or locked. The current state of the snapshot is selected. In this example, the state is locked.

To save changes that were made to a snapshot while it was unlocked, lock the snapshot by using the same process as unlocking. The locking process completes immediately, preventing further modification to the snapshot. In this case, the snapshot might no longer be a point in time copy of the source volume. A modified snapshot can be overwritten by a current point in time copy of the volume by selecting **Actions** → **Snapshot Data** → **Overwrite by Volume**.

Note: A snapshot that is unlocked and then locked again displays Yes under the Modified column in the Snapshots view. You might need to add the Modified column to the default view.

Another option is to create a duplicate or copy of the snapshot to preserve the original snapshot as a backup of the volume while the duplicate snapshot is modified after the duplicate snapshot is unlocked. The option to duplicate or copy a snapshot is under the Snapshot Data option in the Actions menu. Copying a snapshot requires a volume to which to copy. You can specify an existing one or create a volume.

Mapping

Mapping a snapshot is the same as mapping a volume. One difference is that after a snapshot is mapped to a host, a locked snapshot is in a read-only state. To use the snapshot for writing, it must first be unlocked.

Quality of service (QoS)

A performance class might need to be applied to an unlocked snapshot that is used for writing or even for backup operations. Select **Actions** → **QoS**, and either create a QoS definition or assign an existing performance class to the snapshot. The menu option is shown in Figure 2-75.




Figure 2-75 Setting a performance class for a snapshot

Statistics

To view the statistics that are associated with a volume, select **Actions** → **Statistics** → **Show Performance Statistics** as shown in Figure 2-76.

× Actions Snapshot Properties	
Properties >	Show Performance Statistics
Mapping >	
QoS >	
Statistics >	

Figure 2-76 Performance statistics for a snapshot

This action opens the Statistics view under the Snapshot table view. This action is similar to the Volume Statistics view that was shown in Figure 2-61 on page 69.

Snapshot data

Other actions that you can perform on a snapshot are listed under Snapshot Data. These options include formatting a snapshot, overwriting the snapshot from the original volume, duplicating the snapshot (which is taking a snapshot of the snapshot), copying the snapshot, and restoring a volume from a snapshot. These options are shown in Figure 2-77.

× Actions Snapshot Properties	
Properties >	Format
Mapping >	Overwrite by Volume
QoS >	Duplicate to a new Snapshot
Statistics >	Copy this Snapshot
Snapshot Data >	Restore Snapshot

Figure 2-77 Snapshot data options

Consider overwriting a snapshot. For regular backup jobs, you can decide to create snapshots, letting the system delete the old snapshots, or overwriting the existing snapshots with the latest changes to the data. For instance, a backup application requires the latest copy of the data to perform its backup operation.

The overwrite operation modified the pointers to the snapshot data to be reset to the volume. All of the pointers to the original data in the snapshot are lost, and the snapshot appears as new. Storage that was allocated for the data changes between the volume and its snapshot is released.

To perform this operation, select **Snapshot Data** → **Overwrite by Volume**. The GUI displays the request for your confirmation, as shown in Figure 2-78. Click **Apply**.

Over write Snapshot by its Volume?	
• CS_Volume_100GB_003.snapshot_	

Figure 2-78 Overwrite a snapshot

You can restore the data from a snapshot of a snapshot back to the original snapshot. This action can be helpful for operations where data was modified incorrectly and the data needs to be restored to the snapshot. Select **Snapshot Data** → **Restore Snapshot**. The system prompts you to select from a drop-down list the snapshot to restore from, which is useful when more than one snapshot exists to choose from.

You can restore a snapshot from its source snapshot. Select the snapshot in the Snapshot view, and then select **Actions** → **Snapshot Data** → **Restore Snapshot**.

Select **Apply**, as shown in Figure 2-79.




Figure 2-79 Restore a snapshot

After you restore a snapshot from its source snapshot, the process will modify the pointers to the snapshot so that they are now equivalent to the source snapshot pointers. This change occurs only for the partitions that are modified from the time that the source snapshot was taken. The snapshot pointer does not change and continues to point to the original data. The restore process then frees the modified partition space.

You can also restore a volume from a snapshot by using a similar process from the volumes view. If a snapshot is taken and the original volume later increases in size, you can still perform the restore operation. The snapshot has the original volume size, and it restores the original volume correctly.

Deleting a snapshot

When a snapshot is no longer needed, it can easily be deleted. Select the snapshot in the Snapshots view. Then, click **Actions** → **Delete** → **Delete Snapshot(s)**, as shown in Figure 2-80.

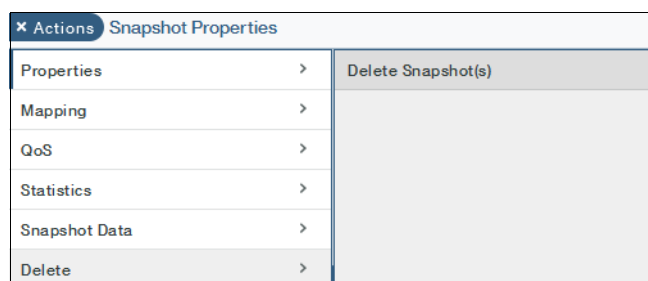


Figure 2-80 Deleting snapshots

The snapshots are displayed on the right side of the window. Confirm that these snapshots are the correct snapshots to delete and click **Apply**. An example can be seen in Figure 2-81.




Figure 2-81 Delete snapshots

Reminder: If a volume is deleted, all of the associated snapshots are also deleted.

2.4.4 Managing snapshots with the CLI

The CLI commands are straightforward in their functions. For a full listing of commands and usage, see *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide*, SC27-8711. The list of basic commands for a volume snapshot are displayed in Figure 2-82. More commands that relate to snapshots can be listed by issuing the following command:

```
help search=snapshot
```

Category	Name	Description
snapshot	snapshot_list	Use the snapshot_list command to list snapshot information.
snapshot	snap_group_format	Use the snap_group_format command to format a snapshot group.
snapshot	mirror_create_snapshot	Use the mirror_create_snapshot command to create a snapshot mirror.
snapshot	mirror_cancel_snapshot	Use the mirror_cancel_snapshot command to cancel all snapshot mirrors ('ad-hoc' sync jobs) of a specified master volume or a master consistency group, that have not run yet.
snapshot	snapshot_delete	Use the snapshot_delete command to delete a snapshot.
snapshot	snapshot_format	Use the snapshot_format command to format a snapshot.
snapshot	snapshot_create	Use the snapshot_create command to create a snapshot of an existing volume.
snapshot	snapshot_create_many	Use the snapshot_create_many command to create multiple snapshots of an existing volume.
snapshot	snapshot_duplicate	Use the snapshot_duplicate command to duplicate an existing snapshot.
snapshot	snapshot_restore	Use the snapshot_restore command to restore a master volume or a snapshot from one of its associated snapshots.
snapshot	snapshot_change_priority	Use the snapshot_change_priority command to change a snapshot's deletion priority.

Figure 2-82 Basic snapshot commands

To create a basic snapshot, use the **snapshot_create** command, as shown in Example 2-12. In this example, we specified the deletion priority. If this option is not used, the default priority is used. The other options for this command are shown:

- ▶ name
- ▶ overwrite
- ▶ ext_id

Example 2-12 Create a snapshot

```
ITS0_2_A9000R>>snapshot_create vol=itso_vol1_001 delete_priority=1
Command run successfully.
```

A snapshot can also be duplicated for testing and backups. Figure 2-83 illustrates an example of duplicating a snapshot, which is followed by listing all of the snapshots that are associated with the same volume.

```
ITS0_2_A9000R>>snapshot_duplicate snapshot=itso_vol1_001.snapshot_00002
Command executed successfully.
ITS0_2_A9000R>>snapshot_list vol=itso_vol1_001
Name                               Size (GB)  Master Name  Consistency Group  Pool      Creator
itso_vol1_001.snapshot_00002      20        itso_vol1_001  itso_vol1_001     itso_pool1  admin
itso_vol1_001.snapshot_00003      20        itso_vol1_001  itso_vol1_001     itso_pool1  admin
ITS0_2_A9000R>>
```

Figure 2-83 Duplicate a snapshot

The CLI also allows a snapshot to be restored from another snapshot. Example 2-13 shows a more recent snapshot (00003) that is restored to a later snapshot (00002). This process results in the pointers for both snapshots that now point to the same data.

Example 2-13 Restoring a snapshot from another snapshot

```
ITS0_2_A9000R>>snapshot_restore snapshot=itso_vol1_001.snapshot_00003
target_snapshot=itso_vol1_001.snapshot_00002
```

Formatting a snapshot is also a simple process with the CLI as shown:

```
snapshot_format snapshot=itso_vol1_001.snapshot_00003
```

Unlock a snapshot by using the **vol_unlock** command:

```
vol_unlock vol="vol_to_unlock"
```

Delete a snapshot by using the **snapshot_delete** command, as shown in Example 2-14.

Example 2-14 Deleting a snapshot

```
ITS0_2_A9000R>>snapshot_delete snapshot=itso_vol1_001.snapshot_00002
```

Command run successfully.

2.4.5 Snapshot Scheduler

The Snapshot Scheduler is an option available with Hyper-Scale Manager starting with V5.3 that allows the system administrator to automate snapshots in a pool on a defined schedule. The schedule is applied to existing and new volumes and consistency groups within the pool. The Snapshot Scheduler is available for the entire IBM Spectrum Accelerate family that is managed by the Hyper-Scale Manager.

To work with the scheduler, select a pool in the pools view. Then, under the **Actions** menu, select **Snapshot Scheduler**. The following options are available:

- ▶ View/Modify Snapshot Schedules: Allows you to make change to existing schedules.
- ▶ Add Snapshot Scheduler: Allows you to create a new schedule.
- ▶ View all the related Snapshots: Lists those snapshots that were created as a result of the scheduler.
- ▶ Volumes missing automated Snapshots
- ▶ CGs missing automated Snapshots
- ▶ Remove all Snapshot Schedules

These options are shown in Figure 2-84.

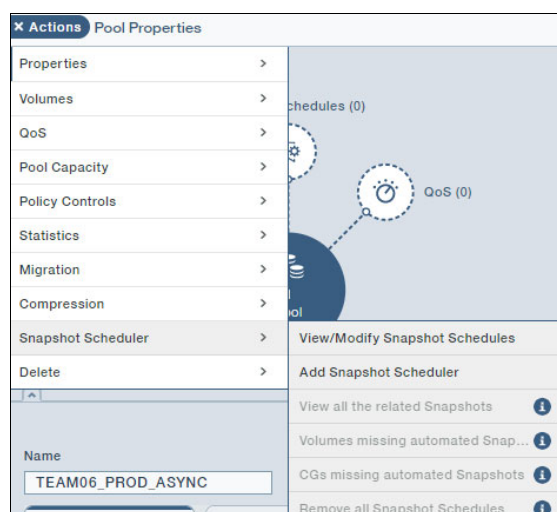


Figure 2-84 Snapshot Scheduler actions

Creating a Snapshot Schedule

This section reviews the steps that are required to create a schedule and some of the options available within adding a schedule.

To create a schedule, select the **Add Snapshot Scheduler** option for the pool. You need to give the schedule a name and define what it applies to from the following options (as shown in Figure 2-85):

- ▶ Volumes & Consistency Groups
- ▶ Volumes (only Volumes not in Consistency Groups)
- ▶ Consistency Groups (only Volumes in Consistency Groups)

After you create a schedule, it affects all existing and new volumes or consistency groups that are created after the schedule is created.



Figure 2-85 Selecting a volume or consistency group

Next, define the time interval (daily, varying weekly intervals for specific days, monthly, and every other month), as shown in Figure 2-86.



Figure 2-86 Selecting a time period

After selecting the time period, select the specific time for the Snapshot Scheduler to run. There can be more than one time slot for a time period. Creating an additional period to run in that time period, click the plus sign (+) to get a new timebox, as shown in Figure 2-87.

Figure 2-87 Schedule multiple time slots

You also need to define the retention policy. This number indicates how many snapshots for the volume or consistency group that are kept on the system. The HSM displays approximately how long the snapshot is retained. If you need snapshots to stick around for longer, creating a schedule for more snapshots to be kept might be appropriate. However, the schedule applies to all volumes or consistency groups within the pool.

Figure 2-88 shows an example where snapshots created by a schedule that is kept for approximately 2 days.

Figure 2-88 Snapshot retention definition

Figure 2-89 shows an example of creating two different types of schedules in the same pool. Notice that the schedule on the left is taken daily for all volumes and consistency groups, and the one on the right is for volumes (not in a consistency group) and is taken Monday through Friday.

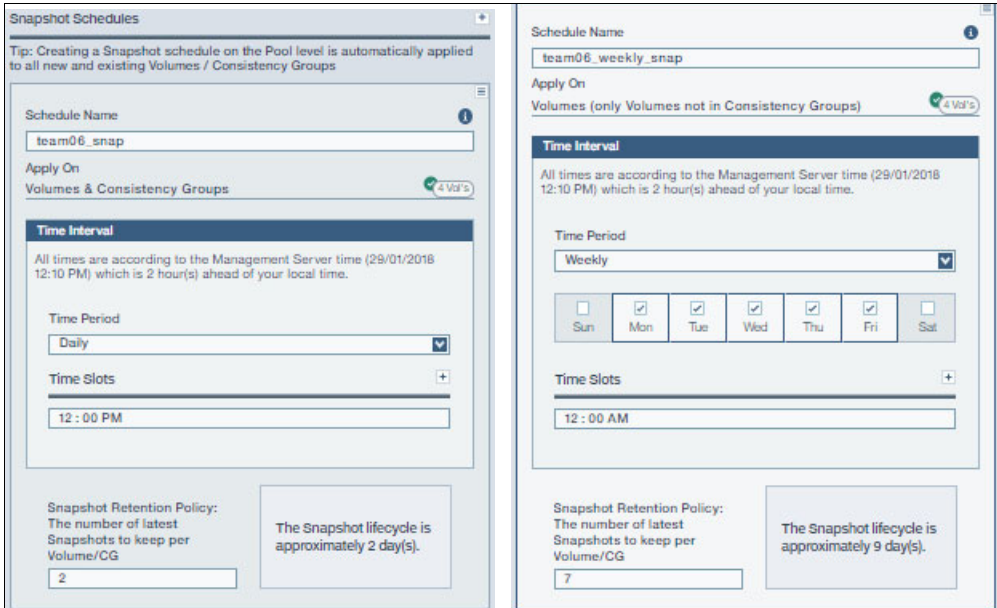


Figure 2-89 Multiple snapshot schedules

The snapshot schedules can be modified, removed, and viewed by using the Actions menu option (Figure 2-90). You can view missing snapshots for volumes and consistency groups if applicable.

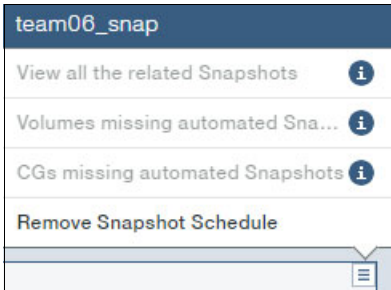


Figure 2-90 Snapshot schedule actions menu

Viewing snapshot schedules

You can view all the pools containing snapshot schedules in all the systems that are managed by the Hyper-Scale Manager by selecting the **Snapshot Scheduled Pools** option under the Pools & Volumes Views menu as shown in Figure 2-91.

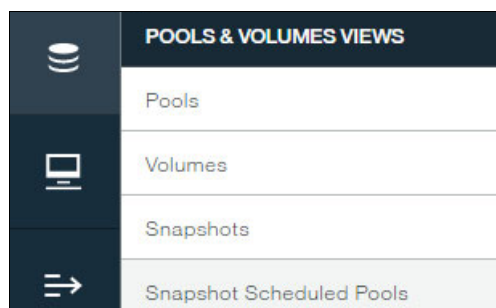


Figure 2-91 Pools with snapshot schedules

Selecting this option opens a customized filter that lists all the pools with schedules. Selecting one of the pools from the list displays the snapshot schedules in the inspector area. Figure 2-92 shows the list of pools.






5	All Systems (5)	▼	POOL	with Snapshot schedule - Yes	 	Click here to adjust filter
1 selected out of 4 Pools						
Pool ^	System	Domain	Pool Size	Snaps...	Snaps...	
team06_pool1	A9000-1322101 	/Global Spa... 	2,067 GB	OK	2	
team19_pool1	A9000R-132...	/Global Spa...	2,067 GB		team19...	
zDemo_dec5	A9000-1331895	/Global Spa...	2,067 GB		zDemo_...	

Figure 2-92 List of pools with snapshot schedules

2.5 Consistency groups

A *consistency group* is a group of volumes of which a snapshot can be made at the same point in time. Therefore, a consistent image is made of all volumes within the group then. The concept of a consistency group is common among storage systems. It is necessary to perform concurrent operations collectively across a set of volumes so that the result of the operation preserves the consistency among the volumes.

For example, effective storage management activities for applications that span multiple volumes, or creating point-in-time backups, is not possible without first employing consistency groups.

This consistency among the volumes in the group is critical to maintaining data integrity from the application perspective. By first grouping the application volumes into a consistency group, it is possible to later capture a consistent state of all volumes within that group at a specified point-in-time by using a special snapshot command for consistency groups.

When you issue this type of command, the following process results:

1. Complete and destage writes across the constituent volumes.
2. Suspend I/O activity simultaneously across all volumes in the consistency group.
3. Create the snapshots.
4. Resume normal I/O activity across all volumes.

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software manage these suspend and resume activities for all volumes within the consistency group.

Restriction: A consistency group cannot span multiple systems across IBM FlashSystem A9000 and A9000R. In IBM XIV Gen3 and IBM Spectrum Accelerate software, the concept of x-Consistency groups exists and can have volumes from multiple systems.

2.5.1 Creating consistency groups

You can use the following methods to create a consistency group within the GUI:

- ▶ Create the consistency group and add the volumes to it.
- ▶ Choose to move a volume to a consistency group and create the consistency group then.

We first look at the second method. Under the Volumes view, select the volumes to move. The selection enables the Actions menu. Select **Actions** → **Consistency Group** → **Move to a Group**, as shown in Figure 2-93.




Figure 2-93 Moving volumes to a consistency group

This action opens the Volume Consistency Group wizard, as shown in Figure 2-94. The storage administrator can either select an existing consistency group or create a consistency group by clicking **CREATE CONSISTENCY GROUP**. In this example, choose to create a consistency group.



Figure 2-94 Create a consistency group from the move volumes to a group option

After the option is selected, a new window opens that changes the view from volumes to consistency groups. The Create Consistency Group wizard is shown in Figure 2-95. Because the option to create was selected from the Volumes view, all of the fields are completed except the name of the consistency group. Also, on any of the filled-in lines (system, domain, and pool), you can optionally go to that object's properties by clicking the arrow. The arrow is displayed only when the cursor is on that line.



Figure 2-95 New consistency group

This action opens a new view. You can get back to the Create Consistency Group view by selecting the CG tab. An example of the tabs is shown in Figure 2-96. The current tab is white. The tab that is labeled Volume is the Volume Consistency Group tab, which is where we started. The number in red indicates the number of required fields in a tab. You can close any of the tabs to cancel the action.



Figure 2-96 Tabs that are open

After you create the consistency group, you are returned to the volumes view. Select **Apply** to move the volumes to the consistency group you created. The view now shows that the volume or volumes that you selected are in a consistency group, similar to Figure 2-97. Four volumes belong to a consistency group in the Hub display.




Figure 2-97 Hub view of volumes that belong to a consistency group

When you look at the properties of the consistency group, you can see the specific items that relate to it, including the four volumes in this example, and any associated snapshot schedules, snapshot groups and mirrors, as shown in Figure 2-98.




Figure 2-98 Hub view of a consistency group that contains volumes

Another option for creating a consistency group is to select the **NEW** menu and choosing **Consistency Group** as shown in Figure 2-99.

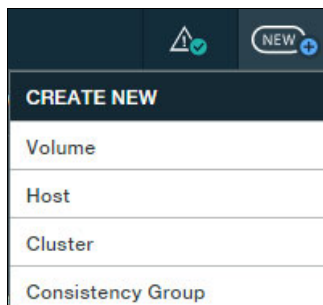


Figure 2-99 Create a consistency group

This action opens the Create Consistency Group window. Enter the required information for the name of the consistency group and select the system, domain, and pool, as shown in Figure 2-100. Doing the action from the Create New menu does not prepopulate the fields as the previous method did.

A screenshot of a 'Create Consistency Group' window. The window has a light blue header with a back arrow icon. Below the header are four sections, each with a label and a text input field or dropdown menu. The first section is 'Name' with the text 'ITSO_RE_cg2'. The second section is 'System' with a dropdown menu showing 'A9000'. The third section is 'Domain' with a dropdown menu showing '/Global Space/'. The fourth section is 'Pool' with a dropdown menu showing 'CS_A9000_Pool_001'. Each section has a small blue arrow icon to its right.

Figure 2-100 Creating a consistency group

After all of the information is entered and applied, the consistency group is created. The consistency group does not contain any volumes yet. The storage administrator can now return to the Volumes view and move volumes to this consistency group.

2.5.2 Adding volumes to a consistency group

To add volumes to a consistency group at any time, select the volumes under the Volumes view and select **Actions** → **Consistency Group** → **Move to a Group**, as shown in Figure 2-93 on page 88.

You can then select an existing consistency group from the drop-down list as shown in Figure 2-101 or create a consistency group as we did in the previous section by using the Create Consistency Group option from the list. In this example, we selected the consistency group, **ITSO_RE_cg2**, from the list. Click **Apply** (not shown), and the action is complete.




Figure 2-101 Select an existing consistency group

Now, the consistency group contains members and you can create *crash consistent* snapshots, either through the GUI or CLI or by using other techniques and application programming interfaces (APIs).

Note: When you add new volumes to a consistency group with snapshot groups, a prompt appears and warns you that the existing snapshot groups will no longer be consistent.

2.5.3 Managing consistency groups by using the GUI

As shown with other configuration tasks, the GUI makes it easy to manage consistency groups. You can set up the view in the consistency group table or continue with the default view (as described at the beginning of this chapter). To access the Consistency Groups view, select **Pools & Volumes Views** → **Consistency Groups** as shown in Figure 2-102.



Figure 2-102 Pools & Volumes Views

After the Consistency Groups option is selected, the list of consistency groups is shown in the table view. To view any of the available actions on a consistency group, a consistency group must be selected from the list.

Tip: You can filter the view or select a system (or subset by type).

The following actions are available:

- ▶ **Properties:** View and modify the properties of a consistency group.
- ▶ **Pool, System:** View and modify the pool that the consistency group belongs to.
- ▶ **Mirror:** Set up and manage a mirrored consistency group.
- ▶ **HyperSwap:** Set up and manage a consistency group in a HyperSwap relationship.
- ▶ **Statistics:** View the consistency group volume and snapshot performance and statistics.
- ▶ **Members:** View and modify the members of a consistency group.

- ▶ Snapshot Groups: Create and modify consistency snapshot groups.
- ▶ Snapshot Scheduler: View an associated snapshot schedule.
- ▶ Delete: Delete a consistency group.

Let us look at some of the options in more detail.

Properties

After a consistency group is selected from the table view, the properties of that consistency group are displayed on the right side of the browser window. Through the Properties tab, you can view and modify the name of the consistency group. This capability is indicated by the box around the editable fields. For consistency groups under properties, the editable field is the name of the group, as shown in Figure 2-103. Editing is a simple task. Place the cursor in the field to edit the name.

Name	ID
ITSO_demo_cg	1f2ae1a600001
CG Written Size (GB)	CG Size (GB)
43	413

Figure 2-103 Consistency group properties

Another part of the Properties view is the graphical Hub display that shows more details about the consistency group. Clicking any of the solid circles displays details about that object. For example, by clicking the **Belonging** circle, the options for viewing systems, domains, and pools are displayed as shown in Figure 2-104.

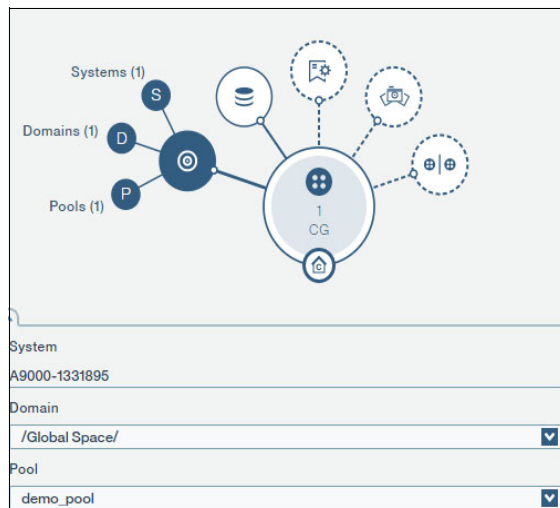


Figure 2-104 Consistency group Hub display

Clicking any of the solid circles displays the Hub details about that object. You can return to the previous window by clicking the back arrow, as shown in Figure 2-105. Starting with HSM V5.4, the back arrow is on the left side of the filter box. Clicking the main CG circle returns the view to the properties of the consistency group. Clicking the lowest level of a circle, such as Pool, changes the table and graphical display to that object.




Figure 2-105 Use the back arrow option after you investigate an object to return to the previous view

Note: A circle with a dashed line indicates that no objects exist. You can click the circle to create an object or association.

Pool and system

To view the Pool, System information that relates to a consistency group, you can also select the option from the Actions menu, as shown in Figure 2-106.




Figure 2-106 View the consistency group pool association

The Consistency Group pool and domain view in Figure 2-107 is then displayed. The storage administrator can view and move the consistency group with all its volumes and snapshots to a different domain or pool, assuming that adequate space is available. Each field has a drop-down list of the available domains and pools in the system.




Figure 2-107 Move a consistency group to another pool

Statistics

You can use the Statistics view for consistency groups to view the performance and capacity statistics for volumes and snapshots within a consistency group. To view the performance statistics, under the Consistency Groups view, select a consistency group and select either **Actions** → **Statistics** → **Show Snapshots Performance Statistics** or **Actions** → **Statistics** → **Show Volumes Performance Statistics**, as shown in Figure 2-108.

x Actions Consistency Group Belonging	
Properties >	Snapshots Performance Statistics ⓘ
Pool, System >	Volumes Performance Statistics
Mirror >	
HyperSwap >	
Statistics >	

Figure 2-108 View statistics that are associated with consistency group members

Members

Members are the volumes and snapshots that make up a consistency group. The storage administrator uses this option under the consistency group Actions menu to view the volumes in a consistency group, add volumes, and remove all members of a consistency group, as shown in Figure 2-109.

x Actions Consistency Group Belonging	
Properties >	
Pool, System >	<div> <div>Scheduler</div> <div>S.Groups</div> <div>Availability</div> </div>
Mirror >	
HyperSwap >	
Statistics >	
Members >	View Volumes
	Add Volume
	Remove all Members

Figure 2-109 Members of a consistency group

Snapshot groups

A consistency group, which consists of multiple volumes, allows a snapshot to be taken of all of the volumes at the same time. This action creates a synchronized snapshot of all of the volumes in the consistency group, and it is suited for applications that span multiple volumes, for example, a database application that stores its data files on multiple volumes.

After the consistency group contains volumes, snapshots can be created by selecting a consistency group in the Consistency Group view and then selecting **Actions** → **Snapshot Groups** → **Create Snapshot Group**, as shown in Figure 2-110. The other available options are viewing, restoring from, deleting, and overwriting a snapshot group.

x Actions Consistency Group Belonging	
Properties >	Create Snapshot Group
Pool, System >	View Snapshot Groups
Mirror >	Restore from Snapshot Group ⓘ
HyperSwap >	Delete all Snapshot Groups ⓘ
Statistics >	Overwrite Snapshot Group ⓘ
Members >	
Snapshot Groups >	

Figure 2-110 Consistency group snapshot groups

After the Create Snapshot Group option is selected, you see a window that is similar to Figure 2-111. The name of the snapshot group and the deletion priority are the only required fields. The name is prepopulated (in this case, ITSO_RE_cg.snap_group_), and can be changed. The default deletion priority is set to 1 - Last, but it can also be changed, depending on your requirements.

Create Snapshot Group

Name

ITSO_RE_cg.snap_group_

Deletion Priority

1 - Last

0 - Golden Snapshot

1 - Last

2

3

4 - First

Figure 2-111 Create a snapshot group

The FlashSystem A9000 and A9000R servers allow a *Golden Snapshot* on a consistency group basis, which means that the system does not automatically delete a snapshot group with priority 0 when it is trying to free space.

Important: The system uses the deletion priority when the system needs to remove snapshots and snapshot groups to clear space.

After the snapshots are created, you can see the snapshots on the Snapshots view. Multiple paths exist in the GUI to perform actions on objects. For example, for snapshot groups, by selecting the snapshot group circle from the Consistency Group Properties view, you can see the Snapshot Group menu, as shown in Figure 2-112.

Actions Consistency Groups' Snapshot Groups ITSO_demo_cg

CG

Belonging

Volumes

Sch

SNAPSHOT GROUPS

Name

ITSO_demo_cg.snap_group_

Rename

Change Deletion Priority

Change Lock State

Restore from Snapshot Group

Restore to Snapshot Group

Overwrite Snapshot Group

Duplicate Snapshot Group

Format Snapshot Group

Disband Snapshot Group

Delete Snapshot Group

Figure 2-112 Snapshot group menu

Tip: Multiple ways to perform the same action on an object are available in the GUI.

Deleting a consistency group

The storage administrator can also remove a consistency group by selecting the consistency group and then by selecting **Actions** → **Delete** → **Delete Consistency Group(s)**, as shown in Figure 2-113. Multiple consistency groups can be selected at the same time for deletion.

✕ Actions Consistency Groups' Snapshot Group	
Properties	> Delete Consistency Group(s)
Pool, System	>
Mirror	>
HyperSwap	>
Statistics	>
Members	>
Snapshot Groups	>
Delete	>

Figure 2-113 Delete a consistency group

Before you delete a consistency group, you must first move all of the volumes and associated snapshots out of the consistency group. If the storage administrator attempts to delete active members, the system prevents the task from completing, as shown in Figure 2-114.




Figure 2-114 Unable to delete a consistency group with members

To remove the members, select the consistency group and select **Actions** → **Members** → **Remove all members**.

The system prompts you for confirmation, as shown in Figure 2-115.




Figure 2-115 Remove members from a consistency group

After all of the members are removed, the consistency group can successfully be deleted.

2.5.4 Managing consistency groups with the CLI

By using the CLI session (or CLI commands), you can also perform the preceding processes. We first look at the commands that relate to consistency groups in the CLI by issuing the following command:

```
help category=consistency-group
```


The output of this command is shown in Figure 2-116.

Tip: Use the search command with help (**help search=cg**) to list the related commands.

ITSO_2_A9000R>>help	category=consistency-group	
Category	Name	Description
consistency-group	cg_list	Lists Consistency Groups.
consistency-group	xcg_list	Lists Cross-system Consistency Group definitions together with contained Consistency Groups.
consistency-group	xcg_create	Creates a Cross-system Consistency Group (XCG) definition.
consistency-group	xcg_add_cg	Associates an existing Consistency Group to a Cross-system Consistency Group definition.
consistency-group	xcg_remove_cg	Removes an existing Consistency Group from a Cross-system Consistency Group definition.
consistency-group	xcg_add_remote_system	Adds a remote system name to the Cross-system Consistency Group definition.
consistency-group	xcg_remove_remote_system	Removes a remote system name from a Cross-system Consistency Group definition.
consistency-group	xcg_delete	Deletes a Cross-system Consistency Group (XCG) definition.
consistency-group	xcg_get_local_cgs	Lists Cross-system Consistency Group definitions together with contained Consistency Groups.
consistency-group	xcg_get_remote_systems	Retrieves the names of remote systems that are a part of the specified Cross-system Consistency Group.
consistency-group	cg_create	Creates a Consistency Group.
consistency-group	cg_add_vol	Adds a volume to a Consistency Group.
consistency-group	cg_remove_vol	Removes a volume from a Consistency Group.
consistency-group	cg_rename	Renames Consistency Groups.
consistency-group	cg_delete	Deletes a Consistency Group.
consistency-group	cg_snapshots_create	Creates a snapshot group of a Consistency Group.
consistency-group	io_pause	Suspends I/O execution on a Consistency Group.
consistency-group	io_resume	Resumes I/Os running on Consistency Groups previously suspended with io_pause.
consistency-group	io_pause_list	Lists Consistency Groups io_pause state.

Figure 2-116 List of consistency group commands in the CLI

To display all of the consistency groups in the system, issue the **cg_list** command, as shown in Example 2-15.

Example 2-15 List the consistency groups

```
ITSO_2_A9000R>>cg_list
Name      Pool Name
itso_cg1  itso_pool1
itso_cg2  itso_pool1
itso_cg3  itso_pool2
```

To create a consistency group, use the **cg_create** command and specify the name of the consistency group with the pool name, as shown in Example 2-16.

Example 2-16 Create a consistency group with the CLI

```
ITSO_2_A9000R>>cg_create cg=itso_cg4 pool=itso_pool2
Command run successfully.
```

The next step, after you create a consistency group, is to add volumes to the group. Use the **cg_vol_add** command. You must specify the consistency group and the volume to add. An example is shown in Example 2-17.

Example 2-17 Add a volume to a consistency group with the CLI

```
ITSO_2_A9000R>>cg_add_vol cg=itso_cg4 vol=itso_vol7
Command run successfully.
```

The storage administrator can also remove a volume from a consistency group by using the command **cg_remove_vol**. In this case, only the volume must be specified. The CLI prompts for confirmation before it removes the volume from the consistency group as shown in Example 2-18.

Example 2-18 Remove a volume from a consistency group with the CLI

```
ITS0_2_A9000R>>cg_remove_vol vol=itso_vol7
```

```
Warning: Are you sure you want to remove volume 'itso_vol7' from its Consistency
Group? y/n: y
Command run successfully.
```

Alternatively, the command can be run by using **-y** to avoid the prompt. This option is useful when the command is run within a script.

To create a snapshot group, issue the command as shown in Example 2-19. More options can be used with this command, including specifying a name for the snapshot group and the deletion priority. By issuing this command, you create all of the snapshots for all of the volumes within the group at the same time.

Example 2-19 Create a snapshot group

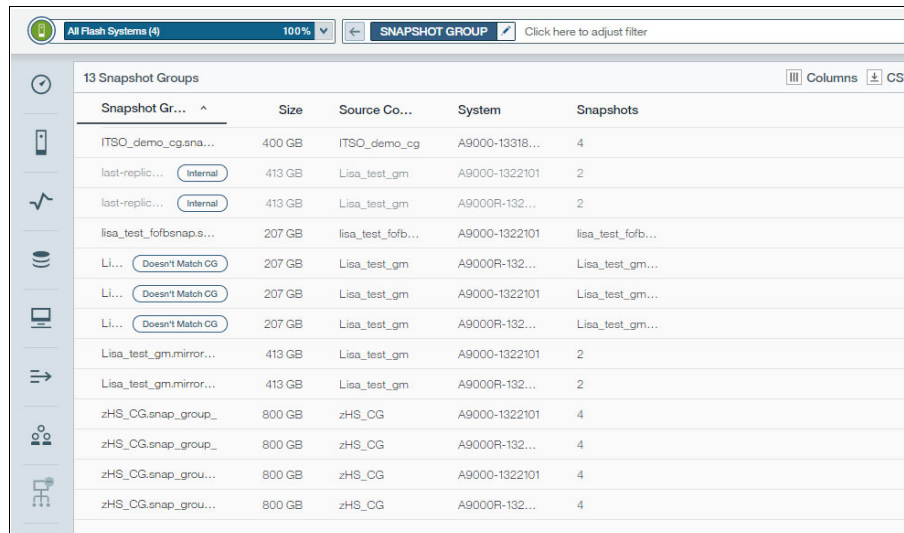
```
ITS0_2_A9000R>>cg_snapshots_create cg=itso_cg1
Command run successfully.
```

To get an application crash consistent snapshot group, first pause I/O from the application before you create the snapshot. Certain applications can perform this action. In the CLI, you can include commands, such as **io_pause** and **io_resume**.

You can also create a single volume snapshot of a volume within a consistency group. However, this single volume snapshot will not be consistent across all of the volumes in the consistency group.

2.5.5 Snapshot Groups

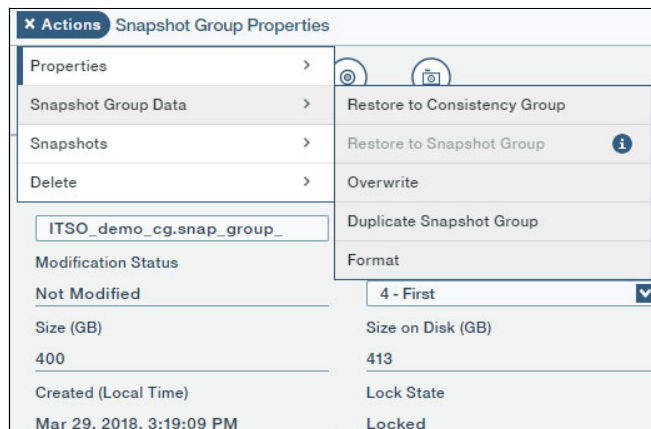
Beginning with HSM V5.4, you can now view and perform actions on the snapshot groups direction by going to the left menu options and selecting **Pools & Volumes View** → **Snapshot Groups**. This action shows all the snapshot groups that exist on the systems for which you filtered. This action also shows the internal snapshot groups that are associated with asynchronously mirrored consistency groups, as shown in Figure 2-117.



Snapshot Gr...	Size	Source Co...	System	Snapshots
ITSO_demo_cg.sna...	400 GB	ITSO_demo_cg	A9000-13318...	4
last-replic... Internal	413 GB	Lisa_test_gm	A9000-1322101	2
last-replic... Internal	413 GB	Lisa_test_gm	A9000R-132...	2
lisa_test_foibnap.s...	207 GB	lisa_test_foib...	A9000-1322101	lisa_test_foib...
Li... Doesn't Match CG	207 GB	Lisa_test_gm	A9000R-132...	Lisa_test_gm...
Li... Doesn't Match CG	207 GB	Lisa_test_gm	A9000-1322101	Lisa_test_gm...
Li... Doesn't Match CG	207 GB	Lisa_test_gm	A9000R-132...	Lisa_test_gm...
Lisa_test_gm.mirror...	413 GB	Lisa_test_gm	A9000-1322101	2
Lisa_test_gm.mirror...	413 GB	Lisa_test_gm	A9000R-132...	2
zHS_CG.snap_group_...	800 GB	zHS_CG	A9000-1322101	4
zHS_CG.snap_group_...	800 GB	zHS_CG	A9000R-132...	4
zHS_CG.snap_grou...	800 GB	zHS_CG	A9000-1322101	4
zHS_CG.snap_grou...	800 GB	zHS_CG	A9000R-132...	4

Figure 2-117 Snapshot groups

You can perform actions on the snapshot group, such as modify, view, and delete, as shown in Figure 2-118.



X Actions Snapshot Group Properties	
Properties	
Snapshot Group Data	Restore to Consistency Group
Snapshots	Restore to Snapshot Group i
Delete	Overwrite
	Duplicate Snapshot Group
	Format
ITSO_demo_cg.snap_group_...	
Modification Status	
Not Modified	4 - First
Size (GB)	Size on Disk (GB)
400	413
Created (Local Time)	Lock State
Mar 29, 2018, 3:19:09 PM	Locked

Figure 2-118 Snapshot group actions



Security

Security is of utmost importance for mission-critical workloads and data. IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software offer many features to ensure that both the access and the data are secure.

This chapter describes local access and authentication that can be configured in IBM FlashSystem A9000 and A9000R and IBM XIV Gen3. For details about LDAP-based authentication, refer to *Enabling LDAP for IBM FlashSystem A9000 and A9000R with Microsoft Active Directory*, REDP-5387. For information regarding encryption, refer to *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.

3.1 Physical access security

When you install IBM XIV Gen3, IBM FlashSystem A9000 and A9000R and IBM Spectrum Accelerate software, the same common sense security practices for any other business-critical IT system apply. These practices include enforcing physical access security to the data center, access doors, the individual systems, and the racks.

Another typical concern is reliable power. Each major component of IBM FlashSystem A9000 and A9000R uses individual and internal redundant battery units to ensure correct operations. In a power loss, these batteries ensure the retention of all caches, destage securely, and perform a graceful shutdown.

However, if someone gains physical access to the equipment, that person might manually shut off or remove components by bypassing the normal process. In this case, the storage system has the potential to lose the contents of several of its caches, resulting in system unavailability or data loss. To eliminate or greatly reduce this risk, IBM FlashSystem A9000R and IBM XIV Gen3 can be equipped with lockable doors. IBM FlashSystem A9000 can be installed in a customer-provided rack that can also provide this feature.

Important: Protect IBM FlashSystem A9000 and A9000R or IBM XIV Gen3 by locking the rack doors and monitoring the physical access to the equipment.

3.2 Native user authentication

To prevent unauthorized access to the configuration of the storage system and ultimately to the information that is stored on its volumes, IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software use password-based user authentication.

Note: Because the graphical user interface (GUI) code is accessed through the IBM Hyper-Scale Manager server, security considerations also apply for that server. It is a preferred practice to regularly change the access code and machine account password. See 3.2.3, “Security considerations for Hyper-Scale Manager” on page 112.

3.2.1 Local credential repository

By default, IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software are configured to use native (locally managed) user authentication. The alternative is to use Lightweight Directory Access Protocol (LDAP), as described in *Enabling LDAP for IBM FlashSystem A9000 and A9000R with Microsoft Active Directory*, REDP-5387.

Native user authentication uses the credential repository that is stored locally on IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software.

The local credential repository maintains the following information:

- ▶ Domain memberships
- ▶ User name
- ▶ User password
- ▶ User role
- ▶ User group
- ▶ Optional account attributes

User name

A user name is a string of 1 - 63 characters that can contain only a - z, A - Z, 0 - 9, . - _ ~, and space symbols. User names are case-sensitive. IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software are configured with a set of predefined user accounts. Predefined user names and corresponding default passwords exist to provide the initial access at the time of installation for system maintenance.

The following user accounts are predefined:

- *technician*

This account is used by the IBM service support representative (SSR) to install and do maintenance.

- *admin*

This account provides the highest level of client access to the system. It can be used for creating new users and changing passwords for existing users in native authentication mode.

Important: Use of the *admin* account must be limited to the initial configuration when no other user accounts are available. Access to the *admin* account needs to be restricted and securely protected.

- *xiv_development* and *xiv_maintenance* user

These IDs are special-case, predefined internal IDs that can be accessed only by qualified IBM development personnel and SSRs.

Predefined user accounts cannot be deleted from the system, and they are always authenticated natively even if the system operates under Lightweight Directory Access Protocol (LDAP) authentication mode. Customers can contact IBM to arrange for the special-case, predefined internal IDs to be removed or changed due to government or intelligence sensitivity, which affects service and support operations.

User accounts can initially be created by the *admin* user only. After the *admin* user creates a user account and assigns it to the *storageadmin* (storage administrator) role, then other user accounts can be created by this *storageadmin* user.

Tip: When you install Hyper-Scale Manager, which is the primary way of managing and monitoring the IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate family of storage systems, a default management *system_management_user* is created.

In native authentication mode, the system is limited to creating up to 128 user accounts. This number includes the predefined users.

User password

The *user password* is a secret word or phrase that is used by the account owner to gain access to the system. The user password is used at the time of authentication to establish the identity of that user. User passwords can be 6 - 12 characters and include these characters:

a - z, A - Z, ~, !, @, #, \$, %, ^, &, *, (,), _ , + , - , =, {, }, |, :, ;, <, >, ?, ., /, \, [, and]

The passwords must not have spaces between characters. In native authentication mode, IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software verify the validity of a password when the password is assigned. Be careful to ensure that strong passwords are used.

Predefined users have the default passwords that are shown in Table 3-1 and that are assigned at the installation.

Table 3-1 Default passwords

Predefined user	Default password
admin	adminadmin
system_management_user	adminadmin

Important: The default *admin* password must be changed at the time of installation to prevent unauthorized access to the system. For instructions, see 3.2.2, “Managing user accounts” on page 108.

The following restrictions apply when you work with passwords in native authentication mode:

- ▶ For security, passwords are not shown in user lists.
- ▶ Passwords are user-changeable. Users can change their own passwords only.
- ▶ Only the predefined user *admin* can change the passwords of other users.
- ▶ Passwords are changeable from both the command-line interface (CLI) and the GUI.
- ▶ Passwords are case-sensitive.
- ▶ User password assignment is mandatory at the time that a user account is created.
- ▶ Creating user accounts with an empty password or removing a password from an existing user account is not permitted.

User roles

Eight predefined user roles exist in the Storage Management GUI and the CLI. In the CLI, roles are also referred to as *categories*, and they are used for day-to-day operation of the storage system. Multiple users can be created for the first four roles that are listed (*storageadmin*, *applicationadmin*, *securityadmin*, and *readonly*). The other four roles are preassigned by the system, and they do not allow additional users to be created with these roles.

The following section describes predefined roles, their level of access, and applicable use:

▶ Storage administrator

The storage administrator (*storageadmin*) role is the user role with highest level of access that is available on the system. A user that is assigned to this role can perform changes on any system resource, except maintaining physical components, changing the status of physical components, or effecting changes that relate to encryption.

▶ Application administrator

The application administrator (*applicationadmin*) role provides flexible access control over volume snapshots. A user that is assigned to the *applicationadmin* role can create snapshots of volumes, perform mapping of their own snapshots to assigned hosts, and delete their own snapshots.

The user group to which an application administrator belongs determines the hosts and their related set of mapped volumes that the application administrator is allowed to manage. If a user group is defined with `access_all="yes"`, application administrators who are members of that group can manage all volumes on the system. For more information about user group membership and group-to-host association, see “User groups” on page 107.

- Security administrator

The security administrator (*securityadmin*) role is to manage certificates, key servers, and other items that relate to encryption. Defining at least two security administrators is a must because two is the minimum number that is required to enable recovery keys.

- Operations administrator

Users that are assigned to the operations role (*readonly*) can view system information only. A typical use for the operations role is a user who is responsible for monitoring system status, system reporting, and message logging. This user must not be permitted to change the system.

- Technician

The technician role has a single predefined user name (*technician*) assigned to it, and it is intended to be used by IBM support personnel for maintaining the physical components of the system. The technician is limited to the following tasks: Physical system maintenance and phasing components in or out of service. The access of the technician to the system is restricted. The technician is unable to perform any configuration changes to pools, volumes, or host definitions on the storage system.

- xiv_development

This single predefined user name is assigned to the xiv_development role, and it is intended to be used by IBM development personnel.

- xiv_maintenance

This single predefined user name is assigned to the xiv_maintenance role, and it is intended to be used by IBM maintenance personnel.

- xiv_hostprofiler

This single predefined user name is assigned to the xiv_hostprofiler role, and it is intended to be used for gathering additional information about hosts that are attached to the IBM XIV Storage System.

User roles: It is not possible to add user roles or to modify predefined roles. In native authentication mode, after a user is assigned a role, the only way to assign a new role is to first delete the user account and then re-create it.

Native authentication mode implements the user role mechanism as a form of *role-based access control* (RBAC). Each predefined user role determines the level of system access and associated functions that a user is allowed to use.

RBAC: IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software use role-based access control (RBAC)-based authentication and authorization mechanisms.

All user accounts must be assigned to a single user role. Assignment to multiple roles is not permitted. Deleting or modifying role assignment of natively authenticated users is also not permitted. Do not confuse this concept with domains. The concept of domains is a separate construct, and multiple domains can be associated with a user's single role.

Global administrator and domain administrator

When you add a user, the user can be assigned to the global space (no domain) or several domains that are defined on IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software, as shown in Figure 3-1. (Select **New** → **Users** from the dashboard.)

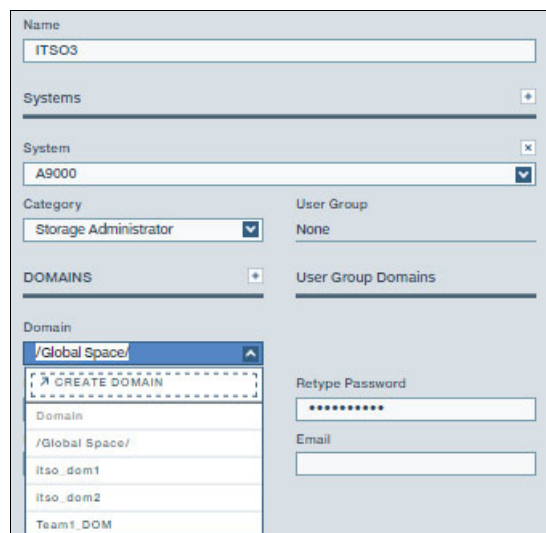
The screenshot shows a web-based user creation form. At the top, there's a 'Name' field with 'ITSO3' entered. Below it is a 'Systems' section with a dropdown menu showing 'A9000'. The 'Category' is set to 'Storage Administrator' and 'User Group' is 'None'. A 'DOMAINS' section contains a list of domains: '/Global Space/', 'itso_dom1', 'itso_dom2', and 'Team1.DOM'. To the right of the domains list are fields for 'Retype Password' and 'Email'. The form has a light blue header and footer.

Figure 3-1 Adding a user to a domain

If a particular user needs to be assigned to multiple domains, the additional association must be made from the Domains view.

With the introduction of domains, two levels of authority exist for a user in a specific role:

- Domain administrator

A user who is associated with one or more domains is a *domain administrator*, and the user has access rights (as implied by the user's role, such as *readonly* or *storageadmin*) to the entities that are uniquely associated with the pertinent domains.

- Global administrator

A user who is not associated with any domain is a *global administrator*, and the user has access rights to all of the entities that are not uniquely associated with a domain (*global space*).

A global administrator who is a *storageadmin* can create, edit, and delete domains, and they can associate resources with a domain.

An open or closed policy can be defined so that a global administrator can (or cannot) extend the reach into a domain.

The intervention of a global administrator who has permission for the global resources of the system is only needed for these situations:

- Initial creation of the domain and the assignment of a domain administrator
- Resolution of hardware issues

User groups

A *user group* is a group of application administrators who share a set of snapshot creation permissions. The permissions are enforced by associating the user groups with hosts or clusters. User groups have the following characteristics:

- ▶ Only users that are assigned to the applicationadmin role can be members of a user group.
- ▶ A user can be a member of a single user group.
- ▶ A maximum of eight user groups can be created.
- ▶ Group names are case-sensitive.
- ▶ In native authentication mode, a user group can contain up to eight members.
- ▶ If a user group is defined with `access_all="yes"`, users that are assigned to the applicationadmin role who are members of that group can manage all snapshots on the system.
- ▶ A user must be assigned to the storageadmin role to be permitted to create and manage user groups.

Important: A user group membership can be defined only for users who are assigned to the applicationadmin role.

User group and host associations

Hosts and clusters can be associated with only one user group. When a user is a member of a user group that is associated with a host, that user is allowed to manage the snapshots of the volumes that are mapped to that host.

User group and host associations have the following properties:

- ▶ User groups can be associated with both hosts and clusters. This configuration enables limiting group member access to specific volumes.
- ▶ A host that is part of a cluster can be associated only with a user group through user group to cluster association. Any attempts to create user group association for that host fail.
- ▶ When a host is added to a cluster, the association of that host is removed. Limitations on the management of volumes that are mapped to the host are controlled by the association of the cluster.
- ▶ When a host is removed from a cluster, the association of that cluster remains unchanged. This configuration enables the continuity of operations so that all scripts that rely on this association continue to work.

Optional account attributes

In this section, we describe optional attributes for email and phone numbers:

- ▶ **Email:** Email is used to notify specific users about events through email messages. Email addresses must follow standard formatting procedures.

Acceptable value: Any valid email address. A default value is not defined.

- ▶ **Phone:** Phone numbers are used to send Short Message Service (SMS) messages to notify specific users about system events. Phone numbers and area codes can be a maximum of 63 digits, hyphens (-), and periods (.).

Acceptable value: Any valid telephone number. A default value is not defined.

3.2.2 Managing user accounts

This section illustrates the procedure for managing user accounts, groups, and group membership in IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software.

Adding users

This process requires that you first log on to IBM Hyper-Scale Manager with storage administrator access rights (storageadmin role). If you are accessing the system for the first time, use the predefined user *admin* (the default password is adminadmin).

Follow these steps to add a user:

1. In the main GUI window, click **NEW+** in the upper right corner of the window. The CREATE NEW menu opens, as shown in Figure 3-2.

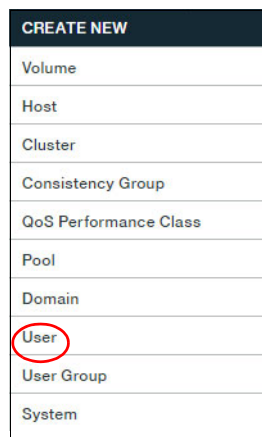


Figure 3-2 Create New menu

2. Click **User**. The Add User window opens. The name, system, and password are required fields. The default category for a new user is Storage Administrator. Use the drop-down list to select the correct category for the user. An optional field is Domains. This field is used if the user must be associated with a domain. For more information about domains, see Chapter 4, “Multitenancy” on page 119.

3. Click **Create** to create the new user. See Figure 3-3.

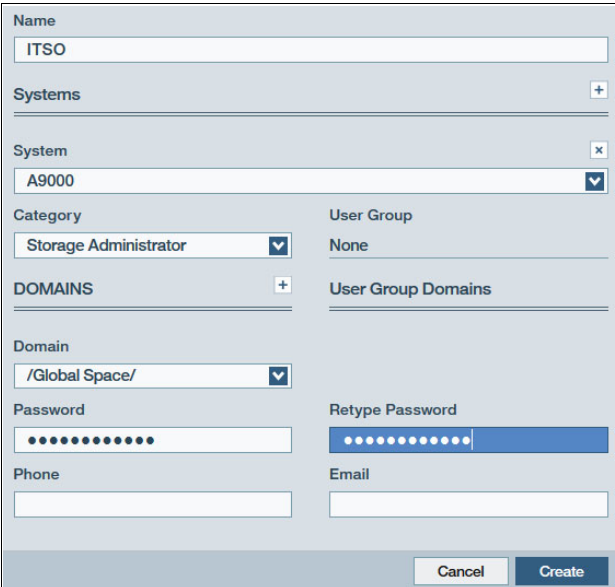
A screenshot of a web-based form for adding a new user. The form is organized into several sections. At the top, there is a 'Name' field with the text 'ITSO' entered. Below this is a 'Systems' section with a '+' icon. The next section is 'System', which contains a dropdown menu currently showing 'A9000'. Below the system dropdown are two columns: 'Category' with a dropdown showing 'Storage Administrator', and 'User Group' with a dropdown showing 'None'. Under these is a 'DOMAINS' section with a '+' icon, followed by 'Domain' with a dropdown showing '/Global Space/'. The 'Password' section has two input fields: 'Password' and 'Retype Password', both filled with dots. Below the password fields are 'Phone' and 'Email' input fields. At the bottom right of the form are two buttons: 'Cancel' and 'Create'.

Figure 3-3 Add user window

Removing users

User maintenance is important for security. It involves removing users when they leave the company or their job changes and they no longer need access to the same systems.

To remove or delete a user, ensure that you log on with the security administrator role and perform the following steps:

1. Click the **ACCESS VIEWS** icon on the left side of the window. The **ACCESS VIEWS** menu is shown in Figure 3-4.




Figure 3-4 ACCESS VIEWS menu

2. Click **Users** from this menu. A list of users, which is similar to Figure 3-5, opens.

USER <small>Click here to adjust filter</small>					
USER ▲	SYSTEM	DOMAIN	CATEGORY	USER GROUP	PHONE
admin	ITSO_1_A9000R	no-domain	Storage Administ...		
admin	ITSO_2_A9000R	no-domain	Storage Administ...		
itso_addadmin2	ITSO_1_A9000R	no-domain	Application Admi...	itso_group	
ITSO_admin	ITSO_1_A9000R	no-domain	Storage Administ...		
itso_appadmin1	ITSO_1_A9000R	no-domain	Application Admi...	itso_group	
itso_dom1_ad...	ITSO_2_A9000R	itso_dom1	Storage Administ...		
itso_secadmin	ITSO_1_A9000R	no-domain	Security Adminis...		

Figure 3-5 List of users

Note: /Global Space/ in the domain field means that the user is not associated with a domain.

3. Click the user name that you want to delete (1 in Figure 3-6). Then, click the **Actions** tab (2 in Figure 3-6). The Actions menu opens as shown in Figure 3-6.

1 selected out of 107 Users				Actions User Properties	
User ^	System	Domain	Cat	Properties >	Delete User(s)
ITSO_Test	A9000	/Global Space/	Sto	Association >	
jbeer	XIV_04_1340008	/Global Space/	Sto	User Group >	
Jens	XIV_04_1340008	/Global Space/	Sto	Security >	
manager server	A9000	/Global Space/	Sto	Delete >	

Figure 3-6 User actions menu

4. From the Actions menu, when you hover over **Delete**, you see the Delete User(s) option that is shown in Figure 3-7. Click **Delete User(s)**. You are prompted for a confirmation.

Actions User Properties	
Properties >	Delete User(s)
Association >	
User Group >	
Security >	
Delete >	

Figure 3-7 Delete a user

5. The confirmation that is shown in Figure 3-8 gives you one last chance to cancel the deletion of this user. If you are certain that you want to delete the user that is named in the message, click **Delete**.




Figure 3-8 Confirm to delete the user

Note: Default users cannot be removed.

User groups

Create a group by selecting the **NEW → User Group** option on the top menu bar. This option displays the Create User Group window as shown in Figure 3-9. In this image, we expanded the explanation for the Limited Access value.

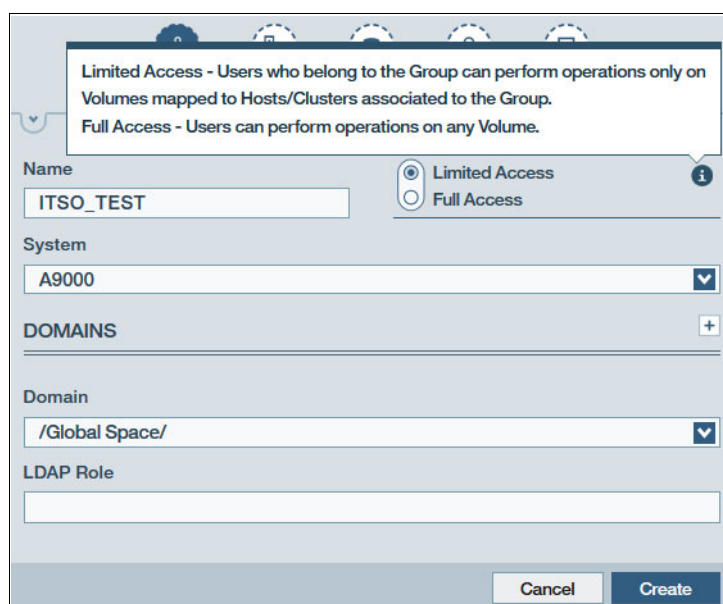


Figure 3-9 System selection for a new user group

Complete the required information for the new users group and select either Limited Access or Full Access. Click **Create** to create the group.

After you create the user group, it is a simple task to create a user in the group by selecting the group from the table view (**Access Views** → **User Groups**) and then selecting **Actions** → **Members** → **Create User Here**. This action prepopulates the system, user group, and domain fields, as shown in Figure 3-10. Complete the remaining required fields (name and password) and select the correct category for the user. We created the group with limited access and selected the user category of **Application Administrator**.

User Group	System
ITSO_appadmin...	A9000
ITSO_appadmin...	A9000
ITSO_RE_group	A9000

Figure 3-10 Add a user to an existing user group

The new group will then be displayed in the table view under **Access Views** → **User Groups**.

To create a user group with the CLI, use the command `user_groupcreate`. An example of this command is shown in Example 3-1.

Example 3-1 Create a user group

```
>>user_group_create user_group=itso_application_group
Command executed successfully.
```

Tip: If you use spaces in user group names, enclose the name between double quotation marks, such as “new group name”.

3.2.3 Security considerations for Hyper-Scale Manager

When you first log on to the system from the GUI by using the default admin credentials, you get unrestricted access to the system. You use this unrestricted access to change the Hyper-Scale Manager settings. This unrestricted or advanced login mode is indicated by the check mark that is displayed next to the Management Server icon as shown in Figure 3-11.




Figure 3-11 Management Server access

To make changes for the Management Server, you must be logged on as the admin user ID (the actual admin user ID, not just any user ID with administrator authority).

Complete the following steps:

1. With the unrestricted access, from the GUI Dashboard, select the **Settings** icon, and then click **Management Server** from the menu as shown in Figure 3-12.




Figure 3-12 Management Server preferences

2. You can then change the access restriction and specify a new access code (a new password) from the **Security** menu, as shown in Figure 3-13. The default Access Code is *adminadmin* and must be changed for security reasons. You can also change the value for the Auto Logout Time.

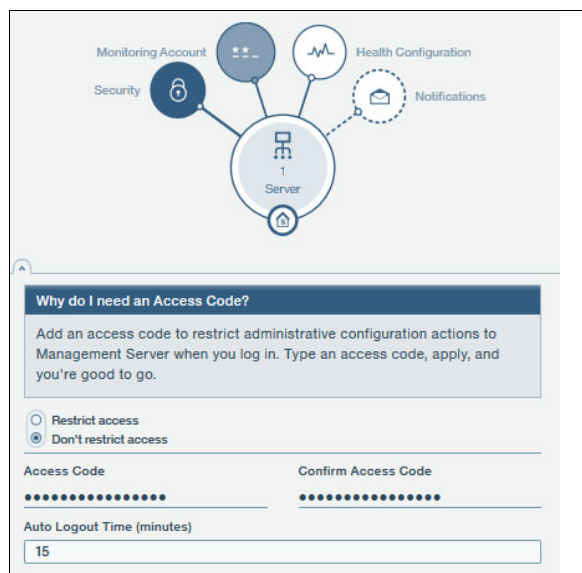


Figure 3-13 Change the access code

3. You can also change the monitoring account password, as shown in Figure 3-14.

BY setting a Monitoring Account, you can monitor the storage systems in your environment. The account can be a local user on the storage system or a user that you configure in LDAP. There can be only one Monitoring Account for all the systems that are monitored in Hyper-Scale Manager. The default Monitoring Account User Name is `manager_server_user`.

Note: If multiple Hyper-Scale Manager servers are installed and used to manage the same storage systems, use the same user credentials for the Monitoring Account on all of your management servers.

Figure 3-14 Change the monitoring account password

The same actions can also be performed directly on the Hyper-Scale Manager server by using the Hyper-Scale Manager configuration menu.

3.2.4 Managing user accounts by using the XCLI

This section summarizes the commands and options that are available to manage user accounts, user roles, user groups, group memberships, and user group to host associations through the XCLI.

Table 3-2 shows the various commands and a brief description for each command. The table also indicates the user role that is required to issue specific commands.

Table 3-2 XCLI access control commands

Command	Description	Role that is required to use the command
<code>access_define</code>	Defines an association between a user group and a host.	storageadmin
<code>access_delete</code>	Deletes an access control definition.	storageadmin
<code>access_list</code>	Lists access control definitions.	storageadmin, readonly, and applicationadmin

Command	Description	Role that is required to use the command
user_define	Defines a new user.	storageadmin and securityadmin
user_delete	Deletes a user.	storageadmin and securityadmin
user_list	Lists all users or a specific user.	storageadmin, readonly, securityadmin, and applicationadmin
user_rename	Renames a user.	storageadmin and securityadmin
user_update	Updates a user. You can rename the user, change a password, modify the Access All setting, modify email, modify the area code, or modify the phone number.	storageadmin, applicationadmin, and securityadmin
user_group_add_user	Adds a user to a user group.	storageadmin
user_group_create	Creates a user group.	storageadmin
user_group_delete	Deletes a user group.	storageadmin
user_group_list	Lists all user groups or a specific one.	storageadmin, readonly, and applicationadmin
user_group_remove_user	Removes a user from a user group.	storageadmin
user_group_rename	Renames a user group.	storageadmin
user_group_update	Updates a user group.	storageadmin
domain_list_users	Lists users that are associated with domains.	storageadmin, securityadmin, applicationadmin, readonly, and technician
domain_add_user	Associate a user with a domain.	storageadmin and securityadmin
domain_remove_user	Disassociate a user from a domain.	storageadmin and securityadmin

Adding users with the XCLI

Before you complete the following steps, you must install the XCLI component on the management workstation. A *storageadmin* user is required. The following examples assume a Microsoft Windows management workstation.

To add users, complete the following steps:

1. Use the **user_list** command to obtain the list of predefined users and categories, as shown in Example 3-2. This example assumes that no users, other than the default users, were added to the system.

Example 3-2 XCLI user_list

```
>>user_list
```

Name	Category	Group	Active...	Access All
xiv_development	xiv_development		yes	
xiv_maintenance	xiv_maintenance		yes	
admin	storageadmin		yes	
technician	technician		yes	
itso	storageadmin		yes	
lab_admin	applicationadmin	Application01_Group	yes	... no

2. If this system is a new system, change the default password for the admin user by running **user_update**, as shown in Example 3-3.

Example 3-3 XCLI user_update

```
>>user_update user=admin password=New-Passw0rd password_verify=New-Passw0rd
Command executed successfully.
```

3. Add a user, as shown in Example 3-4. Define a user by using a unique name, a password, and a role (which is designated here as *category*).

Example 3-4 XCLI user_define

```
>>user_define user=lab_user password=XIV-R3 password_verify=XIV-R3
category=applicationadmin
Command executed successfully.
```

To add user to a domain, complete the following steps:

1. Use the **domain_list_users** command to obtain the list of predefined and domain-associated users and categories, as shown in Example 3-5. This example is based on the assumption that a domain was defined and that one admin associated user is present in the system.

Example 3-5 XCLI domain_list_users

```
XIV_PFE2_1340010>>domain_list_users
```

Domain	User	Category
ITS0_d1	ITS0_User1	storageadmin
no-domain	Residency	applicationadmin
no-domain	xiv_development	xiv_development
no-domain	xiv_maintenance	xiv_maintenance
no-domain	admin	storageadmin
no-domain	technician	technician
no-domain	xiv_hostprofiler	xiv_hostprofiler

2. Add the same user to another domain, as shown in Example 3-6. The example assumes that another domain was defined. Specify the domain and user name. A user can have a domain administrator role in more than one domain.

Example 3-6 XCLI domain_add_user

```
XIV_PFE2_1340010>>domain_add_user domain=pfe.ibm user=ITS0_User1  
Command executed successfully.
```

Defining user groups with the XCLI

To define a user group, complete the following steps:

1. Run **user_group_create**, as shown in Example 3-7, to create a user group that is called Application03_Mainz.

Example 3-7 XCLI user_group_create

```
>>user_group_create user_group=Application03_Mainz  
Command executed successfully.
```

Note: Avoid spaces in user group names. If spaces are required, the group name must be placed between quotation marks, such as “name with spaces”.

The Application03_Mainz user group is empty, and it has no associated host. The next step is to associate a host or a cluster.

2. Associate the Application03_Mainz user group to the Application_host01 host, as shown in Example 3-8.

Example 3-8 XCLI access_define

```
>>access_define user_group=Application03_Mainz host=Application_host01  
Command executed successfully.
```

A host is assigned to the user group. The user group does not have any users that are included in the user group yet.

3. Add the first user, as shown in Example 3-9.

Example 3-9 XCLI user_group_add_user

```
>>user_group_add_user user_group=Application03_Mainz user=lab_user  
Command executed successfully.
```

The user lab_user was assigned to the Application03_Mainz user group. This user is an applicationadmin with the Full Access right set to no.



Multitenancy

Multitenancy provides the option to allocate storage resources to several independent administrative domains. This capability means that any resource that is allocated to a specific domain can be seen and acted upon within that domain only. Other resources are considered to be *global resources* that can be used or seen by all domains, such as the System view, the Fibre Channel (FC) and internet Small Computer System Interface (iSCSI) ports, and events on the system.

Multitenancy provides several advantages:

- ▶ Secure division and isolation of storage resources among numerous tenants
- ▶ Simple, quick delegation of administration tasks and role-based permissions
- ▶ Simple, rapid deployment without the need for extensive planning and tuning, and also field upgradability

In this chapter, we cover the details of domain administration by using the Hyper-Scale Manager with IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software.

4.1 Multitenancy

Today, many organizations are moving toward cloud-based infrastructures, where both compute and storage resources are viewed as commodity services. Users of these resources expect a pay-as-you-go model, where the resources that they require are rapidly available when needed, but they pay only for the resources that they require, nothing more. For storage providers, this shift represents a significant challenge.

From both capital expenditure and resource management (personnel) perspectives, storage consolidation becomes an important consideration. This new approach necessitates storage systems that can easily and securely support the provisioning, administration, and service level management of storage resources for multiple resources to be used concurrently.

The multitenancy of IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software brings flexibility and simplicity to management of tenant data and storage resources across one or more systems. In this chapter, we cover a summary of the domain objects and the details of domain administration by using management tools.

4.1.1 Domains

With IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software, multitenancy is realized through the concept of *domains*. A domain is a logically partitioned container (more than one can exist in a system) with its own administrators and resources. This partitioning enables the secure isolation of users and resources among domains.

Domains enable the ability to control what object a user can perform actions on. A user is able to see and act upon resources only within the domains they are associated with. When a domain user connects to a system using either the graphical user interface (GUI) or command-line interface (CLI), they can access only the resources within their domains.

IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software uses a role-based access control (RBAC) model to control what actions a user can perform. For more information about the available user roles, see Chapter 3, “Security” on page 101.

Domains can be associated with the following objects:

- ▶ Users and user groups
- ▶ Storage pools and all of their content, which can be associated with only a single domain
- ▶ Hosts and clusters
- ▶ Remote mirror targets

Important: The implementation of domains occurs at the management and virtualization level. All physical system resources (IBM MicroLatency® modules, CPUs, memory, and interfaces) are shared among domains.

Domain administrators cannot modify physical system resources and they cannot access resources that belong to another domain, including the “global domain” resources of the system. However, they can be notified of events that relate to physical system attributes because these events might affect the objects within their domain. For example, a domain administrator can be alerted upon the failure of a MicroLatency module or the disconnection of an interface link.

4.1.2 Domain attributes

A domain is assigned resources, including capacity and the maximum number of pools, volumes, consistency groups, mirrors, and data migrations that can be used in the domain. The *global storage administrator* (the storage administrator of the entire IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, or IBM Spectrum Accelerate software defines these numbers when the global storage administrator creates a domain.

The global administrator must remember that these numbers are system-wide numbers. The global administrator must ensure that other domains have the required number of resources available for use. These other domains include resources that are used and that are not in a domain. Resources that are not in a domain are known as “*global space*” resources.

4.1.3 Global administrators and domain administrators

The same user roles in IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software can also be assigned to users within a domain. The main difference between users is with the storage administrator. A super user or overall administrator, which we refer to as the *global administrator*, is required to create domains. The global administrator then can create a new storage administrator user and associate that user with a domain.

The domain storage administrator, after they log in to the system through either the GUI or CLI, can then manage their domain in the same way that the global administrator can manage the domain, including creating users, such as an application administrator and backup storage administrators.

Note: The domain administrator cannot create a domain within a domain.

The domain administrator is also limited in the amount of capacity that is assigned to the domain by the global administrator. This capacity can be in the form of free space and predefined pools. If extra space is needed for the domain, the domain administrator must request more space from the global administrator.

The global administrator also sets how many objects can be created within the domain. To accurately account for the system-wide requirements, the global administrator must plan carefully to ensure that the domain and global space requirements are met. For example, the domain might be allotted four storage pools and 30 volumes. This number contributes to the total for the system-wide allocation.

4.1.4 Domain policy

Note: The security administrator of the system has the authority to set the *domain policy* to closed or open, depending on business needs.

The domain policy can be set to open or closed with the CLI by the security administrator. The command to do this is **domain_policy_set**. Let's look at what these settings mean.

Open policy

If the policy is set to open, the global administrator is able to see several attributes of the domain, such as the number of pools and volumes, but the global administrator cannot see any attributes of those resources, such as the name or size of a pool or volume.

In addition, the domain administrator can add resources and modify attributes for an existing domain if the domain administrator requests additional capacity for the domain or needs to create additional pools, volumes, and other objects but reached the limit for their domain.

With the domain policy set to open, the global administrator can address these requests while still managing those finite system-wide resource limits, such as the maximum number of volumes in a system, to ensure that enough resources exist for all of the users of the system.

Closed policy

If the policy is set to closed, the global administrator has no access to any domain other than the global domain (the resources that are not contained in any domain). The global administrator can still create a domain but after the domain is created, the global administrator has no view or access to the domain.

If any changes need to be made to the domain resources, the security administrator must change the domain resources or change the domain policy back to open for the global administrator. This setting is system-wide so all domains are inaccessible to the global administrator, except for the global space, after the policy is set to closed.

All of the examples in the next section, 4.2, “Managing domains with the GUI” on page 122, were run with the domain policy set to open.

4.2 Managing domains with the GUI

Domains are managed by the domain administrator. The global administrator (which is also called the *storage administrator*) cannot manage the resources within a domain, but they can add capacity, pools, users, hosts, clusters, and targets to a domain if the domain policy is set to open. The day-to-day operations within a domain are performed by the domain administrator.

The domain administrator can add a user who has access only to resources within the domain. The domain administrator can also add a performance class on a pool or volume within the domain. A global administrator can define iSCSI and Fibre Channel interfaces, as these resources are not unique to a domain. A domain administrator cannot perform any actions on these interfaces. The domain administrator can access alerts and events that affect system-wide resources and domain resources.

In this section, we describe the following topics:

- ▶ Creating domains
- ▶ Viewing and editing domain properties
- ▶ Managing domain associations
- ▶ Managing from within a domain
- ▶ Removing domains

Managing domains with the GUI is simple. You can access the Domain view as shown in Figure 4-1. A tabular view is displayed of all of the domains that are configured on the system (or systems if multiple systems are managed by Hyper-Scale Manager) and the global space. After the list of domains is displayed, you can select a domain and then use the Actions menu to manage the domain and resources.




Figure 4-1 SYSTEMS & DOMAINS VIEWS

4.2.1 Creating domains

It is a simple process to create a domain with the Storage Management GUI of IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software:

1. Select **New** from the top navigation pane to open the CREATE NEW menu with a listing of all of the objects that can be created within this menu. **Domain** is the first item in the list.
2. By clicking **Domain**, as shown in Figure 4-2, a new wizard opens with all of the attribute fields for the domain.

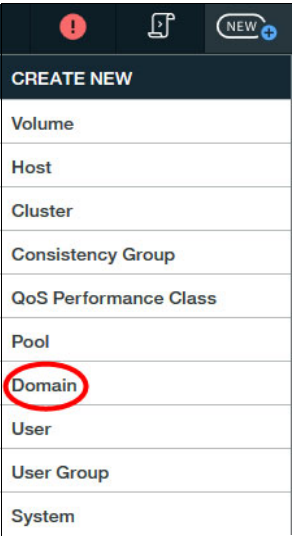


Figure 4-2 Create a domain

3. Fill out the required information for creating a domain, as shown in Figure 4-3.

The screenshot shows the 'Create Domain' window with the following fields and values:

- Name:** ITSO_Domain
- System:** A9000
- Physical Size (GB):** N/A
- Domain Size (GB):** 2000
- Max Pools:** 2
- Max Volumes and Snapshots:** 15
- Max CGs:** 2
- Max mirrors and HyperSwaps:** 10
- Max Data Migrations:** 3
- LDAP Domain ID:** ITSO_Domain

Buttons: Cancel, Create

Figure 4-3 New domain attributes

In this example, we created a domain with a maximum capacity of 2000 GB on an IBM FlashSystem A9000 system.

The fields that are shown in Figure 4-3 are described in detail in Table 4-1.

Table 4-1 Domain attributes

Field name	Description
Name	The logical name of the domain. It must be unique within a system.
System	IBM XIV Gen3, IBM FlashSystem A9000, IBM FlashSystem A9000R or IBM Spectrum Accelerate software where the domain will be defined.
Size (GB)	The capacity that is given to the domain that can be used to create pools within the domain, the minimum capacity is 2000 GB for IBM FlashSystem A9000 and A9000R.
Max Pools	The maximum number of pools that can be configured within this domain.
Max Volumes and Snapshots	The maximum number of volumes and snapshots that can be defined within this domain.
Max CGs	The maximum number of consistency groups that are allowed within this domain.
Max mirrors and HyperSwap	The maximum number of remote mirrors and HyperSwap operations that are allowed within this domain.
Max Data Migrations	The maximum number of active data migrations that can be defined within this domain at any time.
LDAP Domain ID	The ID attribute of the domain to use when you configure multitenancy with Lightweight Directory Access Protocol (LDAP) authentication. The default is set to the name of the domain.

Note: By default, the LDAP Domain ID field is populated with the value of the Domain Name. You can modify this value to a name other than the Domain Name. If you clear the field, the LDAP Domain ID value is automatically set to the value of the Domain Name.

4. You can now verify that the domain was created by using the Domains view, as shown in Figure 4-4. In addition to the new domain, the Global Space (global domain) is also listed. The Global Space is the view of the entire system resources that are not associated with a domain on the system.

3 Domains					
Domain	System ^	Domain Size	Allocated / ...	Written / All...	Stored / Ph
/Global Space/	A9000	1,195,926 GB	3% <div></div>	0% <div></div>	N/A <div></div>
ITSO_Domain	A9000	2,067 GB	0% <div></div>	0% <div></div>	N/A <div></div>

Figure 4-4 Domain list

4.2.2 Customizing the Domain view

First, we describe the customization of this tabular view for domains. Default settings for the view must be changed as needed after the view is closed. After you are in the Domain view, right-click anywhere in the headings, which opens the menu to view the default settings and to customize what you see in this table. An example is shown in Figure 4-5, specifically looking at the domain Associations. In this view, the box next to the fields either shows a check mark or a selection indicator.

The selection indicator means that all of the items within that category were not selected. The check mark indicates that this field or all of the fields within the category are visible in the table. In our example, the Hosts option is the only item that is selected under Associations. You can also choose to see Clusters, Users, User Groups, and Targets from this category.

Figure 4-5 shows you how to customize the Domain view.

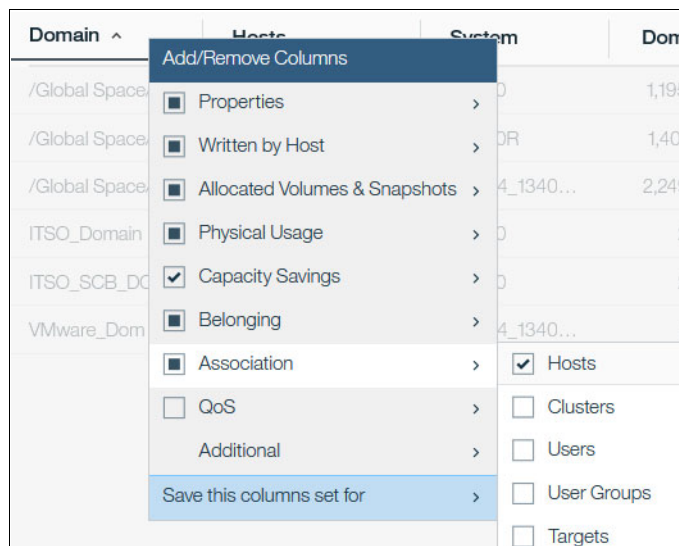


Figure 4-5 Customize the tabular view of domains

By using this technique, you can customize the tabular view for domains to suit your needs. You can then save this view either as a filter or sidebar, as shown in Figure 4-6. After the view is saved, this user will see this view every time they log on to the GUI. Alternatively, you do not have to save the customized view.




Figure 4-6 Save the customized view

4.2.3 Adding a user to a domain

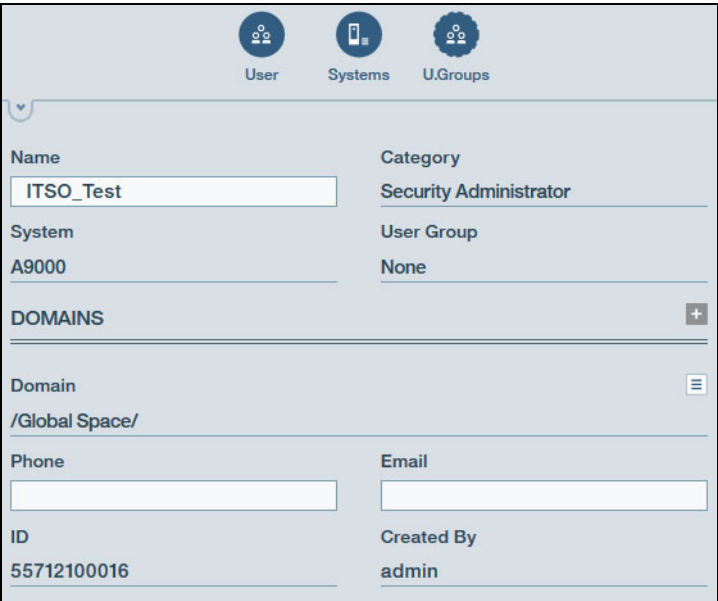
To manage local user associations with a domain select the User view, as shown in Figure 4-7. By selecting one of the users, the properties of that user display to the right of the list, as shown in Figure 4-7.

User ^	System	Domain	Category
admin	A9000	/Global Space/	Storage Admi...
gpieper	A9000	/Global Space/	Storage Admi...
ITSO_SCB	A9000	ITSO_SCB_D...	Storage Admi...
ITSO_Test	A9000	/Global Space/	Security Adm...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
SRM1	A9000	/Global Space/	Storage Admi...

Figure 4-7 User list

We look at the necessary steps to add a user to a domain. In this example, we selected the user ITS0_Test with the storage administrator role. In the Properties view of this user in Figure 4-8, the user is not yet associated with a domain.

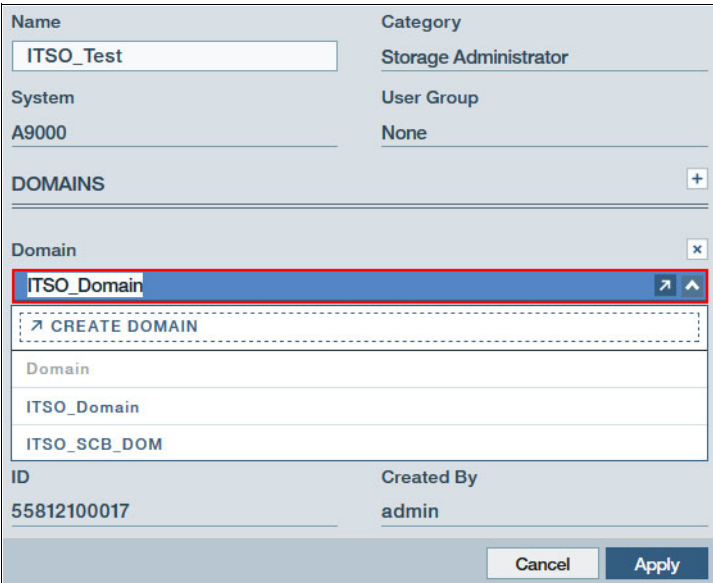
Global Space: A user that is not associated with a specific domain is designated as Global Space in the domain field.



The screenshot shows the 'User Properties' view for a user named 'ITS0_Test'. At the top, there are three tabs: 'User' (selected), 'Systems', and 'U.Groups'. Below the tabs, the user's details are displayed in a form. The 'Name' field contains 'ITS0_Test'. The 'Category' is 'Security Administrator'. The 'System' is 'A9000'. The 'User Group' is 'None'. Below these fields is a section titled 'DOMAINS' with a plus sign (+) to its right. Under 'DOMAINS', the 'Domain' field is set to '/Global Space/'. There are also fields for 'Phone' and 'Email', both of which are empty. At the bottom, the 'ID' is '55712100016' and the 'Created By' is 'admin'.

Figure 4-8 User properties

By clicking the plus sign (+) in the Domains section of the User Properties view, a list of the current domains is displayed (as shown in Figure 4-9). In this example, two domains are shown. Alternatively, you can launch the Create Domain wizard by clicking **Create Domain**. For this example, we selected ITS0_Domain and selected **Apply**.



This screenshot shows the 'User Properties' view for the same user 'ITS0_Test', but with the 'DOMAINS' section expanded. The 'Domain' field now shows a dropdown menu with two options: 'ITS0_Domain' and 'ITS0_SCB_DOM'. The 'ITS0_Domain' option is highlighted with a red border. Below the dropdown, there is a button labeled 'CREATE DOMAIN'. At the bottom of the form, there are 'Cancel' and 'Apply' buttons. The 'ID' is '55812100017' and the 'Created By' is 'admin'.

Figure 4-9 Associate a user with a domain

This action results in a warning message, which is shown in Figure 4-10. The system is warning the storage administrator that the particular user will lose all global privileges of the system. A user that is associated with a domain can access the resources that belong to that domain only.

Warning
The user ITSO_Test will lose global privileges.
Do you want to continue?

Figure 4-10 Warning message

Important: Associating a global user with a domain removes the global privileges from that user. Therefore, that user can only access domains that the user is associated with. For a local user that is associated with one or more domains, removal of that local user from its last domain association results in the deletion of the user itself from the system. A user cannot be converted from a domain user to a global user.

Figure 4-11 shows the user that was created in the **Users** view. The Domain field lists ITSO_Domain as the associated domain for this user.

User ^	System	Domain	Category
ITSO_Test	A9000	ITSO_Domain	Storage Adm...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...
manager_serv...	A9000	/Global Space/	Storage Admi...

Figure 4-11 New user that is listed

Reminder: If the domain policy is set to closed, this user is no longer visible by the global administrator.

4.2.4 Domain properties

After a domain is created, a global administrator with a storage administrator role can manage the system resources that are associated with that domain if the domain policy is set to open. We first look at the domain properties as shown in Figure 4-12. You can see this view by first selecting the **Domains** view, clicking a specific domain, and selecting **Actions**.

* Actions Domain Properties	
Properties	View Properties
Pools, System	Rename
Domain Capacity	Resize
Association	Modify Object Limits

Figure 4-12 Domain properties

Under the Properties section, the global administrator can perform the following tasks:

- ▶ View domain properties
- ▶ Rename a domain
- ▶ Resize a domain
- ▶ Modify object limits, such as pools and volumes

Most of these tasks are self-explanatory. We describe the option Modify Object Limits in more detail. After you select **Actions** → **Modify Object Limits**, the properties of the domain are shown to the right of the tabular view, which is displayed in Figure 4-13. The editable fields are shown in a lighter-colored or dark blue box, such as Size (GB).

The global administrator can change the domain based on requirements or requests from the domain administrator, such as requesting an increase in the number of volumes that the domain can configure and use.

Name		Externally Managed
ITSO_Domain		N/A
Domain Preview Allocation Written		
<div><div><div></div><div>GB</div><div>0</div><div>2,067</div></div><div>Allocated Volumes & Snapshots (GB) 0</div><div>Domain Size (GB) 2067</div></div>		
LDAP Domain ID		ID
ITSO_Domain		4c431a400001
Max Pools		Max Volumes and Snapshots
2		15
Max CGs		Max mirrors and HyperSwaps
2		10
Max Data Migrations		
3		

Figure 4-13 Modify domain objects

After the changes are made, the global administrator clicks **Apply**. The changes take effect immediately, and they are ready for use by the domain administrator. Other changes can be made as needed in this section.

Note: Although we selected 2000 GB for the domain size, the system rounded the value to 2067 GB.

4.2.5 Managing domains

From the **Actions** menu that is shown in Figure 4-12 on page 128, the global administrator can manage various aspects of a domain, as explained in the following sections.

Managing domain pools and systems

The global administrator can globally manage a domain by using the **Pools, Systems** option. The **View Pools/Systems** and **Create Pool Here** options are shown in Figure 4-14.

✕ Actions Domain Properties	
Properties >	View Pools/Systems
Pools, System >	Create Pool Here

Figure 4-14 Domain: Pools and systems

Figure 4-15 shows an example of the **View Pools/Systems** output. This information is displayed to the right of the tabular view under the Hub graphic. This information can also be seen in the tabular view with customization of the view. Although the GUI displays the name of the system to which the domain belongs, the GUI does not return any details about the pools other than the number of pools that are configured within the domain.

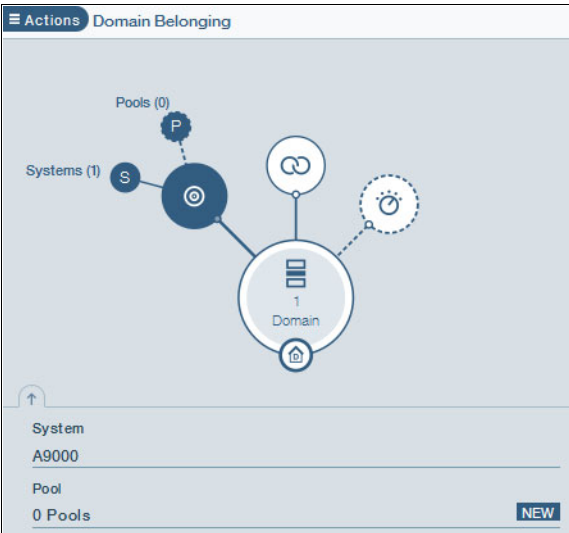


Figure 4-15 Domain: List of pools and systems

When you select **Create Pool Here** (Figure 4-16), the information displays to the right of the tabular view. This selection opens the Create Pool wizard. The details for the new pool need to be completed. The system and domain fields are prepopulated based on where the wizard was launched from.

The screenshot shows the 'Create Pool Here' wizard interface. At the top right, it says 'Units: GB | v'. The form has several sections: 'Name' with a red asterisk and info icon, a text input field, and radio buttons for 'Regular' (selected) and 'Thin'; 'Hard Size (GB)' with an info icon, showing 'N/A'; 'Snapshot Reserve...' with a red asterisk and a text input field; 'Quantity' with a text input field showing '1'; 'Pool Size (GB)' with a red asterisk and info icon, and a text input field; 'System' with a dropdown menu showing 'A9000'; 'Domain' with a dropdown menu showing 'ITSO_Domain'; and two radio button options: 'Read-only (when 100% usage is reached)' (selected) and 'No IO (when 100% usage is reached)'. At the bottom, there are two radio button options: 'Create Volumes as Uncompressed' and 'Create Volumes as Compressed' (selected), with an info icon to the right.

Figure 4-16 Create a pool for a domain

Figure 4-17 on page 132 shows an example of the completed fields. Despite arriving at this wizard from the Domain view, you can change the system and the domain through the drop-down lists as indicated in the example.

Tip: Enough space must be left in the domain for the requested pool size or the system does not allow the creation of the pool.

After the global administrator clicks **Apply**, the new pool is associated with the domain and ready for use. After the pool parameters are set and applied, the pool size is automatically rounded up, as shown in Figure 4-17. To view the pool and change any information, select the **Pool view**. The global administrator uses this selection to change the pool, including the domain which it belongs to, if the domain policy is set to Open.

Figure 4-17 New pool properties

Important:

- ▶ Attempting to add a pool to a domain that is larger than the domain's free capacity will fail.
- ▶ Associating an existing pool that has preexisting resource associations (hosts, clusters, or targets) with a domain, automatically associates those resources with the domain.
- ▶ Associating an existing pool with a domain when it has a preexisting association with another domain removes the association with the other domain.
- ▶ A pool can belong to only one domain at a time.

After the pool is created with this method, the Pool view is displayed, as shown in Figure 4-18.

Pool ^	System	Domain	Pool Size	Vol...	Allocated Volum...
Container_R...	A9000	ITSO_SCB_...	2,067 GB	u_i...	0% <div></div>
ITSO	A9000	/Global Spa...	40,002 ...	4	4% <div></div>
ITSO_p1	A9000	/Global Spa...	2,067 GB	0	0% <div></div>
ITSO_Pool	A9000	ITSO_Domain	2,067 GB	0	0% <div></div>

Figure 4-18 Viewing the new pool and domain association

Figure 4-19 shows the graphical Hub view of the domain associations. We expanded the **Belonging** circle, which shows that the domain is associated with one system and one pool. The global administrator can also create and associate a new pool from this view by selecting **NEW** on the listed pool line.

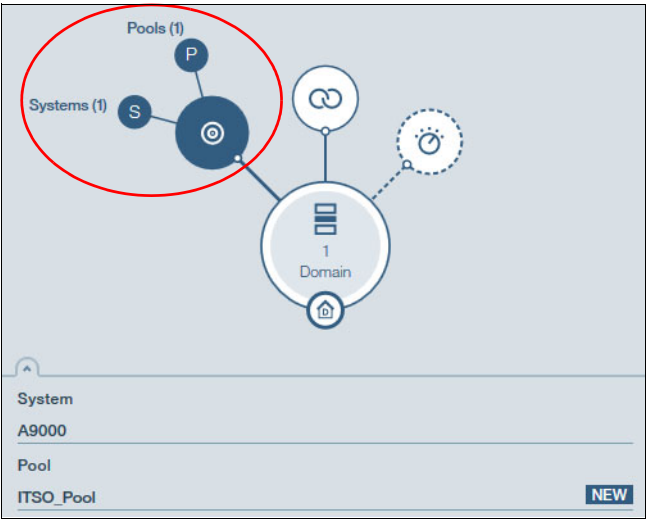


Figure 4-19 Graphical domain pool list

Domain capacity

This action allows the storage administrator to generate a capacity report for just the domain. Figure 4-20 shows the **Domain Capacity** option under the **Actions** menu. The **Capacity Trend/Forecast Report** option is available under **Domain Capacity**.

x Actions Domain Properties	
Properties	Capacity Trend/Forecast Report
Pools, System	
Domain Capacity	

Figure 4-20 Generate domain capacity report

After you select this option, Hyper-Scale Manager generates the report and automatically downloads it through the browser to your workstation, as shown in Figure 4-21.




Figure 4-21 Downloading the capacity planing report

Restriction: The domain administrator can generate a capacity planning report, but the domain administrator can see resources from the domain only, not system-wide resources.

Domain associations

The next section under **Domain Actions** is **Association**. Two options exist: **View Associations** and **Unassociate All**, as shown in Figure 4-22.




Figure 4-22 View domain associations

Looking at the details that are displayed with **View Associations**, we can see the Hub figure that is shown to the right of the tabular view (Figure 4-23).

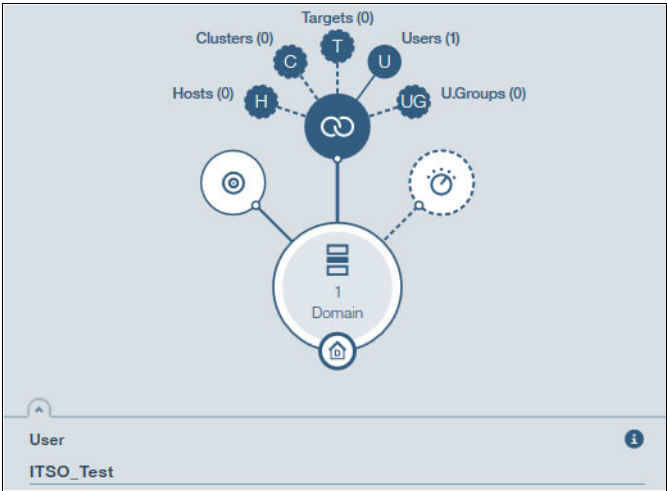


Figure 4-23 Domain users

Here, you can see the number of objects that are associated with the domain under each category in a graphical view. The association categories are listed:

- Hosts
- Clusters
- Targets (for mirroring, HyperSwap or migration)
- Users
- User Groups

This domain is new, so the only association is a single user. Again, no details about the associations are shown, just the quantity for each category.

On Figure 4-22 on page 134, by selecting **Unassociate All**, you remove the associations to the domain.

Reminder: Unassociating all objects from a domain deletes the users from the system if they are not associated with another domain. The user running the command will not be removed from the system. “Users cannot remove themselves”.

Domain quality of service

The global administrator can also assign a performance class to a domain by setting a quality of service for the domain. For information about quality of service, see 4.3, “Performance classes (quality of service)” on page 143.

Domain statistics

The next option on the **Actions** menu is the **Statistics** option. Two options are available: **Show Performance Statistics** and **Show Capacity Statistics**, as shown in Figure 4-24.

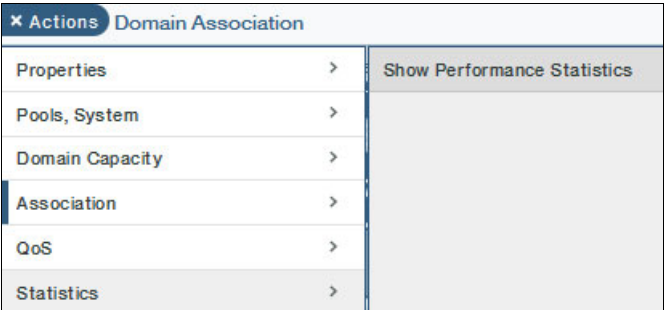


Figure 4-24 Domain statistics

In our example, we select **Show Performance Statistics** for the domain from the **Actions** menu. This action causes the Statistics view to expand, as shown in Figure 4-25.

From this view, you can select a specific interface and other statistical views as shown in the drop-down list. Two buttons are available on the upper-right side of this view. If the plus sign (+) is selected, a new statistics window opens above the current statistics window. The domain administrator can use this new window to see multiple options in the same view. By clicking x, the domain administrator closes this expanded view and returns to the summary, as shown in Figure 4-25.

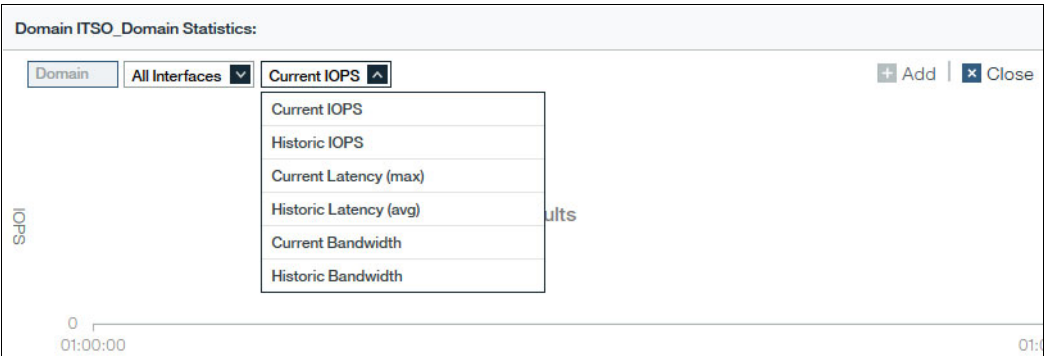


Figure 4-25 Expanded statistics view

4.2.6 Managing within a domain

Up to this point, we focused on what the global administrator can do with a domain if the domain policy is set to Open. We briefly look at the view from the perspective of the domain administrator with the storage administrator role.

After the domain administrator connects to the GUI, the Dashboard view is shown, as with all other users. However, the amount of capacity and the amount of potential capacity that are seen by this user are restricted. We look at the All Systems table view, as shown in Figure 4-26. The state of the system and version are the same as with the global administrator but the amount of capacity that is seen by this user is less.


System ^	Status	Hardware T...	Version	Written ...	Projecte...	Free to ...
A9000	 Fully Protected	A9000:5.7TB...	12.2.0	0 TB	N/A	1.76 TB

Figure 4-26 All Systems view for the domain administrator

We can see a similar view when we look at the Hub and System Properties, as shown in Figure 4-27.



Figure 4-27 System Properties from the domain administrator's viewpoint

We first look at creating a pool. The domain administrator clicks **NEW** in the upper-right corner of the menu and selects **Pool** from the list, as shown in Figure 4-30. Selecting this option opens the new pool wizard.

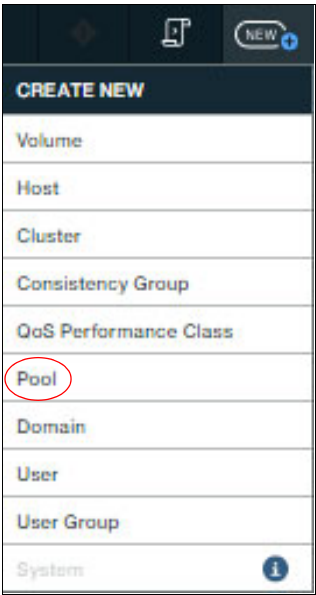


Figure 4-30 Domain administrator creates a pool

Restriction: A domain administrator is unable to create a domain. Attempting to create a new domain results in failure.

After the wizard is displayed, the domain administrator completes the required fields. The size of the pool cannot exceed the size of the domain. If the domain administrator tries to create a pool that is larger than the domain, the Hyper-Scale Manager GUI displays a message that indicates insufficient free space. In addition, the domain administrator must select their domain from the list and their system. An example is shown in Figure 4-31.



Figure 4-31 Create a pool within a domain

Also, only the domains that the domain administrator has authority to access display in the list of available domains (Figure 4-32).




Figure 4-32 List of domains for a new pool

After the pool is created, the domain administrator can create volumes within the pool container. Although the new volumes wizard can be launched from the **NEW** option as shown in Figure 4-30 on page 138, we show another path to launch the wizard, which automatically prepopulates several of the fields. From the pools table view, select a pool and then click **Actions**. From the list of actions, select **Volumes** and click **Create volumes here** as shown in Figure 4-33.




Figure 4-33 Create volumes from the pool view

After the Create Volume wizard opens, the system and domain fields are prepopulated based on the pool that was selected in the table view. Even though these fields can be changed if needed, the wizard takes care of this step for you. An example is shown in Figure 4-34.

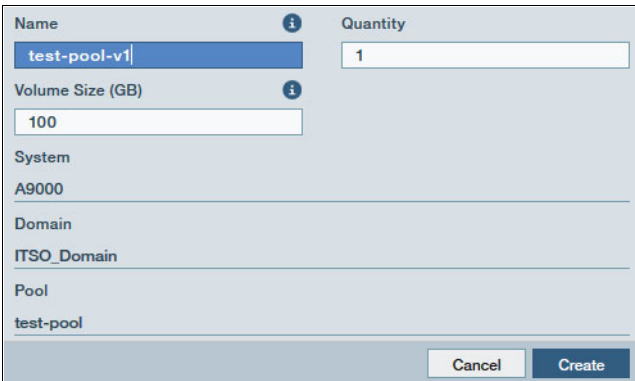


Figure 4-34 Create a volume within a domain

The domain administrator can list only the volumes within the domain that they are associated with. Although many more volumes might be on the system, only those volumes in the domain or domains are listed. An example is shown in Figure 4-35. The volume in this list is associated with a domain, unlike the volumes that the global administrator sees if the domain policy is set to Open (all global space volumes and domain volumes).

1 Volume						
	Volume ^	System	Pool	Volume...	Size (Di...	Snap...
	test-pool-v1	A9000	test-pool	100 GB	103 GB	0

Figure 4-35 Volumes within a domain

An example of what the global administrator sees is shown in Figure 4-36.

1 selected out of 30 Volumes						
	Volume ^	System	Domain	Pool	Volume...	Size (Di...
	ITSO_SVC_vol_014	A9000	/Global Space/	ITSO_S...	29,000 ...	29,045 ...
	ITSO_SVC_vol_015	A9000	/Global Space/	ITSO_S...	29,000 ...	29,045 ...
	ITSO_SVC_vol_016	A9000	/Global Space/	ITSO_S...	29,000 ...	29,045 ...
	ITSO_SVC_vol_017	A9000	/Global Space/	ITSO_S...	29,000 ...	29,045 ...
	ITSO_SVC_vol_018	A9000	/Global Space/	ITSO_S...	29,000 ...	29,045 ...
	test-pool-v1	A9000	ITSO_Domain	test-pool	100 GB	103 GB
	u_instance1_Docker...	A9000	ITSO_SCB_DOM	Contain...	5 GB	5 GB

Figure 4-36 Global administrator's view of volumes

The domain administrator can create as many volumes and pools as are allocated for the domain by the global administrator. If the domain administrator runs out of capacity to create pools or volumes, a request for more capacity can be made. Only the global administrator can increase the capacity and the number of pools and volumes that can be created in the domain. An example of running into the volume limit is shown in Figure 4-37. The system displays a message about the limit. In this example, the message states that the maximum number of volumes was created in this domain.

The destination Domain reached its max number of Volumes.

Name

vol_{001}

Quantity

15

Figure 4-37 Number of volumes exceeds the domain allocation

The domain administrator (and other authorized users) also can set up data migrations, mirroring, consistency groups, and snapshots, like users in the non-domain environment. The number of mirroring pairs, data migration pairs, and consistency groups depends on what the global administrator authorized for the domain. However, the number of snapshots does not need to be specified by the global administrator.

If snapshot space was reserved in the pool, volume snapshots can be created. One of the ways to create a snapshot is to select the **Pools & Volumes Views**, click a volume, and select **Actions** → **Snapshots** → **Create Snapshot(s)**, as shown in Figure 4-38. This task is performed the same way whether or not the volume is part of a domain. For more information about snapshots, see Chapter 2, “Storage management and configuration” on page 37.

x Actions Volume Properties	
Properties	Create Snapshot(s)
Snapshots	View Snapshots
Mirror	Restore from Snapshot

Figure 4-38 Create a snapshot of a volume within a domain

We look more closely at what the domain administrator will see when they look at the properties and Hub view for a volume (Figure 4-39).

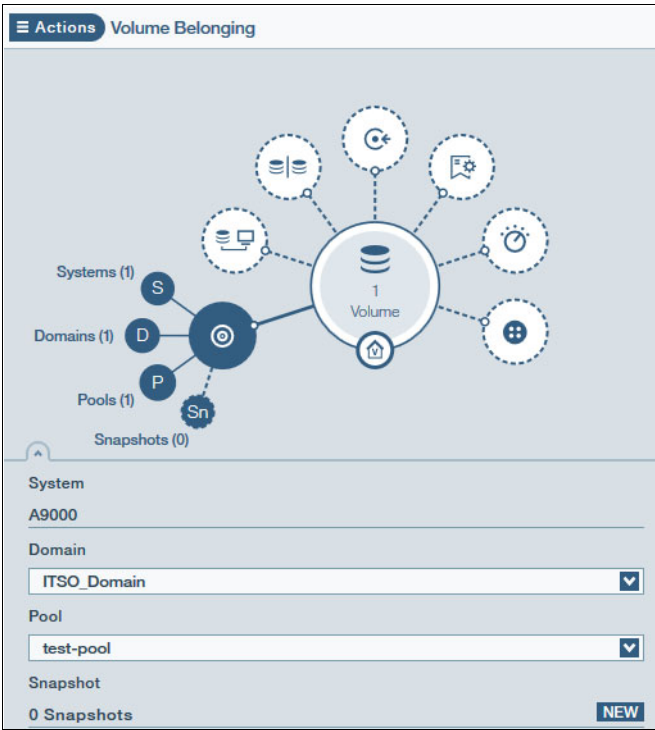


Figure 4-39 Volume properties within a domain

4.2.7 Removing a domain

To remove a domain, login as global administrator and select the **Systems & Domains Views** → **Domains**. Select the domain to remove. Click **Actions** → **Delete** → **Delete Domain(s)** as shown in Figure 4-40.

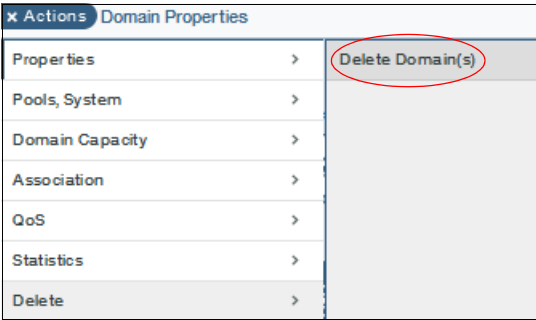


Figure 4-40 Domain: Delete a domain

This option deletes a domain if users and pools are not associated with the domain. If the pool has either of these associations, the system prevents the deletion and displays a message similar to the one shown in Figure 4-41. The global administrator must first remove or move the pool or pools from the domain and all of the users that are associated with the domain.



Figure 4-41 Domain: Preventing a deletion

We look at how the global administrator can move a pool and all of its volumes from a domain. Alternatively, the pool can be deleted.

The global administrator can move a pool from the Pools view by selecting the pool and clicking **Actions** → **Properties** → **View/Modify Domain**.

The Pool Belonging window opens, as shown in Figure 4-42. Alternatively, the global administrator can click the **Belonging** circle in the Hub view to display the same information.

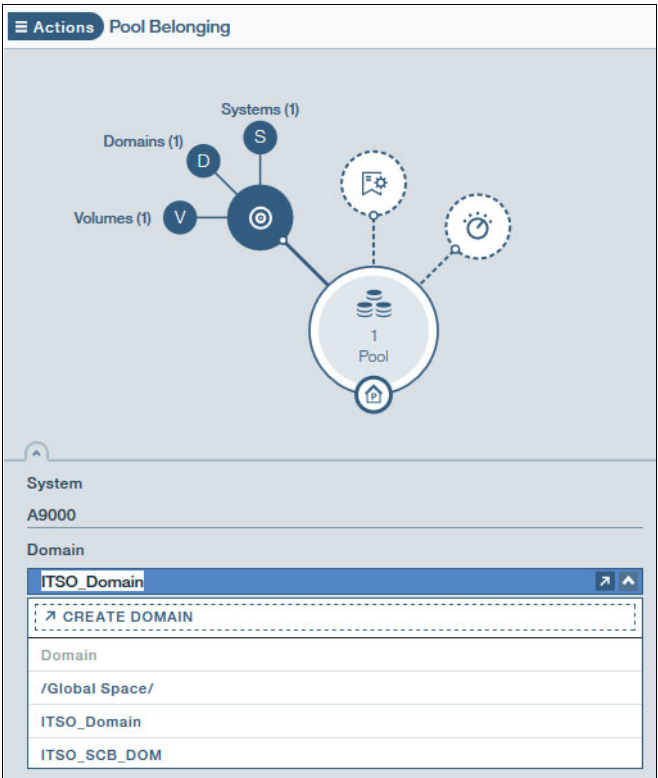


Figure 4-42 Pool Belonging window with system domains list

The global administrator can select the domain or Global Space to move the pool and all its volumes to. After a pool is selected, a new symbol (an asterisk (*)) is displayed next to the domain as shown in Figure 4-43. By hovering over this new symbol, additional information is displayed that indicates that the pool will be moved from one domain to another domain (or to the global space).




Figure 4-43 Domain attribute for a pool will change

After all of the pools are removed from a domain, the global administrator must disassociate the users from the domain. The domain can then be deleted from the system.

Important: Removing a user from the user's last domain association removes the user from the system.

4.3 Performance classes (quality of service)

The quality of service (QoS) feature allows IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate software to deliver different service levels to hosts that are connected to the same system. The QoS feature favors the performance of critical business applications that run concurrently with noncritical applications. Because resources are shared among all applications and all hosts are attached to the same resources, division of these resources among both critical and noncritical applications might negatively affect the performance of critical applications.

QoS can address this issue by limiting the rate, based on bandwidth and input/output operations per second (IOPS), for noncritical applications. Limiting performance resources for non-critical applications means that the remaining resources are available without limitation for the business-critical applications.

The QoS feature is managed through the definition of performance classes and then by associating hosts with a performance class. The feature can also be set by domains and storage pools. Each performance class is now implicitly one of two types: A host type or a pool/domain type.

The possibilities and limitations of the QoS feature can be summarized in the following manner:

- ▶ Up to 500 performance classes are configurable.
- ▶ QoS is applicable to host, domain, pool, volume, and restricted combinations of these entities. For instance, hosts cannot be specified for a performance class that already contains a domain or pool.
- ▶ Limits can be defined as *Total*, which means for an IBM XIV Gen3, IBM FlashSystem A9000 and A9000R, or IBM Spectrum Accelerate software system as a whole, or *Per Interface*.

- Limits are specified as IOPS or bandwidth, as shown in Figure 4-44.
- The limit calculation is based on preferred practices for setup and zoning.

The limited IOPS processes are expected to always come through all active interface nodes (equal to active interface modules). For example, on an FlashSystem A9000 with four grid controllers, the total I/O or bandwidth rate is divided by 4 (the number of active interface modules). If a limit total of 3,000 IOPS is specified, it results in a limitation of 750 IOPS per interface module.

In addition, if the limited IOPS come through only two of the four interface modules (as a result of the storage area network (SAN) zoning), the effective limitation is $2 \times 750 \text{ IOPS} = 1,500 \text{ IOPS}$ rather than the expected 3,000 I/O limitation.

The calculation-related information is visible in the GUI, as shown in Figure 4-44.



Figure 4-44 QoS can be set by bandwidth or IOPS

4.3.1 Creating a performance class

The management of the QoS feature is easy. To define a new performance class, click **NEW+** and then select **QoS Performance Class** from the menu (Figure 4-45), which opens the menu that is shown in Figure 4-46 on page 145.




Figure 4-45 QoS: Create a performance class

Figure 4-46 shows the window to create a QoS performance class.

Figure 4-46 Fields to create a QoS performance class

Enter a suitable performance class name and an IOPS limit, a bandwidth limit, or a combination of both, based on your business needs. You can set this limit so that your reporting systems are not allowed to use excessive IOPS to the detriment of other systems or other tenants of your storage system.

Two choices are available when you enter the limitation settings: *Shared* or *Independent*. Shared indicates that the limits will be shared across the QoS members. Independent means that the limit is specific to each QoS member.

4.3.2 Associating a domain with an existing performance class

After we create a QoS performance class, we can easily use this class to apply these limits to an existing domain. First, we navigate to the Domain view and select the domain that we are interested in on the left side of the window. Then, we click **Actions** → **QoS** → **Limit traffic by QoS Perf Class** as shown in Figure 4-47.

Figure 4-47 QoS: Limit a domain with a performance class

From the Domain QoS Performance Class window, we choose an existing QoS Performance Class, as in Figure 4-48 and click **Apply**. Alternatively, we can create a QoS performance class by selecting that option from the menu that is shown in Figure 4-48.

Figure 4-48 QoS: Choose a performance class to apply to this domain

Tip: You can also use the CLI. With the CLI, look for domain commands by typing `domain` followed by pressing the Tab key twice. To look at QoS commands, type `perf_class` followed by pressing the Tab key twice. For the user, type `user`, followed by pressing the Tab key twice. Typing `help` lists all of the commands.

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM FlashSystem A9000 and A9000R Architecture and Implementation (Version 12.3.1)*, SG24-8345
- ▶ *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401
- ▶ *IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM XIV Storage System: Host Attachment and Interoperability*, SG24-8368
- ▶ *IBM HyperSwap and Multi-site HA/DR for IBM FlashSystem A9000 and A9000R*, REDP-5434
- ▶ *Using the IBM Spectrum Accelerate Family in VMware Environments: IBM XIV, IBM FlashSystem A9000 and IBM FlashSystem A9000R, and IBM Spectrum Accelerate*, REDP-5425

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Other publications and online resources

These publications are also relevant as further information sources:

- ▶ IBM Fix Central:
<http://www.ibm.com/support/fixcentral/>
- ▶ IBM FlashSystem A9000 on IBM Knowledge Center:
<http://www.ibm.com/support/knowledgecenter/STJKN5>

The following publications are at this website:

- *Hyper-Scale Manager 5.3 REST API Specification*, SC27-6440
- *Hyper-Scale Manager 5.3 User Guide*, SC27-8560
- *IBM FlashSystem A9000 Command-Line Interface (CLI) Reference Guide*, SC27-8559
- *IBM FlashSystem A9000 Product Overview*, GC27-8583
- *IBM FlashSystem A9000 Models 9836-415, 9838-415, 9836-425, and 9838-425 Deployment Guide*, GC27-8564

- ▶ IBM FlashSystem A9000R on IBM Knowledge Center:
<http://www.ibm.com/support/knowledgecenter/STJKMM>
 The following publications are at this website:
 - *Hyper-Scale Manager 5.3 REST API Specification*, SC27-6440
 - *Hyper-Scale Manager 5.3 User Guide*, SC27-8560
 - *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide*, SC27-8711
 - *IBM FlashSystem A9000R Models 9835-415, 9837-415, 9835-425, and 9837-425 Deployment Guide*, GC27-8565
 - *IBM FlashSystem A9000R Product Overview*, GC27-8558
- ▶ IBM FlashSystem A9000 product page:
<http://www.ibm.com/systems/storage/flash/a9000>
- ▶ IBM FlashSystem A9000R product page:
<http://www.ibm.com/systems/storage/flash/a9000r>
- ▶ *IBM Hyper-Scale and Its Implementation in XIV Storage*, TSL03121USEN
<http://public.dhe.ibm.com/common/ssi/ecm/ts/en/ts103121usen/TSL03121USEN.PDF>
- ▶ *IBM Hyper-Scale Manager Installation as application*, GC27-5984
- ▶ *IBM Hyper-Scale Manager Installation as a virtual appliance*, GC27-5985
- ▶ *IBM Hyper-Scale: A powerful new approach to scaling storage management*, TSW03281USEN
<http://public.dhe.ibm.com/common/ssi/ecm/ts/en/tsw03281usen/TSW03281USEN.PDF>
- ▶ IBM Offering Information page (announcement letters and sales manuals):
http://www.ibm.com/common/ssi/index.wss?request_locale=en
 On this page, enter A9000, select the information type, and click **Search**. On the next page, narrow your search results by geography and language.
 - ▶ *IBM Spectrum Accelerate Command-Line Interface (CLI) Reference Guide*, SC27-6697
 - ▶ *IBM Spectrum Accelerate Planning, Deployment, and Operation Guide*, SC27-6695
 - ▶ *IBM Spectrum Accelerate Product Overview*, SC27-6696
- ▶ For the latest hardware and software requirements, see the IBM System Storage Interoperation Center (SSIC) website:
<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>
- ▶ *IBM XIV Storage System Application Programming Interface*, GC27-3916
- ▶ *IBM XIV Storage System Management Tools Operations Guide*, SC27-5986
- ▶ *IBM XIV Storage System: Product Overview*, GC27-3912
- ▶ *IBM XIV Storage System User Manual*, GC27-3914

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



SG24-8376-03

ISBN 0738457582

Printed in U.S.A.

Get connected

