

# Unleash the Power of Flash: Getting Started with IBM Storage Virtualize Version 8.7 on IBM Storage FlashSystem and IBM SAN Volume Controller

Vasfi Gucer

Andy Harchen

Jon Herd

Hartmut Lonzer

Jonathan Wilkie



Storage





IBM Redbooks

**Unleash the Power of Flash: Getting Started with IBM  
Storage Virtualize Version 8.7**

October 2024

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

**First (October 2024)**

This edition applies to IBM Storage Virtualize Version 8.7.0.

# Contents

<b>Figures</b> .....	vii
<b>Tables</b> .....	xi
<b>Examples</b> .....	xiii
<b>Notices</b> .....	xv
Trademarks .....	xvi
<b>Preface</b> .....	xvii
Authors .....	xvii
Now you can become a published author, too! .....	xix
Comments welcome .....	xix
Stay connected to IBM Redbooks .....	xix
<b>Chapter 1. Introduction and system overview</b> .....	1
1.1 IBM Storage Virtualize .....	2
1.1.1 Overview .....	2
1.1.2 Benefits of IBM Storage Virtualize .....	2
1.2 IBM Storage Virtualize V8.7 supported products .....	3
1.2.1 Status of supported products .....	3
1.2.2 IBM Storage Virtualize V8.7.0 current product features .....	4
1.3 Changes and enhancements in IBM Storage Virtualize V8.7.0 .....	6
1.3.1 File system awareness for ransomware detection .....	6
1.3.2 GUI updates .....	7
1.3.3 Flash Grids and storage partition migration .....	9
1.3.4 Automatic Patch Updating and Automatic Drive Downloading .....	10
1.3.5 Management IP changes .....	12
1.3.6 Remote Copy support on IBM Storage Virtualize V8.7 .....	12
1.4 Preparation and upgrading to IBM Storage Virtualize V8.7.0 .....	13
1.4.1 Confirm hardware support .....	13
1.4.2 Correct any issues as needed .....	13
1.4.3 Schedule the upgrade .....	14
<b>Chapter 2. Initial configuration</b> .....	15
2.1 Prerequisites .....	16
2.2 System initialization .....	17
2.2.1 Before starting initialization .....	17
2.2.2 System initialization process .....	20
2.3 System setup .....	24
2.3.1 System Setup wizard .....	24
2.3.2 Adding an enclosure in IBM FlashSystem .....	43
2.3.3 Adding a node or hot spare node in IBM SAN Volume Controller systems .....	45
2.3.4 Business continuity with policy-based high availability .....	49
2.3.5 Configuring quorum disks or applications .....	50
2.3.6 Configuring the local Fibre Channel port masking .....	53
2.3.7 Automatic configuration for IBM SAN Volume Controller back-end storage .....	55
<b>Chapter 3. Step-by-step configuration</b> .....	61
3.1 The Storage Virtualize GUI .....	62

3.1.1	Accessing the GUI	62
3.1.2	Brief introduction to the GUI	62
3.2	Network configuration	63
3.2.1	Management IP addresses	63
3.2.2	Service IP addresses	64
3.2.3	Additional Ethernet ports	64
3.2.4	Portsets	65
3.3	Pools and managed disks configuration	65
3.3.1	Provisioning policies	66
3.3.2	Types of pools	66
3.3.3	Ransomware threat detection	66
3.3.4	Creating storage pools	67
3.3.5	Creating RAID array managed disks in a storage pool	68
3.3.6	Adding external managed disks into a storage pool	69
3.3.7	Child pools	70
3.4	Configuring volumes	72
3.4.1	Creating volume groups	72
3.4.2	Creating volumes	72
3.4.3	Virtual volumes	73
3.5	Configuring hosts	73
3.5.1	Host attachment overview	73
3.5.2	Fibre channel host connectivity	74
3.5.3	Ethernet host connectivity	75
3.5.4	Host objects	77
3.5.5	Mapping volumes for host access	77
3.6	Snapshots and replication	78
3.6.1	Volume group snapshots	78
3.6.2	Asynchronous policy-based replication	80
	<b>Chapter 4. Verifying configuration and basic operations</b>	<b>83</b>
4.1	Verifying the configuration	84
4.1.1	System Health Dashboard	84
4.1.2	Verifying network configuration	84
4.1.3	Verifying storage configuration	85
4.1.4	Verifying volume configuration	85
4.1.5	Verifying host configuration	86
4.2	Additional settings and basic operations	87
4.2.1	Security settings	87
4.2.2	Audit log	89
4.2.3	Support settings	89
	<b>Chapter 5. IBM Storage Insights and IBM Storage Insights Pro</b>	<b>91</b>
5.1	IBM Storage Insights overview	92
5.2	IBM Storage Insights monitoring	93
5.2.1	Component health	94
5.2.2	Capacity monitoring	95
5.2.3	Performance monitoring	96
5.2.4	Logging support tickets by using IBM Storage Insights	98
5.2.5	Managing existing support tickets by using IBM Storage Insights	102
5.2.6	Enhancements to IBM Storage Insights Pro	104
	<b>Chapter 6. Storage Virtualize troubleshooting and diagnostics</b>	<b>107</b>
6.1	Troubleshooting	108
6.1.1	Storage Insights	109

6.1.2 Using the GUI . . . . .	109
6.1.3 Recommended actions and fix procedure. . . . .	110
6.1.4 Storage Virtualize failure recovery . . . . .	111
6.1.5 Using the command-line interface. . . . .	114
6.2 Collecting diagnostic data . . . . .	115
6.2.1 IBM Storage Virtualize systems data collection . . . . .	115
6.2.2 Drive data collection: drivedumps . . . . .	120
6.2.3 Host multipath software . . . . .	121
6.2.4 More data collection . . . . .	127
<b>Abbreviations and acronyms . . . . .</b>	<b>129</b>
<b>Related publications . . . . .</b>	<b>131</b>
IBM Redbooks . . . . .	131
Online resources . . . . .	131
Help from IBM . . . . .	131





# Figures

1-1 IBM FlashSystems and SVC Family . . . . .	3
1-2 Async policy-based replication and partition base HA user experience improvements . . . . .	7
1-3 Volume group tile and assigning ownership groups to volume groups . . . . .	8
1-4 GUI performance panel. . . . .	9
1-5 IBM Storage Virtualize upgrade support matrix . . . . .	13
2-1 Technician port FlashSystem 9500 . . . . .	18
2-2 Technician port FlashSystem 7300 . . . . .	18
2-3 Technician port FlashSystem 5200 . . . . .	18
2-4 Technician port FlashSystem 5300 . . . . .	18
2-5 Technician port FlashSystem 5045 . . . . .	18
2-6 Technician port FlashSystem 5015 . . . . .	19
2-7 Technician port IBM SAN Volume Controller 2145-SV3 . . . . .	19
2-8 Technician port IBM SAN Volume Controller 2145-SV2 . . . . .	19
2-9 Logging in to Service Assistant by way of the technician port . . . . .	21
2-10 System Initialization: Canister detection . . . . .	21
2-11 System Initialization: Initialize the first enclosure . . . . .	22
2-12 System Initialization: Initialize the first IBM SAN Volume Controller node . . . . .	22
2-13 System Initialization: Enter Management IP . . . . .	23
2-14 System Initialization: Web-server restart timer counting down from 5 minutes. . . . .	23
2-15 System Initialization completed. . . . .	24
2-16 Logging in for the first time . . . . .	25
2-17 Initial Setup Window . . . . .	25
2-18 Setup Call Home . . . . .	26
2-19 Transmission Types for Call Home . . . . .	27
2-20 Setup Internal Proxy Server . . . . .	28
2-21 Connection Test to the Support Center . . . . .	28
2-22 System Location . . . . .	29
2-23 Summary page . . . . .	30
2-24 System Setup Welcome page. . . . .	30
2-25 Accept License Agreement. . . . .	31
2-26 Change password . . . . .	32
2-27 System Name . . . . .	33
2-28 License Functions . . . . .	34
2-29 DNS Server setup . . . . .	35
2-30 Date and Time . . . . .	35
2-31 Activate Encryption License . . . . .	36
2-32 Encryption Licensed . . . . .	37
2-33 Change Call Home settings . . . . .	37
2-34 Setup Support Assistance. . . . .	38
2-35 System communicating with named IBM Support servers . . . . .	39
2-36 Remote support access settings. . . . .	39
2-37 Automatic Configuration for Virtualization. . . . .	40
2-38 Summary Page . . . . .	41
2-39 System Initialization . . . . .	42
2-40 Setup completed . . . . .	42
2-41 Dashboard . . . . .	43
2-42 Add Enclosure button . . . . .	43
2-43 Selecting the control enclosure to add . . . . .	44

2-44	Add Node button . . . . .	46
2-45	Adding a node . . . . .	46
2-46	IBM SAN Volume Controller is adding node to the cluster . . . . .	47
2-47	Node added . . . . .	47
2-48	Download IPv4 quorum button . . . . .	51
2-49	Download IP quorum application window . . . . .	51
2-50	IP quorum application that is deployed and connected . . . . .	52
2-51	Changing the quorum mode . . . . .	53
2-52	Applying a port mask by using a GUI . . . . .	54
2-53	Modify Connection dialog box . . . . .	54
2-54	Automatic Configuration wizard enablement . . . . .	56
2-55	Automatic configuration: Add Enclosure . . . . .	56
2-56	Defining a host cluster . . . . .	57
2-57	Hosts inside an IBM SAN Volume Controller host cluster . . . . .	58
2-58	Begin the automatic configuration process . . . . .	58
2-59	Automatic pool configuration . . . . .	59
2-60	Pools configuration . . . . .	59
2-61	Automatic configuration running commands . . . . .	60
2-62	Automatic configuration complete . . . . .	60
3-1	Welcome page with the dashboard . . . . .	62
3-2	Ethernet ports . . . . .	64
3-3	Add IP address . . . . .	65
3-4	Portsets . . . . .	65
3-5	Create Pool . . . . .	67
3-6	Create Pool panel . . . . .	67
3-7	Add Storage . . . . .	69
3-8	RAID array . . . . .	69
3-9	Child pools with different purposes . . . . .	70
3-10	Create Child Pool . . . . .	71
3-11	Add Host . . . . .	78
3-12	Snapshot policies . . . . .	80
3-13	Suspend Policy . . . . .	80
3-14	Create Partnership . . . . .	81
3-15	Partnership Created . . . . .	82
3-16	Setup policy-based replication wizard . . . . .	82
4-1	System Health . . . . .	84
4-2	Portset Mappings . . . . .	84
4-3	Portset IP addresses . . . . .	85
4-4	Host with asymmetrical logins . . . . .	86
4-5	Audit log . . . . .	89
5-1	IBM Storage Insights System overview (classic view) . . . . .	93
5-2	IBM Storage Insights System overview (Carbon enhanced view) . . . . .	94
5-3	Component Health overview . . . . .	94
5-4	Ports in error . . . . .	95
5-5	Capacity area of the IBM Storage Insights system overview . . . . .	95
5-6	Capacity planning for one system . . . . .	96
5-7	System overview: Performance . . . . .	96
5-8	IBM Storage Insights: Performance view . . . . .	97
5-9	Filtered performance graph . . . . .	97
5-10	Performance List View . . . . .	98
5-11	Get Support (see highlighted area) . . . . .	98
5-12	Get Support window . . . . .	99
5-13	Create Ticket wizard . . . . .	99

5-14	Add a note or attachment window . . . . .	100
5-15	Selecting a Severity Level window . . . . .	100
5-16	Review the ticket window . . . . .	101
5-17	Update ticket . . . . .	102
5-18	View tickets . . . . .	102
5-19	Adding a log package to the ticket . . . . .	103
5-20	Confirming the log upload . . . . .	103
5-21	Log upload completed and processing . . . . .	104
5-22	IBM Storage insights Pro and IBM Flash Grid integration. . . . .	105
6-1	Events icon in the GUI . . . . .	109
6-2	System Health expanded section in the dashboard . . . . .	110
6-3	Recommended actions . . . . .	111
6-4	Monitoring → Events window . . . . .	113
6-5	Properties and Sense Data for an event . . . . .	114
6-6	Upload Support Package details. . . . .	118



# Tables

1-1 IBM Storage Virtualize V8.7 supported product list . . . . .	4
1-2 IBM FlashSystem current products feature summary comparison chart . . . . .	4
6-1 Useful AIX lspath commands . . . . .	122
6-2 Useful AIX lsmpio commands . . . . .	122
6-3 Useful Windows mpclaim.exe commands . . . . .	123
6-4 Useful Windows PowerShell cmdlets . . . . .	123



# Examples

2-1	Reenabling the onboard Ethernet port 2 as the technician port . . . . .	17
2-2	Listing the I/O groups . . . . .	44
2-3	Listing the candidate control enclosures . . . . .	45
2-4	Adding a control enclosure . . . . .	45
2-5	Adding an expansion enclosure . . . . .	45
2-6	Listing I/O groups . . . . .	48
2-7	Listing the candidate nodes . . . . .	48
2-8	Adding a node as a spare . . . . .	48
2-9	Adding a node to an I/O group . . . . .	48
2-10	Single IO-group (two nodes) and one spare . . . . .	49
2-11	Two IO-groups (four nodes) configured- no spare . . . . .	49
2-12	Starting the IP quorum application on the Windows operating system . . . . .	51
2-13	Viewing the local port mask . . . . .	55
2-14	Setting a local port mask by running the chsystem command . . . . .	55
6-1	The svc_livedump command . . . . .	119
6-2	preplivedump and lslivedump commands . . . . .	120
6-3	The triggerdrivedump command . . . . .	120
6-4	Output for the multipath -ll command . . . . .	121
6-5	Output of esxcli storage core path list command . . . . .	124
6-6	Output of esxcli storage core path list -d <naaID> . . . . .	125
6-7	Output for esxcli storage nmp device list . . . . .	126





# 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 <https://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.

Redbooks (logo)  ®	HyperSwap®	IBM Research®
AIX®	IBM®	IBM Spectrum®
DS8000®	IBM Cloud®	PowerHA®
Easy Tier®	IBM FlashCore®	Redbooks®
FlashCopy®	IBM FlashSystem®	

The following terms are trademarks of other companies:

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

ITIL is a Registered Trade Mark of AXELOS Limited.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Red Hat, OpenShift, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, VMware vSphere, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

# Preface

IBM® Storage Virtualize (formerly IBM Spectrum® Virtualize) can simplify infrastructure management for block storage across diverse workloads on-premises, off-premises, or in hybrid cloud environments. This core offering of the IBM Storage portfolio enables rapid deployment and streamlines management for SAN Volume Controller and IBM FlashSystem® systems, including support for hybrid multicloud deployments.

This IBM Redbooks® publication focuses on IBM Storage Virtualize Version 8.7, guides users through new features, upgrades, and configuration for both new and existing systems. It is intended for pre-sales and post-sales technical support personnel and storage administrators.

## Authors

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



**Vasfi Gucer** leads projects for the IBM Redbooks team, leveraging his 20+ years of experience in systems management, networking, and software. A prolific writer and global IBM instructor, his focus has shifted to storage and cloud computing in the past eight years. Vasfi holds multiple certifications, including IBM Certified Senior IT Specialist, PMP, ITIL V2 Manager, and ITIL V3 Expert.



**Andy Harchen** is a subject matter expert (SME) who is based at the EMEA TLS Remote Technical Support Storage, IBM Germany. He has over 25 years of experience in onsite and remote technical support with a focus on disk storage and virtualization solutions. In his current role, he delivers technical support for IBM Storage Virtualize Products (IBM SAN Volume Controller and IBM FlashSystem). He is a member of the Predictive Support Team, which provide proactive support based on cloud technology such as Storage Insights, supported by AI to detect and prevent impacts and outages of systems.



**Jon Herd** is an IBM Senior Executive Advocate working for the TLS EMEA Remote Technical Support and Client Care team based in IBM Germany. He covers the United Kingdom, Ireland and beyond, advising customers on a portfolio of IBM storage products, including FlashSystem products. He also works as a senior advisor to the TLS EMEA RTS/CC management on new products, strategy, and new technologies that might affect the TLS business. Jon has been with IBM for more than 45 years, and has held various technical roles, including Europe, Middle East, and Africa (EMEA) level 2 support on mainframe servers and technical education development. He has written many IBM Redbooks on the FlashSystem products and is an IBM Redbooks Platinum level author. He holds IBM certifications in Product Services profession at a thought leader L3 level, and is a Technical Specialist at an experienced L1 level. He also is a certified Chartered Member of the British Computer Society (MBCS - CITP), a Certified Member of the Institution of Engineering and Technology (MIET), and a Certified Technical Specialist of the Open Group (TOG).



**Hartmut Lonzer** brings 45 years of technical and sales expertise at IBM to his role as a Storage Advisory Partner Technical Specialist for DACH. Previously, he served as an OEM Alliance Manager for Lenovo at IBM Germany. Based at the company's headquarters in Ehningen, his focus lies on the IBM FlashSystem Family and IBM SAN Volume Controller products, with experience dating back to their introduction.



**Jonathan Wilkie** is an Advanced Subject Matter Expert/L3 support representative for IBM Spectrum® Virtualize and IBM FlashSystem. He has more than 20 years of experience in IBM storage technical support. Over his career, he has provided technical support for Shark, DS4000, DS6000, and IBM DS8000® products. He has been supporting IBM Storage Virtualize-based products since 2010.

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

**Elias Luna, Andrew Greenfield**  
IBM USA

**Lucy Harris, Evelyn Perez, Chris Bulmer, Chris Canto, Daniel Dent, Bill Passingham, Nolan Rogers, David Seager, Russell Kinmond**  
IBM UK

**Sushil H Sharma, Ramakrishna Vadla**  
IBM India

**Diana Laura Silva Gallardo**  
IBM Mexico

## 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 IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at: [ibm.com/redbooks/residencies.html](http://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 found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- ▶ Send your comments in an email to:  
[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)
- ▶ Mail your comments to:  
IBM Corporation, IBM Redbooks  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:  
<https://www.linkedin.com/groups/2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:  
<https://www.redbooks.ibm.com/subscribe>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:  
<https://www.redbooks.ibm.com/rss.html>





# Introduction and system overview

This chapter dives into IBM Storage Virtualize V8.7, covering its new features, supported systems, and upgrade considerations.

This chapter has the following sections:

- ▶ 1.1, “IBM Storage Virtualize” on page 2
- ▶ 1.2, “IBM Storage Virtualize V8.7 supported products” on page 3
- ▶ 1.3, “Changes and enhancements in IBM Storage Virtualize V8.7.0” on page 6
- ▶ 1.4, “Preparation and upgrading to IBM Storage Virtualize V8.7.0” on page 13

# 1.1 IBM Storage Virtualize

## 1.1.1 Overview

IBM Storage Virtualize (formerly IBM Spectrum Virtualize) can simplify managing block storage for various workloads, on-premises or in the cloud. It runs on IBM FlashSystem and SAN Volume Controller hardware, offering data protection, rapid cloud deployment, and performance for analytics. IBM Storage Virtualize provides a way to manage and protect huge volumes of data from mobile and social applications. IBM Storage Virtualize enables rapid and flexible cloud services deployments and delivers the performance and scalability that is needed to gain insights from the latest analytics technologies.

**Note:** For more information, see [IBM Storage FlashSystem](#) and [IBM SAN Volume Controller](#).

With the introduction of the IBM Storage family, the *software* that runs on IBM SAN Volume Controller and on IBM Storage FlashSystem (IBM FlashSystem) products is called IBM Storage Virtualize. The name of the underlying hardware platform is not changed.

## 1.1.2 Benefits of IBM Storage Virtualize

IBM Storage Virtualize delivers benefits that improve storage infrastructure in many ways, including the following examples:

- ▶ Reduces the cost of storing data by increasing the use and accelerating applications to speed business insights. To achieve this goal, the solution provides the following functions:
  - Uses data reduction technologies to increase the amount of data that you can store in the same space
  - Enables rapid deployment of cloud storage for disaster recovery (DR) along with the ability to store copies of local data
  - Moves data to the most suitable type of storage based on policies that you define by using IBM Storage Control to optimize storage
  - Improves storage migration performance so that you can do more with your data
- ▶ Protects data from theft or inappropriate disclosure and enables a high availability (HA) strategy that includes protection for data and application mobility and DR. To achieve this goal, the solution provides the following functions:
  - Uses software-based encryption to improve data security
  - Provides fully duplexed copies of data and automatic switchover across data centers to improve data availability
  - Eliminates storage downtime with nondisruptive movement of data from one type of storage to another type.
- ▶ Simplifies data by providing a data strategy that is independent of your choice of infrastructure, which delivers tightly integrated functions and consistent management across heterogeneous storage. To achieve this goal, the solution provides the following functions:
  - Integrates with virtualization tools, such as VMware vCenter to improve agility with automated provisioning of storage and easier deployment of new storage technologies



- Enables supported storage to be deployed with Kubernetes and Docker container environments, including Red Hat OpenShift
- Consolidates storage, regardless of the hardware vendor for simplified management, consistent functions, and greater efficiency
- Supports common capabilities across storage types, which provide flexibility in storage acquisition by allowing a mix of vendors in the storage infrastructure

**Note:** These benefits are a subset of the list of features and functions that are available with IBM Storage Virtualize software.

Figure 1-1 shows the current IBM FlashSystem and IBM SAN Volume Controller Family.








<h2>FlashSystem Family</h2>  <p><b>FlashSystem 5015 &amp; 5045</b></p> <p>SAS based entry to the world of IBM Storage Virtualize and IBM Storage FlashSystem.</p>	<p>Released May 2024</p>  <p><b>FlashSystem 5300</b></p> <p>Next-gen award-winning NVMe efficiency leader. Densest 1U system unlocks the FlashCore Module value proposition along with significant data resilience.</p>	 <p><b>FlashSystem 7300</b></p> <p>The 5<sup>th</sup> generation of the award-winning 7000 series, striking the right performance, scaling and value balance.</p>	 <p><b>FlashSystem 9500</b></p> <p>Extreme performance for large, mixed and consolidated workloads</p>
<h2>Powered by IBM Storage Virtualize</h2>			
 <p><b>Storage Virtualize for Public Cloud</b></p> <p>Hybrid Cloud integration with all FlashSystem products allows seamless data migration from on premise to your chosen cloud provider.</p>	 <p><b>SAN Volume Controller (SVC) SV3 Storage Engine</b></p> <p>Celebrating 20 years in market, SVC supports more than 500 models of storage controller with seamless SVC and storage replacement.</p>	 <p><b>IBM FlashCore Modules</b></p> <p>The 4<sup>th</sup> generation of one of FlashSystem key value propositions with:</p> <ul style="list-style-type: none"> <li>• AI-enabled Ransomware Threat Detection embedded in hardware without any performance degradation</li> <li>• Data Reduction with hardware offload</li> <li>• Densest Tier-0 flash drive</li> </ul>	

Figure 1-1 IBM FlashSystems and SVC Family

**Note:** IBM Storage Virtualize for Public Cloud is not currently supported on IBM Storage Virtualize V8.7. This function is planned for a future release.

## 1.2 IBM Storage Virtualize V8.7 supported products

### 1.2.1 Status of supported products

Table 1-1 on page 4 shows the IBM Storage Virtualize V8.7 supported product list and whether the product is still currently sold or is designated as End-of-Marketing (EOM).

Table 1-1 IBM Storage Virtualize V8.7 supported product list

Product	Machine Type	Model	Comment
FS9500/R	4666, 4983	AH8, UH8	Current Product
FS7300	4657	924, U7D	Current Product
FS5300	4662	7H2	Current Product
FS5200	4662	6H2, UH6	Current Product
FS5045	4680	3P2, 3P4	Current Product
FS5015	4680	2P2, 2P4	Current Product
SVC	2145, 2147	SA2, SV3	Current Product
SVC	2145, 2147	SV2	EOM 01/2023
FS9200/R	9846, 9848, 4666	AG8, UG8	EOM 07/2022
FS7200	2076, 4664	824, U7C	EOM 07/2022
FS9100	9846, 9848	AF8, UF8	EOM 07/2022
FS5000 (FS5015, FS5035)	2072	2N2, 2N4, 3N2, 3N4	EOM 12/2023

**Note:** This version of the IBM Redbooks includes systems that can run IBM Storage Virtualize V8.7. Some products that are listed in the book are no longer sold by IBM but can still run the V8.7 software. Where this is applicable, it is mentioned in the text.

## 1.2.2 IBM Storage Virtualize V8.7.0 current product features

This is a brief summary of the technical specifications for this solution. It provides a concise overview of the essential features and specifications.

Table 1-2 shows the IBM Storage FlashSystem Family feature summary and comparison, for currently marketed products.

Table 1-2 IBM FlashSystem current products feature summary comparison chart

	SVC	5015	5045	5300	7300	9500
<b>Machine Type</b>	2145 2147	4680	468	4662	4657	4983 4666
<b>Controller Models</b>	SA2 (No Drives) SV3 (No Drives)	2P2 (12-drive) 2P4 (24-drive)	3P2 (12-drive) 3P4 (24-drive)	7H2 (12-drive)	924 (24-drive)	AH8 (48-drive)
<b>Expansion Models</b>	N/A	12H (12-drive) 24H (24-drive) 92H (92-drive)	12H (12-drive) 24H (24-drive) 92H (92-drive)	12G (12-drive) 24G (24-drive) 92G (92-drive)	12G (12-drive) 24G (24-drive) 92G (92-drive)	AFF (24-drive) A9F (92-drive)

	<b>SVC</b>	<b>5015</b>	<b>5045</b>	<b>5300</b>	<b>7300</b>	<b>9500</b>
<b>Processors</b>	2 Intel Xeon CPUs SV3 24 cores each SA2 8 Cores each	2 Intel Xeon CPUs 2 cores each	2 Intel Xeon CPUs 6 cores each	2 Intel Xeon CPUs 12 cores each	2 Intel Xeon CPUs 10 cores each	4 Intel Xeon CPUs 24 cores each
<b>Memory</b>	SA2 128 GB 386 GB 764 GB SV3 512 GB 1 TB 1.5 TB	32 GB 64 GB	32 GB 64 GB	64 GB 256 GB 512 GB	256 GB 756 GB 1.5 TB	1 TB 2 TB 3 TB
<b>Height</b>	2U	2U	2U	1U	2U	4U
<b>Connectivity (standard)</b>	N/A	1 Gb/s iSCSI	10 Gb/s iSCSI	25/10 Gb/s iSCSI or NVMe/TCP	10 Gb/s iSCSI	N/A
<b>Connectivity (optional)</b>	SV3 only: 64 Gb/s FC or NVMe/FC SV3 and SA2: 32 Gb/s FC or NVMe/FC 100 Gb/s iSCSI or NVMe /TCP 25/10 Gb/s iSCSI or NVMe /TCP	16 Gb/s FC 10 Gb/s iSCSI 12 Gb/s SAS	16 Gb/s FC 10 Gb/s iSCSI 12 Gb/s SAS	64 Gb/s FC or NVMe/FC 32 Gb/s FC or NVMe/FC 10 Gb/s iSCSI or NVMe /TCP	32 Gb/s FC or NVMe/FC 25/10 Gb/s iSCSI or NVMe /TCP	64 Gb/s FC or NVMe/FC 32 Gb/s FC or NVMe/FC 100 Gb/s iSCSI or NVMe /TCP 25/10 Gb/s iSCSI or NVMe /TCP
<b>Max ports</b>	12	8	8	16	24	48
<b>Max IOPS (4K read miss)</b>	SA2 1 M SV3 2.5 M	140k	400k	700k	1M	2.5M
<b>Max bandwidth (256K read miss)</b>	SA2 45 GB/s SV3 100 GB/s	8 GB/s	12 GB/s	28 GB/s	45 GB/s	100 GB/s
<b>Warranty and Support</b>	2145 Enterprise Class Support and a one-year warranty. 2147 Enterprise Class Support and a three-year warranty.	One year 9x5 standard 1–5 Expert Care Basic, Advanced, or Premium	One year 9x5 standard 1–5 Expert Care Basic, Advanced, or Premium	One year 9x5 standard 1–5 Expert Care Basic, Advanced, or Premium	One year 9x5 standard 1–5 Expert Care Basic, Advanced, or Premium	One year 24x7 standard 1–5 Expert Care Advanced or Premium

	SVC	5015	5045	5300	7300	9500
<b>Dimensions</b>	Control Enclosure: Height: 8.7 cm (3.4 in.) Width 44.6 cm (17.6 in.) Depth 82.6 cm (32.5 in.)	Control enclosure: Height: 8.7 cm (3.4 in.) Width: 48.3 cm (19.0 in.) Depth: 55.6 cm (21.9 in.)	Control enclosure Height: 8.7 cm (3.4 in.) Width: 48.3 cm (19.0 in.) Depth: 55.6 cm (21.9 in.)	Control enclosure Height: 4.3 cm (1.7 in.) Width: 44.6 cm (17.5 in.) Depth: 77 cm (30.3 in.)	Control enclosure Height: 8.8 cm (3.5 in.) Width: 48.3 cm (19.0 in.) Depth: 85 cm (33.5 in.)	Control enclosure Height: 17.43 cm (6.8 in.) Width: 44.6 cm (17.6 in.) Depth: 82.6 cm (32.6 in.)
<b>Weight</b>	25 kg (55 lb) to 30 kg (65 lb) depending on configuration	Fully configured (12 drives): 28.3 Kg Fully configured (24 drives): 27.3Kg	Fully configured (12 drives): 28.3 Kg Fully configured (24 drives): 27.3 Kg	Fully configured (12 drives): 19.5 Kg	Fully configured (24 drives): 46.6 Kg	Fully configured (48 drives): 70.5 Kg

## 1.3 Changes and enhancements in IBM Storage Virtualize V8.7.0

IBM Storage Virtualize V8.7.0 provides more features and updates to the IBM Storage Virtualize family of products, which include IBM FlashSystems and the IBM SAN Volume Controller.

### 1.3.1 File system awareness for ransomware detection

Volumes can be used by many different applications, operating systems and file systems. This might pose a challenge for AI-powered ransomware detection. Knowing the specific file system that is used on each volume can improve I/O pattern analysis, a key element in ransomware detection.

However, storage administrators often lack complete visibility into volume usage. Applications and different teams might employ volumes for diverse purposes, often resorting to cryptic volume names that don't reflect the actual use case. This lack of clear information hinders efficient storage management and ransomware detection strategies. IBM Support can also benefit from understanding which file systems are in each volume in some recovery scenarios.

IBM Storage Virtualize V8.7.0 provides the following file level awareness for ransomware detection:

1. Every 12 hours, the file system is automatically updated for each volume. The file system can also be updated by `analyzevdisk` or `analyzevdiskbysystem` CLI commands.
2. Background reads are sent to a volume.
3. Open-source libraries are used to determine file system.
4. Output is displayed in the `file_system` field of the `lsvdiskanalysis` command:
  - 15 character max for field
  - Can display multiple file systems

5. File system used by inferencing engine to improve ransomware detection.

### Enhancement to IBM Storage Insights for threat detection

A key part of monitoring your system includes the detection of potential ransomware attacks. To ensure that you have the latest storage metadata for detecting those types of attacks, compression and cyber resiliency statistics for volumes are collected every 5 minutes.

With these statistics, IBM Storage Insights builds a historical model of a storage system and uses its built-in intelligence and formulas to identify when and where ransomware attacks might be occurring. For more information about statistics, see [IBM Storage Insights](#).

**Note:** For more information, see the blog post [IBM Storage Virtualize 8.7.0 including Flash Grid](#) by *Barry Whyte and Andrew Martin*, which provides a good overview of the new features in Version 8.7.

### 1.3.2 GUI updates

IBM Storage Virtualize V8.7.0 provides these three areas of GUI enhancements:

#### Async policy-based replication and partition base HA user experience improvements

Figure 1-2 shows the async policy-based replication and partition base HA user experience improvements.

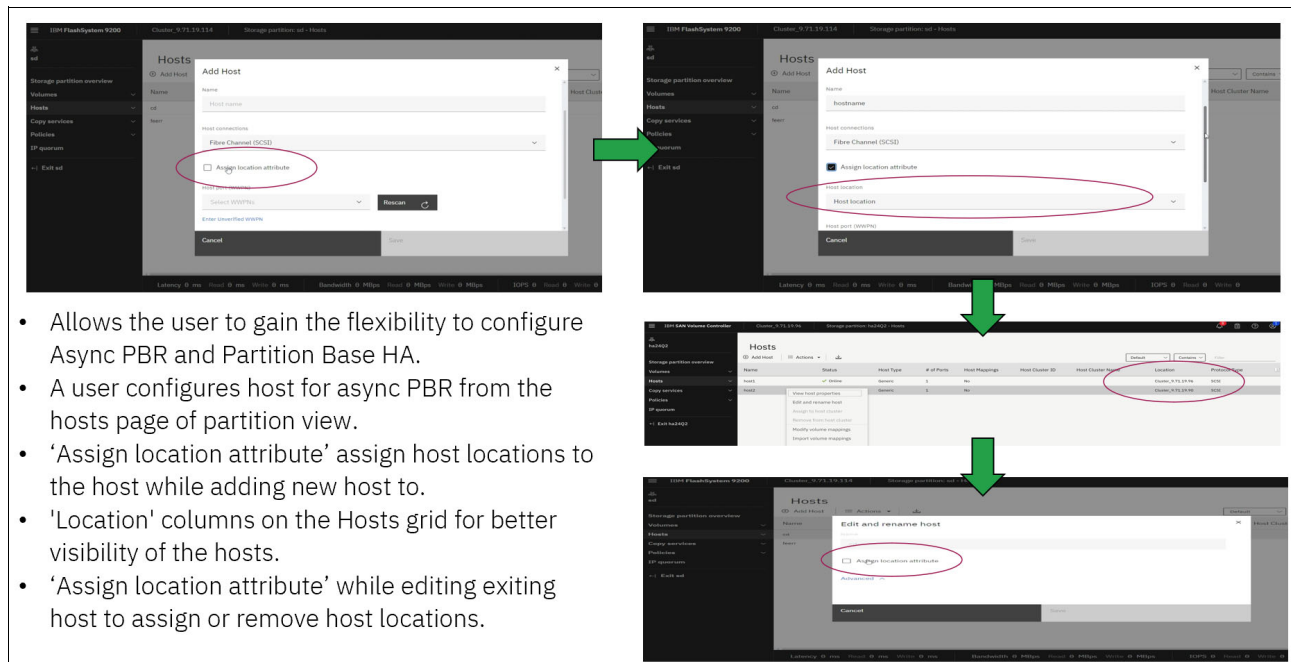


Figure 1-2 Async policy-based replication and partition base HA user experience improvements

## Volume group tile and assigning ownership groups to volume groups

Figure 1-3 shows the enhancement of a volume group tile and assigning ownership groups to volume groups.

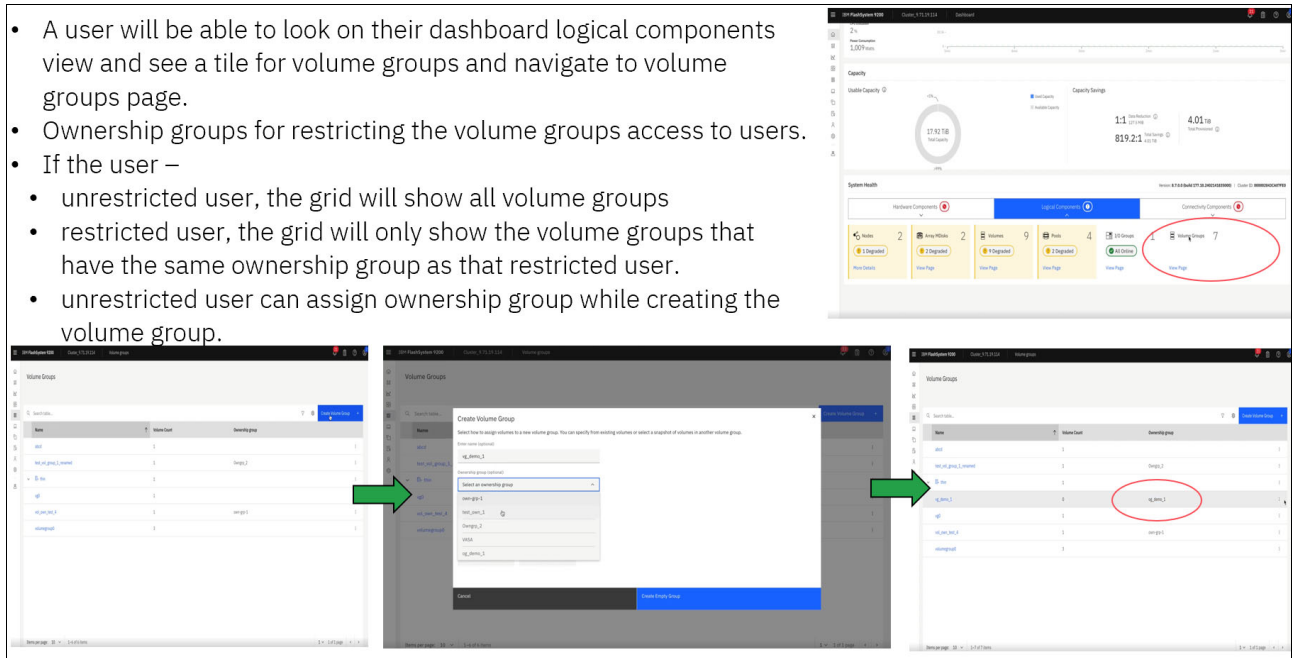


Figure 1-3 Volume group tile and assigning ownership groups to volume groups

## Performance panel carbonization

Storage Virtualize introduces GUI performance panel modernization by using Carbon v11 components and Carbon Charts:

- Tabs-based implementation for future scalability of the charts to incorporate a growing number of statistics
- Charts by user customization across the user sessions
- Responsive flexible layout of 1 column, 2 column or 3 columns
- Improved usability and accessibility through drag and drop feature
- New option for power and temperature charts to display statistics at systems, node, or enclosure level
- Ability to compare the IOPS, bandwidth, and latency through a single view
- *Restore to default* feature to remove all user customization and restore default view

Figure 1-4 shows the newly designed GUI performance panel.

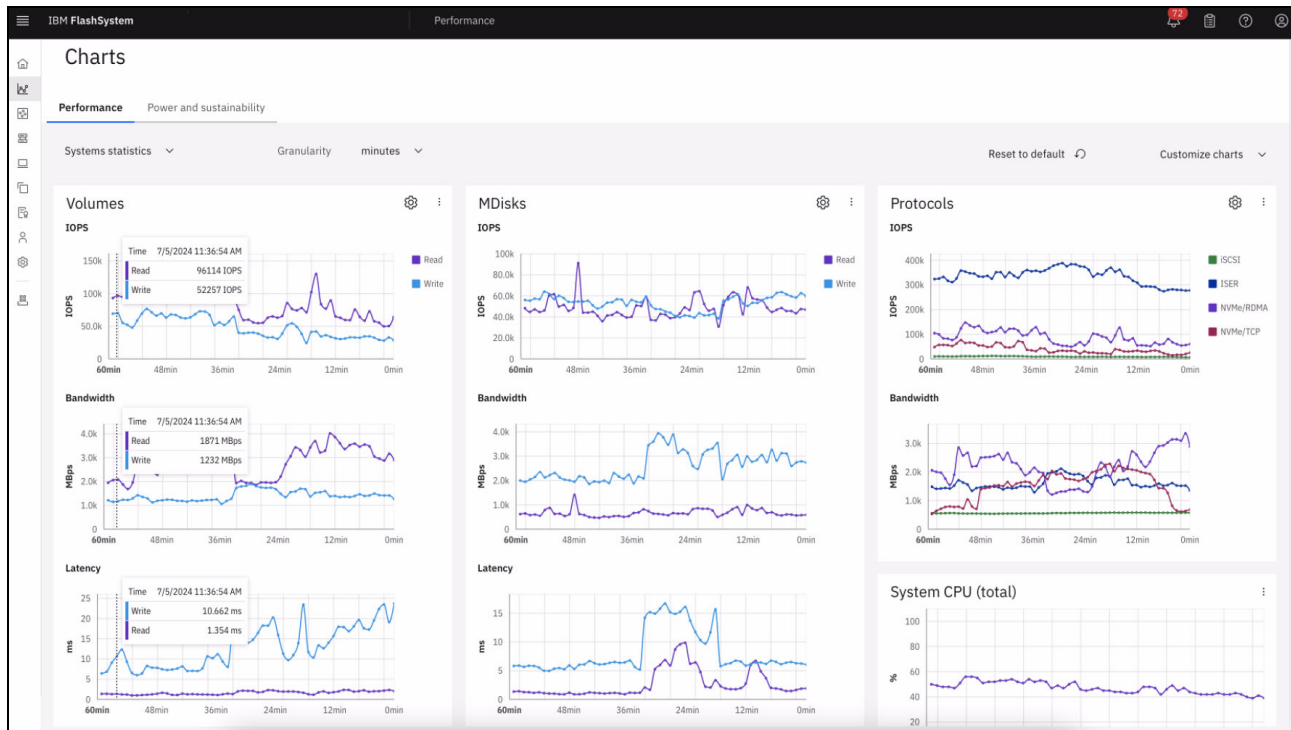


Figure 1-4 GUI performance panel

### 1.3.3 Flash Grids and storage partition migration

A *Flash Grid* is a collection of FlashSystems or SVC systems that you can administer as a single storage solution. You can manage it from a single browser window and use it to run nondisruptive migration.

The current I/O group structure presents several limitations that hinder performance scalability and flexibility:

- ▶ Limited scalability. A maximum of four I/O groups restricts the overall performance achievable by a single system.
- ▶ Hardware compatibility challenges. Compatibility requirements between I/O groups complicate hardware upgrades.
- ▶ Disruptive upgrades. System-wide upgrades are needed for both software and hardware, leading to downtime and complexity.
- ▶ Nonlinear object limits. Volume, snapshot, and host counts are limited per system, not per I/O group, hindering scalability.
- ▶ Feature restrictions. Several advanced features, such as policy-based HA, vVol replication, and storage partitions, are only available on single I/O group systems.

Flash Grid addresses these limitations by offering a more granular and flexible approach to storage management. Storage Virtualize 8.7.0 is the first phase of this implementation and includes the following key features:

- ▶ CLI-driven Flash Grid management. Initial configuration and management are primarily done through the command-line interface (CLI). The CLI uses AI-assisted storage partition

migration, which uses artificial intelligence to optimize the migration of storage partitions across Flash Grid member systems.

- ▶ Scalability for performance and flexibility. Flash Grid allows grouping up to eight systems, providing scalability for I/O and storage capacity.
- ▶ Broad hardware compatibility. All NVMe-based FlashSystem products and SVC models are supported within a Flash Grid, promoting hardware choice and future upgrades.
- ▶ Centralized management and ownership. A single system, which is designated as the Flash Grid owner, manages the membership of other systems within the grid.
- ▶ Tiering for optimized performance. By using Flash Grid, you can combine different hardware within a single grid, and you can tier your storage based on performance and capacity needs.
- ▶ Independent software updates. Each member system in a Flash Grid can receive software updates independently, which offers greater flexibility and reduced downtime for maintenance.

**Note:** IBM plans to include the Flash Grid implementation, monitoring, and management in the GUI in a future release. There is also a plan to more closely integrate with IBM Storage Insights to give AI capable operations to storage partitions migration across systems in the Flash Grid.

### 1.3.4 Automatic Patch Updating and Automatic Drive Downloading

Storage Virtualize V8.7 helps to simplify patching with two new features.

#### Automatic Patch Updating

*Automatic Patch Updating* is a mechanism that enables security, or other patches, to be scheduled to automatically update on a user's system as new patch versions are published.

A patch is a small update to a function or service that can be installed on a user's system. A patch install never requires a node reboot or reset.

**Note:** A patch installation might restart a Linux service when installed. It can be installed on all platform types and is small in size.

A process for creating and publishing patches is already in place on Storage Virtualize 8.6.0. When developers identify an issue, they create patches such as bug fixes and security updates to address issues in Storage Virtualize. Patches are published on [IBM FixCentral](#). IBM Cloud® Call Home is used to access patches. Newer versions of IBM Storage Virtualize code can include older patches that were released in previous Storage Virtualize versions.

The Automatic Patch Updating has the following benefits to the clients' systems:

- ▶ You can use enhanced Patching framework to schedule automatic patch updates for your Storage Virtualize systems. This eliminates the need for complex full Program Temporary Fix (PTF) or concurrent code upgrades and can save you time and effort.
- ▶ Benefits users whose systems have patches that might need frequent updates.

**Note:** An example might be Ransomware Threat Detection, where the inference data files might be regularly changed.



- ▶ Users can configure their systems and know that updates of vital patches happen in the background.
- ▶ Automatic Patch Updating can be configured on a user's system by using either the GUI or CLI commands.

After it is configured, automatic patching performs daily checks on IBM Fix Central. If any selected patches are available for download, they are automatically downloaded and applied to your system.

**Important:** Automatic Patch Updating uses IBM Cloud Call Home to access patch information and lists. Therefore, a functioning IBM Cloud Call Home is a prerequisite before you configure automatic updates.

## Automatic Drive Firmware Download

*Automatic Drive Firmware Download* is a mechanism that uses the new patch infrastructure to enable drive firmwares to be stored on a cluster and ensure that a standard drive firmware level is maintained.

**Note:** This change is for FCM drives only.

The Automatic Drive Firmware Download process is as follows:

- ▶ Building on the patching infrastructure, FCM drive firmware updates are now built as patches.
- ▶ You can use the CLI command **applysoftware <firmware\_patch>** to copy the firmware patch to all cluster nodes, which can streamline the update process.
- ▶ When a new drive is added to the cluster, it automatically checks for and applies the latest available firmware, helping ensure your storage remains up to date.
- ▶ Firmware patches are at IBM FixCentral with other software updates for Storage Virtualize.
- ▶ Newer versions of Storage Virtualize can include older firmware patches, which eliminates the need to search for and apply them individually.

Because of Automatic Drive Firmware Download, any FCM Field Replaceable Unit (FRU) replacements or additional drives that are added to your array are automatically updated, so they are compatible with your system's firmware.

## Example scenario

Consider the following scenario:

1. A user has an FCM4 array that uses firmware version 4.1.4. A drive fails and requires replacement.
2. A FRU arrives with version 4.0.4. The user performs the Dynamic Drive Pool operation and replaces the failed drive with the replacement FRU.
3. As the drive attempts to rejoin the array, Automatic Drive Download automatically verifies that version 4.1.4 is available and upgrades the drive.

### 1.3.5 Management IP changes

IBM Storage Virtualize 8.7 gives more flexibility in assigning and managing the system management ports on the IBM FlashSystem and SVC machines.

The following changes are included:

- ▶ Ability to configure system IP address on any port and with VLAN
- ▶ An increase to 4 routable data IP addresses per port per node
- ▶ System defined default management portset for system IP addresses
- ▶ Limiting the number of system IP addresses to 2
- ▶ System IP address on any port other than default port 1 and 2
- ▶ VLAN support.
- ▶ Changes in GUI for system IP addresses:
  - Management IP address panel is changed.
  - Adding new management IP addresses.

#### Reclaim space of thin-provisioned volumes in standard pools

This feature supplies a new mechanism to automatically reclaim thin-provisioned volumes in standard pools. It automatically recovers space in standard pools after hosts submit unmap I/O (or overwrite with zeros). This feature was developed in response to feedback from customers who previously relied on manual processes and close monitoring.

### 1.3.6 Remote Copy support on IBM Storage Virtualize V8.7

IBM Storage Virtualize V8.7.0 is the final version that supports all Remote Copy based features, including IBM HyperSwap®, Metro Mirror, Global Mirror, and Global Mirror with Change Volumes. When you plan new deployments, consider using policy-based replication and policy-based high availability to avoid causing an unnecessary migration.

**Note:** Remote Copy is supported on V8.7.0 if the hardware has a valid support contract. This also includes the following functions:

- ▶ Global Mirror
- ▶ Global Mirror with Change Volumes
- ▶ HyperSwap
- ▶ Metro Mirror
- ▶ Migration relationships
- ▶ HyperSwap and Metro Mirror 3-site solutions

**Important:** Entry-level IBM FlashSystem 5015 and 5035 do have replication capabilities if upgraded beyond V8.7.0.

Storage Virtualize V8.7.0 and later includes the following changes to remote copy support:

- ▶ Global Mirror and Global Mirror with Change Volumes are replaced by policy-based replication. For more information, see [Migrating to Safeguarded snapshots](#).
- ▶ HyperSwap is replaced by policy-based HA.
- ▶ Migration is using storage partition migration.

# 1.4 Preparation and upgrading to IBM Storage Virtualize V8.7.0

To run IBM Storage Virtualize V8.7.0 on your selected hardware, there are some tasks and checks that need to be done before implementing this level of IBM Storage Virtualize software.

## 1.4.1 Confirm hardware support

Confirm that your current hardware is able to support IBM Storage Virtualize V8.7.0.

Figure 1-5 shows the matrix of supported hardware versus the IBM Storage Virtualize software levels.

The “from” level is your current IBM Storage Virtualize software level and the “to” level is IBM Storage Virtualize 8.7.0.

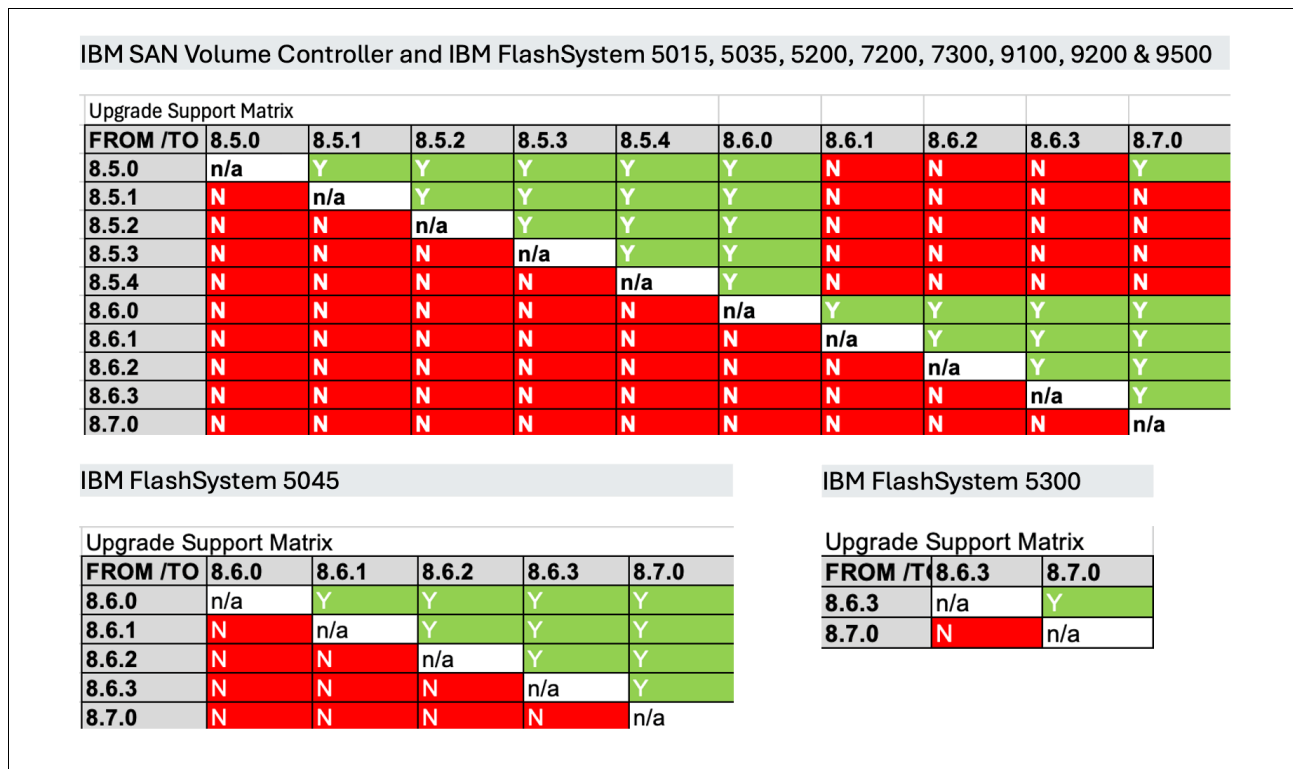


Figure 1-5 IBM Storage Virtualize upgrade support matrix

Examine the matrix in Figure 1-5 and confirm your IBM Storage Virtualize hardware can upgrade to the IBM Storage Virtualize 8.7.0 level.

## 1.4.2 Correct any issues as needed

Verify that any outstanding issues or errors are corrected or fixed before the upgrade. If required, place a service call with IBM for assistance with getting these resolved.

Also, there are limits with some features in IBM Storage Virtualize 8.7.0 that might be applicable to your configuration. Ensure that you understand these limits before upgrading to

IBM Storage Virtualize 8.7.0. For more information, see [V8.7.0.x Configuration Limits for IBM FlashSystem and SAN Volume Controller](#).

### 1.4.3 Schedule the upgrade

Schedule time to upgrade to IBM Storage Virtualize V8.7.0 software. The upgrade is concurrent. However, it is recommended to upgrade the systems during times of low I/O activity. During an upgrade each node goes offline and the partner node in the I/O group handles all the I/O activity.

For the specific steps of the upgrade process including the prechecks and code download, see [Software update](#).

Also, verify that the drive firmware on your IBM FlashSystem is at the latest level. Updates to drive firmware are not done during system updates. For more information, see [Drive update](#).



# Initial configuration

This chapter describes the initial configuration of the following systems:

- ▶ IBM FlashSystem 9500
- ▶ IBM FlashSystem 7300
- ▶ IBM FlashSystem 5300
- ▶ IBM FlashSystem 5200
- ▶ IBM FlashSystem 5045
- ▶ IBM FlashSystem 5015
- ▶ IBM SAN Volume Controller

It also provides step-by-step instructions for the initial setup process and defines the baseline system settings. These settings are typically applied during the implementation phase, which is before volume creation and provisioning.

This chapter includes the following topics:

- ▶ 2.1, “Prerequisites” on page 16
- ▶ 2.2, “System initialization” on page 17
- ▶ 2.3, “System setup” on page 24

## 2.1 Prerequisites

Before you initialize and set up the system, ensure that the following prerequisites are met:

- ▶ The physical components fulfill all the requirements and are correctly installed:
  - The FlashSystem control enclosures or IBM SAN Volume Controller nodes are physically installed in the racks.
  - The Ethernet and Fibre Channel (FC) cables are connected.
  - Existing expansion enclosures are physically installed and attached to the control enclosures that use them.
  - The control enclosures or IBM SAN Volume Controller nodes and optional expansion enclosures are powered on.

**Note:** IBM SAN Volume Controller nodes need enough time to charge the batteries. How long it takes to recharge depends on how long it was waiting idle in stock and not in production. You cannot start the nodes without a fully charged battery.

- ▶ The web browser that is used for managing the system is supported by the management GUI. For the list of supported browsers, see [Management GUI](#).
- ▶ The required information for remote management of the system is available:
  - The IPv4 (or IPv6) addresses that are assigned for the system's management interfaces:
    - The unique cluster IP address, which is the address that is used for the management of the system.
    - Unique service IP addresses, which are used to access node service interfaces. You need one address for each IBM SAN Volume Controller node or IBM FlashSystem node (two per control enclosure).
    - The IP subnet mask for each subnet that is used.
    - The IP gateway for each subnet that is used.
  - The licenses that might be required to use specific functions. Whether these licenses are required depends on the hardware that is used. For more information, see [Licensed functions](#).
  - Information that is used by a system when performing Call Home functions:
    - The company name and system installation address.
    - The name, email address, and phone number of the storage administrator whom IBM can contact if necessary.
  - The following information is optional:
    - The Network Time Protocol (NTP) server IP address
    - The Simple Mail Transfer Protocol (SMTP) server IP address, which is necessary if you want to enable Call Home or want to be notified about system events through email
    - The IP addresses for Remote Support Proxy Servers, which are required only if you want to use them with the Remote Support Assistance feature

**Note:** IBM FlashSystem 9500 and IBM SAN Volume Controller are installed by an IBM System Services Representative (IBM SSR). Provide all the necessary information to the IBM SSR by completing the following planning worksheets:

- ▶ [Planning worksheets for IBM FlashSystems](#)
- ▶ [Planning worksheets for IBM SAN Volume Controller](#)

After the IBM SSR completes their portion of the setup, see 2.3, “System setup” on page 24 to continue the setup process.

## 2.2 System initialization

This section provides step-by-step instructions to create the system cluster.

You can view the following demonstration videos. Although the videos are based on IBM Storage Virtualize V8.6, they are still applicable to V8.7.

- ▶ [IBM Storage Virtualize V8.6 Initial setup: SSR configuration tasks](#)
- ▶ [IBM Storage Virtualize V8.6 Initial setup: Customer configuration tasks](#)
- ▶ [IBM Storage Virtualize V8.6 Initial setup: Setting up a cluster from the service IP](#)

### 2.2.1 Before starting initialization

To start the initialization procedure, connect a workstation directly to the *technician port*. The technician port is a dedicated 1 Gb Ethernet (GbE) port that is located on the rear of each storage node canister. On all platforms except IBM FlashSystem 5015, it can be used only for initializing or servicing the system. It cannot be connected to an Ethernet switch because it supports only a direct connection to a workstation, and it remains disconnected after the initial setup is complete.

On IBM FlashSystem 5015, the technician port is enabled initially. However, the port is switched to internet Small Computer Systems Interface (iSCSI) host attachment mode after the setup wizard is complete.

To re-enable an onboard Ethernet port on a system to be used as the technician port, refer to the command shown in Example 2-1.

*Example 2-1 Reenabling the onboard Ethernet port 2 as the technician port*

---

```
IBM_IBM FlashSystem 9100:superuser>satask chserviceip -techport enable -force
```

---

The location of the technician port on an IBM FlashSystem 9500 is shown in Figure 2-1.




Figure 2-1 Technician port FlashSystem 9500

The location of the technician port of an IBM FlashSystem 7300 is shown in Figure 2-2.




Figure 2-2 Technician port FlashSystem 7300

The location of the technician port of an IBM FlashSystem 5200 is shown in Figure 2-3



Figure 2-3 Technician port FlashSystem 5200

The location of the technician port of an IBM FlashSystem 5300 is shown in Figure 2-4



Figure 2-4 Technician port FlashSystem 5300

The location of the technician port of an IBM FlashSystem 5045 is shown in Figure 2-5.



Figure 2-5 Technician port FlashSystem 5045



The location of the technician port of an IBM FlashSystem 5015 is shown in Figure 2-6.




Figure 2-6 Technician port FlashSystem 5015

The location of a technician port on the IBM SAN Volume Controller 2145-SV3 is shown in Figure 2-7.




Figure 2-7 Technician port IBM SAN Volume Controller 2145-SV3

The location of a technician port on the IBM SAN Volume Controller 2145-SV2 is shown in Figure 2-8.



Figure 2-8 Technician port IBM SAN Volume Controller 2145-SV2

The technician port runs an IPv4 DHCP server, and it can assign an address to any device that is connected to this port. Ensure that your workstation Ethernet adapter is configured to use a DHCP client if you want the IP to be assigned automatically.

If you prefer not to use DHCP, you can set a static IP on the Ethernet port from the 192.168.0.x/24 subnet; for example, 192.168.0.2 with the netmask 255.255.255.0.

The default IP address of a technician port on a node canister is 192.168.0.1. Do *not* use this IP address for your workstation.

**Note:** Ensure that the technician port is not connected to the organization's network. No Ethernet switches or hubs are supported on this port.

## 2.2.2 System initialization process

Before initialization, each node canister in a new system remains in the candidate state and cannot process I/O.

During initialization, the nodes within a single control enclosure are joined into a cluster. This cluster is later configured to process data. For an IBM SAN Volume Controller system, the cluster initially consists of only one node.

If your system has multiple control enclosures or IBM SAN Volume Controller nodes, initialize only the first enclosure or node. The remaining enclosures or nodes can be added to the cluster later by using the cluster management interface (GUI or CLI) after the initial setup.

During initialization, you must specify an IPv4 or IPv6 system management address. This address is assigned to Ethernet port 1 on each node and is used to access the management GUI and CLI. You can configure additional IP addresses after the system is initialized.

**Note:** Do not perform the system initialization procedure on more than one node canister of one control enclosure. After initialization is done, use the management GUI or CLI to add control enclosures to the system.

To initialize a new system, complete the following steps:

1. Connect your workstation to a technician port of any canister of the control enclosure or the IBM SAN Volume Controller system. Ensure that you obtained a valid IPv4 address with DHCP.
2. Open a supported web browser and browse to `https://install`. The browser is automatically redirected to the System Initialization wizard. You can also use the IP address `https://192.168.0.1` if you are not automatically redirected.

**Warnings about untrusted certificates:** During system initialization, you might see warnings about untrusted certificates. This happens because the system uses self-signed certificates, which are not verified by a well-known authority.

However, if you are directly connected to the service interface, there is no intermediary that might impersonate the system with a fake certificate. Therefore, you can safely accept the certificates in this scenario.

If the system is not in a state that allows initialization, the system does not start the System Initialization wizard, and you are redirected to the Service Assistant interface. Use the displayed error codes to troubleshoot the problem.

3. If the system is *not* in a state that allows initialization, the window that is used to log in to Service Assistant opens (Figure 2-9). Otherwise, the System Initialization wizard opens (Figure 2-10). Enter the default superuser password of `passw0rd` and click **Log in**.

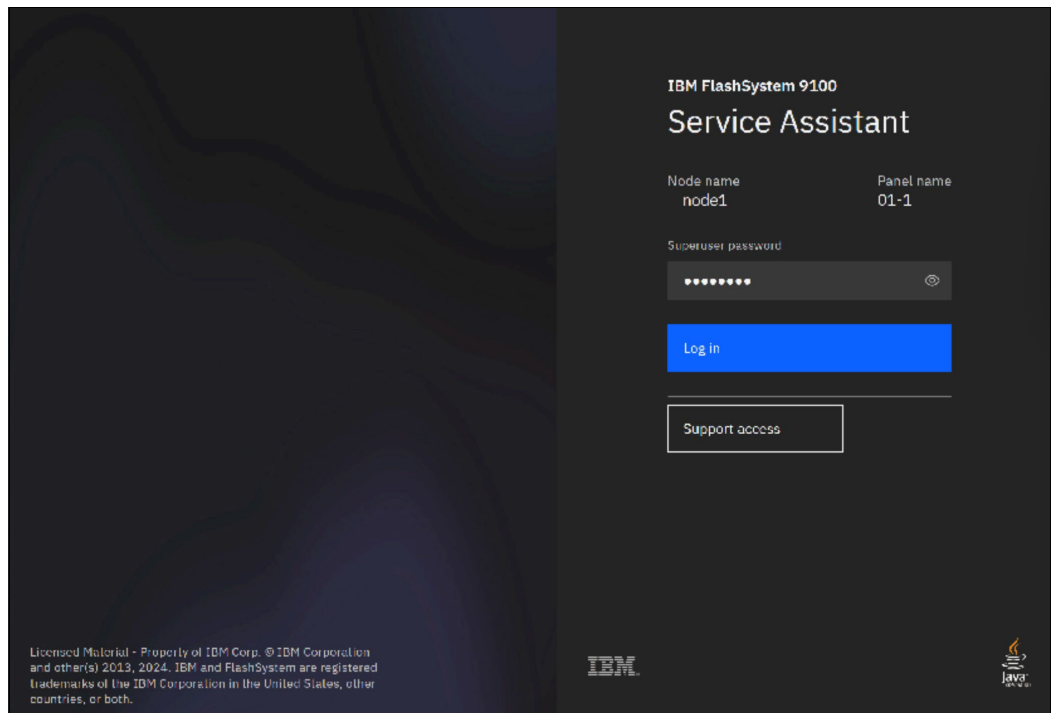


Figure 2-9 Logging in to Service Assistant by way of the technician port

4. The System Initialization wizard shows the detected canisters, as shown in Figure 2-10. Click **Proceed** to continue. This window is not shown for IBM SAN Volume Controller nodes.

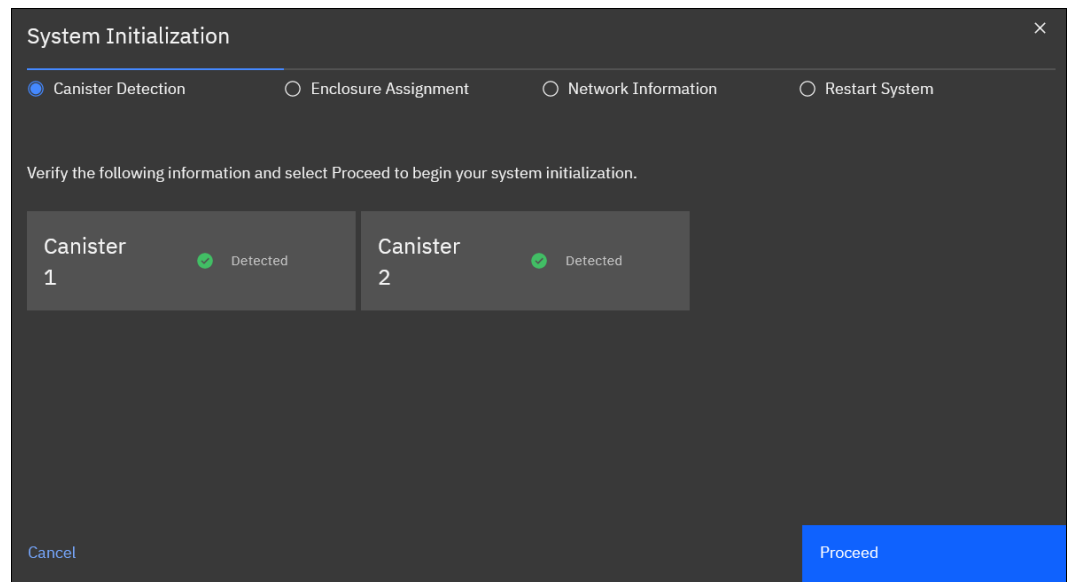


Figure 2-10 System Initialization: Canister detection

5. The System Initialization wizard shows the enclosure assignment. Select **As the first enclosure in a new system**, as shown in Figure 2-11. Click **Next**.




Figure 2-11 System Initialization: Initialize the first enclosure

For IBM SAN Volume Controller systems, the initialization window might differ (see Figure 2-12). You are likely to be prompted to add nodes directly, rather than enclosures.

If you select **As an additional node in an existing system**, you are directed to disconnect from the technician port and use the system's GUI for further configuration.

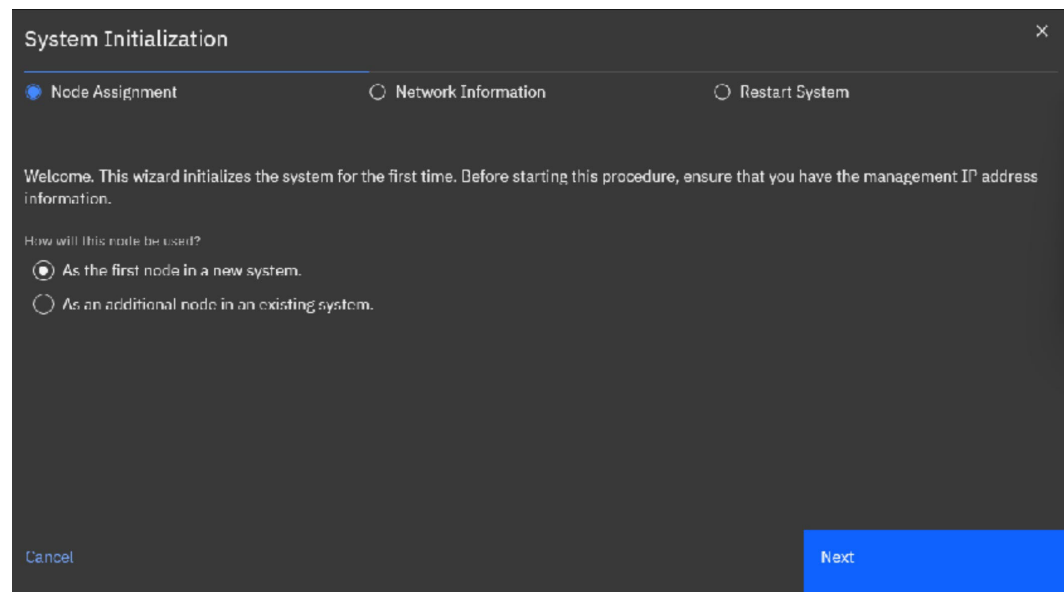


Figure 2-12 System Initialization: Initialize the first IBM SAN Volume Controller node

6. Enter the management IP address information for the new system as shown in Figure 2-13. Set the IP address, network mask, and gateway. Click **Next**.

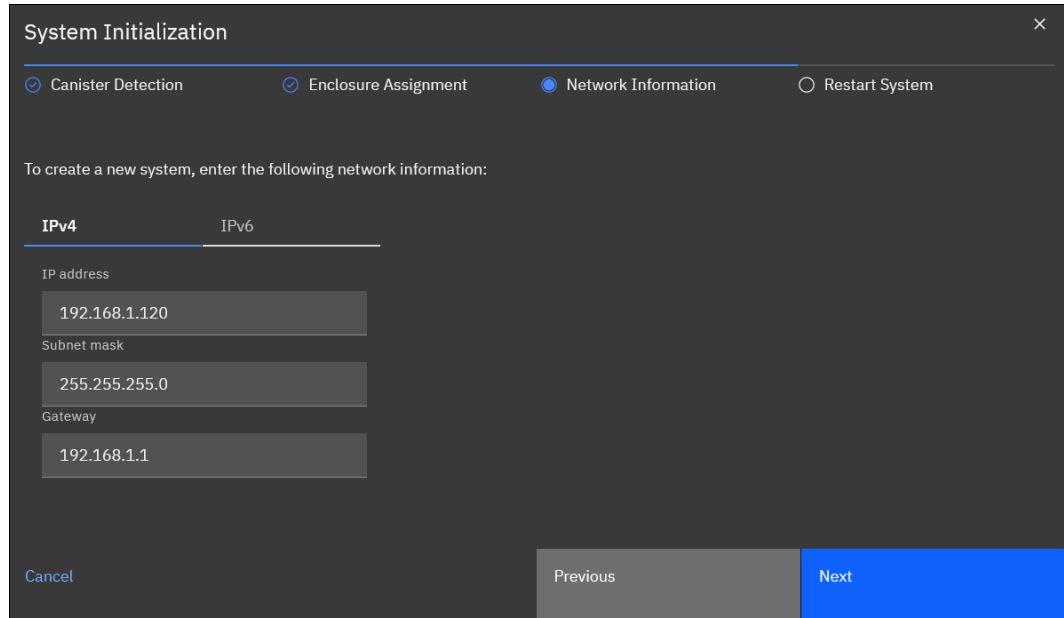


Figure 2-13 System Initialization: Enter Management IP

7. A window that includes a restart timer opens (Figure 2-14). When the timeout is reached, the window is updated to reflect success or failure. Failure occurs if the system is disconnected from the network, which prevents the browser from updating with the IBM FlashSystem web server.




Figure 2-14 System Initialization: Web-server restart timer counting down from 5 minutes

8. The System Initialization completed message is shown in Figure 2-15. Click **Finish**.

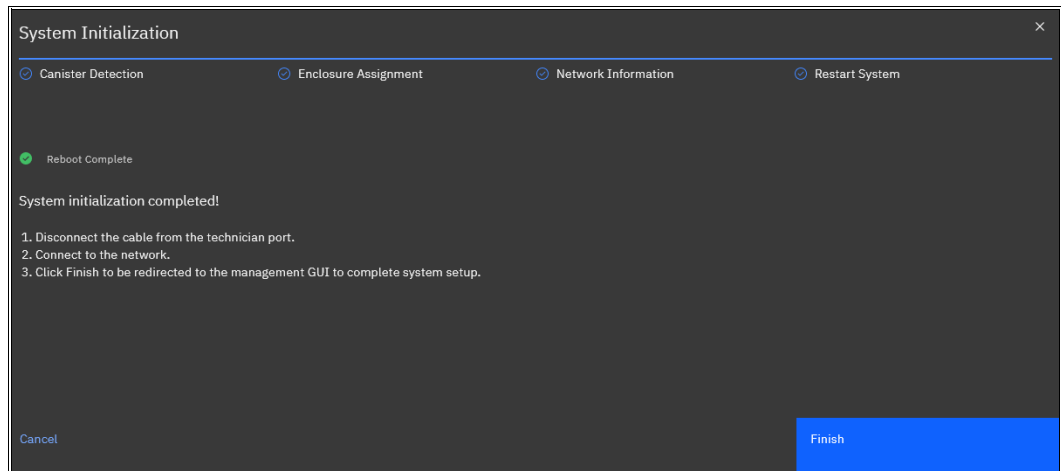


Figure 2-15 System Initialization completed

Follow the instructions, and direct your browser to the management IP address to access the system GUI after you click **Finish**.

System Setup is also available directly from the technician port. The System Setup wizard is available through both the management IP address and the technician port.

## 2.3 System setup

This section provides instructions about how to define the basic settings of the system by using the System Setup wizard.

### 2.3.1 System Setup wizard

Complete the System Setup wizard to define the basic settings of the system. After the initialization is complete, you are redirected to a management GUI from your workstation, or you can browse to the management IP address of a freshly initialized system from another workstation.

The first time that you connect to the management GUI, you can be prompted to accept untrusted certificates because the system certificates are self-signed. If your company policy requests certificates that are signed by a trusted certificate authority (CA), you can install them after you complete the System Setup.

To finish the System Setup wizard, complete the following steps:

1. Log in to system GUI. Until the wizard is complete, you can use only the *superuser* account, as shown in Figure 2-16 on page 25. Click **Sign in**.

**Note:** The default password for the *superuser* account is passw0rd (with the number zero, *not* the uppercase letter O). The default password must be changed by using the System Setup wizard or after the first CLI login. The new password cannot be set to the default password.

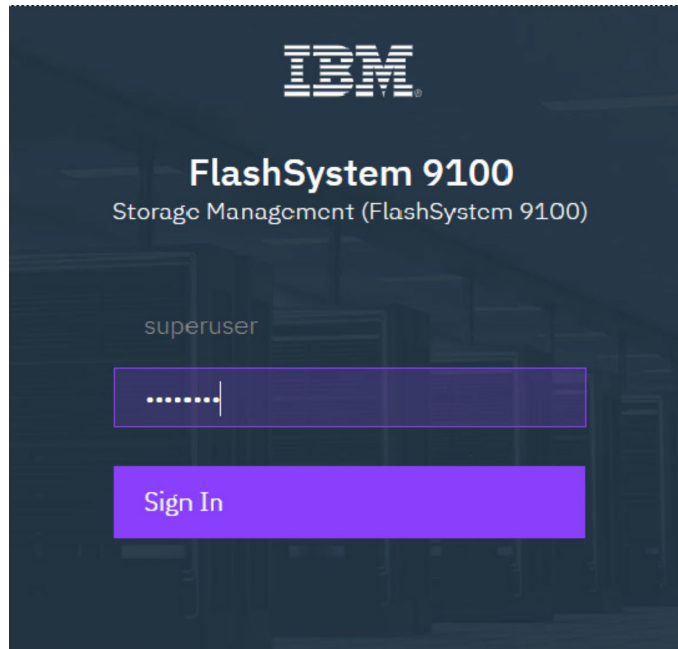


Figure 2-16 Logging in for the first time

2. The Initial Setup starts with the Welcome page, as shown in Figure 2-17. Click **Next**.

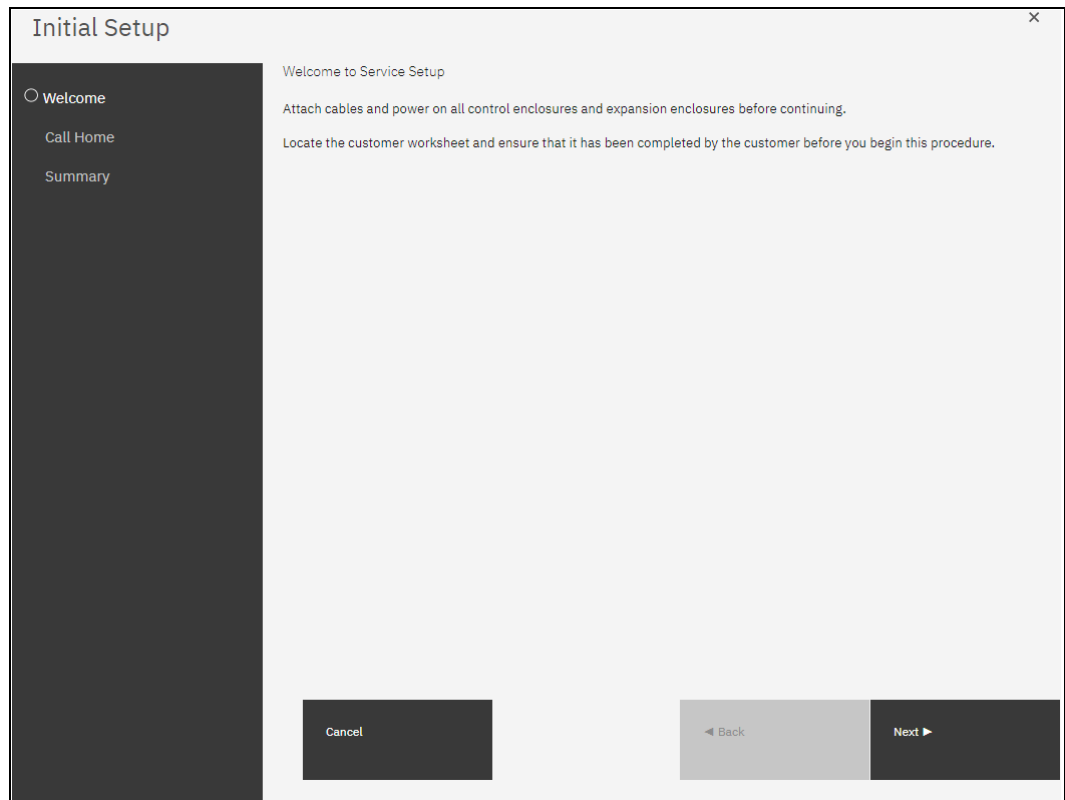


Figure 2-17 Initial Setup Window

3. Configure Call Home. See Figure 2-18. Call Home provides multiple benefits. It enables the automatic creation of support tickets at IBM if errors occur, which can help improve the speed and efficiency of problem resolution. Call Home also enables Storage Insights and Remote Support.

**Note:** For more information, see “Call Home” on page 89.

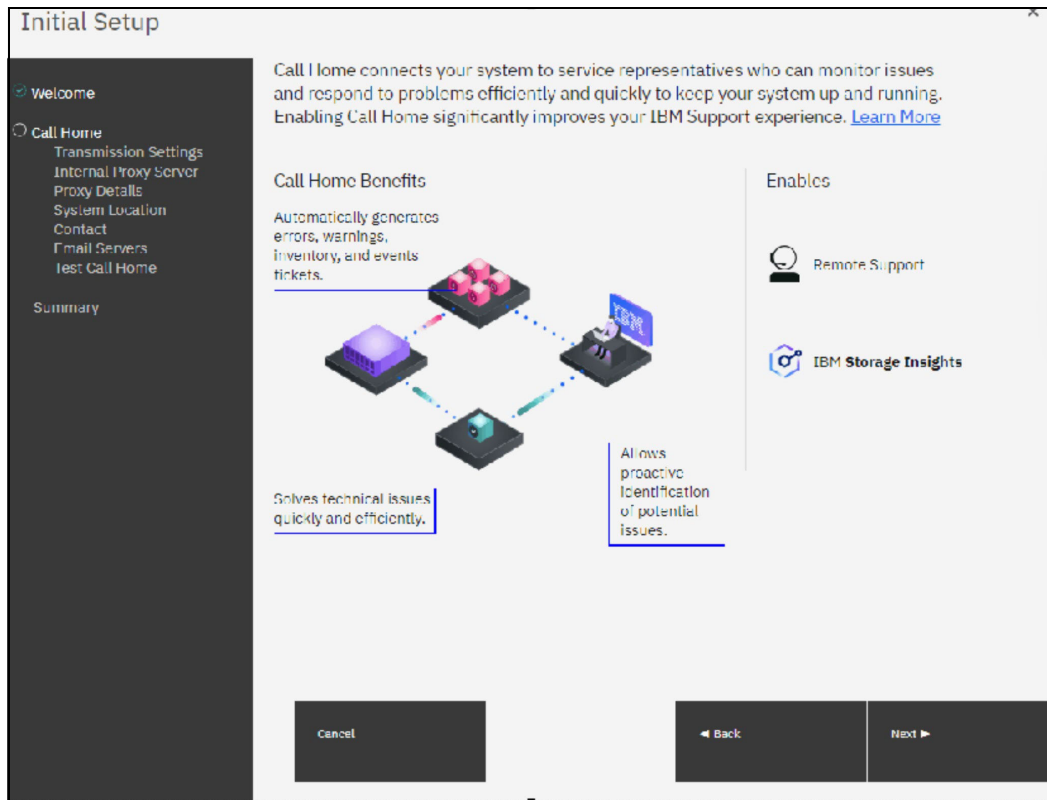


Figure 2-18 Setup Call Home

On IBM FlashSystem 9500 systems and IBM SAN Volume Controller systems, an IBM SSR configures Call Home during installation. Verify that all the entered data is correct.

All IBM FlashSystem products and IBM SAN Volume Controller systems support the following methods of sending Call Home notifications to IBM:

- Cloud Call Home
- Call Home with email notifications

Cloud Call Home is the default and preferred option for a system to report event notifications to IBM Support. With this method, the system uses RESTful application programming interfaces (APIs) to connect to an IBM centralized file repository that contains troubleshooting information that is gathered from customers. This method requires no extra configuration.

The system can also be configured to use email notifications for this purpose. If this method is selected, you are prompted to enter the SMTP server IP address.

If both methods are enabled, Cloud Call Home is used, and the email notifications method is kept as a backup.



If either of these methods is selected, the system location and contact information must be entered. This information is used by IBM to provide technical support. All fields in the form must be completed. In this step, the system also verifies that it can contact the Cloud Call Home servers.

4. Click **Next** to enter the Transmission Type for Call Home.
5. Select which transmission types to use for Call Home. See Figure 2-19.

**Note:** It is *not* recommended to select **I don't want to use Call Home**.

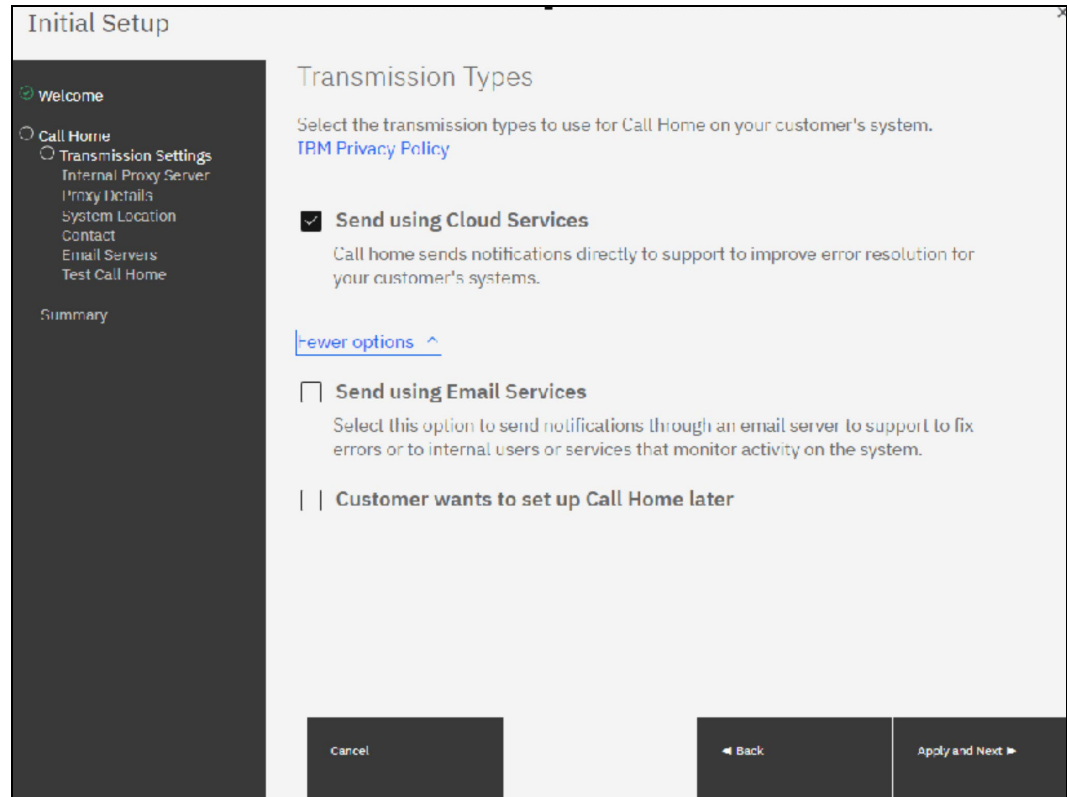


Figure 2-19 Transmission Types for Call Home

6. Select your choice. In the example **Send using Cloud services** is selected. Click **Apply and Next** to setup the Internal Proxy Server. See Figure 2-20 on page 28.
7. Enter the requested information. After you set up the Proxy Server, the system checks the connection to the Support Center. See Figure 2-21 on page 28.

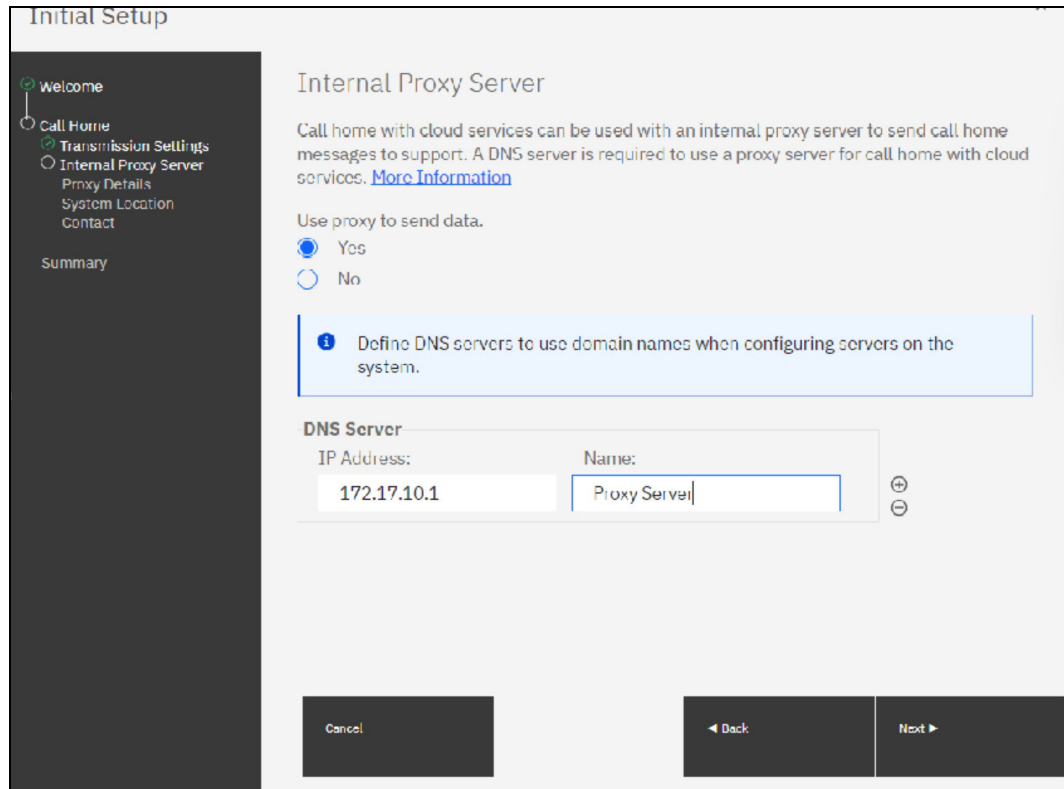


Figure 2-20 Setup Internal Proxy Server

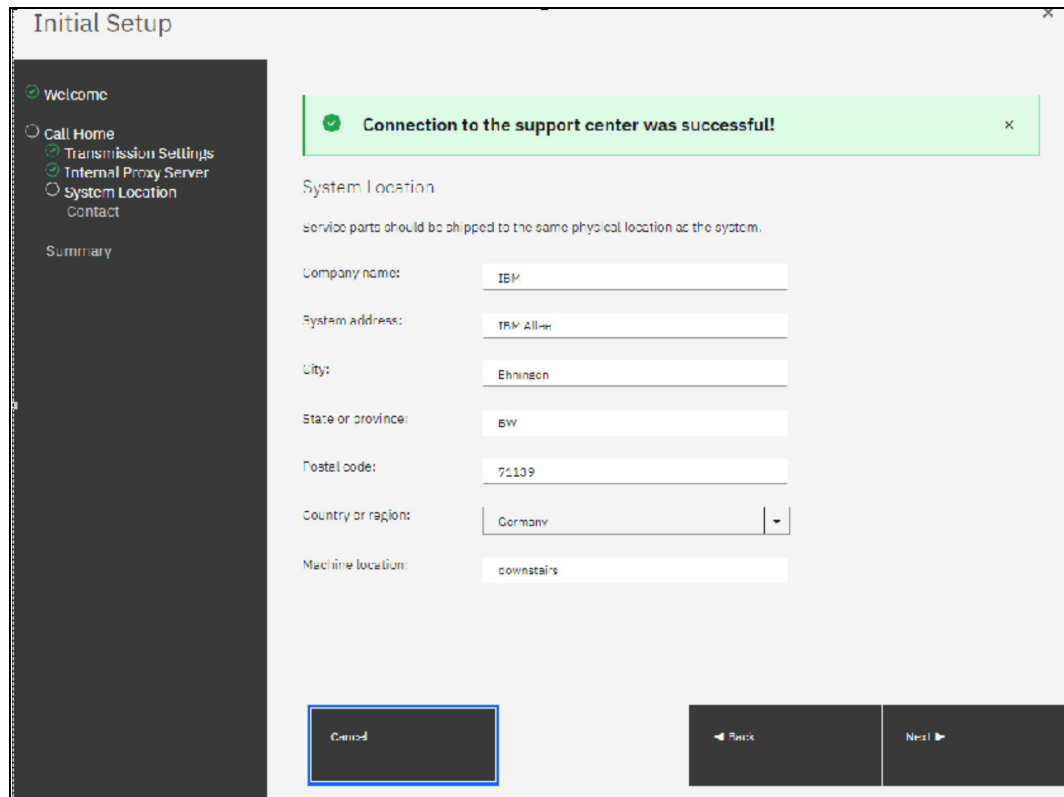


Figure 2-21 Connection Test to the Support Center

8. Enter all required information for the System Location.
9. Click **Next** and enter the Contact information

Figure 2-22 shows the panel for the contact information. Use the Company contact information to comply with privacy regulations. IBM might use the contact data if you allow it.

10. To complete the registration, click **Apply and Next**.

The screenshot shows the 'Initial Setup' window. The left sidebar contains a navigation menu with the following items: 'Welcome' (checked), 'Call Home', 'Transmission Settings' (checked), 'Internal Proxy Server' (checked), 'System Location' (checked), 'Contact' (selected), and 'Summary'. The main content area is titled 'Contact' and includes the following elements: a sub-header 'Enter business-to-business contact information' with a note: 'To comply with privacy regulations, personal contact information for individuals with your organization is not recommended.'; four input fields: 'Name' (Lorzer), 'Email' (lorzcr@dc.ibm.com), 'Phone (primary)' (+49-1234567), and 'Phone (alternate)' (+49-1234568); a toggle switch for 'TDM may use my contact data to keep me informed of Storage related products, services and offerings.' which is currently turned 'Off'; and three buttons at the bottom: 'Cancel', 'Back', and 'Apply and Next'.

Figure 2-22 System Location

11. Review the Summary information. If all is correct, click **Finish**. See Figure 2-23 on page 30.
12. The system saves the entered information, and you are prompted to log in again. After login, you are guided to the System Setup page. See Figure 2-24 on page 30.

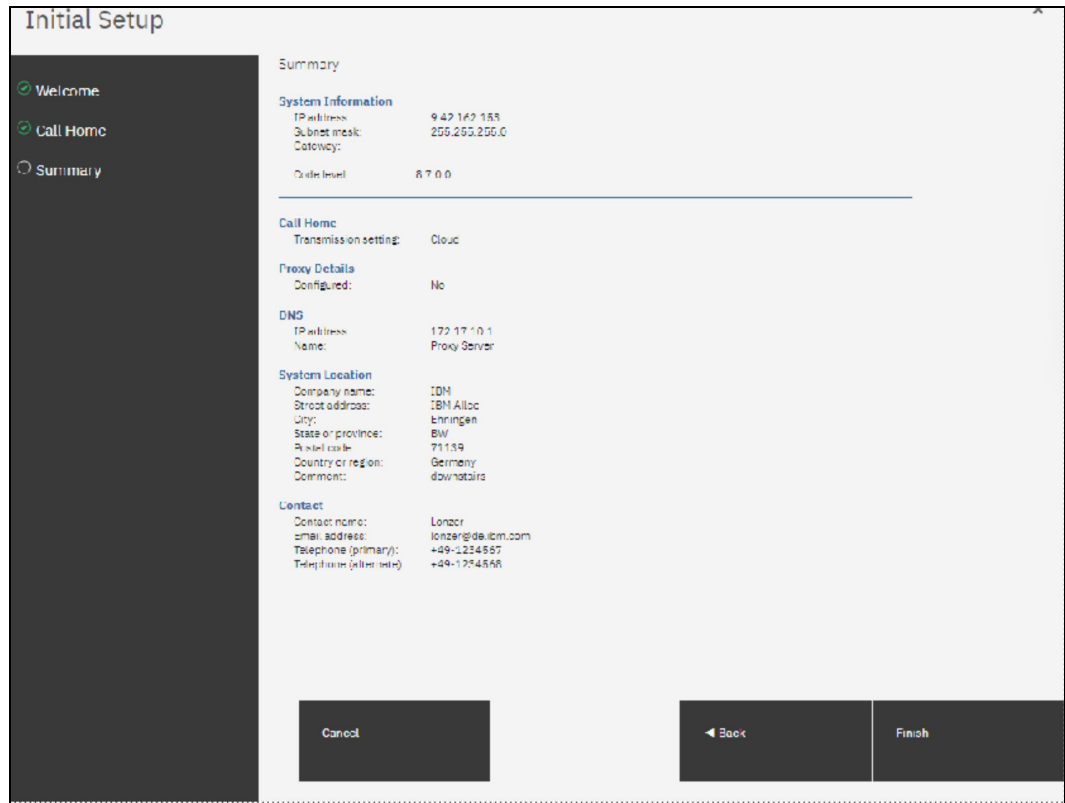


Figure 2-23 Summary page

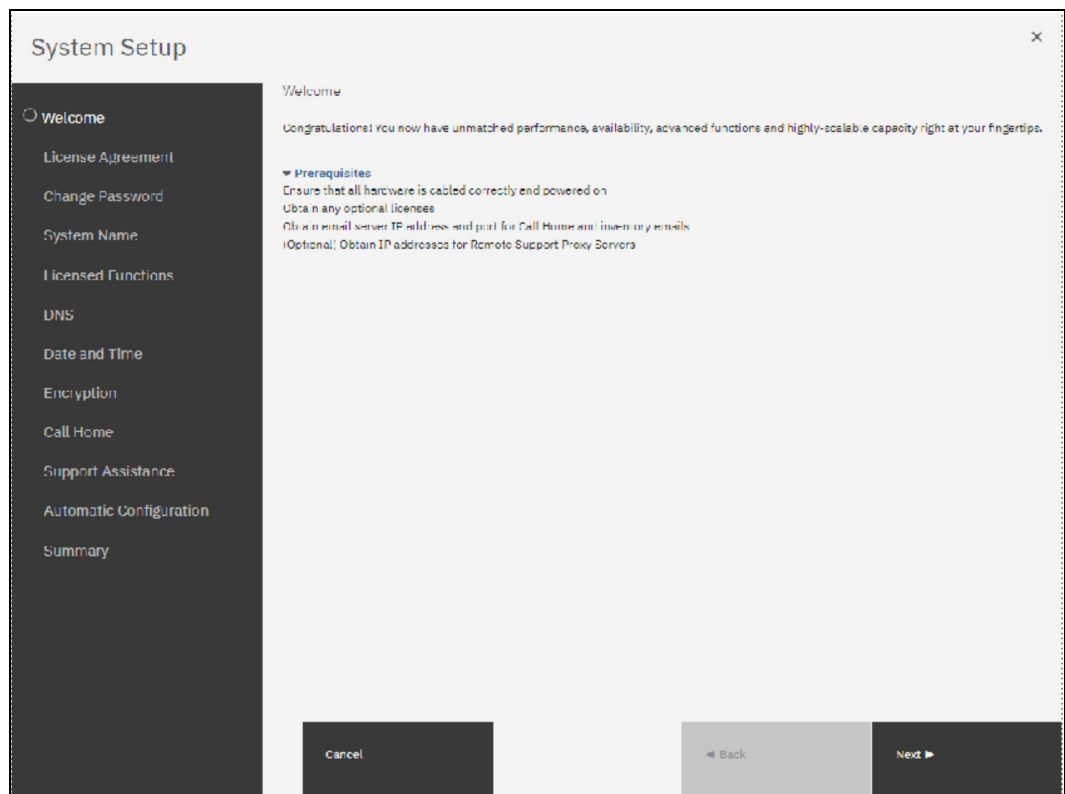


Figure 2-24 System Setup Welcome page

13. Click **Next** to view the License Agreement page. Read the license agreement. Select **I agree with the terms in the license agreement** if you want to continue the setup. Otherwise, the system stops the setup. See Figure 2-25.

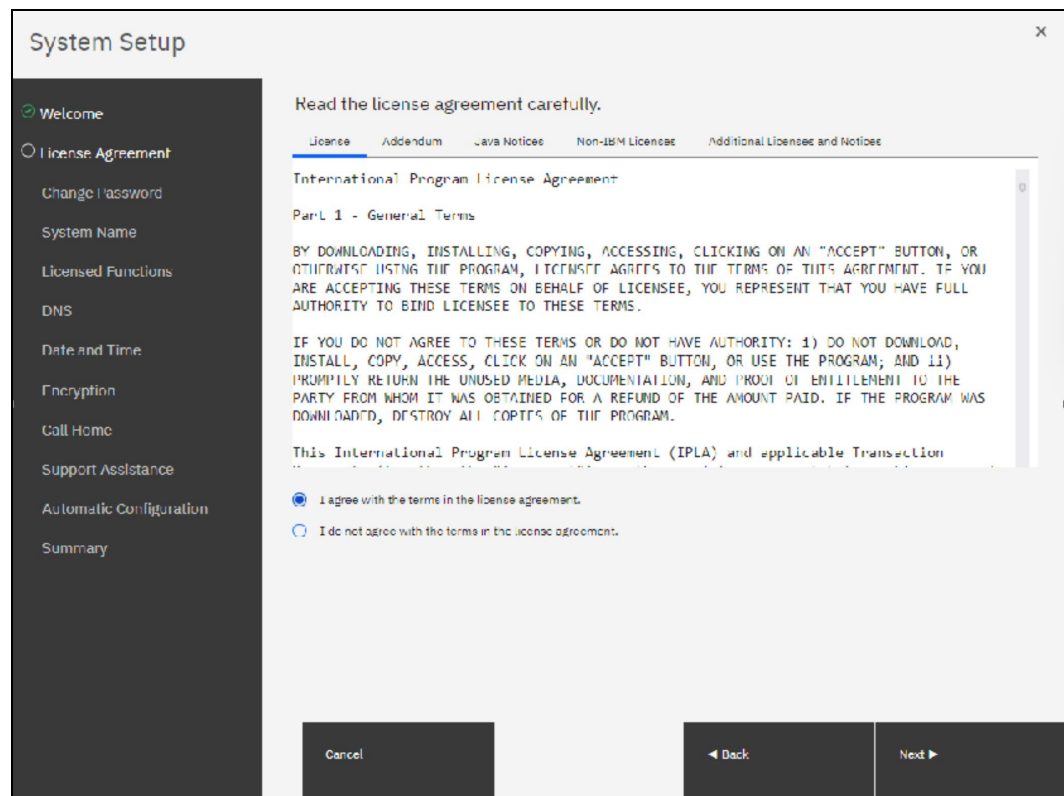


Figure 2-25 Accept License Agreement

14. You are prompted to change the password as shown in Figure 2-26 on page 32. Enter a new password for *superuser*. A valid password is 8–64 characters and cannot begin or end with a space. Also, the password cannot be set to match the default password.

**Note:** All configuration changes that are made by using the System Setup wizard are applied immediately, including the password change. The user sees the system running commands during the System Setup wizard.

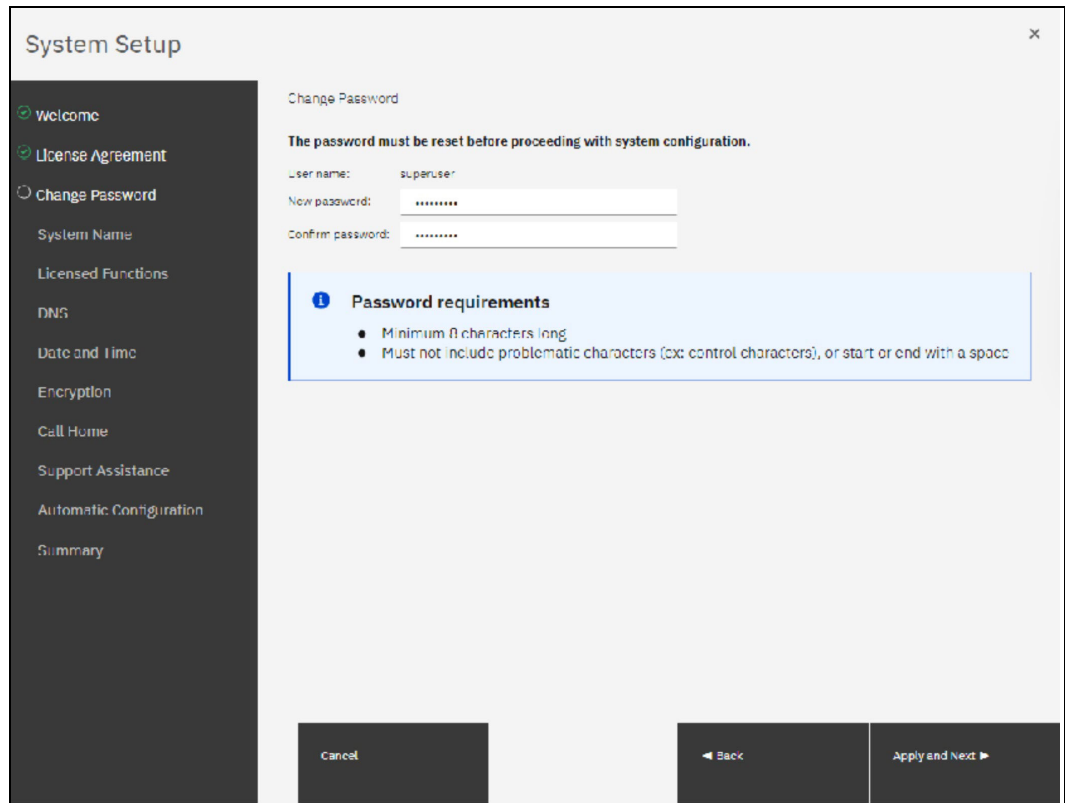


Figure 2-26 Change password

15. To confirm press **Apply and Next**.

16. Provide a System Name as shown in Figure 2-27 on page 33. Avoid the use of an underscore ( **\_** ) in a system name because underscores cannot be used in domain name server (DNS) shortnames and fully qualified domain names (FQDNs). Therefore, such naming might cause confusion and access issues. The following characters can be used: A - Z, a - z, 0 - 9, and - (hyphen).

**Note:** In a 3-Site Replication solution, ensure that the system name is unique for all three clusters when you prepare the IBM Storage Virtualize clusters at Master, AuxNear, and AuxFar sites to work. The system names must remain different for the life of the 3-site configuration.

For more information about 3-Site Replication, see *IBM Spectrum Virtualize 3-Site Replication*, SG24-8504.

If required, the system name can be changed by running the **chsystem -name <new\_system\_name>** command. The system can also be renamed in the management GUI by clicking **Monitoring** → **System Hardware** and selecting **System Actions** → **Rename System**.

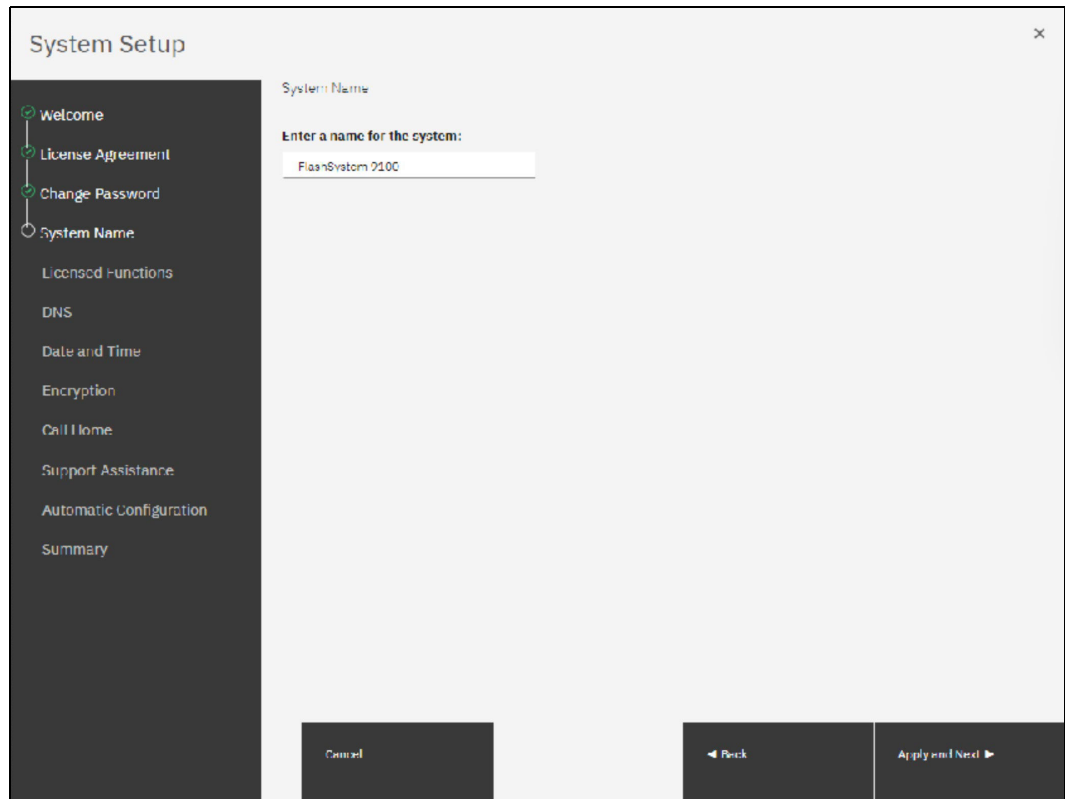


Figure 2-27 System Name

17. Click **Apply and Next**.

18. Provide the System Name and enter the required licenses for each function. Figure 2-28 on page 34 shows you an example.

The window for this step in the system setup might look different, depending on the systems that are used. Also, the way the license is enabled depends on the system that is used

**Note:** Encryption uses a key-based licensing scheme.

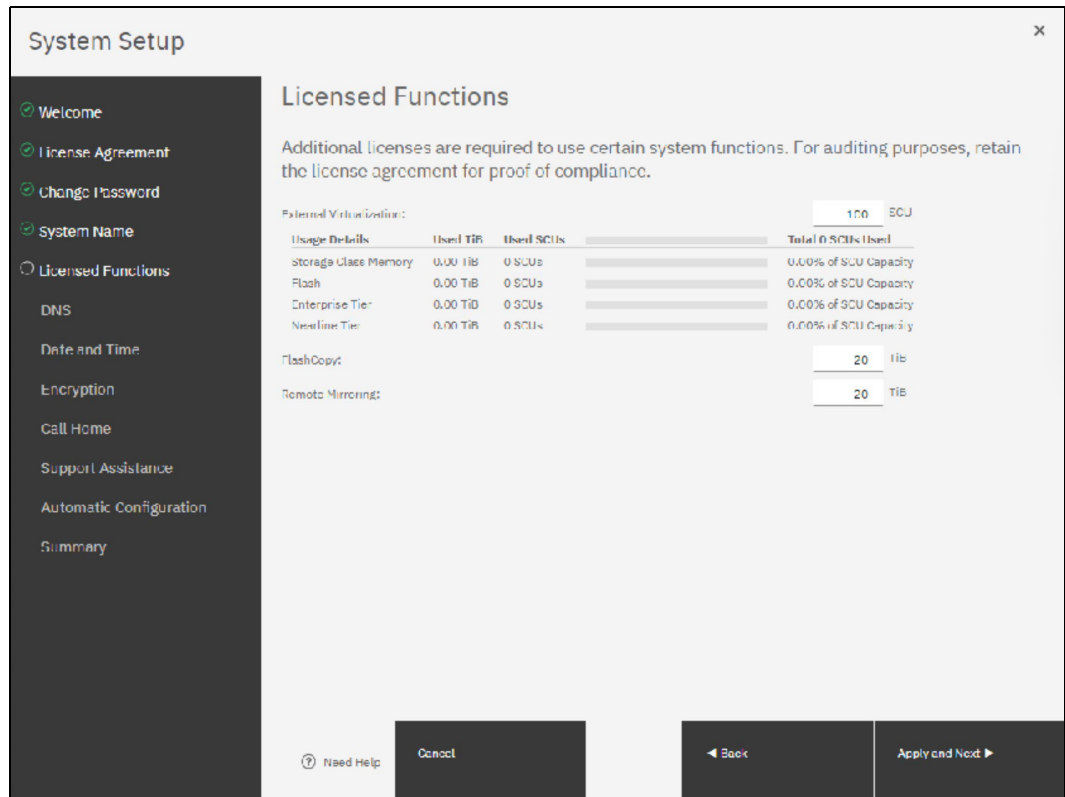


Figure 2-28 License Functions

19. When done, click **Apply and Next**.
20. You can configure DNS on the system as shown in Figure 2-29 on page 35. DNS helps the system to resolve the names of the computer resources that are in the external network if they are not identified by an IP address.
21. When you finish the DNS configuration, click **Apply and Next**.
22. Set the date and time. In the example in Figure 2-30 on page 35, the date and time are set by using manual settings. You can use an NTP server so that all of your storage area network (SAN) and storage devices have a common timestamp. This practice facilitates troubleshooting and prevents time stamp-related errors.
23. When done, click **Apply and Next**.



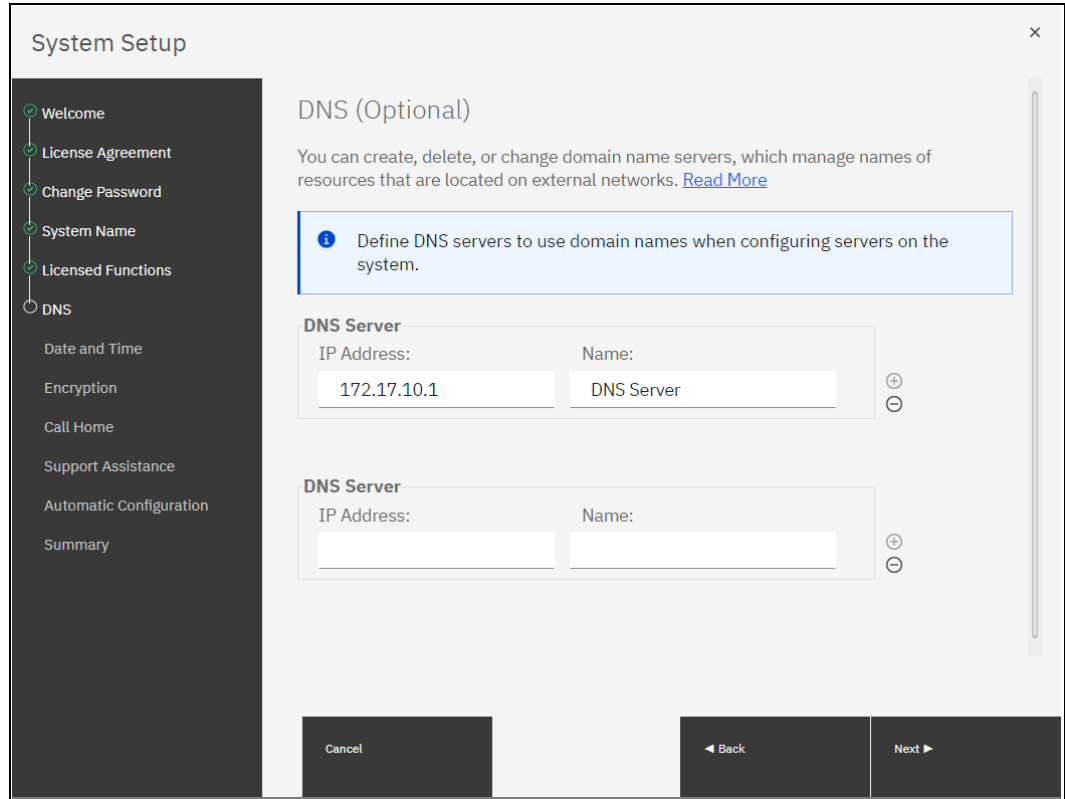


Figure 2-29 DNS Server setup

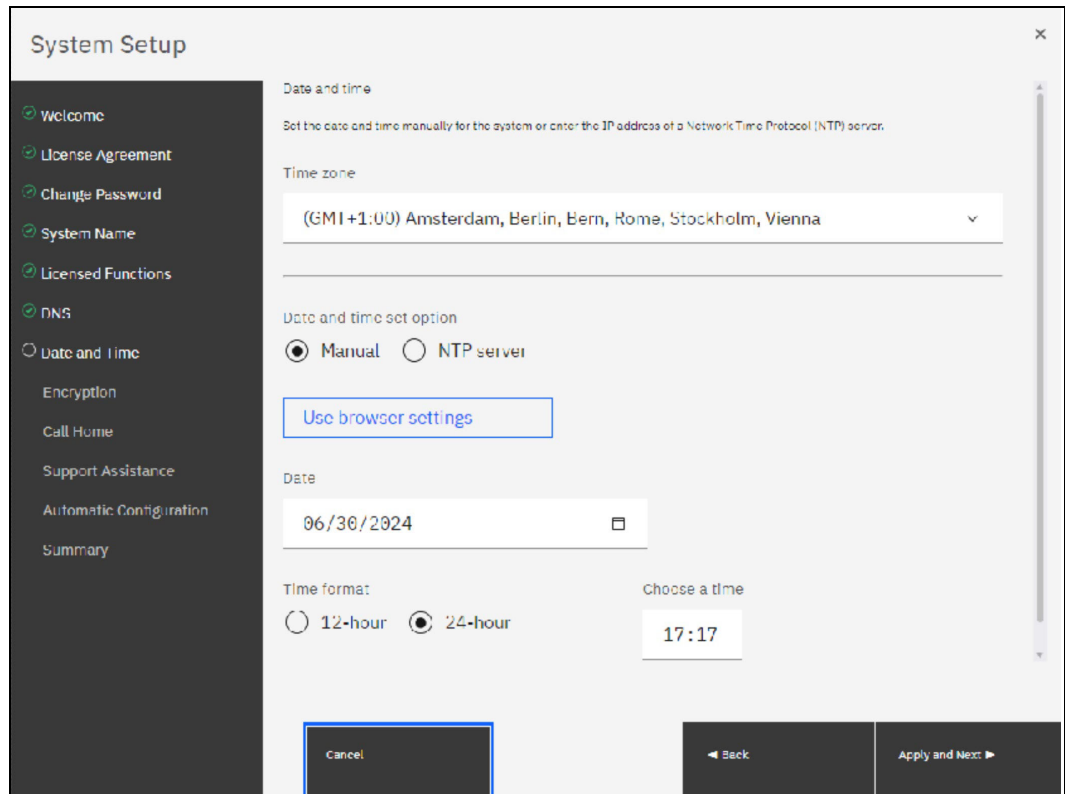


Figure 2-30 Date and Time

24. If you have purchased an Encryption License, activate it here. See Figure 2-31.

If encryption is not planned now, select **No** and then click **Next**. You can enable this feature later.

**Note:** When encryption is enabled on the system, encrypted storage pools can be created. If the system is a single control enclosure system where all FCM-drives should be in the same storage pool, encryption must be enabled before creating the storage pool. If a storage pool is created before encryption is enabled, any data in that pool must be migrated to an encrypted storage pool, if the data must be encrypted.

If you purchased the encryption feature, you are prompted to activate your license manually or automatically. The encryption license is key-based and required for each control enclosure.

You can use automatic activation if the workstation that you use to connect to the GUI and run the System Setup wizard has Internet access. If no Internet connection is available, use manual activation and follow the instructions.

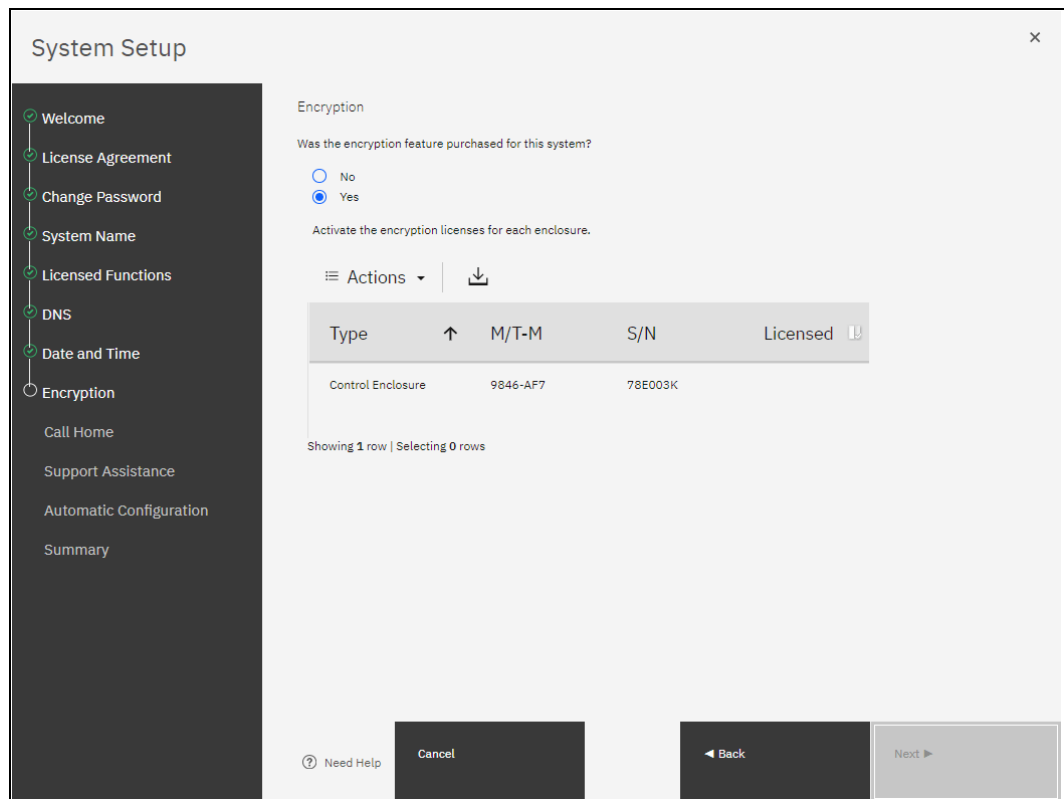


Figure 2-31 Activate Encryption License

25. After the encryption license is activated, you see a green checkmark for each enclosure, as shown in Figure 2-32 on page 37. After all the control enclosures show that encryption is licensed, click **Next**.

26. If you want to modify your previously entered Call Home settings, you can do so here. See Figure 2-33 on page 37.

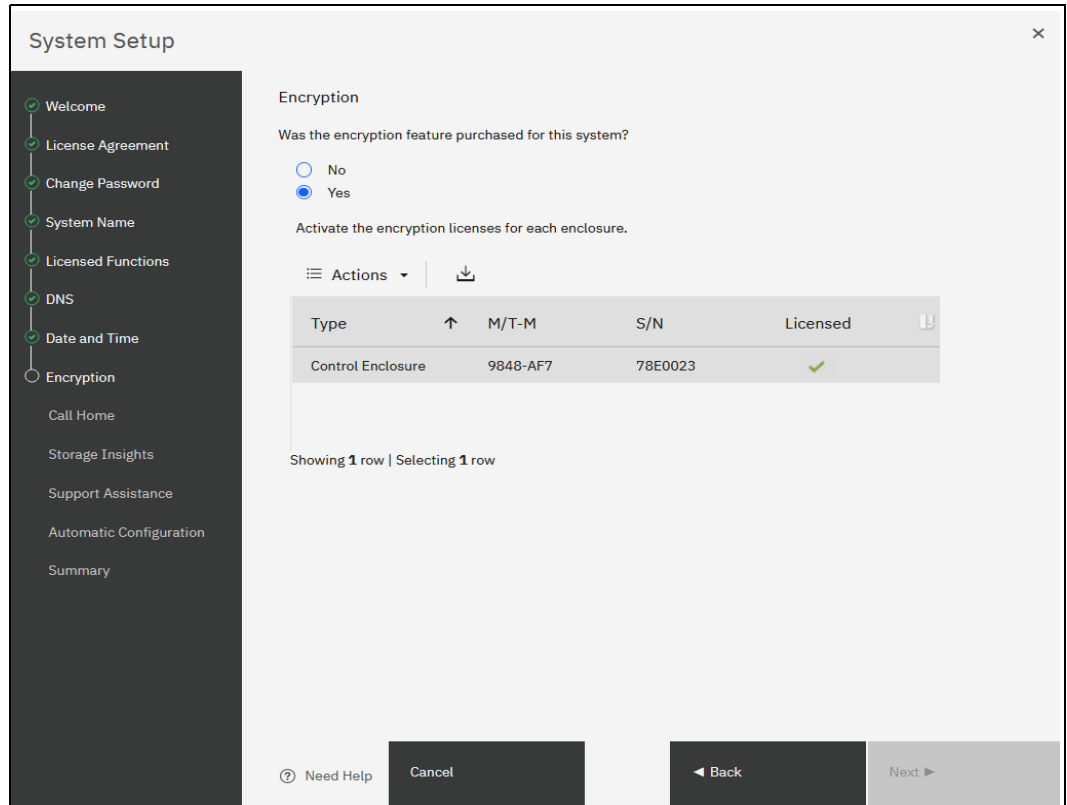


Figure 2-32 Encryption Licensed

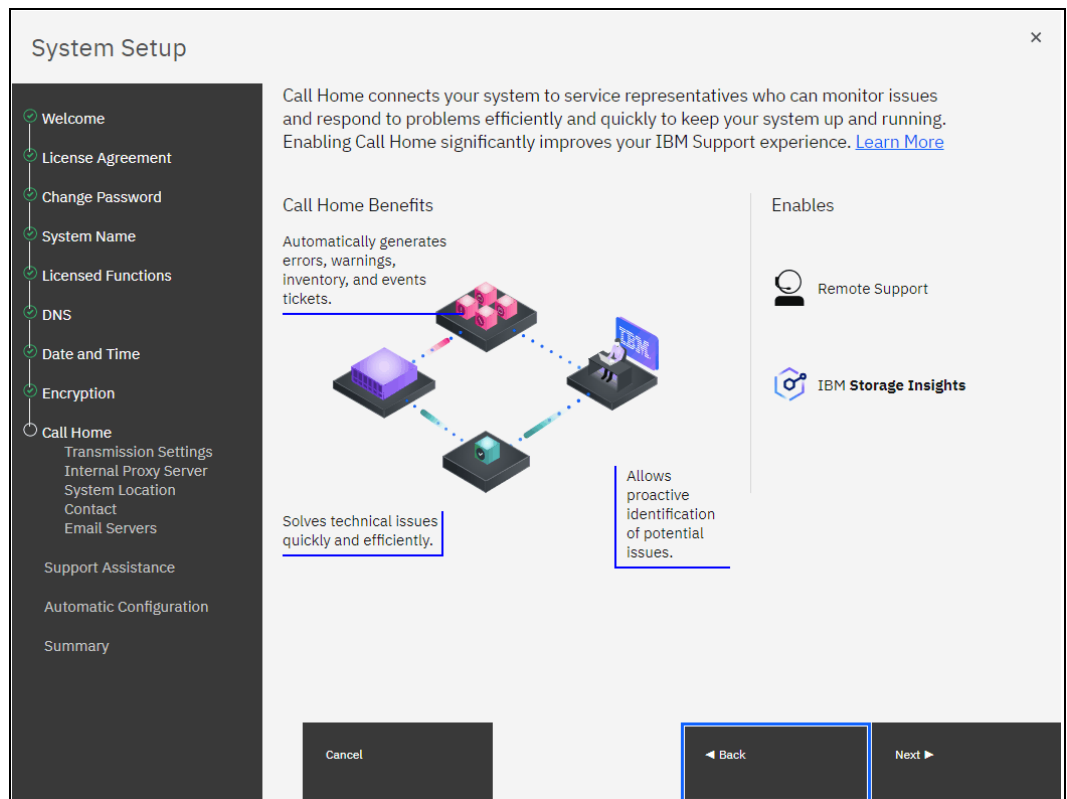


Figure 2-33 Change Call Home settings

27. If you want to use the **Support Assistance** offering from IBM, you can choose your options here (Figure 2-34). If you enabled at least one Call Home method, the Support Assistance configuration window opens. The Support Assistance function requires Call Home; therefore, if it is disabled, Support Assistance cannot be used. Click **Next** to continue.

**Note:** Refer to “Remote Support Assistance” on page 90 for more information.

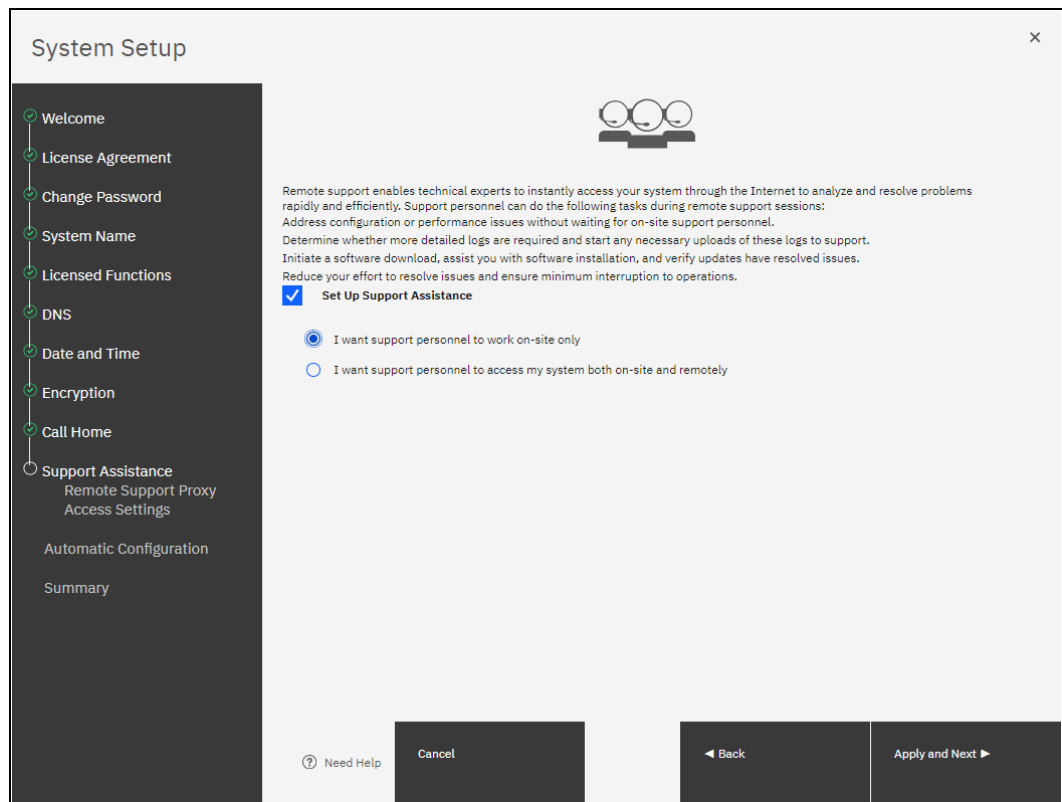


Figure 2-34 Setup Support Assistance

With the Support Assistance feature, you allow IBM Support to perform maintenance tasks on your system with support personnel onsite or remotely.

If an IBM SSR is onsite, the SSR can log in locally with your permission and a special user ID and password so that a superuser password does not need to be shared with the IBM SSR.

You can also enable Support Assistance with remote support to allow IBM Support personnel to log in remotely to the machine with your permission through a secure tunnel over the Internet.

If you allow remote support, you are provided with the IP addresses and ports of the remote support centers and an opportunity to provide proxy server details (if required) to allow the connectivity, as shown in Figure 2-35 on page 39. Click **Apply and Next**.

28. You can also allow remote connectivity at any time or only after obtaining permission from the storage administrator, as shown in Figure 2-36 on page 39.

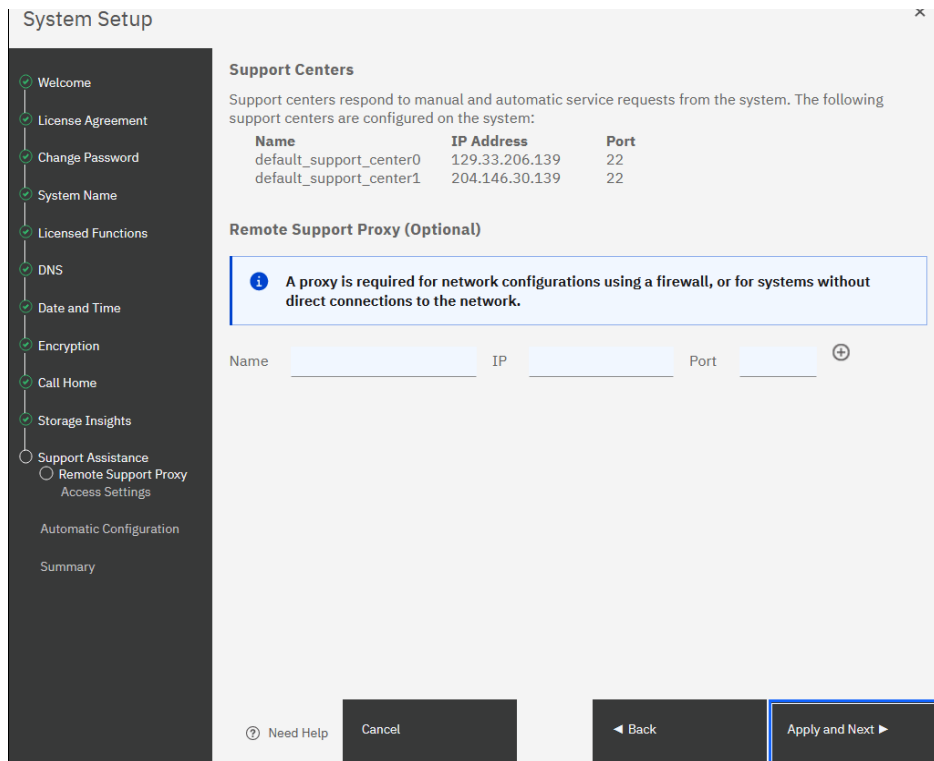


Figure 2-35 System communicating with named IBM Support servers

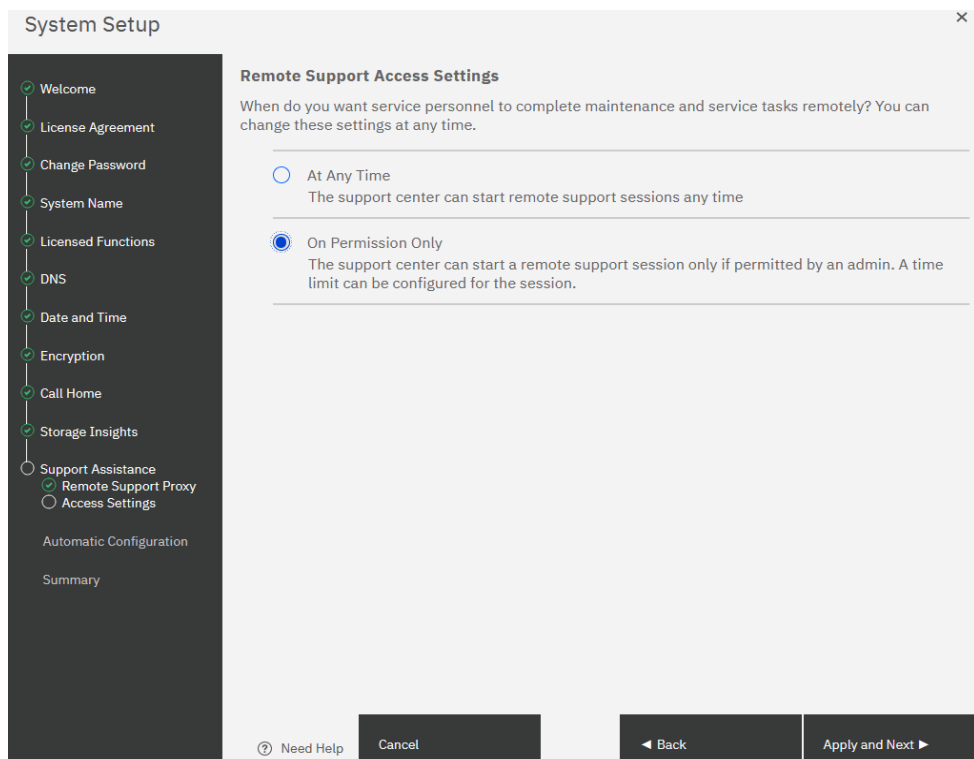


Figure 2-36 Remote support access settings

29. Click **Apply and Next**.

30. If you install your system below an IBM San Volume Controller, you can use **Automatic Configuration for Virtualization**. System Setup offers this option on IBM FlashSystem products only to automatically configure the system if it is used as FC-attached, back-end storage for IBM SAN Volume Controller. If you plan to use the system in stand-alone mode that is, not behind an IBM SAN Volume Controller, leave Automatic Configuration turned off, as shown Figure 2-37. Click **Next** to continue.

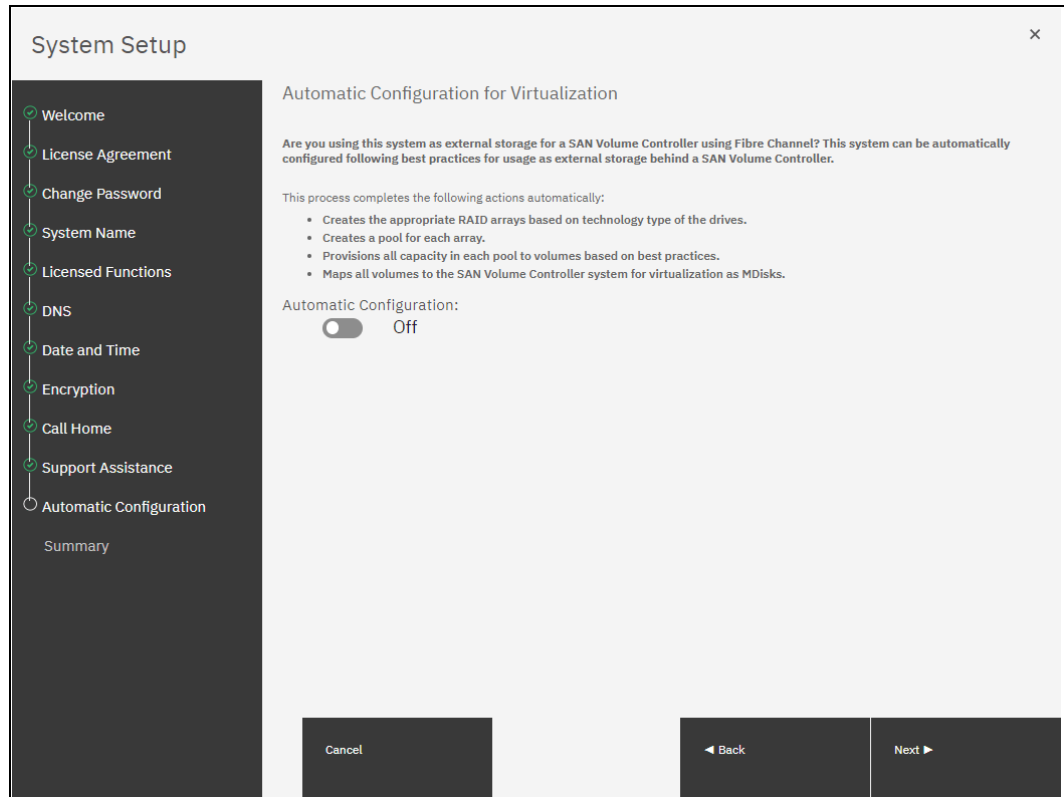


Figure 2-37 Automatic Configuration for Virtualization

For more information about how to enable Automatic configuration for IBM SAN Volume Controller on a running system after the System Setup wizard, see 2.3.7, “Automatic configuration for IBM SAN Volume Controller back-end storage” on page 55.

31. On the Summary page, the settings that were selected by the System Setup wizard are shown. If corrections are needed, you can return to a previous step by clicking **Back**. Otherwise, click **Finish** to complete the system setup wizard shown Figure 2-38 on page 41.

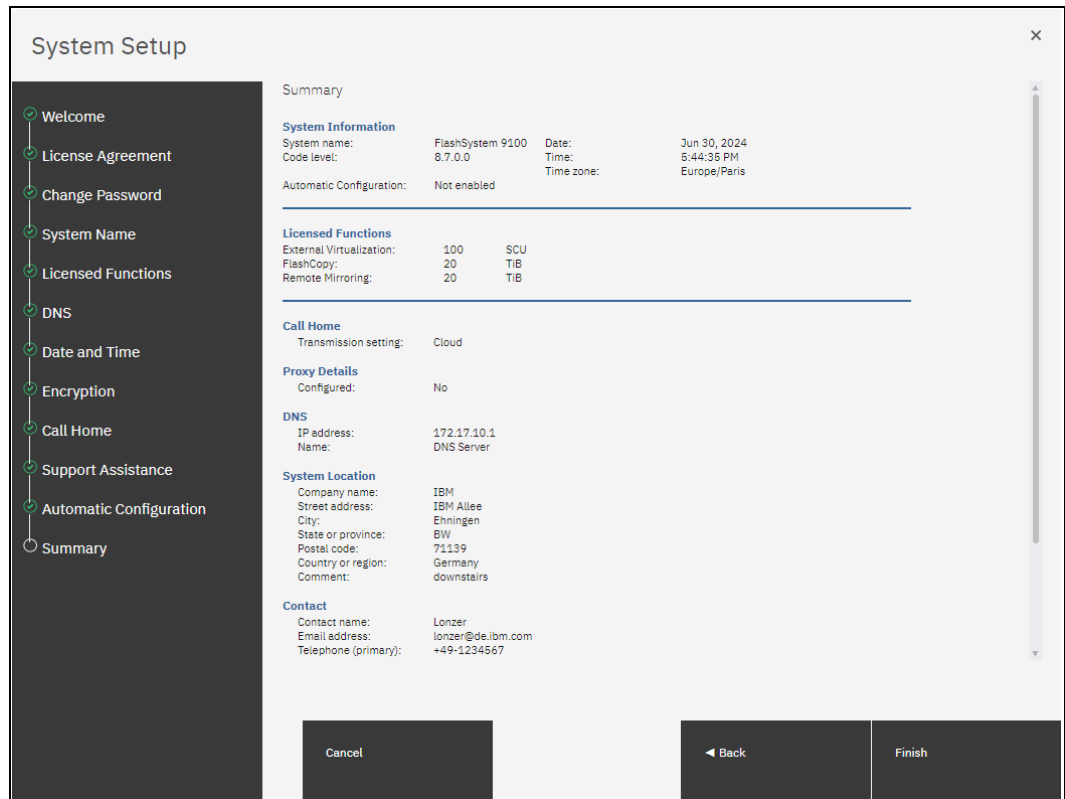


Figure 2-38 Summary Page

When the system setup wizard is done, your IBM FlashSystem consists of only the control enclosure that includes the node canister that you used to initialize the system and its partner, and the expansion enclosures that are attached to them.

When you set up an IBM SAN Volume Controller, your system consists of only one node in the cluster, which might see other candidate nodes in the service GUI if they are connected to SAN and zoned together.

If you have other control and expansion enclosures or IBM SAN Volume Controller nodes, you must add them to complete the System Setup.

For more information about how to add a control or expansion enclosure, see 2.3.2, “Adding an enclosure in IBM FlashSystem” on page 43.

For more information about how to add a node or hot spare node, see 2.3.3, “Adding a node or hot spare node in IBM SAN Volume Controller systems” on page 45.

If no other enclosures or nodes are to be added to this system, the System Setup process is complete and you can click **Finish** to be returned to the login window of the IBM FlashSystem.

All the required steps of the initial configuration are complete. If needed, you can configure other global functions, such as system topology, user authentication, or local port masking before configuring the volumes and provisioning them to hosts.

32. When you click **Finish**, you confirm your choices. See Figure 2-39.

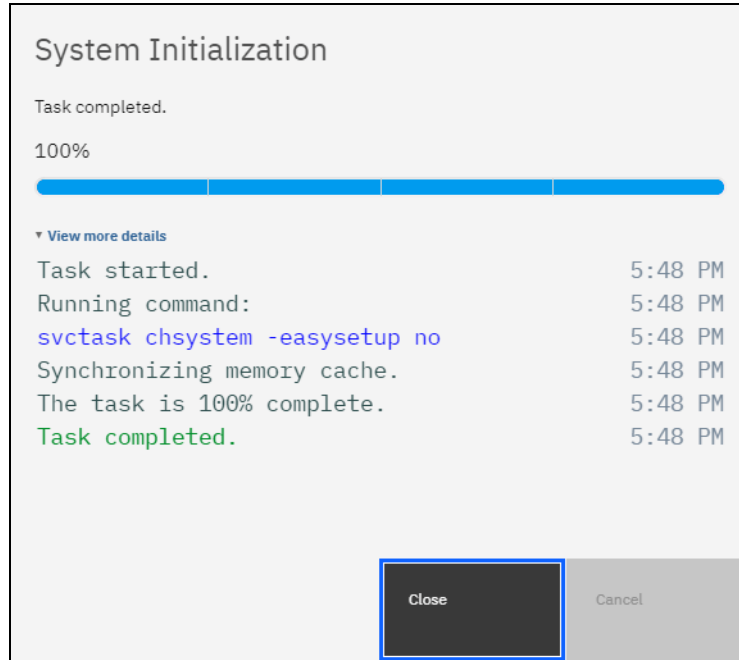


Figure 2-39 System Initialization

Click **Close** as shown in Figure 2-39.

33. In the Setup Completed window, click Close to be redirected to the Management GUI. See Figure 2-40.

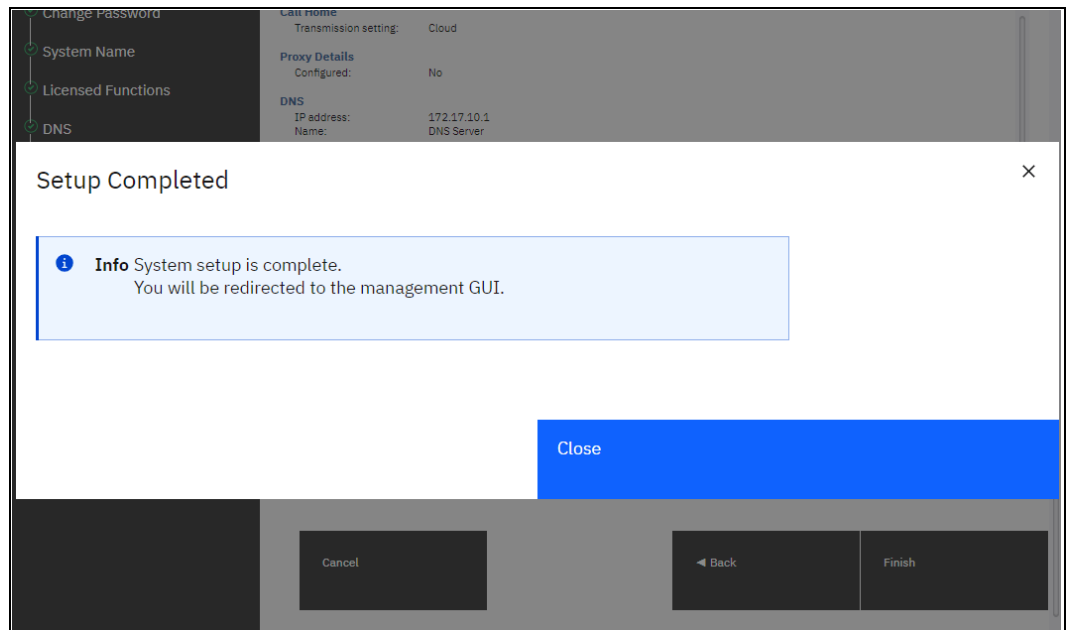


Figure 2-40 Setup completed



34. Clicking **Close and Finish** takes you to the Dashboard (see Figure 2-41).

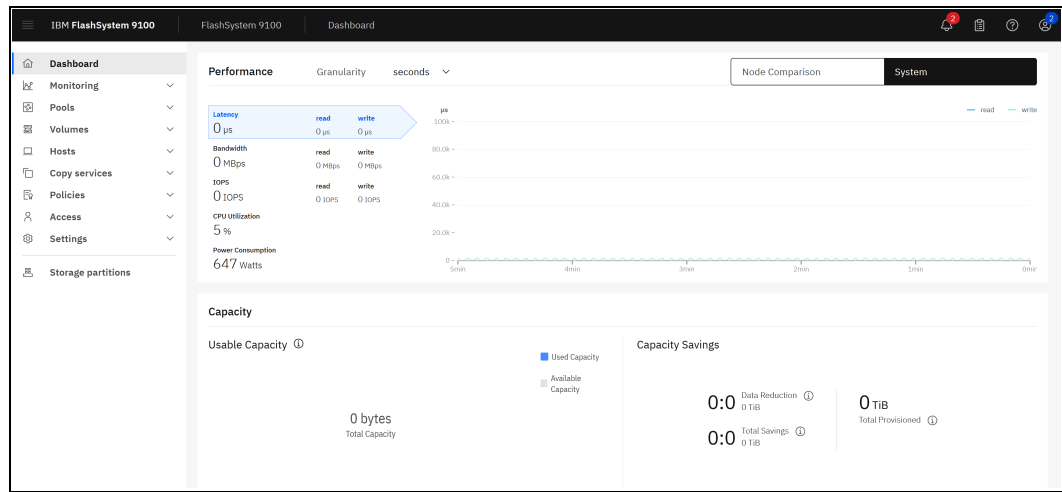


Figure 2-41 Dashboard

The tasks that are described next are used to define global system configuration settings. Often, they are performed during the System Setup process. However, they can also be performed later, such as when the system is expanded or the system environment is reconfigured.

### 2.3.2 Adding an enclosure in IBM FlashSystem

This procedure is the same whether you are configuring the system for the first time or expanding it. When the system GUI is used to add an enclosure, the same steps are used for adding expansion or control enclosures.

Before beginning this process, ensure that the new control enclosure is correctly installed and cabled to the system.

For FC node-to-node communication, verify that the correct SAN zoning is set.

For node-to-node communication over RDMA-capable Ethernet ports, ensure that the IP addresses are configured and a connection between nodes can be established.

#### Using the GUI to add an enclosure

To add an enclosure to the system from the GUI, complete the following steps:

1. In the GUI, select **Monitoring** → **System Hardware**. When a new enclosure is detected by a system, the Add Enclosure button appears next to System Actions, as shown in Figure 2-42.



Figure 2-42 Add Enclosure button

**Note:** If the Add Enclosure button does not appear, review the installation instructions to verify that the new enclosure is connected and set up correctly.

2. Click **Add Enclosure**, and a list of available candidate enclosures opens, as shown in Figure 2-43. To light the Identify light-emitting diode (LED) on a selected enclosure, select **Actions** → **Identify**. When the required enclosure (or enclosures) is chosen, click **Next**.

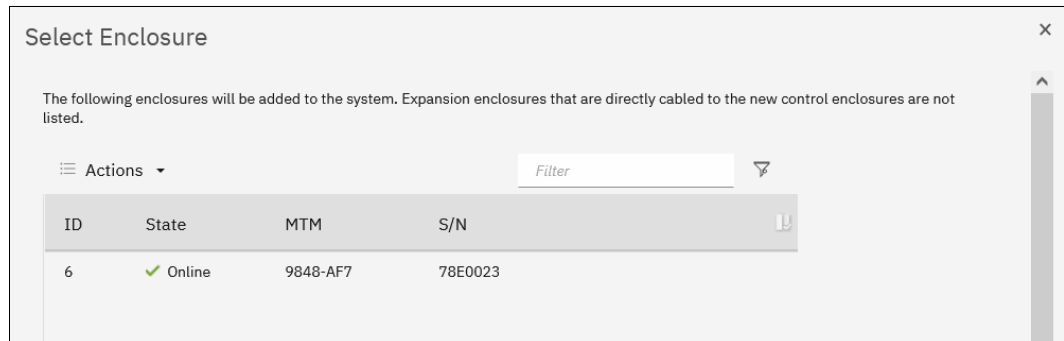


Figure 2-43 Selecting the control enclosure to add

3. Review the summary in the next window and click **Finish** to add the expansion enclosure or the control enclosure and all expansions that are attached to it to the system.

**Note:** When a new control enclosure is added, the software version that is running on its nodes is upgraded or rolled back to match the system software version. This process can take up to 30 minutes or more, and the enclosure is added only when this process completes.

4. After the control enclosure is successfully added to the system, a success message appears. Click **Close** to return to the System Overview window and check that the new enclosure is visible and available for management.

## Using the CLI to add an enclosure

To perform the same procedure by using a CLI, complete the following steps. For more information about the detailed syntax for each command, see [Command-line interface](#).

1. When you add control enclosures, check for unpopulated I/O groups by running the `lsiogrp` command. Because each control enclosure includes two nodes, it forms an I/O group.

Example 2-2 shows that only `io_grp0` has nodes. Therefore, a new control enclosure can be added to `io_grp1`.

Example 2-2 Listing the I/O groups

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>lsiogrp
id name          node_count vdisk_count host_count site_id site_name
0  io_grp0        2          0          0          0
1  io_grp1        0          0          0          0
2  io_grp2        0          0          0          0
3  io_grp3        0          0          0          0
4  recovery_io_grp 0          0          0          0
```

- To list control enclosures that are available to add, run the `lscontrolenclosurecandidate` command, as shown in Example 2-3. To list the expansion enclosures, run the `lsenclosure` command. Expansions that have the `managed` parameter set to no can be added.

*Example 2-3 Listing the candidate control enclosures*

---

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>lscontrolenclosurecandidate
serial_number product_MTM machine_signature
78E005D        9848-AF8      4AD2-EA69-8B5E-D0C0
```

---

- Add a control enclosure by running the `addcontrolenclosure` command, as shown in Example 2-4. The command triggers only the process, which starts in the background and can take 30 minutes or more.

*Example 2-4 Adding a control enclosure*

---

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>addcontrolenclosure -iogrp 1 -sernum
78E005D
```

---

- To add an expansion enclosure, change its `managed` status to yes by running the `chenclosure` command, as shown in Example 2-5.

*Example 2-5 Adding an expansion enclosure*

---

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>lsenclosure
id status type      managed IO_group_id IO_group_name product_MTM serial_number
1  online control  yes     0          io_grp0      9848-AF8    78E006A
2  online expansion no       0          io_grp0      9848-AFF    78CBVF5
```

---

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>chenclosure -managed yes 2
```

---

### 2.3.3 Adding a node or hot spare node in IBM SAN Volume Controller systems

This procedure is the same whether you are configuring the system for the first time or expanding it later. The same process is used to add a node to an I/O group, or a hot spare node.

Before beginning this process, ensure that the new control enclosure is correctly installed and cabled to the system.

For FC node-to-node communication, verify that the correct SAN zoning is set.

For node-to-node communication over RDMA-capable Ethernet ports, ensure that the IP addresses are configured and a connection between nodes can be established.

To add a node to the system, complete the following steps:

- In the GUI, select **Monitoring** → **System Hardware**. When a new enclosure is detected by a system, the **Add Node** button appears on the System - Overview window next to System Actions, as shown in Figure 2-44 on page 46.




Figure 2-44 Add Node button

**Note:** If the Add Node button does not appear, review the installation instructions to verify that the new node is connected and set up correctly.

2. Click **Add Node**. A form that you can use to assign nodes to I/O groups opens, as shown in Figure 2-45. To illuminate the Identify LED on a node, click the LED icon that is next to a node name. When the required node or nodes is selected, click **Finish**.




Figure 2-45 Adding a node

The Monitoring → Systems Hardware window changes and shows that the node is added, as shown in Figure 2-46 on page 47.

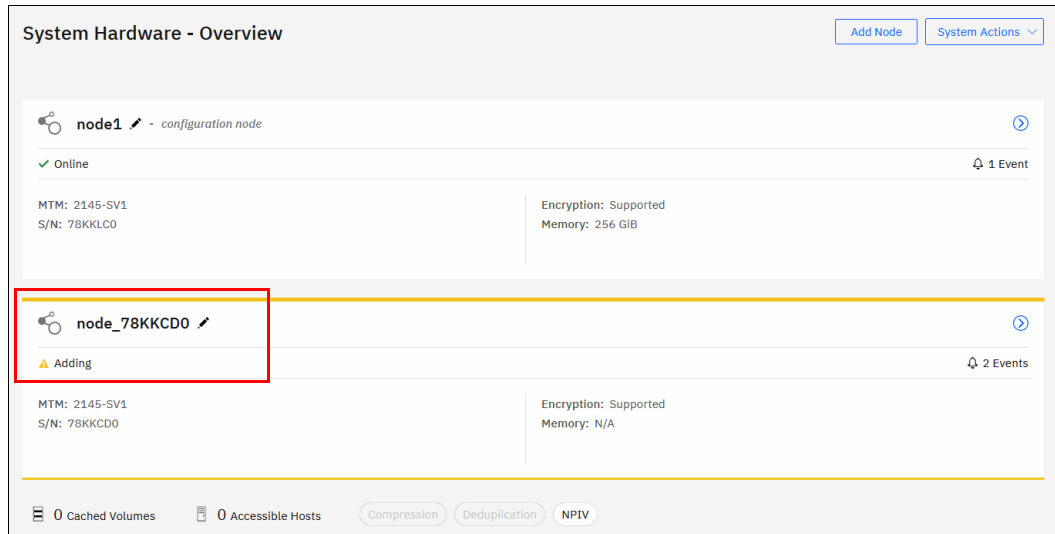


Figure 2-46 IBM SAN Volume Controller is adding node to the cluster

The node is added, as shown in Figure 2-47.

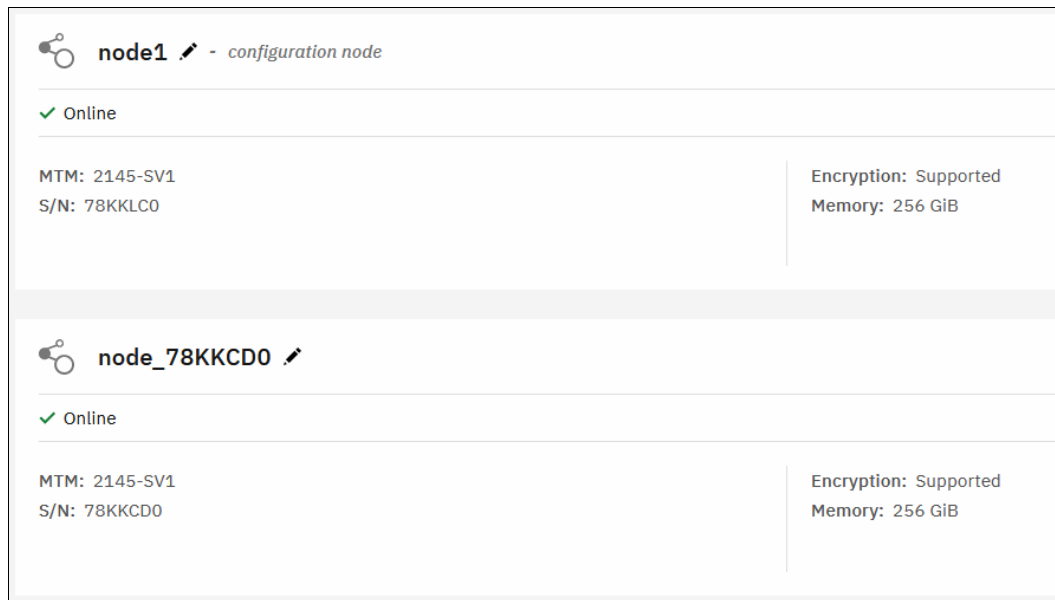


Figure 2-47 Node added

**Note:** When a node is added, the software version that is running is upgraded or rolled back to match the cluster software version. This process can take 30 minutes or more to complete. The node is added only after this process finishes.

To perform the same procedure by using a CLI, complete the following steps. For more information about the detailed syntax for each command, see [Command-line interface](#).

1. When you add nodes, check for unpopulated I/O groups by running **lsiogrp**. Each complete I/O group has two nodes. Example 2-6 shows that only `io_grp0` has nodes; therefore, a new control enclosure can be added to `io_grp1`.

*Example 2-6 Listing I/O groups*

---

```
IBM_2145:ITS0-SVC:superuser>lsiogrp
id name          node_count vdisk_count host_count site_id site_name
0  io_grp0        2          0          0          0
1  io_grp1        0          0          0          0
2  io_grp2        0          0          0          0
3  io_grp3        0          0          0          0
4  recovery_io_grp 0          0          0          0
```

---

2. To list nodes that are available to add to the I/O group, run the **lnodecandidate** command, as shown in Example 2-7.

*Example 2-7 Listing the candidate nodes*

---

```
BM_2145:ITS0-SVC:superuser>lnodecandidate
id          panel_name UPS_serial_number UPS_unique_id hardware
serial_number product_mtm machine_signature
500507680C00D98F 78KKLD0          500507680C00D98F SV1      78KKLD0
2145-SV1      3F25-557E-21E6-2B7D
500507680C00D98A 78KKCH0          500507680C00D98A SV1      78KKCH0
2145-SV1      702D-D5FE-76AA-4034
```

---

3. Add a node by running the **addnode** command. The command in Example 2-8 adds a node as a spare. The command starts in the background and can take 30 minutes or more.

*Example 2-8 Adding a node as a spare*

---

```
IBM_2145:ITS0-SVC:superuser>addnode -panelname 78KKLD0 -spare
Node, id [3], successfully added
```

---

In Example 2-9 the **addnode** command is used to add a node to I/O group `io_grp1`.

*Example 2-9 Adding a node to an I/O group*

---

```
IBM_2145:ITS0-SVC:superuser>addnode -panelname 78KKCH0 -name node3 -iogrp 1
Node, id [4], successfully added
```

---

4. List the nodes in the system by using CLI. As shown in Example 2-10 on page 49, the IBM SAN Volume Controller is configured with two nodes, which forms one IO-group. A spare node is configured for the IO-group.

### Example 2-10 Single IO-group (two nodes) and one spare

```
IBM_2145:ITSO-SVC:superuser>lsnode
id name          UPS_serial_number WWNN          status IO_group_id IO_group_name config_node
UPS_unique_id hardware iscsi_name          iscsi_alias panel_name
enclosure_id canister_id enclosure_serial_number site_id site_name
1 node1_78KKLCO 500507680C00D990 online 0          io_grp0      yes
SV1 iqn.1986-03.com.ibm:2145.itso-svc.node178kklc0 78KKLCO
2 node2_78KKCDO 500507680C00D982 online 0          io_grp0      no
SV1 iqn.1986-03.com.ibm:2145.itso-svc.node278kkcd0 78KKCDO
3 spare1        500507680C00D98F spare   0          io_grp0      no
SV1                                     78KKLDO
```

A two-IO-group system with no spare is shown in Example 2-11.

### Example 2-11 Two IO-groups (four nodes) configured- no spare

```
IBM_2145:ITSO-SVC:superuser>lsnode
id name          UPS_serial_number WWNN          status IO_group_id IO_group_name config_node
UPS_unique_id hardware iscsi_name          iscsi_alias panel_name
enclosure_id canister_id enclosure_serial_number site_id site_name
1 node1_78KKLCO 500507680C00D990 online 0          io_grp0      yes
SV1 iqn.1986-03.com.ibm:2145.itso-svc.node178kklc0 78KKLCO
2 node2_78KKCDO 500507680C00D982 online 0          io_grp0      no
SV1 iqn.1986-03.com.ibm:2145.itso-svc.node278kkcd0 78KKCDO
3 node3_78KKCHO 500507680C00D98A online 1          io_grp1      no
SV1 iqn.1986-03.com.ibm:2145.itso-svc.node378kkch0 78KKCHO
4 node4_78KKLDO 500507680C00D98F online 1          io_grp1      no
SV1 iqn.1986-03.com.ibm:2145.itso-svc.node478kkld0 78KKLDO
```

The administrator might want to rename the nodes to feature consistent names. This process can be done by clicking **Monitoring** → **System Hardware** → **Node Actions** → **Rename**.

## 2.3.4 Business continuity with policy-based high availability

Business continuity ensures an organization can deliver services even during disruptions. Although some applications might tolerate temporary outages, major disasters can cause significant downtime and data loss, leading to immense recovery costs. Organizations should minimize data loss and downtime to lessen business impact and financial strain.

From a storage perspective, business continuity involves maintaining data consistency and availability for uninterrupted application access, achieved through two key concepts: disaster recovery (DR) and high availability (HA). DR focuses on replicating data to remote locations for recovery. HA prioritizes continuous data accessibility.

Disasters can range from entire site outages to data corruption or theft. Data protection relies on local or remote backups. IBM Storage Virtualize offers functionalities to safeguard your data against various threats, such as hardware failures, software errors, or cyberattacks. Policy-based replication and policy-based high availability protect against site failures by automatically failing over to a secondary site, helping ensure business continuity. Although it is not covered here, Storage Virtualize offers additional features such as snapshots and Safeguarded snapshots to protect against data corruption or cyberattacks.

**Note:** Policy-based high availability is *not* supported by the IBM FlashSystem 5015.

For more information about this topic, refer to IBM Redbooks *Ensuring Business Continuity: A Practical Guide to Policy-Based Replication and Policy-Based High Availability for IBM Storage Virtualize Systems*, SG24-8569.

## 2.3.5 Configuring quorum disks or applications

Quorum devices are required for a system to hold a copy of important system configuration data. An internal drive of an IBM FlashSystem, a managed disk (MDisk) from FC-attached external back-end storage, or a special application that is connected over an IP network can work as a quorum device.

One of these items is selected for the *active quorum* role, which is used to resolve failure scenarios where half the nodes on the system become unavailable or a link between enclosures is disrupted. The active quorum determines which nodes can continue processing host operations. It also avoids a “split brain” condition, which occurs when both halves of the system continue I/O processing independently of each other.

For IBM FlashSystem products with a single control enclosure and IBM SAN Volume Controller systems with a standard topology, quorum devices are automatically selected from the internal drives or assigned from an MDisk, respectively. No special configuration actions are required. This function also applies for IBM FlashSystem products with multiple control enclosures, a standard topology, and virtualizing external storage.

For policy-based HA or for HyperSwap or Enhanced Stretched Cluster topology systems, configure an active quorum device on a third, independent site. Because of the costs that are associated with deploying a separate FC-attached storage device on a third site, an IP-based quorum device can be used for this purpose.

Without a third arbitration site (quorum server), a tie-breaker mechanism must be chosen for the two existing sites. During a network outage between the sites, the pre-configured *winner* continues operating and processing I/O requests. The *loser* site is unavailable until the connection is restored. IP quorum settings, within the configuration options, determine the preferred site for handling these scenarios. If a site outage occurs at the winning site, the system stops processing I/O requests until this site is recovered or the manual quorum override procedure is used.

On IBM FlashSystem products in a standard topology system with two or more control enclosures and no external storage, none of the internal drives can be the active quorum device. For such configurations, it is a best practice to deploy an IP-based quorum application to avoid a “split brain” condition.

### Creating and installing an IP quorum application

To create and install an IP quorum application, complete the following steps:

1. Select **Settings** → **System** → **IP Quorum** to download the IP quorum application, as shown in Figure 2-48 on page 51. If you use IPv6 for management IP addresses, the **Download IPv6 Application** button is available and the IPv4 option is disabled. In the example, **Download IPv4 Application** is selected.






Figure 2-48 Download IPv4 quorum button

2. Click **Download...** and a window opens, as shown in Figure 2-49. It provides an option to create an IP application that is used for tie-breaking only, or an application that can be used as a tie-breaker and to store recovery metadata.

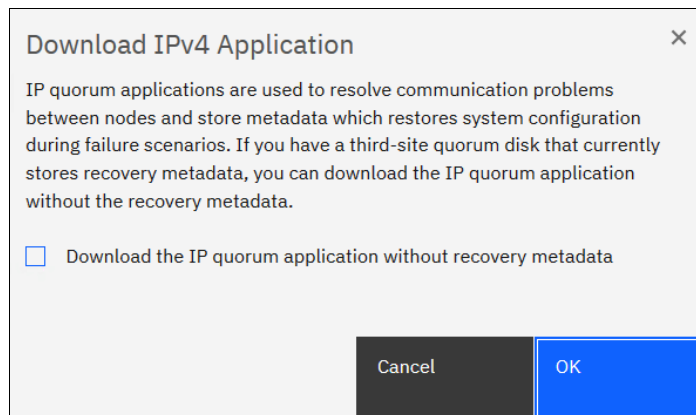


Figure 2-49 Download IP quorum application window

An application that does not store recovery metadata requires less channel bandwidth for a link between the system and the quorum app, which might be a decision-making factor for using a multi-site HA system.

For a full list of IP quorum app requirements, see [IP quorum application configuration](#).

3. Click **OK**. The `ip_quorum.jar` file is created. Save the file and transfer it to a supported AIX®, Linux, or Windows host that can establish an IP connection to the service IP address of each system node. Move it to a separate directory and start the application, as shown in Example 2-12.

*Example 2-12 Starting the IP quorum application on the Windows operating system*

```
C:\IPQuorum>java -jar ip_quorum.jar
=== IP quorum ===
Name set to null.
Successfully parsed the configuration, found 2 nodes.
Trying to open socket
Trying to open socket
```

Handshaking  
Handshaking  
Waiting for UID  
Creating UID  
\*Connecting  
Connected to 10.0.0.42  
Connected to 10.0.0.41

**Note:** Add the IP quorum application to the list of auto-started applications at each start or restart or configure your operating system to run it as an auto-started service in the background. The server hosting the IP quorum application must reside within the same network subnet as the IBM FlashSystem for proper communication. Up to five IP quorums can be deployed in your environment.

The IP quorum log file and recovery metadata are stored in the same directory with the ip\_quorum.jar file.

4. Check that the IP quorum application is successfully connected and running by verifying its online status by selecting **Settings** → **System** → **IP Quorum**, as shown in Figure 2-50.

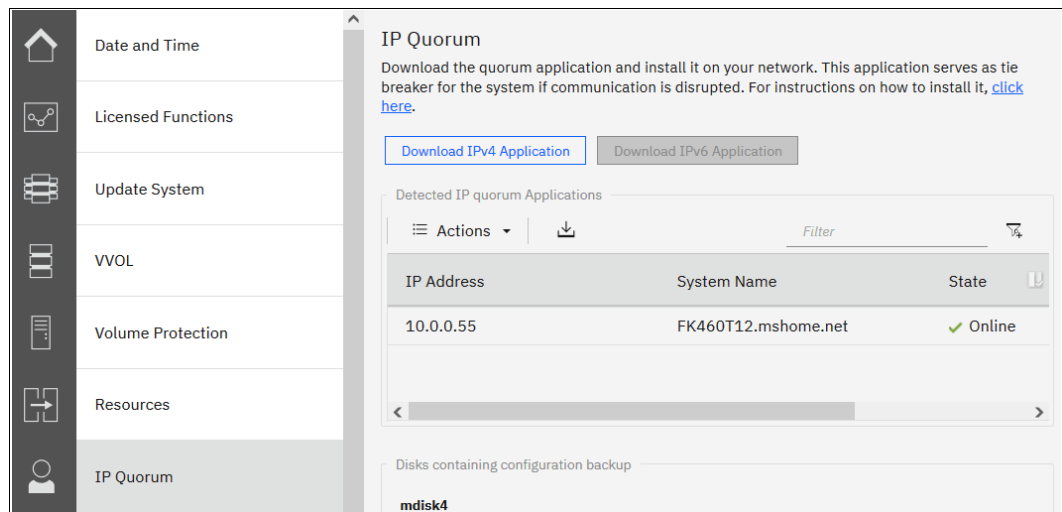


Figure 2-50 IP quorum application that is deployed and connected

## Configuring the IP quorum mode

On a standard topology system, only the Standard quorum mode is supported. No other configuration is required. On a policy-based HA topology, a tie occurs when half of the nodes that were a member of the system are present. You can configure the following tie-breaker scenarios:

- ▶ If the quorum mode is set to Standard, both sites have an equal chance to continue working after the tie breaker.
- ▶ If the quorum mode is set to Preferred, then during a disruption the system delays processing tie-breaker operations on non-preferred sites, which leaves more time for the preferred site to win. If during an extended period a preferred site cannot contact the IP quorum application (for example, if it is destroyed), a non-preferred site continues working.
- ▶ If the quorum mode is set to Winner, the selected site is always the tie-breaker winner. If the winner site becomes unavailable, the remaining site can continue operating only after manual intervention.

The Preferred quorum mode is supported by an IP quorum only.

To set a quorum mode, select **Settings** → **System** → **IP Quorum** and then click **Quorum Setting**. The Quorum Setting window opens, as shown in Figure 2-51.

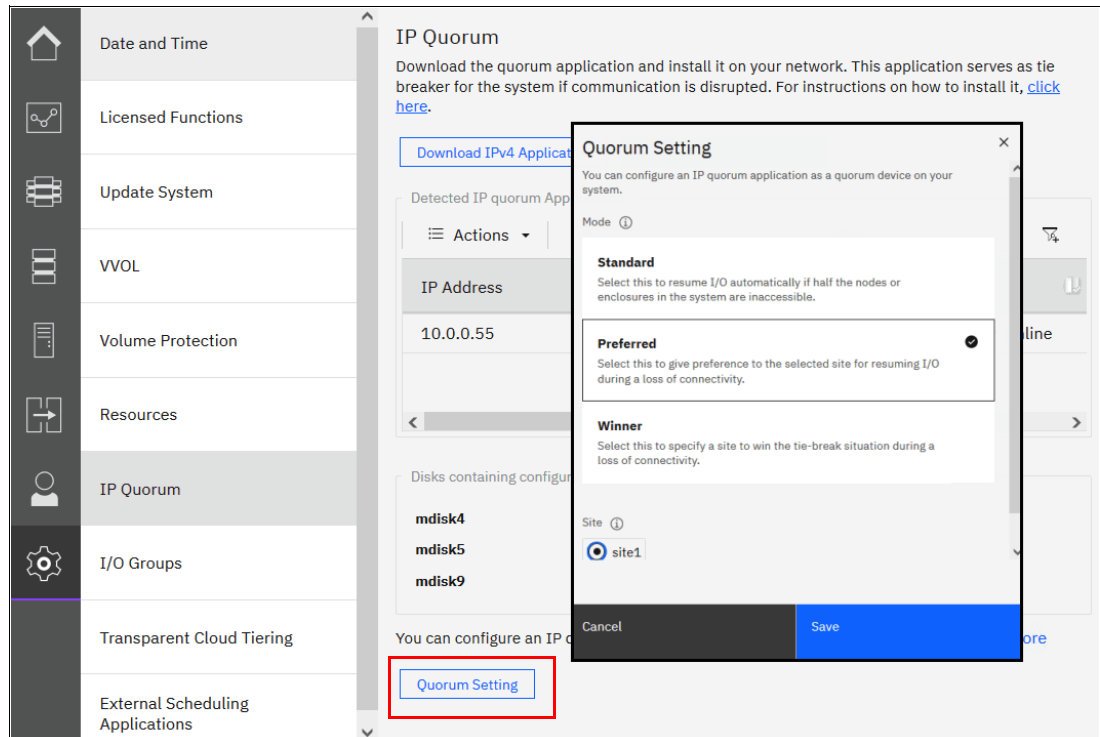


Figure 2-51 Changing the quorum mode

### 2.3.6 Configuring the local Fibre Channel port masking

With FC port masking, you control the use of FC ports. By applying a mask, you restrict node-to-node communication or replication traffic on selected ports.

To set the FC port mask by using the GUI, complete the following steps:

1. Select **Settings** → **Network** → **Fibre Channel Ports**. In a displayed list of FC ports, the ports are grouped by a system port ID. Each port is configured identically across all nodes in the system. You can click the arrow next to the port ID to expand a list and see which node ports (N\_Port) belong to the selected system port ID and their worldwide port names (WWPNs).
2. Right-click a system port ID that you want to change and select **Modify Connection**, as shown in Figure 2-52 on page 54.

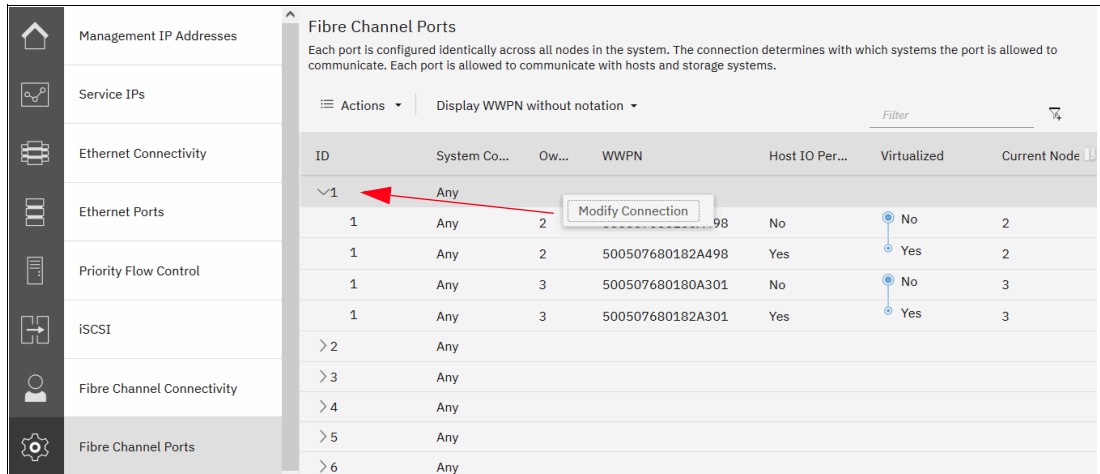


Figure 2-52 Applying a port mask by using a GUI

By default, all system ports can send and receive traffic of any kind, including the following examples:

- Host traffic
- Traffic to virtualized back-end storage systems
- Local system traffic (node to node)
- Partner system (remote replication) traffic

The first two types are always allowed, and you can control them only with SAN zoning. The other two types can be blocked by port masking.

3. In the Modify Connection dialog box (Figure 2-53), you can choose which type of traffic a port can send. For example, Remote if the port is dedicated to Remote Replication traffic. Click **Modify** when done.

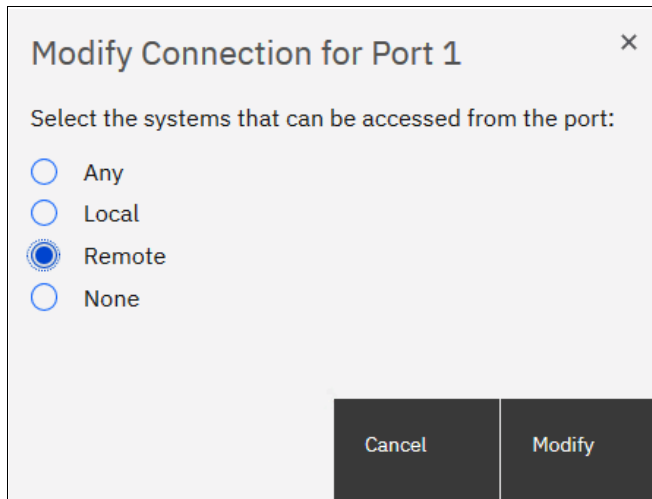


Figure 2-53 Modify Connection dialog box

The following types of traffic are allowed for each choice:

- ▶ Any. A port can work with all types of traffic.
- ▶ Local. Local node-to-node traffic is allowed. Remote replication traffic is blocked on this port.
- ▶ Remote. Remote replication traffic is allowed. Local node-to-node traffic is blocked.

- ▶ None. Local and remote systems traffic is allowed, but system-to-host and system-to-back-end storage communication still exists.

Port masks can also be set by using the CLI. Local and remote partner port masks are internally represented as a string of zeros and ones. The last digit in the string represents port one. The previous digits represent ports two, three, and so on.

If the digit for a port is set to 1, the port is enabled for the specific type of communication. If it is set to 0, the system does not send or receive traffic that is controlled by a mask on the port.

To view the current port mask settings, run the `lssystem` command, as shown in Example 2-13. The output shows that all system ports allow all types of traffic.

*Example 2-13 Viewing the local port mask*

---

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>lssystem |grep mask
local_fc_port_mask 1111111111111111111111111111111111111111111111111111111111111111
partner_fc_port_mask 1111111111111111111111111111111111111111111111111111111111111111
```

---

To set the `localfcportmask` for node to node traffic or the `partnerfcportmask` for remote replication traffic, run the `chsystem` command. Example 2-14 shows the mask setting for a system with four FC ports on each node and that has FC relationships. Masks are applied to allow local node-to-node traffic only on ports 1 and 2, and replication traffic only on ports 3 and 4.

*Example 2-14 Setting a local port mask by running the chsystem command*

---

```
IBM_IBM FlashSystem:ITS0-FS9500:superuser>chsystem -localfcportmask 0011
IBM_IBM FlashSystem:ITS0-FS9500:superuser>chsystem -partnerfcportmask 1100
IBM_IBM FlashSystem:ITS0-FS9500:superuser>lssystem |grep mask
local_fc_port_mask 0000000000000000000000000000000000000000000000000000000000000011
partner_fc_port_mask 0000000000000000000000000000000000000000000000000000000000001100
```

---

The mask is extended with zeros, and all ports that are not set in a mask have the selected type of traffic blocked.

**Note:** When replacing or upgrading your node hardware, consider that the number of FC ports and their arrangement might be changed. If so, make sure that any configured port masks are still valid for the new configuration.

### 2.3.7 Automatic configuration for IBM SAN Volume Controller back-end storage

If a system is to be used as FC-attached back-end storage for IBM SAN Volume Controller, you can enable Automatic Configuration for Virtualization during the initial System Setup or anytime later by selecting **Settings** → **System** → **Automatic Configuration**.

Automatic Configuration for Virtualization is intended for a new system. If host, pool, or volume objects are configured, all the user data must be migrated out of the system and those objects must be deleted.

The Automatic Configuration for Virtualization wizard starts immediately after you complete the initial setup wizard if you set Automatic Configuration to 0n.

Complete the following steps:

1. Click **Settings** → **System** → **Automatic Configuration**. Then, select Automatic Configuration **ON** and click **Save**, as shown in Figure 2-54.

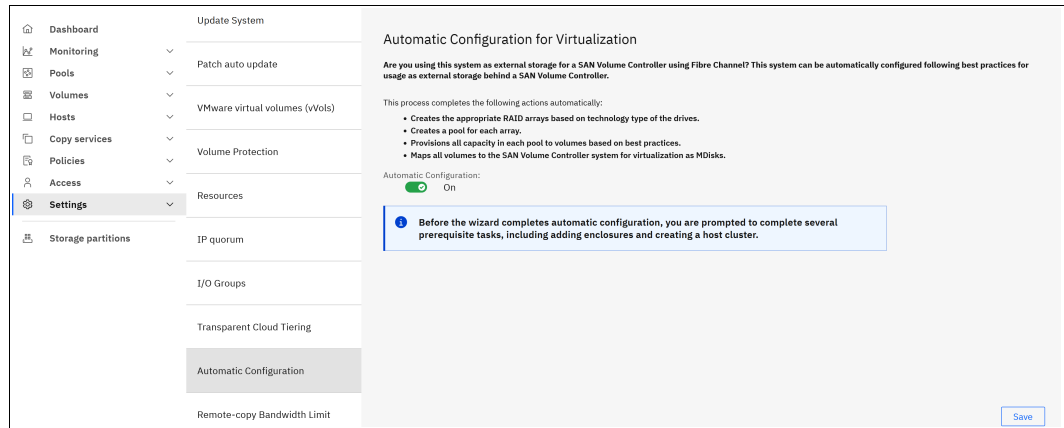


Figure 2-54 Automatic Configuration wizard enablement

2. You can add any control or expansion enclosures as part of the external storage to be virtualized. If you do not have more enclosures to add, this part of the prerequisite steps can be skipped.

Click **Add Enclosure** to add the enclosures, or click **Skip** to move to the next step (see Figure 2-55).

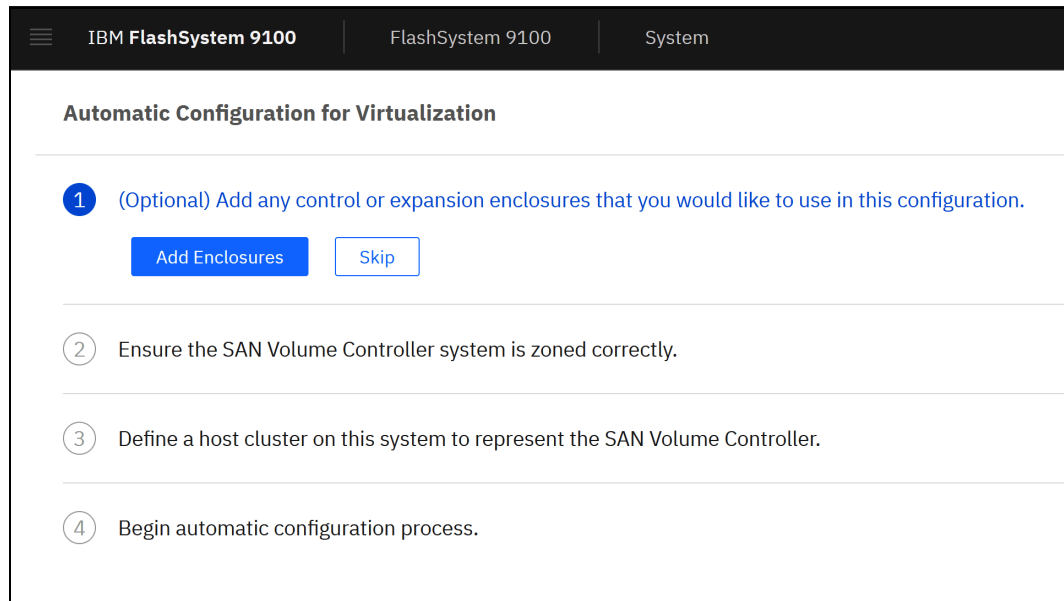


Figure 2-55 Automatic configuration: Add Enclosure

**Note:** You can turn off the Automatic Configuration for Virtualization wizard at any step by clicking the dotted symbol in the upper right corner.

- The wizard checks whether the IBM SAN Volume Controller is correctly zoned to the system. By default, newly installed systems run in N\_Port ID Virtualization (NPIV) mode (Target Port Mode). The system's virtual (host) WWPNs must be zoned for IBM SAN Volume Controller. On the IBM SAN Volume Controller side, physical WWPNs must be zoned to a back-end system independently of the NPIV mode setting.
- Create a host cluster object for IBM SAN Volume Controller. Each IBM SAN Volume Controller node has its own worldwide node name (WWNN). Make sure to select all WWNNs that belong to nodes of the same IBM SAN Volume Controller cluster.

Figure 2-56 shows that because the system detected an IBM SAN Volume Controller cluster with dual I/O groups, four WWNNs are selected.

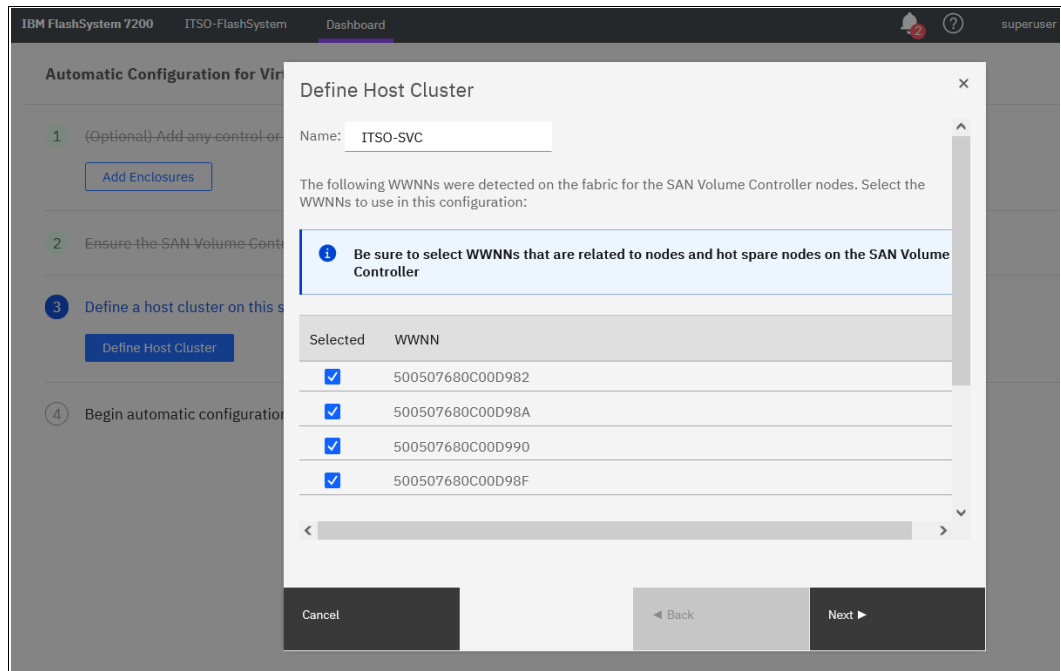


Figure 2-56 Defining a host cluster

- When all nodes of an IBM SAN Volume Controller cluster (including the spare cluster) are selected, you can change the host object name for each one, as shown in Figure 2-57 on page 58. For convenience, name the host objects to match the IBM SAN Volume Controller node names or serial numbers.




Figure 2-57 Hosts inside an IBM SAN Volume Controller host cluster

6. Click **Automatic Configuration** and check the list of internal resources that are used, as shown in Figure 2-58.

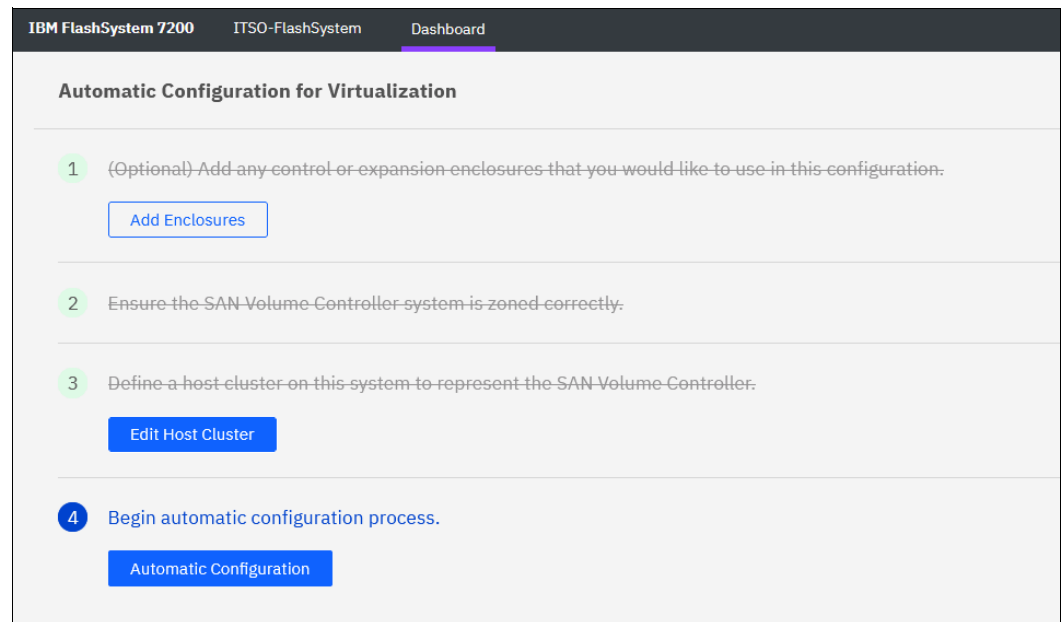


Figure 2-58 Begin the automatic configuration process



- If the system uses compressed drives (FCM drives), you are prompted to enter your expected compression ratio or the total capacity that is to be provisioned to IBM SAN Volume Controller (Figure 2-59). If IBM SAN Volume Controller uses encryption or writes data that is not compressible, set the ratio to 1:1 and then click **Next**.

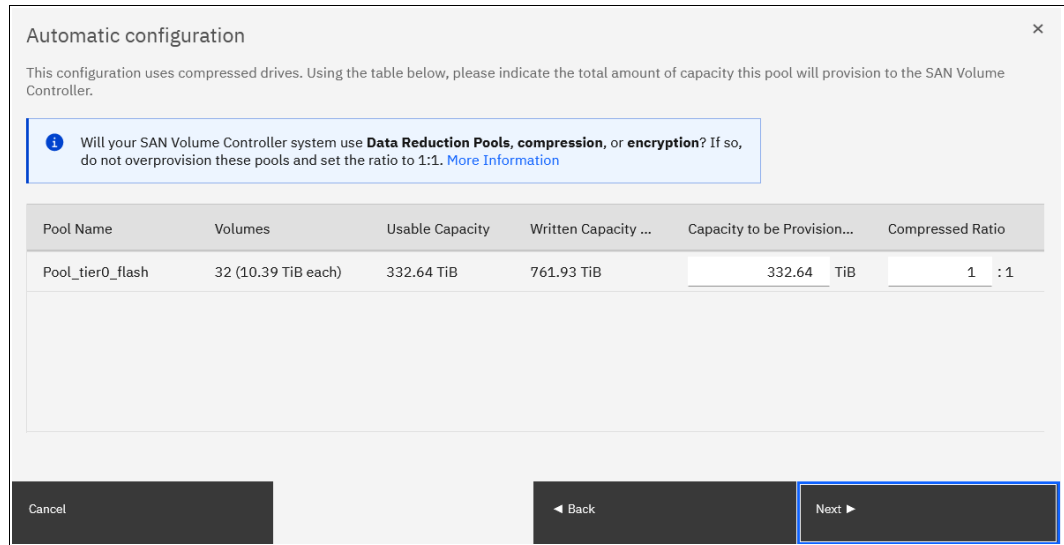


Figure 2-59 Automatic pool configuration

- Review the configuration of the pools, (Figure 2-60), and click **Proceed** to apply the configuration.

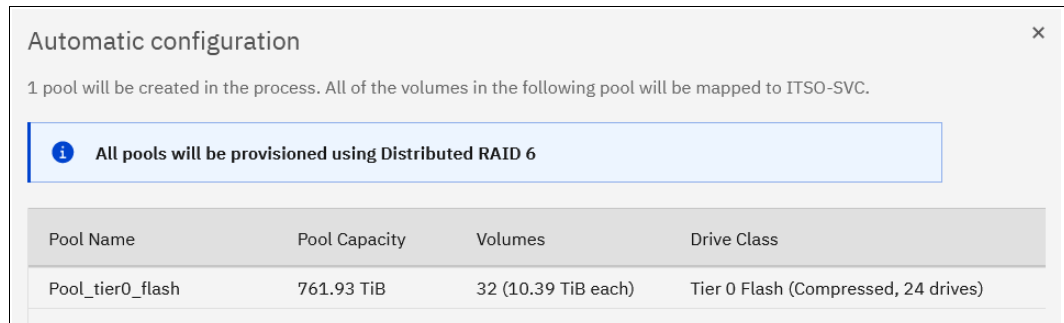


Figure 2-60 Pools configuration

- When the Automatic Configuration for Virtualization wizard completes, you see the window that is shown in Figure 2-61. After clicking **Close**, you can proceed to the IBM SAN Volume Controller GUI and configure a new provisioned storage.

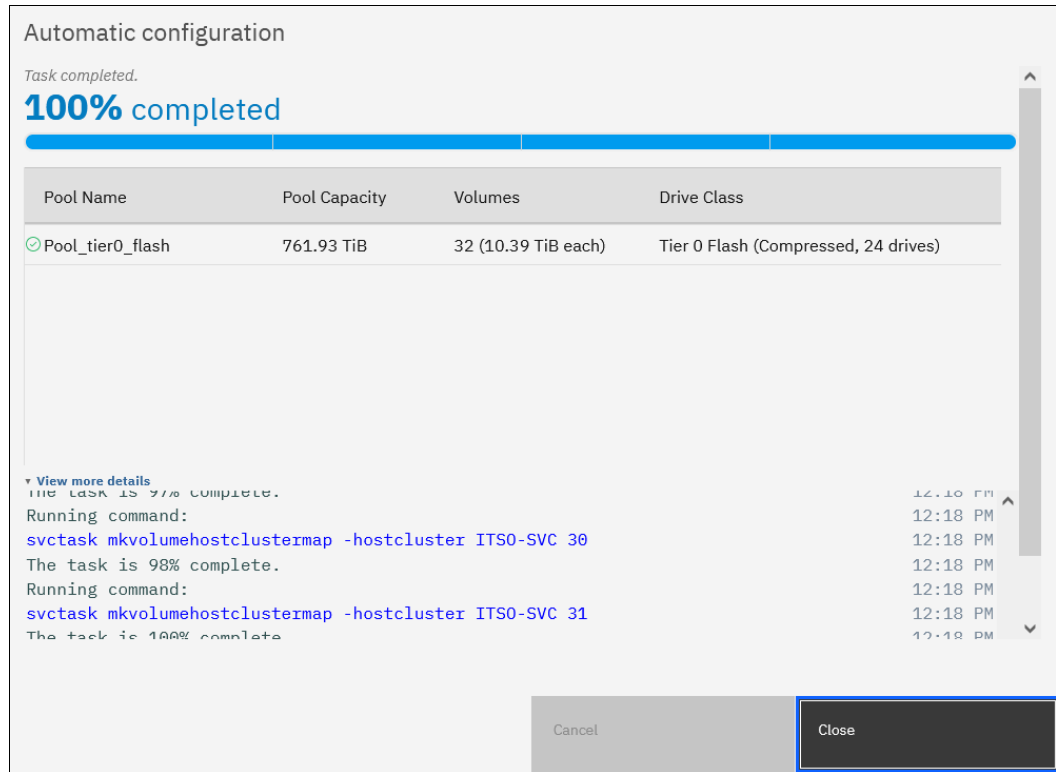


Figure 2-61 Automatic configuration running commands

- You can export the system volume configuration data in .csv format by using this window or anytime by selecting **Settings** → **System** → **Automatic Configuration**.

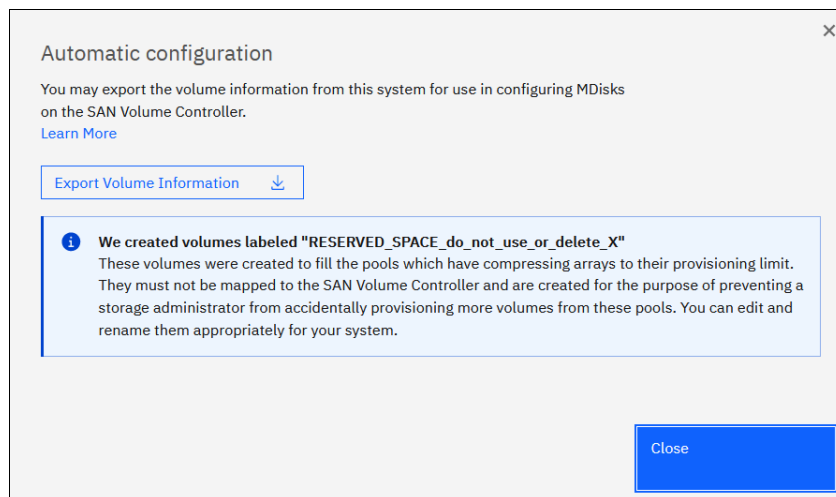


Figure 2-62 Automatic configuration complete



## Step-by-step configuration

This chapter describes the steps used to quickly implement a Storage Virtualize system by using the most commonly used functions. Where applicable links are provided for additional reading about advanced functions.

This chapter describes the Storage Virtualize GUI, the steps needed for network configuration, creating pools and assigning storage, configuring hosts, basic snapshots, and asynchronous replication configuration.

The following topics are included:

- ▶ 3.1, “The Storage Virtualize GUI” on page 62
- ▶ 3.2, “Network configuration” on page 63
- ▶ 3.3, “Pools and managed disks configuration” on page 65
- ▶ 3.4, “Configuring volumes” on page 72
- ▶ 3.5, “Configuring hosts” on page 73
- ▶ 3.6, “Snapshots and replication” on page 78

## 3.1 The Storage Virtualize GUI

The Storage Virtualize GUI is a built-in software component within the IBM Storage Virtualize Software. Multiple users can be logged in to the GUI. However, because no locking mechanism exists, if two users change the same object simultaneously, then the last action that is entered from the GUI is the action that takes effect.

### 3.1.1 Accessing the GUI

To access the IBM GUI, enter the system IP address that was defined during the initial setup process into your web browser. You can connect from any workstation that can communicate with the system.

**Recommendation:** It is a recommended practice for each user to have their own unique account

The default user accounts can be disabled for use or their passwords changed and kept secured for emergency purposes only. This approach helps to identify any personnel who are working on the systems and track all important changes that are done by them. The *superuser* account is for initial configuration and servicing the system only. For more information on user accounts, see [Users](#).

### 3.1.2 Brief introduction to the GUI

After a successful login, the Welcome window opens and displays the system dashboard. See Figure 3-1.

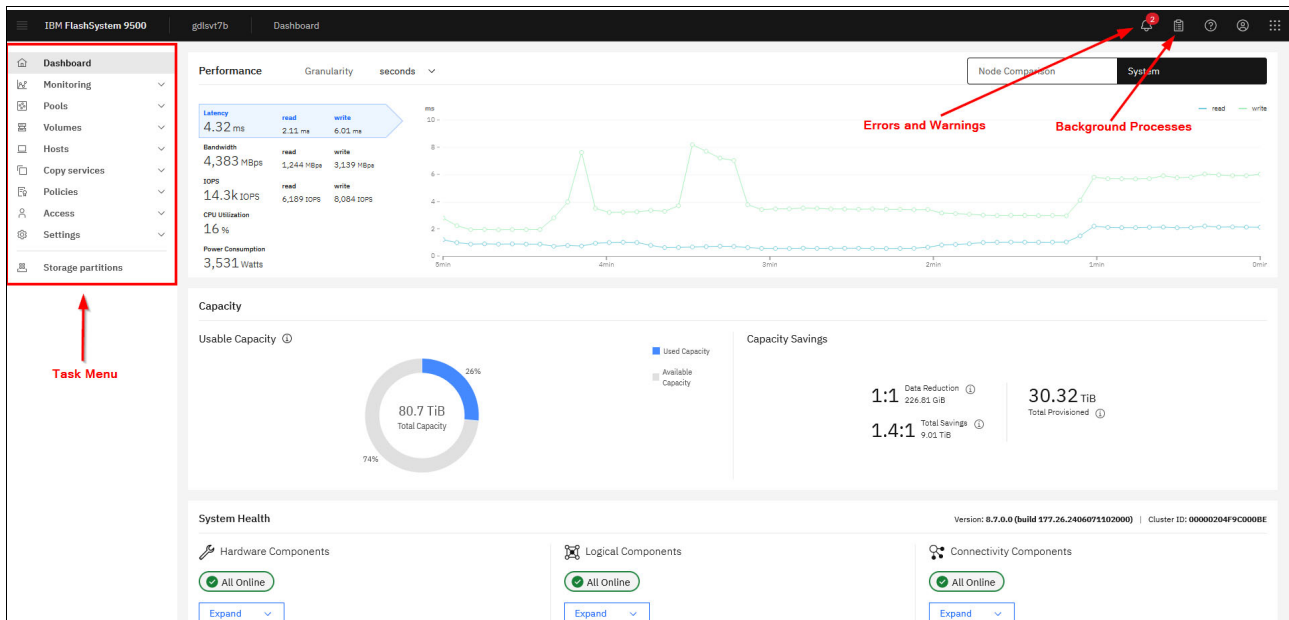


Figure 3-1 Welcome page with the dashboard

## Task menu

The IBM Storage Virtualize GUI task menu is always available on the left panel of the GUI window. To browse by using this menu, click the action and choose a task that you want to display.

## Notifications icons and help

Three notification icons are in the upper navigation area of the GUI.

- ▶ The left icon indicates warning and error alerts that were recorded in the Event log.
- ▶ The middle icon shows running jobs and suggested tasks.
- ▶ The third rightmost icon offers a help menu with content that is associated with the current tasks and the currently opened GUI menu.

## Performance

This section provides important information about latency, bandwidth, input/output operations per second (IOPS), and CPU usage. All this information can be viewed at the system or canister levels. A Node comparison view shows the differences in characteristics of each node. The performance graph is updated with new data every 5 seconds. The granularity of the metrics can be adjusted from seconds to days. For more detailed performance charts, select **Monitoring** → **Performance**.

## Capacity

This section shows the current usage of attached storage. It also shows provisioned capacity and capacity savings.

## System Health

This section indicates the status of all critical system components, which are grouped in three categories: Hardware Components, Logical Components, and Connectivity Components. When you click **Expand**, each component is listed as a subgroup. You can then go directly to the section of GUI where the component in which you are interested is managed.

## 3.2 Network configuration

The network configuration panel is accessed by selecting **Settings** → **Network**. Here you can configure or change configuration for the management IP, service IP, host attachment IPs, back-end storage IP addresses, replication IP addresses, priority flow control, iSCSI, DNS, internal proxy server, and portsets.

You can also view node to node Ethernet connectivity, fibre channel connectivity, NVMe connectivity, and fibre channel ports.

### 3.2.1 Management IP addresses

During the system initialization, one management IP address is set. A second management IP can be added. Starting in Storage Virtualize 8.7.0, the two management IPs can be configured on any ports and are not required to use ports 1 and 2. The management IPs are included in a default management portset. The management IPs are not tied to a single node. Whichever node is the config node is the node that responds to requests to the management IP.

**Note:** The system always uses the management IP on the lowest numbered port for outbound communication, for example, Cloud Call Home, email notifications, DNS lookup.

For more information on configuring ports in a FlashSystem storage unit, refer to the IBM Redpaper *The Definitive Guide to FlashSystem 5300 Port Configuration*, REDP-5734.

### 3.2.2 Service IP addresses

On each node, port id 1 is assigned a default service IP. The first node in an enclosure is assigned 192.168.70.121 and the second node in an enclosure is assigned 192.168.70.122. Change the default service IP addresses to addresses that are accessible on the network.

By connecting to a service IP address with a browser or SSH client, you can access the Service Assistant Interface, which can be used for maintenance and service tasks. The service IPs are also used for some system functions, for example to access a key server or IP quorum or for remote support assistance.

### 3.2.3 Additional Ethernet ports

The Ethernet ports menu is for configuring ports for host attachment, for replication, and for virtualizing back-end storage with iSCSI.

To configure an IP address on a port, select a port and go to **Actions** → **Manage IP addresses** (Figure 3-2).

			↑ Link State	Speed	Host Attach	Storage	Replication
			Active	1Gb/s	No	No	No
			Active	1Gb/s	No	No	No
node2	2	Inactive	No	No	No	No	
node1	2	Inactive	No	No	No	No	
node2	3	Active	Yes	Yes	Yes	Yes	
node1	3	Active	Yes	Yes	Yes	Yes	
node1	4	Active	Yes	Yes	Yes	Yes	
node2	4	Active	Yes	Yes	Yes	Yes	
node1	5	Active	Yes	Yes	Yes	Yes	
node2	5	Active	Yes	Yes	Yes	Yes	
node1	6	Active	Yes	Yes	Yes	Yes	
node2	6	Active	Yes	Yes	Yes	Yes	

Figure 3-2 Ethernet ports

On the next screen, select **Add IP address** to configure the IP address and add to a portset. See Figure 3-3 on page 65.

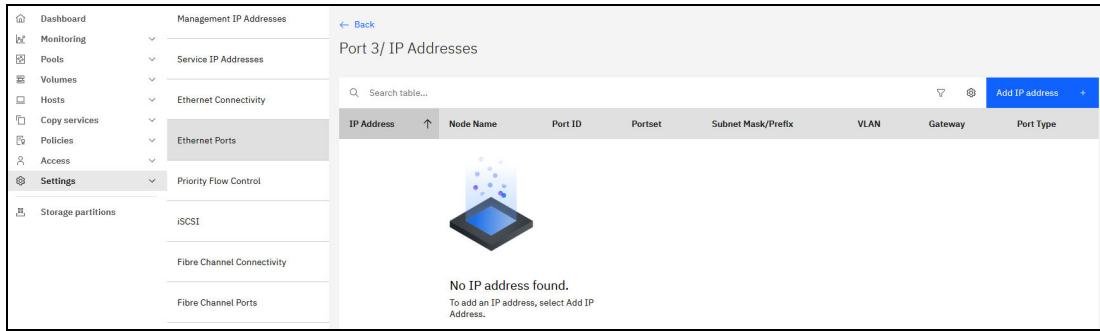


Figure 3-3 Add IP address

### 3.2.4 Portsets

Portsets are groupings of logical addresses that are associated with the specific traffic types. The system comes with one Fibre Channel and five Ethernet portsets defined. They are used for host attachment, system management, remote copy, and back-end storage virtualization. For more information, see [Portsets](#).

Figure 3-4 shows how to create or modify portsets.

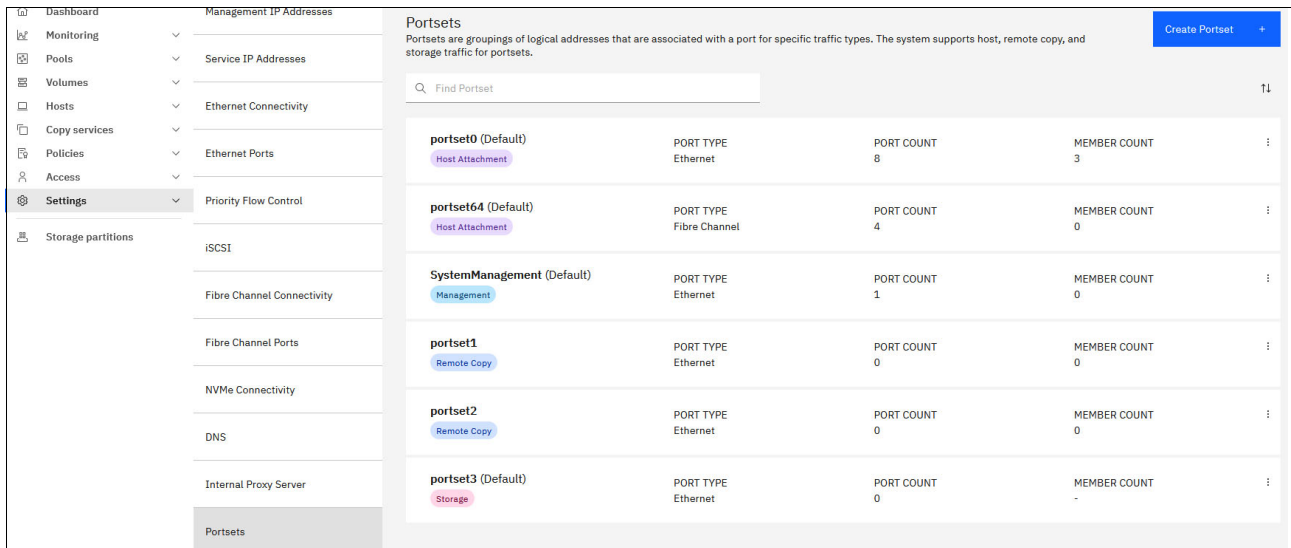


Figure 3-4 Portsets

## 3.3 Pools and managed disks configuration

This section describes how the storage system manages physical storage resources. All storage resources under system control are managed by using storage pools.

Storage pools aggregate internal and external capacity as managed disks (Mdisks) and provide containers in which you can create volumes that can be mapped to host systems. Storage pools help make it easier to dynamically allocate resources, maximize productivity, and reduce costs.

A storage pool is created as an empty container with no storage assigned to it. Storage is then added in the form of MDisks. MDisks can be a redundant array of independent disks

(RAID) arrays that are created by using internal storage, such as drives and flash modules, or logical units (LUs) that are provided by external storage systems. A single storage pool can contain both types of MDisks, but a single MDisk can be part of only one storage pool. MDisks are not visible to host systems.

Arrays are assigned to storage pools at creation time. Arrays cannot exist outside of a storage pool and they cannot be moved between storage pools. It is possible to delete an array by removing it from a pool and re-create it within a new pool.

External MDisks can exist within or outside of a pool. The MDisk object remains on a system if it is visible from external storage, but its access mode changes depending on whether it is assigned to a pool.

### 3.3.1 Provisioning policies

Consider using a provisioning policy for pools. If a provisioning policy is assigned, any volumes that are created from the pool are provisioned based on the capacity savings method that is defined in the policy. A policy can be created and assigned or unassigned to a pool at any time, but it effects only volumes that were created while the policy was active. To create a provisioning policy, select **Policies** → **Provisioning policies** → **Create policy**. For more information, see [Provisioning policy](#).

**Note:** Provisioning policy does not change any parameters of volumes that already exist in the pool when a policy is assigned. If you already have volumes in the pool, then after assigning a provisioning policy you might need to change volumes capacity savings settings manually.

### 3.3.2 Types of pools

The system supports standard pools and data reduction pools (DRPs). Both support parent pools and child pools.

Child pools are created from capacity that is assigned to a parent pool instead of created directly from MDisks. When a child pool is created from a standard pool, the capacity for a child pool is reserved from the parent pool. This capacity is no longer reported as available capacity of the parent pool. In terms of volume creation and management, child pools are similar to parent pools. Child pools that are created from DRPs are quota-less. Their capacity is not reserved but is shared with a parent pool.

DRPs use a set of techniques, such as compression and deduplication, that can reduce the required amount of usable capacity to store data. Data reduction can increase storage efficiency and performance, and reduce storage costs, especially for flash storage. These techniques can be used in addition to compression on Flash Core Modules (FCMs).

In standard pools, there can be no compression on a pool layer, but data is still compressed on the FCM layer if the pool contains drives with this technology. For more information, see [Pools](#).

### 3.3.3 Ransomware threat detection

Ransomware threat detection is automatically enabled if the following requirements are met:

- ▶ The pool must be created at Storage Virtualize code level 8.6.2 or higher.



- ▶ The pool consists of only FCM4 drives with firmware 4.1 or higher configured in a single DRAID6 array.
- ▶ Each node contains at least 128 GB RAM.
- ▶ Volumes are in a standard pool or fully allocated within a DRP.

### 3.3.4 Creating storage pools

If you want to create an encrypted pool, the encryption license must be installed and encryption enabled before creating the pool. A pool cannot be changed to encrypted after creation. For more information, see [Encryption documentation](#).

To create a storage pool, complete the following steps:

1. Select **Pools** → **MDisks by Pools** and click **Create Pool** or select **Pools** → **Pools** and click **Create** → **Create Pool**. Figure 3-5 shows the Create Pool menu.

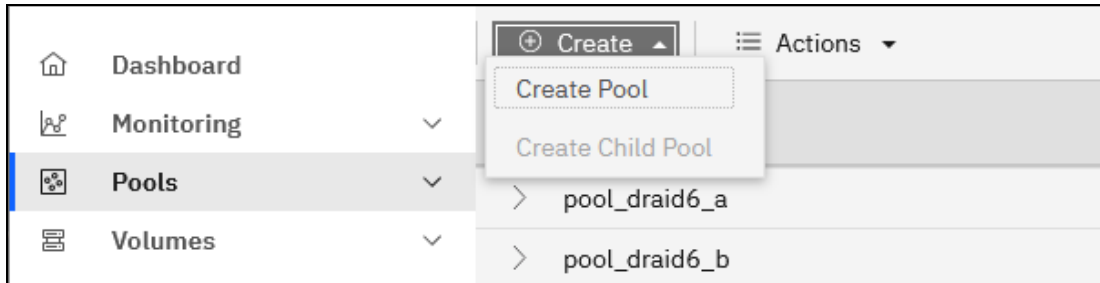


Figure 3-5 Create Pool

Both alternatives open the dialog box that is shown in Figure 3-6.




Figure 3-6 Create Pool panel

2. Select the **Data reduction** option if you want to create a DRP. Leaving it clear creates a standard storage pool.

**Note:** Limitations, capacity requirements, and performance characteristics of DRPs are different from standard pools. Verify with your system architect or IBM representative that your system was sized to be used with DRP and with its reduction features before creating a DRP.

The size of the extents is selected at creation time and cannot be changed later. The extent size controls the maximum total storage capacity that is manageable per system (across all pools). For DRPs, the extent size also controls the maximum pool stored capacity per IO group. For more information, see [V8.7.0.x Configuration Limits for IBM FlashSystem and SAN Volume Controller](#).

**Important:** Do not create DRPs with small extent sizes. For more information, see this [IBM Support alert](#).

If an encryption license is installed and enabled, you can select whether the storage pool is encrypted. The encryption setting of a storage pool is selected at creation time and cannot be changed later. By default, if encryption is licensed and enabled, the encryption check-box is selected.

Enter the name for the pool and click **Create**.

**Naming rules:** When you choose a name for a pool, the following rules apply:

- ▶ Names must begin with a letter.
- ▶ The first character cannot be numerical.
- ▶ The name can be a maximum of 63 characters.
- ▶ Valid characters are uppercase letters (A - Z), lowercase letters (a - z), digits (0 - 9), underscore (\_), period (.), hyphen (-), and space.
- ▶ Names must not begin or end with a space.
- ▶ Object names must be unique within the object type. For example, you can have a volume that is named ABC and a storage pool that is called ABC, but not two storage pools that are both called ABC.
- ▶ The default object name is valid (object prefix with an integer).
- ▶ Objects can be renamed at a later stage.

The new pool is created and is included in the list of storage pools. It has no storage in it, so its capacity is zero. Storage in a form of disk arrays or externally-virtualized MDisk must be assigned to the pool before volumes can be created.

### 3.3.5 Creating RAID array managed disks in a storage pool

To create a RAID array and assign it to a pool, select **Pools** → **Pools** → **Select an already created pool** → **Actions** → **Add Storage**. The Add Storage menu is shown in Figure 3-7 on page 69.

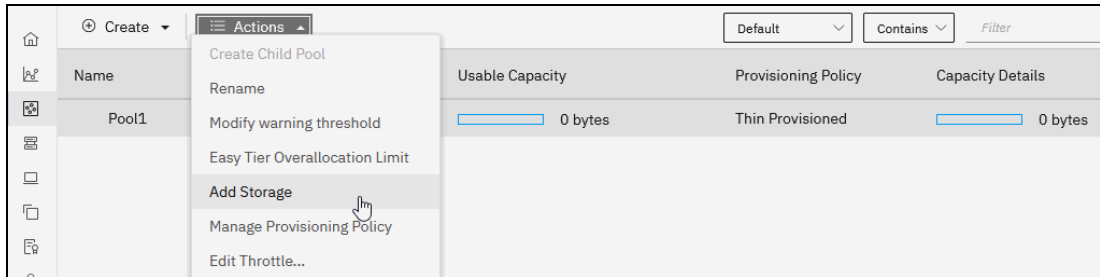


Figure 3-7 Add Storage

This opens a new pane with a suggested RAID array configuration based on the installed drives. See Figure 3-8

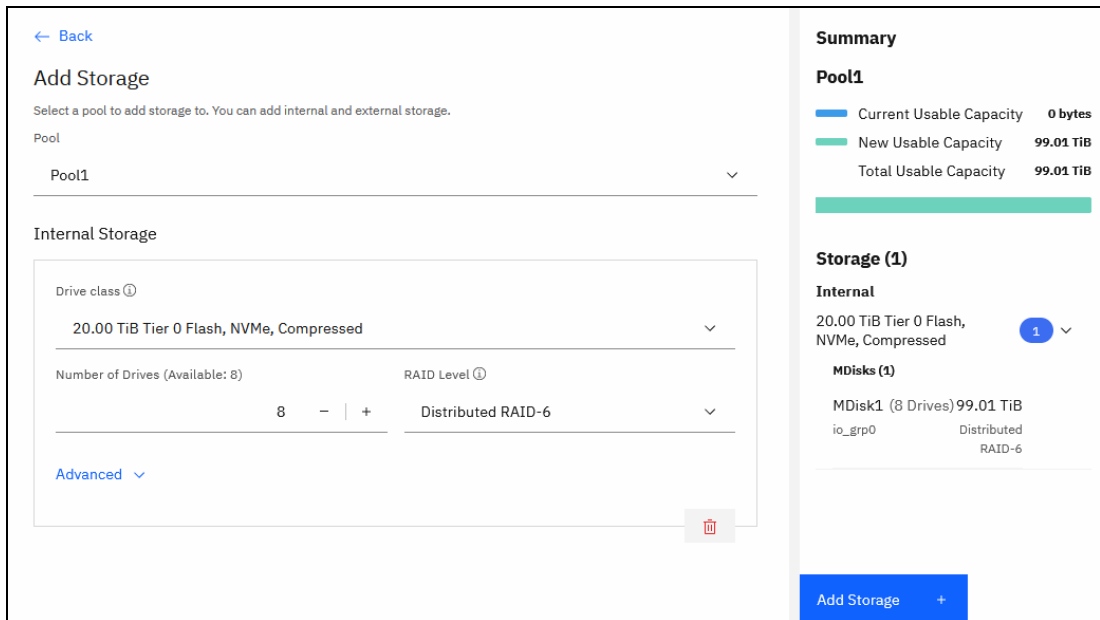


Figure 3-8 RAID array

### 3.3.6 Adding external managed disks into a storage pool

Controllers are external storage systems that provide storage resources that are used as MDisks. The system supports external storage controllers that are attached through internet Small Computer Systems Interface (iSCSI) and through Fibre Channel (FC).

A key feature of the system is its ability to consolidate disk controllers from various vendors into storage pools. The storage administrator can manage and provision storage to applications from a single user interface and use a common set of advanced functions across all of the storage systems under the control of the system.

This concept is called *External Virtualization*, which can make your storage environment more flexible, more cost-effective, and easier to manage.

#### System layers

The system layer affects how the system interacts with a system and other external systems that run IBM Storage Virtualize software. To virtualize another system using Storage

Virtualize software, one system must be in the replication layer and one system must be in the storage layer. For more information, see [System layers documentation](#).

### External storage systems

IBM Storage Virtualize based systems support a wide range of storage controllers. They can be attached through Fibre channel or iSCSI. To check the compatibility of a system use the [IBM System Storage Interoperation Center \(SSIC\)](#).

For detailed instructions on configuring an external storage system, review the [External storage documentation](#).

When external LUs are discovered by the IBM Storage Virtualize system, they are visible in **Pools** → **MDisks** by pools under Unassigned MDisks. Select the MDisks that are to be added to a pool and select **Actions** → **Assign**.

When you add MDisks to pools, you must assign them to the correct storage tiers. It is important to set the tiers correctly if you plan to use the IBM Easy Tier® feature. The use of an incorrect tier can mean that the Easy Tier algorithm might make wrong decisions and thus affect system performance.

The storage tier setting can also be changed after the MDisk is assigned to the pool. For more information, see [Easy tier](#).

### 3.3.7 Child pools

A *child pool* is a storage pool that is created within another storage pool. The storage pool in which the child storage pool is created is called the *parent storage pool*. Unlike a parent pool, a child pool does not contain MDisks. Its capacity is provided by the parent pool.

A child pool cannot be created within another child pool. Multiple child pools can be created within a single parent pool.

Multiple child pools can be created from a single parent pool for different uses. Each child pool can use a different provisioning policy. Child pools can also be linked to a remote pool for policy-based replication. See Figure 3-9.

Name	Properties	State	Provisioning Policy	Capacity Details
Pool0		Online	-	30.00 TiB / 59.22 TiB (51%)
Child_capacity		Online	capacity_optimized	0 bytes / 10.00 TiB (0%)
Child_performance		Online	performance_optimized	0 bytes / 15.00 TiB (0%)
Child_PBR		Online	-	0 bytes / 5.00 TiB (0%)

Figure 3-9 Child pools with different purposes

Child pools created from standard pools and child pools that are created from data reduction pools have a significant difference:

- ▶ A child pool with a standard pool as a parent has a type child\_thick. Child pools of Standard pools have a fixed capacity, which is taken, or reserved, from the parent pool. Free capacity of a parent pool reduces when a child pool is created. Volumes in a child pool of a standard pool cannot occupy more capacity that is assigned to the child.

- ▶ A child pool with DRP as a parent, has type `child_quotaless`. Quotaless child pools share its free and used capacity with the parent pool and do not have their own capacity limit. Free capacity of a DRP does not change when a new quotaless child pool is created.

The capacity of a `child_thick` type pool is set at creation time, but can be modified later nondisruptively. The capacity must be a multiple of the parent pool extent size and must be smaller than the free capacity of the parent pool.

Child pools of a `child_thick` type can be used to implement the following configurations:

- ▶ Limit the capacity that is allocated to a specific set of volumes
 

It can also be useful when strict control over thin-provisioned volume expansion is needed. For example, you might create a child pool with no volumes in it to act as an emergency set of extents so that if the parent pool uses all its free extents, you can use the ones from the child pool.
- ▶ As a container for VMware vSphere virtual volumes (VVOLs)
 

Data reduction pools are *not* supported as parent pools for VVOL storage.
- ▶ Migrate volumes from nonencrypted parent storage pool to encrypted child pools
 

When you create a child pool of type `child_thick` after encryption is enabled, an encryption key is created for the child pool, even when the parent pool is not encrypted. You can then use volume mirroring to migrate the volumes from the nonencrypted parent pool to the encrypted child pool.

Encrypted `child_quotaless` type child pools can be created only if the parent pool is encrypted. The data reduction child pool inherits an encryption key from the parent pool.

## Creating a child storage pool

To create a child pool, complete the following steps:

Select **Pools** → **Pools**. Right-click the parent pool that you want to create a child pool from and select **Create Child Pool**. The Create Child Pool pane opens. See Figure 3-10.

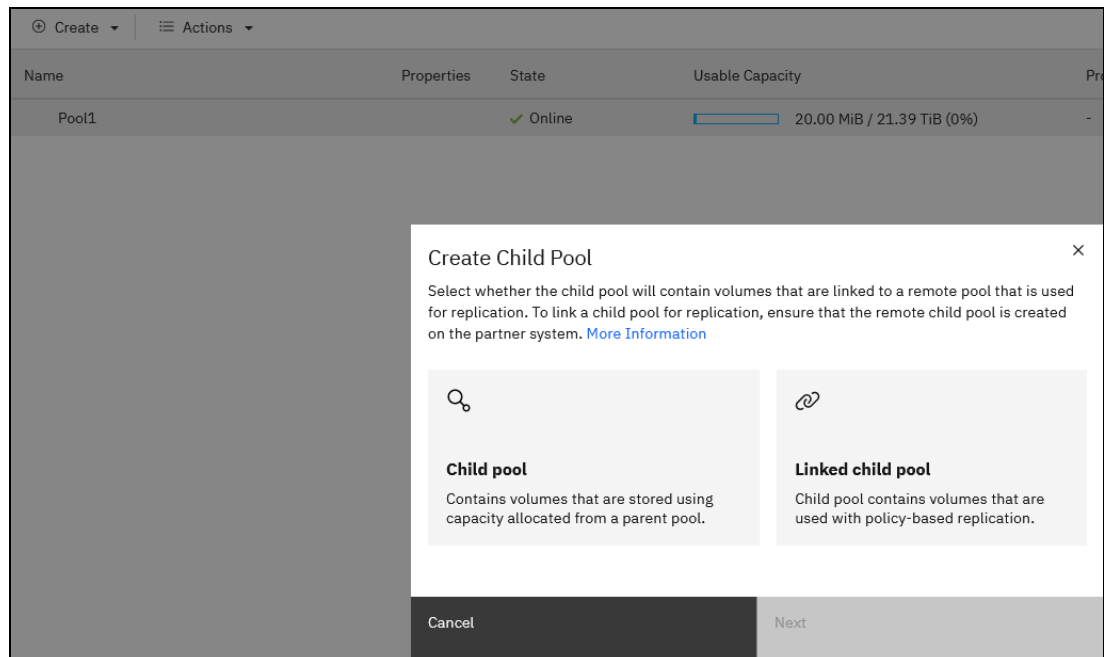


Figure 3-10 Create Child Pool

## 3.4 Configuring volumes

In IBM Storage Virtualize, a volume is storage space that is provisioned out of a storage pool and presented to a host as a logical unit (LU), which is also called a logical disk.

This section describes how to create and provision volumes on IBM Storage Virtualize systems. For more information on volumes and the various volume types, see [Volumes](#).

### 3.4.1 Creating volume groups

Volume groups refer to a collection or grouping of volumes that share common characteristics or are organized together for specific reasons. A volume group is a group of volumes. Volume groups can be formed based on various factors:

- ▶ Volumes with similar SLA requirements, such as performance targets, availability, or data protection policies can be grouped within a volume group. This helps ensure that the volumes within the group are managed and treated according to the same service level guidelines.
- ▶ Some applications or data sets might require mutual consistency among multiple volumes. In such cases, these volumes can be grouped within a volume group to ensure that they are synchronized and maintain consistency in terms of data updates or access.
- ▶ Volumes residing on the same server can be grouped within a volume group. This grouping facilitates efficient management and administration of the volumes within the server environment, which allows for streamlined operations and centralized control.

It is important to note that volume groups are distinct from consistency groups. Although in some cases, the underlying system might use a consistency group concept internally when managing volume groups

Volume groups are used with the following functions:

- ▶ Safeguarded Copy function
- ▶ Policy-based replication
- ▶ Snapshot function

To create a volume group select **Volumes** → **Volume groups** → **Create Volume Group**.

**Note:** If you plan to use policy-based replication, then if you configure it and assign a replication policy to the volume group before creating volumes within the group, you can skip the initial copy of replicated data to the remote site.

### 3.4.2 Creating volumes

If a volume group was created, select the volume group and select **Actions** → **Create New Volumes**. If you are creating volumes outside a volume group, select **Volumes** → **Volumes** → **Create Volumes**.

On the next panel, you define the volume properties. However, if the pool has a pre-assigned provisioning policy, the capacity savings option is locked and reflects the policy's settings.

There is a toggle on the screen for Advanced settings mode, which allows manual selection of I/O group and preferred node parameters.

The newly created volumes automatically start formatting. This is a background process and the volume is immediately available for host access. The default format speed is 2 MiB/s per volume. If you want to increase the format rate for a volume, right-click the volume and select **Modify Mirror Sync Rate**. Then choose the preferred rate.

It is possible to overuse the system's resources by formatting too many volumes at too high a rate. If you experience a system performance problem after you increase the mirror sync rate, you can reduce it in the same manner.

**Note:** If you are using a system with IBM FlashCore® Modules, the data that is written to the system is compressed automatically. There is no requirement to also create the volumes as compressed.

### 3.4.3 Virtual volumes

The system provides native support for VMware vSphere APIs for Storage Awareness (VASA) through a VASA Provider (also known as a Storage Provider). The VASA provider sends and receives information about storage that is used by VMware vSphere to the vCenter Server. Through VASA, the system also supports VMware Virtual Volumes (vVols), which allows VMware vCenter to automate the creation, deletion and mapping of volumes.

For more information about configuring vVols with IBM Storage Virtualize, see *IBM Storage Virtualize and VMware: Integrations, Implementation and Best Practices*, SG24-8549.

Also see [VMWare virtual volumes \(vVols\)](#).

## 3.5 Configuring hosts

A host system can be defined as any networked computer server, virtual, or physical that provides workloads and services to the storage.

This section describes the processes that are required to attach a supported host system to an IBM Storage Virtualize storage system through various supported interconnect protocols.

For more information, see [Hosts](#).

### 3.5.1 Host attachment overview

IBM Storage Virtualize family supports various open system host types from IBM and non-IBM vendors.

These hosts can connect to the storage systems through any of the following protocols:

- ▶ Fibre Channel Protocol (FCP)
- ▶ Fibre Channel over Ethernet (FCoE)
- ▶ iSCSI
- ▶ SAS
- ▶ iSCSI Extensions for Remote Direct Memory Access (RDMA) (iSER)
- ▶ Non-Volatile Memory Express (NVMe) over Fibre Channel (FC-NVMe)
- ▶ NVMe over Remote Direct Memory Access (NVMe over RDMA)
- ▶ NVMe over Transmission Control Protocol (NVMe over TCP)

**Note:** Specific host operating systems can be connected directly to the IBM Storage Virtualize storage system without the use of SAN switches. For more information, see the [IBM System Storage Interoperation Center \(SSIC\)](#).

To enable multiple access paths and enable correct volume presentation, a host system must have a multipathing driver installed.

In addition, the multipathing driver serves the following purposes:

- ▶ Protection from:
  - Fabric path failures, including port failures on IBM Storage Virtualize system nodes
  - A host bus adapter (HBA) failure (if two HBAs are used)
  - Failures if the host is connected through two HBAs across two separate fabrics
- ▶ Load balancing across the host HBAs.

For more information about the native operating system multipath drivers that are supported for IBM Storage Virtualize systems, see the [SSIC](#).

For more information about how to attach specific supported host operating systems to the storage systems, see [Host attachment](#).

**Note:** If a specific host operating system is not mentioned in the SSIC, contact your IBM representative or IBM Business Partner to submit a special request for support.

## 3.5.2 Fibre channel host connectivity

This section includes a discussion of Fibre channel host connectivity.

### N\_Port ID Virtualization

IBM Storage Virtualize systems use N\_Port ID Virtualization (NPIV), which is a method for virtualizing a physical FC port that is used for host I/O.

NPIV mode creates a virtual worldwide port name (WWPN) for every physical system FC port. This WWPN is available for host connection only. During node maintenance, restart, or failure, the virtual WWPN from that node is transferred to the same port of the other node in the I/O group.

Ensure that the FC switches support the ability to create four more NPIV ports on each physically connected system port.

When performing zoning configuration, virtual WWPNs are used for host communication only. That is, system-to-host zones must include virtual WWPNs. Internode, intersystem, and back-end storage zones must use the WWPNs of physical ports. Ensure that equivalent ports with the same port ID are on the same fabric and in the same zone.

**Important:** IBM i Systems that are attached to FlashSystem or SVC must be converted to use FlashSystem NPIV before upgrading to 8.7 or higher. For FlashSystem or SVC systems that have NPIV in a state of disabled or transitional and that have any IBM i hosts, a modified procedure must be used when enabling NPIV, to avoid loss of host access to data. For more information, see [IBM i Systems attached to FlashSystem or SVC must be converted to use FlashSystem NPIV before upgrading to 8.7. or higher](#).



To view the virtual WWPNs to be used in system to host select **Settings** → **Network** → **Fibre Channel Ports**. Expand the section for each port. Columns indicate WWPN, Host IO Permitted, and Protocol type. SCSI is for Fibre Channel Protocol (FCP).

**Note:** The NPIV WWPNs do not become active until there is at least one online volume.

### Host zones

A host must be zoned to an I/O group to access volumes that are presented by this I/O group.

The preferred zoning policy is *single initiator zoning*. To implement it, create a separate zone for each host bus adapter (HBA) port, and place one port from each node in each I/O group that the host accesses in this zone. A typical Fibre Channel host has two ports zoned to each IO group, which creates a total of four paths. For deployments with more than 64 hosts that are defined in the system, this host zoning scheme must be used.

**Note:** Cisco Smart Zoning and Brocade Peer Zoning are supported. You can use either to insert target ports and multiple initiator ports in a single zone for ease of management. However, either acts as though each initiator and target are configured in isolated zones. The use of these zoning techniques is supported for host attachment and storage virtualization. As a best practice, use normal zones when configuring ports for clustering or for replication because these functions require the port to be an initiator and a target.

Consider the following rules for zoning hosts over SCSI or FC-NVMe:

- ▶ For any volume, the number of paths through the SAN from the host to a system must not exceed eight. For most configurations, four paths to an I/O group are sufficient.
- ▶ Balance the host load across the system's ports. For example, zone the first host with ports 1 and 3 of each node in the I/O group, zone the second host with ports 2 and 4, and so on. To obtain the best overall performance of the system, the load of each port must be equal. Assuming that a similar load is generated by each host, you can achieve this balance by zoning approximately the same number of host ports to each port.
- ▶ Spread the load across all system ports. Use all ports that are available on your machine.
- ▶ Balance the host load across HBA ports. If the host has more than one HBA port per fabric, zone each host port with a separate group of system ports.

All paths must be managed by the multipath driver on the host side. Make sure that the multipath driver on each server can handle the number of paths that is required to access all volumes that are mapped to the host.

### 3.5.3 Ethernet host connectivity

You can attach your IBM Storage Virtualize system to iSCSI, iSER, NVMe over RDMA, and NVMe over TCP hosts by using the Ethernet ports of the system.

The same ports can be used for iSCSI and iSER host attachment concurrently. However, a single host can establish an iSCSI or session, but not both

Hosts connect to the system through IP addresses, which are assigned to the Ethernet ports of the node. If the node fails, the address becomes unavailable and the host loses communication with the system through that node.

To allow hosts to maintain access to data, the node-port IP addresses for the failed node are transferred to the partner node in the I/O group. The partner node handles requests for its

own node-port IP addresses and for node-port IP addresses on the failed node. This process is known as *node-port IP failover*.

In addition to node-port IP addresses, the iSCSI name and iSCSI alias for the failed node are transferred to the partner node. After the failed node recovers, the node-port IP address and the iSCSI name and alias are returned to the original node.

## iSCSI

iSCSI is a protocol that uses the Transmission Control Protocol and Internet Protocol (TCP/IP) to encapsulate and send SCSI commands to storage devices that are connected to a network. iSCSI is used to deliver SCSI commands from a client interface, which is called an iSCSI Initiator, to the server interface, which is known as the iSCSI Target. The iSCSI payload contains the SCSI CDB and optionally, data. The target carries out the SCSI commands and sends the response back to the initiator.

## NVMe over Remote Direct Memory Access

IBM Storage Virtualize can be attached to an NVMe host through NVMe over Remote Direct Memory Access (RDMA). NVMe over RDMA uses RDMA over Converged Ethernet (RoCE) v2 as the transport protocol. RoCE v2 is based on user datagram protocol (UDP).

RDMA is a host-offload, host-bypass technology that allows an application (including storage) to make data transfers directly to and from another application's memory space. The RDMA-capable Ethernet NICs (RNICs), and not the host, manage reliable data transfers between source and destination.

RNICs can use RDMA over Ethernet by way of RoCE encapsulation. RoCE wraps standard InfiniBand payloads with Ethernet or IP over Ethernet frames, and is sometimes called *InfiniBand over Ethernet*. The following main RoCE encapsulation types are available:

- ▶ RoCE V1

This type uses dedicated Ethernet Protocol Encapsulation (Ethernet packets between source and destination MAC addresses by using EtherType 0x8915).

- ▶ RoCE V2

This type uses dedicated UDP over Ethernet Protocol Encapsulation, IP UDP packets that use port 4791 between source and destination IP addresses. UDP packets are sent over Ethernet by using source and destination MAC addresses. This type is *not* compatible with other Ethernet options, such as RoCE v1.

**Note:** Unlike RoCE V1, RoCE V2 is routable.

## NVMe over TCP

IBM Storage Virtualize can be attached to an NVMe host through NVMe over Transmission Control Protocol (TCP). NVMe over TCP is a transport that allows NVMe performance without any constraint to the data center infrastructure.

NVMe over TCP needs more CPU resources than protocols using RDMA. Each NVMe/TCP port on FlashSystem supports multiple IP addresses and multiple VLANs. Generally, NVMe-TCP runs on all switches and is routable.

For operating system support and multipathing, see [IBM System Storage Interoperation Center \(SSIC\)](#).

## 3.5.4 Host objects

Before a host can access the storage capacity, it must be presented to the storage system as a *host object*.

A host object is configured by using the GUI or command-line interface (CLI) and must contain the necessary credentials for host-to-storage communications. After this process is completed, storage capacity can be mapped to that host in the form of a volume.

IBM Storage Virtualize supports configuring the following host objects:

- ▶ Host
- ▶ Host cluster

A host cluster object groups clustered servers and treats them as a single entity. This configuration allows multiple hosts to access the same volumes through one shared mapping.

**Note:** Any volume that is mapped to a host cluster is automatically assigned to all of the members in that cluster with the same SCSI ID.

A typical use case for a host cluster object is to group multiple clustered servers with a common operating system, such as IBM PowerHA® and Microsoft Cluster Server, and enable them to have shared access to common volumes.

To create a host object select **Hosts** → **Hosts** → **Add Host**. The Add Host page opens. See Figure 3-11 on page 78.

**Tip:** The Host port drop-down menu shows FCP initiator WWPNs that are currently logged in to the system. If an expected WWPN is missing, examine switch zoning and rescan the storage from the hosts. Some operating systems log out if no LUNs are mapped to the host. If an expected host is not listed, then select **Enter Unverified WWPN** and enter the host WWPNs manually.

## 3.5.5 Mapping volumes for host access

For the host or host cluster to access the volumes, the volumes must be mapped to the host or host cluster. To perform the mappings, select **Volumes that you wish to map** → **Actions** → **Map to Host or Host Cluster** → **Select the host or host cluster** → **Next**. In most cases, leave the radio button for the system to assign SCSI LU IDs.

**Note:** Usually, all volumes within a volume group are mapped to the same host or host cluster, and the mapping can be done within the volume group view.

Figure 3-11 Add Host

## 3.6 Snapshots and replication

You can create copies of data by using snapshots and replication.

### 3.6.1 Volume group snapshots

Before the introduction of volume group snapshots, when volumes were dependent on each other, creating point-in-time copies required taking snapshots at the same time. This was achieved by using consistency groups (CG), which consisted of a group of mappings that had to be started simultaneously. Configuring the mappings and targets for Flashcopies within a consistency group was a complex process that had to be repeated each time a new point-in-time copy was created. This complexity posed limitations on usage and made it challenging to incorporate new functions.

The purpose of the volume group snapshot management model is to simplify the implementation of standard IBM FlashCopy® operations. It achieves this by offering a more straightforward setup process and separating the snapshot and clone features. By using volume group snapshots, administrators can create snapshots of volume groups with more ease and efficiency, without the need for complex consistency group configurations.

Snapshots cannot be mapped to a host. To access the data on a snapshot, create a thin-clone of the snapshot and map it to a host.

**Demonstration videos:** The following demonstration videos are available:

- ▶ [IBM Storage Virtualize V8.6: Handling Snapshots using the GUI.](#)
- ▶ [IBM Storage Virtualize V8.6: Handling snapshots using the command line interface.](#)

For more information, see [Snapshots](#).

## Triggering volume group snapshots

The process for triggering a volume group snapshot involves a streamlined version of the FlashCopy mapping trigger process. They can be triggered by using the GUI, on a schedule by using a snapshot policy, or by an external application such as Copy Services Manager (CSM).

## Volume group snapshot policy

A volume group snapshot policy is designed to automate the creation and deletion of snapshots based on predefined schedules, which eliminates the need for external applications:

- ▶ Users can choose from predefined snapshot policies or create custom policies that are tailored to their specific needs.
- ▶ The snapshot policies that are configured with the volume group snapshot scheduler are reusable and can be applied to multiple volume groups as needed.
- ▶ The volume group snapshot scheduler comes with default snapshot policy parameters, offering convenient options for most use cases.
- ▶ Users can specify the creation frequency of snapshots in minutes, hours, weeks, days, or months. The minimum creation frequency allowed is 60 minutes. Also, the retention of snapshots can be specified in terms of days.
- ▶ When a snapshot policy is assigned to a volume group, there is an option to select Safeguarded. Safeguarded snapshots can be deleted before their expiration time by a security administrator only.

**Note:** When using a snapshot policy, after the initial snapshot, snapshots are triggered based on the frequency defined. This means that the time of day the snapshot is triggered might shift forward and backward with Daylight Saving Time changes.

To create, view, or assign a snapshot policy select **Policies** → **Snapshot policies**. See Figure 3-12 on page 80.

**Snapshot Policies**  
Create policies and assign them to volume groups to automate the creation and retention of your snapshots.

Search table... Filter Settings [Create snapshot policy](#)

Name	Target	Frequency	Retention	Volume group count
predefinedsspolicy0	Local	Every 6 hours	7 Days	0
predefinedsspolicy1	Local	Every week on Saturday at 11:00 PM	30 Days	0
predefinedsspolicy2	Local	Every month on the 2nd at 11:00 PM	365 Days	0
predefinedsspolicy38	Cloud	Every day at 11:00 PM	30 Days	0
predefinedsspolicy39	Local and cloud	Multiple	Multiple	0

Figure 3-12 Snapshot policies

You can also suspend or unassign a policy from within the volume group. See Figure 3-13.

Volume groups / VG1

**VG1** 500.00 GiB Total Group Capacity

[Snapshot Policy](#)

Volumes (5) **Policies (1)** Local snapshots (0) Cloud Snapshots

Policy-based replication is not configured.

[Configure policy-based replication](#) +

**Snapshot Policy**

Name: predefinedsspolicy0 Target: Local

Local schedule: Every 6 hours, retained for 7 days

Next at: 2024-6-13 11:00 AM

[Suspend Policy](#)

[Unassign Policy](#)

Figure 3-13 Suspend Policy

### 3.6.2 Asynchronous policy-based replication

Asynchronous policy-based replication provides a variable, greater-than-zero recovery point that aims to achieve the best possible recovery point for the current conditions. This type of replication helps ensure mutual consistency between all volumes in the volume group.

To configure policy-based replication between two systems, both require at least one IP address that is created and assigned to a replication portset with at least one pool with storage created. Multiple IP addresses can be added to a replication portset. If there is a second independent inter-site link between the systems, a second portset can be used and added to the partnership.

1. On the primary system, select **Copy Services** → **Partnerships** → **Create Partnership**. If using IP, select **IP** and **enter the partner IP address** then select **Test Connection**. If the partner meets requirements for policy-based replication the **Use policy-based**

**replication** checkbox can be selected. Enter the requested information and select **Create**. Repeat these steps on the partner system. See Figure 3-14.

Create Partnership

Create a partnership to establish a connection to a remote system for replication.

Type

Fibre Channel

IP (long distances using TCP)

IP (short distances using RDMA)

Partner IP Address or Domain

Test Connection ✓

Secured IP partnerships

Use policy-based replication

**i** IP Partnerships cannot use high-availability  
Only Fibre Channel partnerships support high-availability

**View certificate**

**!** The remote system is using a CA-signed certificate. Review the certificate to ensure that it matches what you expect

Certificate from [IP Address]

[See details](#)

**Link specification**

Link bandwidth is available between systems, in megabits per second (Mbps)

Cancel Create

Figure 3-14 Create Partnership

2. When the partnership shows a green dot and configured select **Setup policy-based replication**. See Figure 3-15 on page 82 and Figure 3-16 on page 82

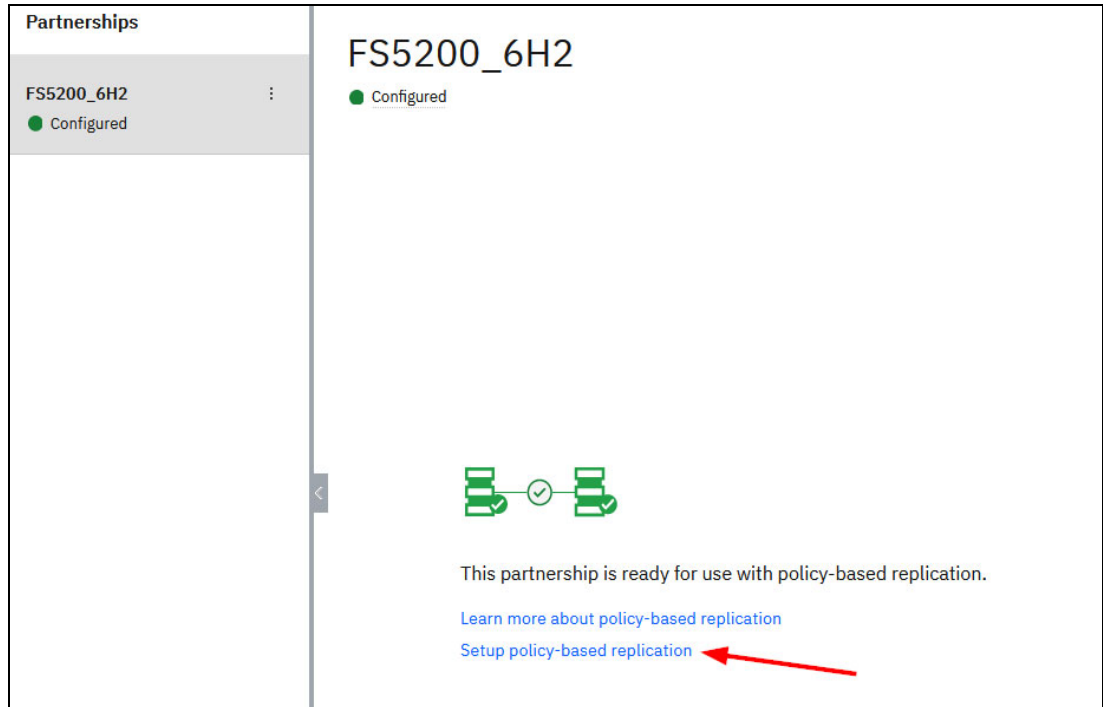


Figure 3-15 Partnership Created

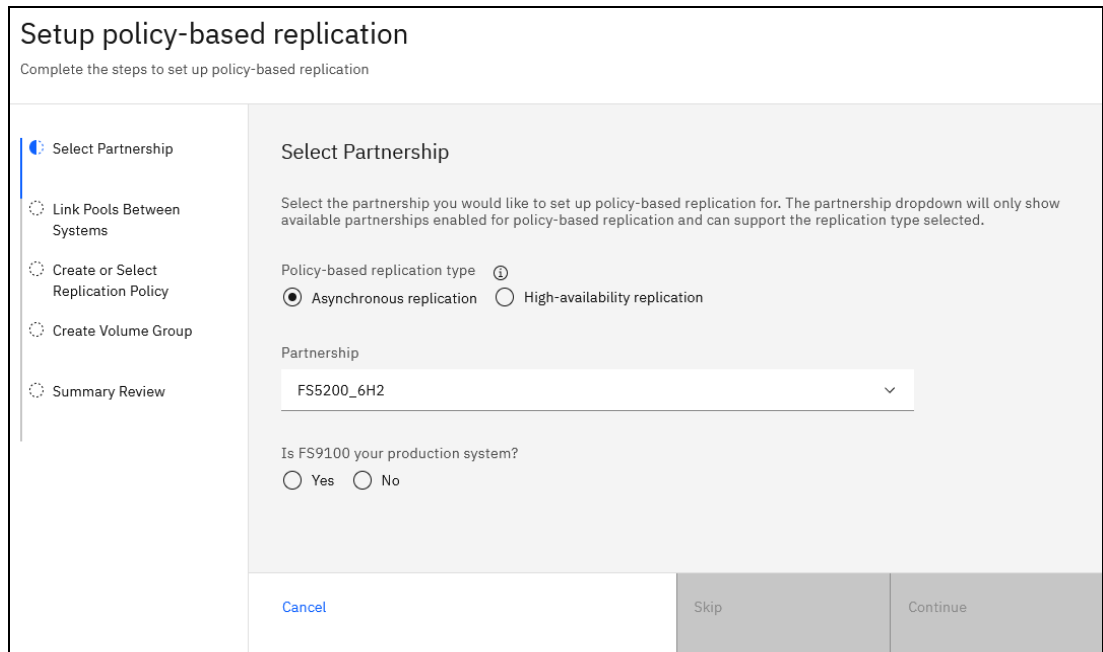


Figure 3-16 Setup policy-based replication wizard

For more information, see IBM Documentation [Asynchronous disaster recovery replication](#) and see *Policy-Based Replication with IBM Storage FlashSystem, IBM SAN Volume Controller and IBM Storage Virtualize*, REDP-5704.

Synchronous replication is provided by policy-based high availability, which is beyond the scope of this document. For more information, see [High Availability](#).





# Verifying configuration and basic operations

This chapter provides steps to verify the configuration and includes tips to resolve common implementation problems. It also provides additional information about functions to consider when you implement a new system.

This chapter includes discussions of the system health dashboard, verifying configuration of objects configured in Chapter 3, “Step-by-step configuration” on page 61, system security, getting support from IBM, and data migration.

This chapter has the following sections:

- ▶ 4.1, “Verifying the configuration” on page 84
- ▶ 4.2, “Additional settings and basic operations” on page 87

## 4.1 Verifying the configuration

Use the various GUI pages to verify health, network storage, volume, and host configurations.

### 4.1.1 System Health Dashboard

The system dashboard provides a way to quickly assess the overall condition of the system and to view notifications of any critical issues that require immediate action. The bottom third of the dashboard provides system health details including tiles for hardware components, logical components, and connectivity components. Each tile has a link to access relevant information. See Figure 4-1.

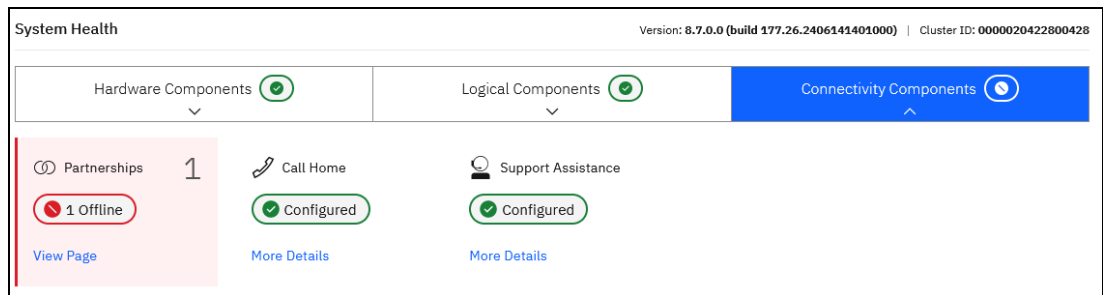


Figure 4-1 System Health

### 4.1.2 Verifying network configuration

You can verify the network IP addresses and WWPNs by selecting **Settings** → **Network** → **Portsets**. Select any of the portsets to view details of the ports and IP addresses assigned to that portset. See Figure 4-2.



Figure 4-2 Portset Mappings

Select **IP Addresses** or **Partnerships** for additional details. See Figure 4-3

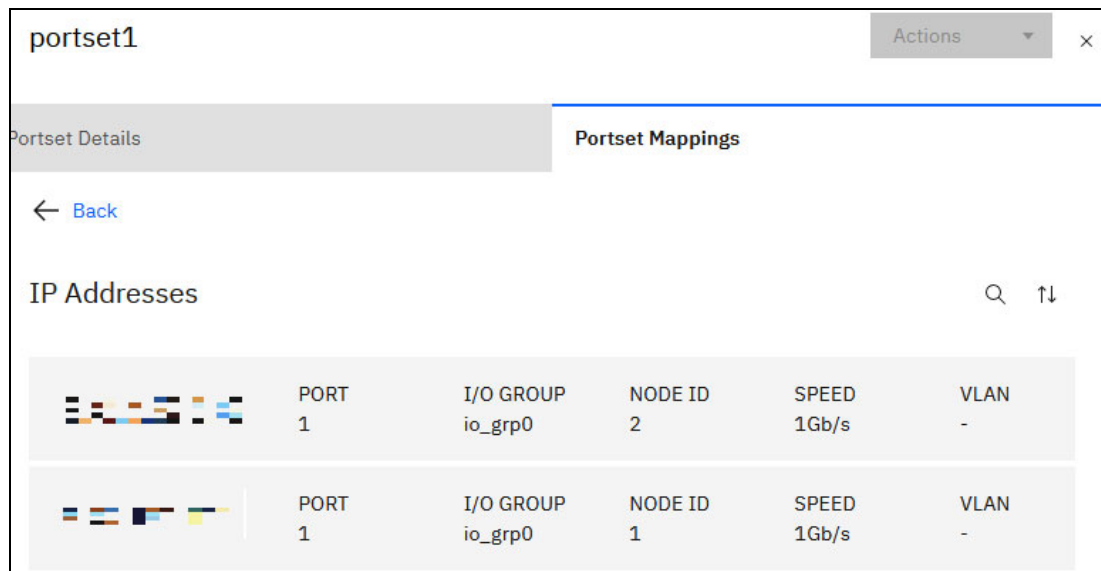


Figure 4-3 Portset IP addresses

### 4.1.3 Verifying storage configuration

Select **Pools** → **MDisks** by pools to view the storage configuration. All pools and Mdisks are expected to be in an online state. If any are degraded, then view **Monitoring** → **Events** for any related events.

There is a column for Usable Capacity and a column for Written Capacity Limit. If you are using over-provisioned storage, for example, FCMs the values are different. The usable capacity represents the physical capacity available after the data is reduced because of compression and deduplication. The Written Capacity Limit is the effective capacity of data that is written to the system before its size is reduced.

### 4.1.4 Verifying volume configuration

Select **Volumes** → **Volumes** to view the volume configuration. All volumes are expected to be online. If any are degraded, check **Monitoring** → **Events** for any related events.

You can change which columns are displayed by right-clicking the column titles bar. Some useful capacity related columns can be displayed.

If a volume is thin provisioned in a standard pool, adding the columns Real Capacity and Used Capacity can provide useful information. Used capacity is the capacity used by the data written to the volume. Real Capacity is the Used Capacity plus a contingency capacity that is used for new writes. These values are effective capacity.

Adding the column Compression Savings lists information on how compressible the data is.

#### Capacity Savings report

*IBM Comprestimator* is a utility that estimates the capacity savings that can be achieved when compression is used for storage volumes. The utility is integrated into the system and results can be viewed by using the GUI and the CLI. The integrated Comprestimator is always

enabled and running continuously to provide up-to-date compression estimation over the entire cluster, both in GUI and IBM Storage Insights.

IBM Comprestimator provides a quick and accurate estimation of compression and thin-provisioning benefits. The utility performs read-only operations, so it does not affect the data that is stored on the volume.

To view the results and the date of the latest estimation cycle, under the volumes view, right-click the volume then select **Capacity Savings** → **Estimate Compression Savings**.

To download a capacity savings reports, under the volumes view, select **Actions** → **Capacity Savings** → **Download Savings report**.

The report is also useful for determining the physical capacity used by each volume when the volume is compressed by FCMs or when the volumes are compressed in a Data Reduction Pool.

A stand-alone comprestimator utility can be installed and used on host systems to estimate savings before you move data to a Storage Virtualize system. To download the Comprestimator that can be installed on a server, see [IBM FlashSystem Comprestimator](#).

### 4.1.5 Verifying host configuration

Select **Hosts** → **Hosts** to view the hosts status. The status of all hosts is expected to be online. A host with a degraded status is typically caused by the host being partially connected to the storage. For example, both host WWPNs are logged in to node1 and one host WWPN is logged in to node2.

To review the host connectivity, select **Settings** → **Network** → **Fibre Channel or NVMe Connectivity**. The results can be filtered by the host. In the following example Host1 is degraded because each WWPN is logged in to node1 twice and node 2 once. See Figure 4-4.

Fibre Channel Connectivity									
Display the connectivity between nodes and other storage systems and hosts that are attached through the Fibre Channel network.									
View connectivity for: Hosts ▾ Host1 ▾ <a href="#">Show Results</a>									
<span>≡ Actions ▾</span>   <span>Display WWPN with colon (:) ▾</span>   <span>⬇️</span>   <span>Default ▾</span>   <span>Contains ▾</span>									
Name	Remote WWPN	↑	Remote ...	Local WWPN	Local Port	Local NP...	State	Node Na...	Type
Host1	10:00:00:90:FA:A0:36:88		121200	50:05:07:68:10:18:02:14	4	120001	✓ Active	node1	Host
Host1	10:00:00:90:FA:A0:36:88		121200	50:05:07:68:10:17:02:16	3	120301	✓ Active	node2	Host
Host1	10:00:00:90:FA:A0:36:88		121200	50:05:07:68:10:17:02:14	3	120101	✓ Active	node1	Host
Host1	10:00:00:90:FA:A0:36:89		081200	50:05:07:68:10:18:02:16	4	080201	✓ Active	node2	Host
Host1	10:00:00:90:FA:A0:36:89		081200	50:05:07:68:10:16:02:14	2	080001	✓ Active	node1	Host
Host1	10:00:00:90:FA:A0:36:89		081200	50:05:07:68:10:15:02:14	1	080101	✓ Active	node1	Host

Figure 4-4 Host with asymmetrical logins

When a host shows as degraded but there is no hardware failure on the host or storage, verify the fibre channel switch zoning and rescan the storage from the host.

**Tip:** If you are using Broadcom Fibre Channel switches, the `fcping` command that is run from the switch CLI can be used to verify zoning and WWPN connectivity.

## 4.2 Additional settings and basic operations

The following sections discuss additional settings that an administrator can use when implementing a new system along with basic operations.

### 4.2.1 Security settings

Storage Virtualize systems implement various security-related features. To configure or modify security functions select **Settings** → **Security**. Some of the more commonly used settings are in this section.

For more information, see [Security](#) and see [IBM Storage Virtualize, IBM Storage FlashSystem, and IBM SAN Volume Controller Security Feature Checklist](#), REDP-5716.

#### Remote authentication

You can use remote authentication to authenticate to the system by using credentials that are stored on an external authentication service. When you configure remote authentication, you do not need to configure users on the system or assign more passwords. Instead, you can use your existing passwords and user groups that are defined on the remote service to help simplify user management and access, to enforce password policies, and to separate user management from storage management.

A remote user is authenticated on a remote LDAP server. It is not required to add a remote user to the list of users on the system, although they can be added to configure optional SSH keys. For remote users, an equivalent user group must be created on the system with the same name and role as the group on the remote LDAP server.

For more information, see [Remote Authentication Documentation](#).

#### Ownership groups

An *ownership group* defines a subset of users and objects within the system. You can create ownership groups to further restrict access to specific resources that are defined in the ownership group. Only users with Administrator or Security Administrator roles can configure and manage ownership groups.

Ownership groups restrict access to only those objects that are defined within that ownership group. An owned object can belong to one ownership group.

An *owner* is a user with an ownership group that can view and manipulate objects within that group.

The system supports the following resources that you assign to ownership groups:

- ▶ Child pools
- ▶ Volumes
- ▶ Volume groups
- ▶ Hosts
- ▶ Host clusters
- ▶ Host mappings
- ▶ FlashCopy mappings
- ▶ FlashCopy consistency groups
- ▶ User groups
- ▶ Portsets

The following basic use cases can be applied to the use of ownership groups on the system:

- ▶ Objects are created within the ownership group. Other objects can be on the system that are not in the ownership group.
- ▶ On a system where these supported objects are configured, and you want to migrate these objects to use ownership groups.

When a user group is assigned to an ownership group, the users in that user group retain their role, but are restricted to only those resources within the same ownership group. User groups can define the access to operations on the system, and the ownership group can further limit access to individual resources.

For example, you can configure a user group with the Copy Operator role, which limits access of the user to Copy Services functions, such as FlashCopy and Remote Copy operations. Access to individual resources, such as a specific FlashCopy consistency group, can be further restricted by adding it to an ownership group.

When the user logs on to the management GUI, only resources that they can access through the ownership group are displayed. Also, only events and commands that are related to the ownership group to which a user belongs are viewable by those users.

For more information, see [Ownership groups](#).

## System certificates

SSL certificates are used to establish secure communications for many services. The system uses a certificate to identify itself when authenticating with other devices. Depending on the scenario, the system might be acting as either the client or the server.

The system has a root certificate authority (CA) that can be used to create internally signed system certificates. System setup creates a certificate that is signed by the root CA to secure connections between the management GUI and the browser. The root certificate can be exported from the system and added to truststores on other systems, browsers, or devices to establish trust. Internally signed certificates can be renewed automatically before they expire. Automatic renewal can simplify the certificate renewal process and can prevent security warnings from expired certificates. Automatic renewal is only supported by using an internally signed certificate.

Externally signed certificates are issued and signed by a trusted third-party provider of certificates, called an external certificate authority (CA). This CA can be a public CA or your own organization's CA. Most web browsers trust well-known public CAs and include the root certificate for these CAs in the device or application. Externally signed certificates cannot be renewed automatically because they must be issued by the external CA. Externally signed certificates must be manually updated before they expire by creating a new certificate signing request (CSR) on the system and supplying it to the CA. The CA signs the request and issues a certificate that must be installed on the system. The system raises a warning in the event log 30 days before the certificate expires.

An externally signed certificate must meet the following requirements:

- ▶ X.509v3 Key Usage Extensions. Include a Digital Signature.
- ▶ X.509v3 Extended Key Usage Extensions. Include TLS Web Server Authentication and TLS Web Client Authentication. Also, Any Extended Key Usage (anyEKU) can be included.

Ensure that the Certificate Authority (CA) used to sign the certificate includes these extensions.

## Security protocol levels

Security administrators can change the security protocol level for either SSL or SSH protocols. When you change the security level for either of these security protocols, you can control which encryption algorithms, ciphers, and version of the protocol are permitted on the system.

The GUI gives a high-level description of each level. For a more detailed description including the ciphers supported with each level, see [Security protocol levels](#).

## 4.2.2 Audit log

The audit log is useful when analyzing past configuration events, especially when trying to determine, for example, how a volume became shared by two hosts or why the volume was overwritten. The audit log is also included in the `svc_snap` support data to aid in problem determination.

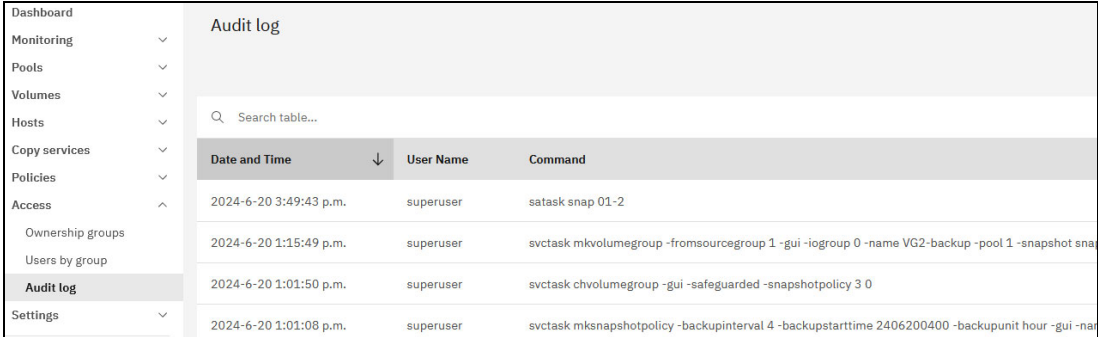
The audit log tracks action commands that are issued through an SSH session, management GUI, or Remote Support Assistance. It provides the following entries:

- ▶ Identity of the user who ran the action command.
- ▶ Name of the actionable command.
- ▶ Timestamp of when the actionable command ran on the configuration node.
- ▶ Parameters that ran with the actionable command.

The following items are not documented in the audit log:

- ▶ Commands that fail are not logged.
- ▶ A result code of 0 (success) or 1 (success in progress) is not logged.
- ▶ Result object ID of node type (for the `addnode` command) is not logged.
- ▶ View commands are not logged.

The audit log is accessed by selecting **Access** → **Audit Log** see Figure 4-5.



The screenshot shows a web interface with a sidebar on the left containing navigation items: Dashboard, Monitoring, Pools, Volumes, Hosts, Copy services, Policies, Access, Ownership groups, Users by group, Audit log (highlighted), and Settings. The main content area is titled 'Audit log' and contains a search bar and a table. The table has three columns: 'Date and Time', 'User Name', and 'Command'. It lists four entries, all performed by 'superuser'.

Date and Time	User Name	Command
2024-6-20 3:49:43 p.m.	superuser	satask snap 01-2
2024-6-20 1:15:49 p.m.	superuser	svctask mkvolumegroup -fromsourcegroup 1 -gui -iogroup 0 -name VG2-backup -pool 1 -snapshot sna
2024-6-20 1:01:50 p.m.	superuser	svctask chvolumegroup -gui -safeguarded -snapshotpolicy 3 0
2024-6-20 1:01:08 p.m.	superuser	svctask mksnapshotpolicy -backupinterval 4 -backupstarttime 2406200400 -backupunit hour -gui -na

Figure 4-5 Audit log

## 4.2.3 Support settings

Support settings can be configured and modified by selecting **Settings** → **Support**.

### Call Home

IBM Call Home is a support function that is embedded in all IBM Storage Virtualize storage products. By enabling call home, the health and functionality of your system is constantly monitored by IBM. If a software or hardware error occurs, the call home function notifies IBM support of the event and then automatically opens a service request. By obtaining information

in this way, IBM support is quickly informed about the issue and can develop an action plan for problem resolution.

There are two methods available for a system to call home and both can be enabled simultaneously:

1. Cloud Services uses HTTPS to connect directly to IBM from the management IP address assigned to the lowest physical port ID over port 443
2. Email services require an SMTP server to forward the email to IBM. Email services can also send alerts to local administrators.

## Remote Support Assistance

Remote Support Assistance enables IBM Support to remotely connect to an IBM FlashSystem system through a secure tunnel to perform analysis, log collection, and software updates. The tunnel can be enabled as needed by the client or as a permanent connection.

For detailed information on call home or remote support assistance, see the white paper [IBM Storage Virtualize Products Call Home and Remote Support Overview](#).

**Note:** Remote support assistance uses the *service IP addresses* to make an outbound connection to IBM on port 22.

The connections for both Call Home and Remote Support Assistance can be routed through a client-supplied web proxy.

## Support package

If you encounter a problem and contact the IBM Support Center, you are asked to provide a support package, which is often referred to as a *snap*.

You can use two methods to collect and upload the support package from the GUI of your Storage Virtualize system:

- ▶ Upload Support Package  
Use this feature if your system is connected to the internet to upload the Support Package directly from the storage system.
- ▶ Download Support Package  
Use this feature if your system is not connected to the internet to upload the Support Package manually.

The support agent provides the type of support package to collect based on the problem. For general guidelines and the differences between the different support package types, see [What Data Should You Collect for a Problem on Spectrum Virtualize systems](#).





# IBM Storage Insights and IBM Storage Insights Pro

Managing storage systems can be complex. You must monitor performance, capacity, and overall health to ensure that your data is always accessible and secure. IBM Storage Insights offers a solution to help you monitor, manage, and optimize your storage resources.

This chapter has the following sections:

- ▶ 5.1, “IBM Storage Insights overview” on page 92
- ▶ 5.2, “IBM Storage Insights monitoring” on page 93

## 5.1 IBM Storage Insights overview

IBM Storage Insights is another part of the monitoring capability of the IBM FlashSystems and IBM SAN Volume Controller systems running IBM Storage Virtualize software and supplements the views that are available in the product GUI.

IBM strongly recommends that all customers install and use this no-charge, cloud-based IBM application because it provides a single dashboard that provides a clear view of all your IBM block storage. You can make better decisions by seeing trends in performance and capacity.

**Note:** IBM Storage Insights is available at no cost to clients who have IBM Storage Systems on either IBM warranty or maintenance. The more fully featured IBM Storage Insights Pro is a chargeable product, which can be purchased separately and can also be included in certain levels of IBM Storage Expert Care and IBM Storage Control.

With storage health information, you can focus on areas that need attention. When IBM support is needed, IBM Storage Insights simplifies uploading logs, speeds resolution with online configuration data, and provides an overview of open tickets all in one place.

IBM Storage Insights includes the following features:

- ▶ A unified view of IBM systems:
  - Provides a single view to see all your system's characteristics.
  - Displays all of your IBM storage inventory.
  - Provides a live event feed so that you know in real time what is going on with your storage so that you can act quickly.
- ▶ IBM Storage Insights collects telemetry data and Call Home data, and provides real-time system reporting of capacity and performance.
- ▶ Overall storage monitoring, which provides the following information:
  - The overall health of the system.
  - Monitoring of the configuration to see whether it meets preferred practices.
  - System resource management to determine which system is overtaxed and provides proactive recommendations to fix it.
- ▶ IBM Storage Insights provides advanced customer service with an event filter that you can use to accomplish the following tasks:
  - You and IBM Support can view, open, and close support tickets, and track trends.
  - You can use the autolog collection capability to collect the logs and send them to IBM before IBM Support looks into the problem. This capability can save as much as 50% of the time to resolve the case.

In addition to the no-charge version of IBM Storage Insights, IBM offers IBM Storage Insights Pro. IBM Storage Insights Pro is a subscription service that provides longer historical views of data, more reporting and optimization options, and supports IBM file and block storage with EMC VNX and VMAX.

**Note:** For a comparison of the features in the IBM Storage Insights and Insights Pro editions, see [IBM Storage Insights vs IBM Storage Insights Pro](#).

- ▶ For more information about IBM Storage Insights and for registration, see the following resources: [IBM Storage Insights Fact Sheet](#)
- ▶ *IBM Storage Insights Security Guide*, [SC27-8774](#)
- ▶ [IBM Storage Insights](#)
- ▶ [Product registration](#), which is used to sign up and register for this no-charge service

## 5.2 IBM Storage Insights monitoring

By using IBM Storage Insights, you can optimize your storage infrastructure by using this cloud-based storage management and support platform with predictive analytics.

The monitoring capabilities that IBM Storage Insights provides are useful for things like capacity planning, workload optimization, and managing support tickets for ongoing issues.

For a live demo of IBM Storage Insights, see [Storage Insights Demo](#) (requires login).

**Demonstration videos:** To view videos about Storage Insights, see [Videos for IBM Storage Insights](#). The videos include new features and enhancements of IBM Storage Insights.

After you add your systems to IBM Storage Insights, you see the Dashboard, where you can select a system that you want to see the overview for.

There are two versions of the dashboard, the classic version and the new Carbon enhanced version.

Figure 5-1 shows the classic version view of the IBM Storage Insights dashboard.

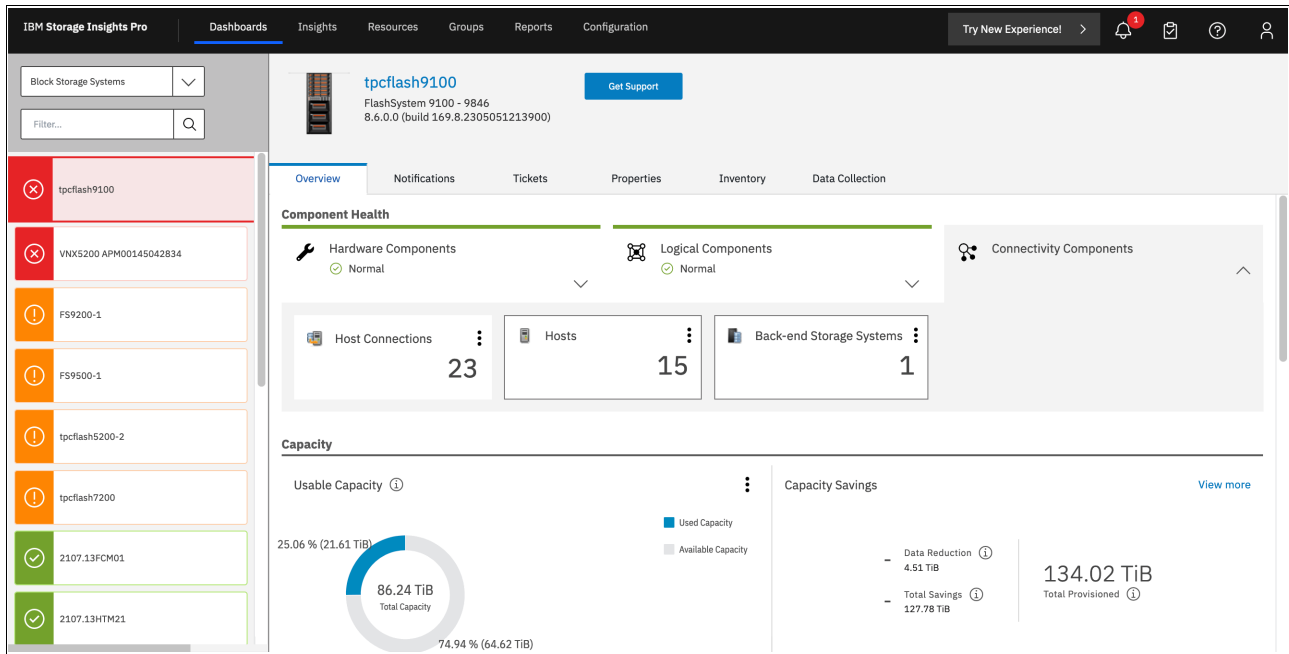


Figure 5-1 IBM Storage Insights System overview (classic view)

Figure 5-2 shows the newer Carbon enhanced view of the IBM Storage Insights dashboard.

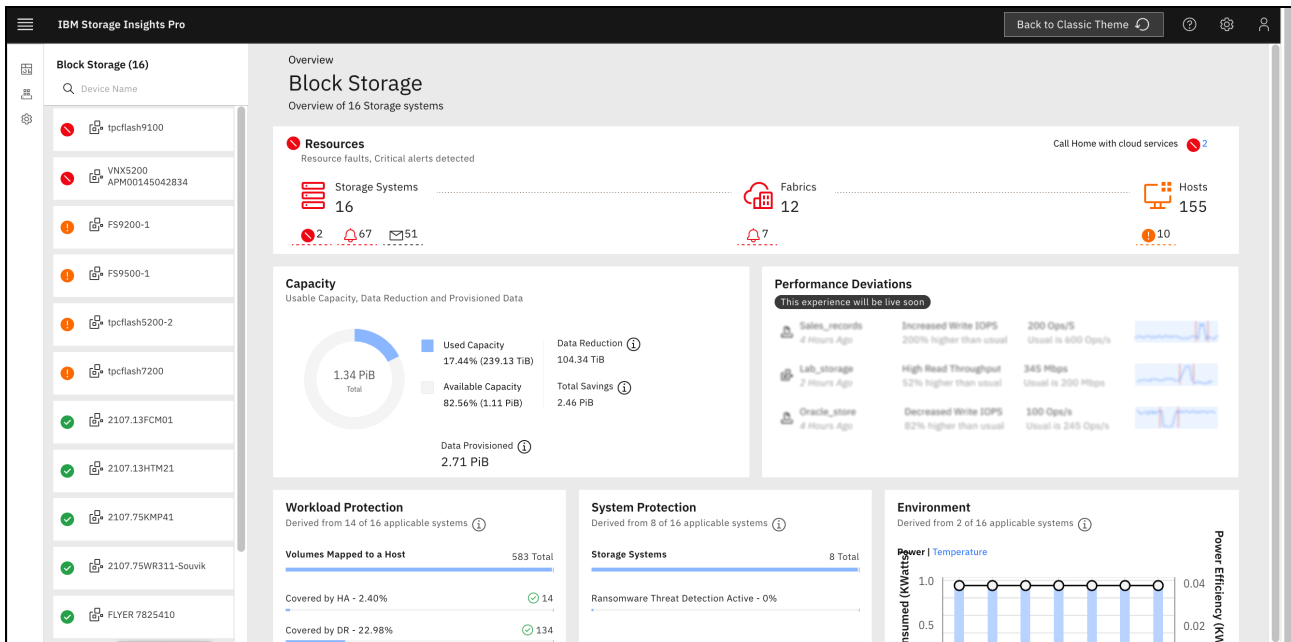


Figure 5-2 IBM Storage Insights System overview (Carbon enhanced view)

The next few examples use the classic view screens.

## 5.2.1 Component health

Component health is shown at the upper center of the window. If there is a problem with one of the Hardware, Logical, or Connectivity components, then errors are shown here, as shown in Figure 5-3.

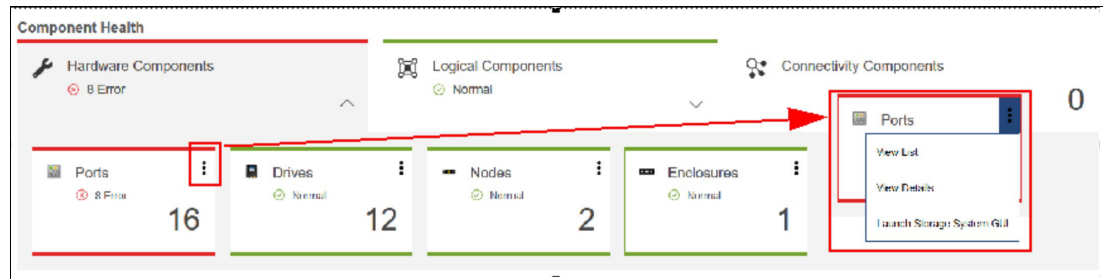


Figure 5-3 Component Health overview

The error entries can be expanded to obtain more details by selecting the three dots at the upper-right corner of the component that has an error and then selecting **View Details**. The relevant part of the more detailed System View opens, and what you see depends on which component has the error, as shown in Figure 5-4 on page 95.

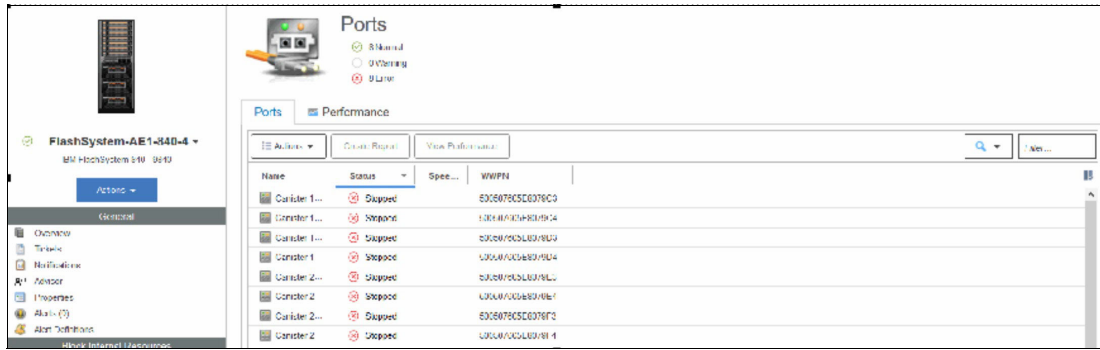


Figure 5-4 Ports in error

In Figure 5-4, the GUI lists which components have the problem and exactly what is wrong with them. You can use that information to open a support ticket with IBM if necessary.

## 5.2.2 Capacity monitoring

You can see key statistics such as Usable and Provisioned Capacity and Capacity Savings as shown in Figure 5-5. Capacity can be viewed by volume or pool and the **View More** button shows a trend curve.




Figure 5-5 Capacity area of the IBM Storage Insights system overview

In the Capacity view, the user can select the required system. Clicking any of these items takes the user to the detailed system view for the selection option. From there, you can get a historical view of how the system capacity changed over time, as shown in Figure 5-6 on page 96. At any time, the user can select the timescale, resources, and metrics to be displayed on the graph by clicking any options around the graph.

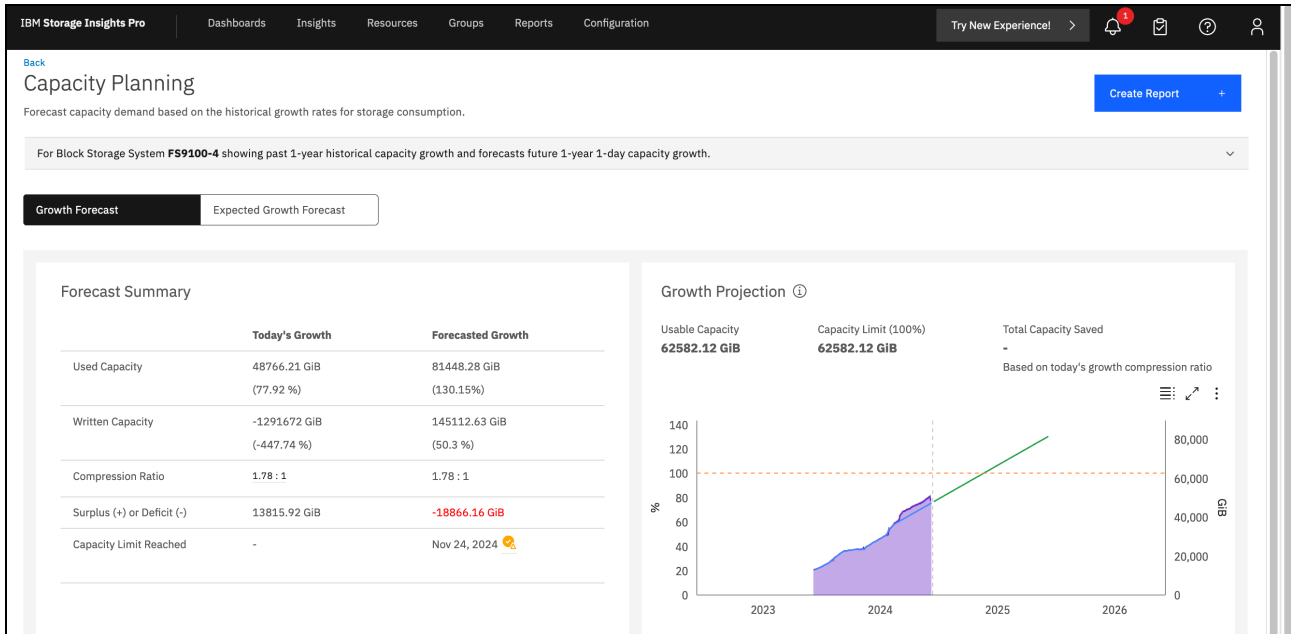


Figure 5-6 Capacity planning for one system

## 5.2.3 Performance monitoring

From the system overview, you can scroll down and see the three key performance statistics for your system, as shown in Figure 5-7. For the Performance overview, these statistics are aggregated across the whole system, and you cannot drill down by Pool, Volume, or other items.




Figure 5-7 System overview: Performance

To view more detailed performance statistics, enter the system view again, as described in 5.2.2, “Capacity monitoring” on page 95.

For this performance example, select **View Pools**, and then select **Performance** from the System View pane, as shown in Figure 5-8 on page 97.

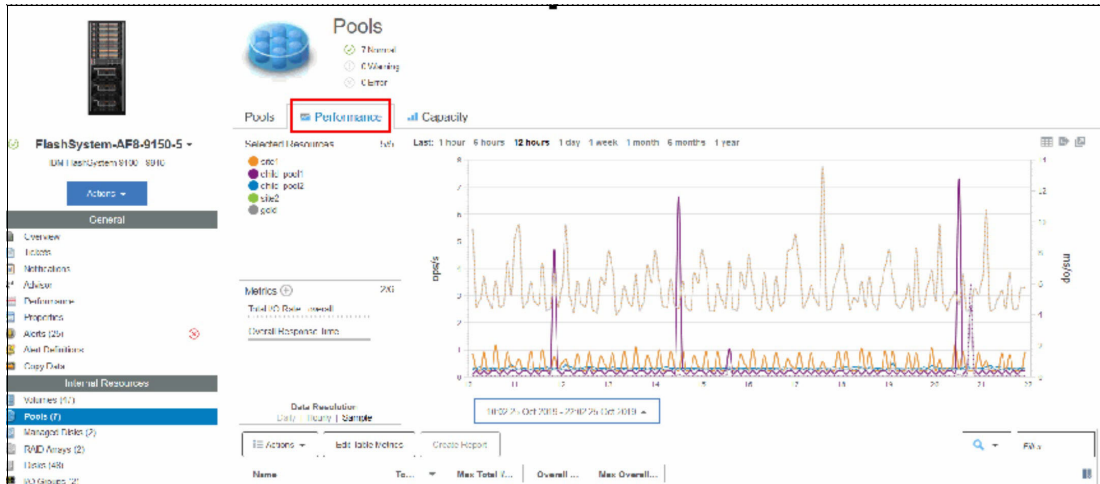


Figure 5-8 IBM Storage Insights: Performance view

It is possible to customize what can be seen on the graph by selecting the metrics and resources. In Figure 5-9, the Overall Response Time for one IBM FlashSystem over a 12-hour period is displayed.




Figure 5-9 Filtered performance graph

Scrolling down the graph, the Performance List view is visible, as shown in Figure 5-10 on page 98. Metrics can be selected by clicking the filter button at the right of the column headers. If you select a row, the graph is filtered for that selection only. Multiple rows can be selected by holding down the Shift or Ctrl keys.



Figure 5-10 Performance List View

## 5.2.4 Logging support tickets by using IBM Storage Insights

With IBM Storage Insights, you can log existing support tickets that complement the enhanced monitoring opportunities that the software provides. When an issue is detected and you want to engage IBM Support, complete the following steps:

1. Select the system to open the System Overview window and click **Get Support**, as shown in Figure 5-11.

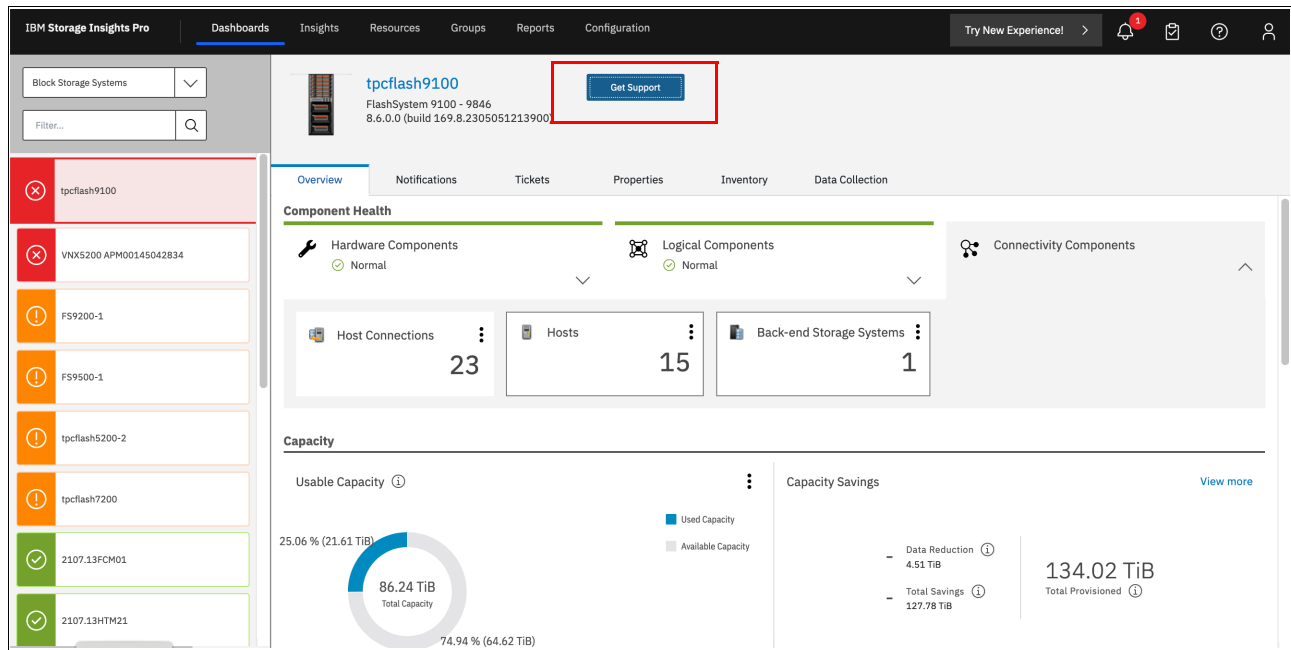


Figure 5-11 Get Support (see highlighted area)

A window opens where you can create a ticket or update an existing ticket, as shown in Figure 5-12 on page 99.



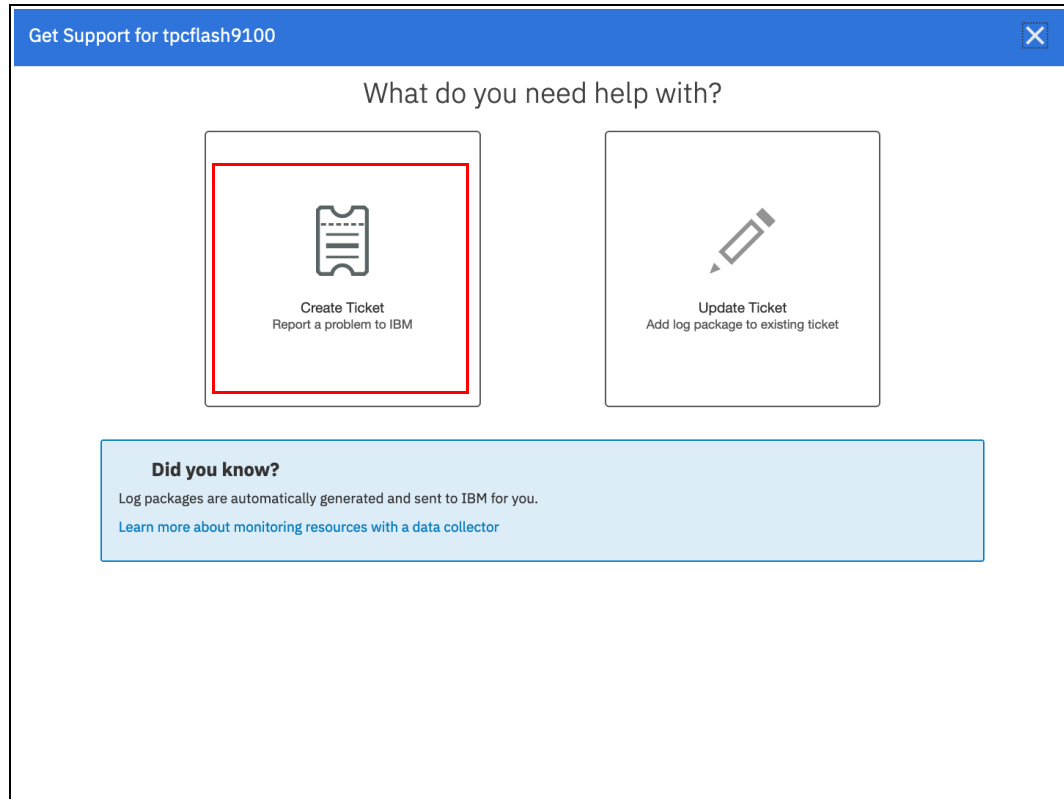


Figure 5-12 Get Support window

2. Select **Create Ticket**, and the ticket creation wizard opens. Details of the system are automatically populated, including the customer number, as shown in Figure 5-13. Select **Next**.

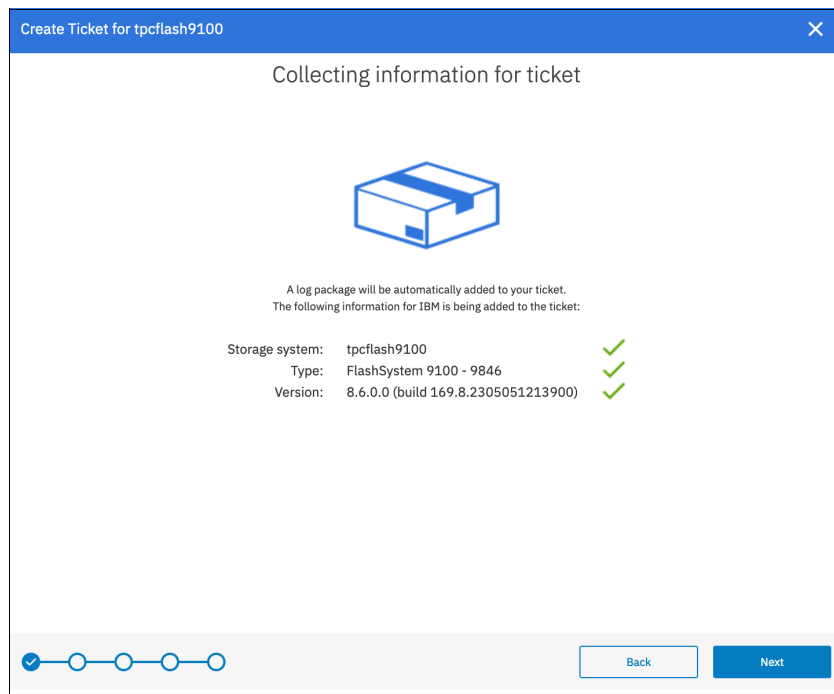


Figure 5-13 Create Ticket wizard

3. You can enter relevant details about your problem to the ticket, as shown in Figure 5-14. It is also possible to attach images or files to the ticket, such as PuTTY logs and screen captures. After you add any images or files, select **Next**.

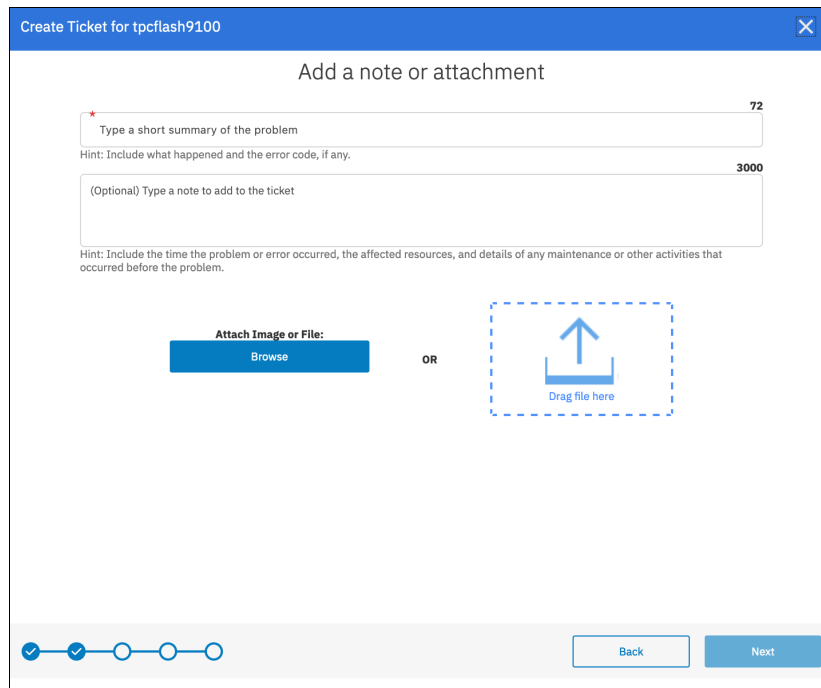


Figure 5-14 Add a note or attachment window

4. You can select a severity for the ticket. Examples of what severity you to select are shown in Figure 5-15. Because in the example there are storage ports offline with no impact, select **severity 3** because there is only minor impact.

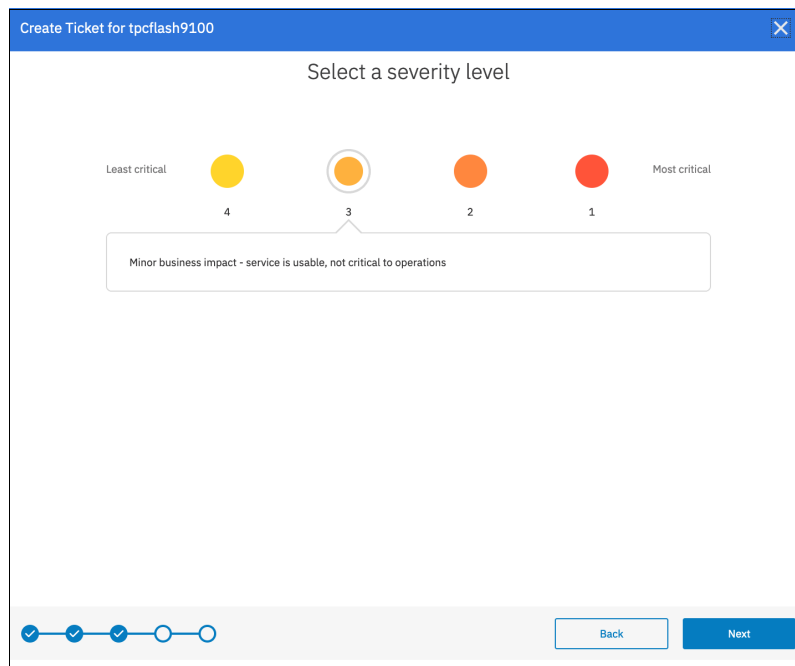


Figure 5-15 Selecting a Severity Level window

5. Choose whether this is a hardware or a software problem. For this example, the offline ports are likely caused by a physical layer hardware problem. Click **Next**.
6. Review the details of the ticket to be logged with IBM, as shown in Figure 5-16. Contact details must be entered so that IBM Support can respond to the correct person. You must also choose which type of logs to attach to the ticket. For more information about the types of snaps, see Figure 5-16. Click **Create Ticket**.

Problem summary:

Description:

Severity level: 3 Minor business impact - service is usable, not critical to operations

Log package: Type 1: Standard logs

Type of problem: Hardware

Contact name: \*

Contact email: \*

Contact phone: \*

Country: United States

Storage system: tpcflash9100

Type: FlashSystem 9100 - 9846

Version: 8.6.0.0 (build 169.8.2305051213900)

Enclosure: Control enclosure: 78E034L (78E034L)

Back Create Ticket

Figure 5-16 Review the ticket window

7. A confirmation window opens, as shown in Figure 5-17 on page 102, and IBM Storage Insights automatically uploads the snap to the ticket when it is collected. Click **Close**.

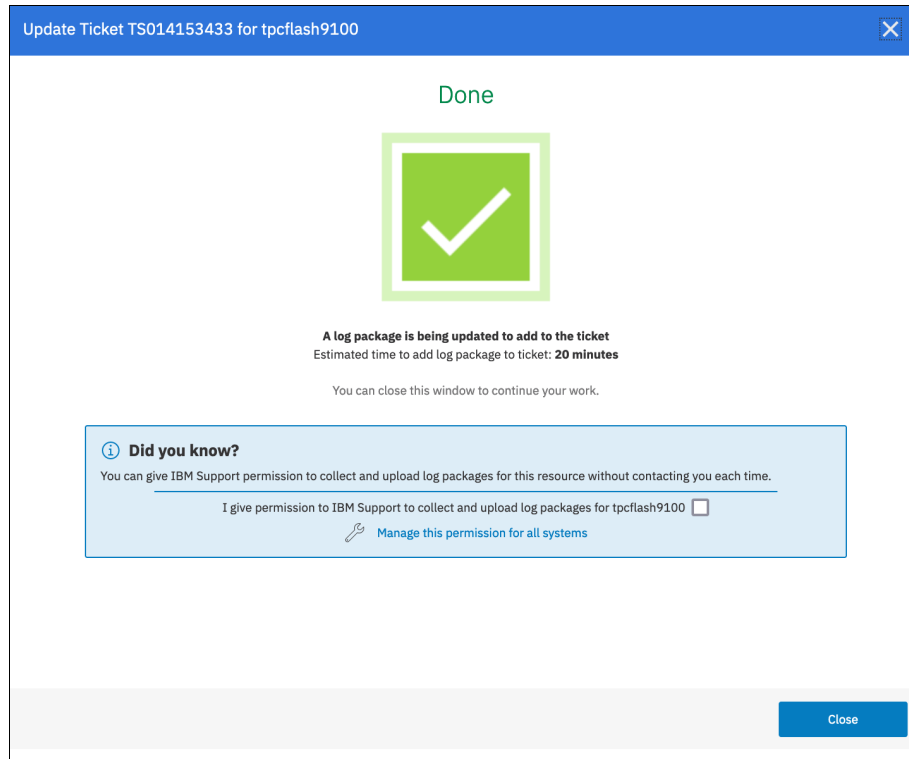


Figure 5-17 Update ticket

## 5.2.5 Managing existing support tickets by using IBM Storage Insights

With IBM Storage Insights, you can track existing support tickets and upload logs to them. To do so, complete the following steps:

1. From the System Overview window, select **Tickets**, as shown in Figure 5-18.

This window shows the newly created ticket number and a history of support tickets that were logged through IBM Storage Insights for the system. Tickets that are not currently open are listed under **Closed Tickets**, and currently open tickets are listed under **Open Tickets**.

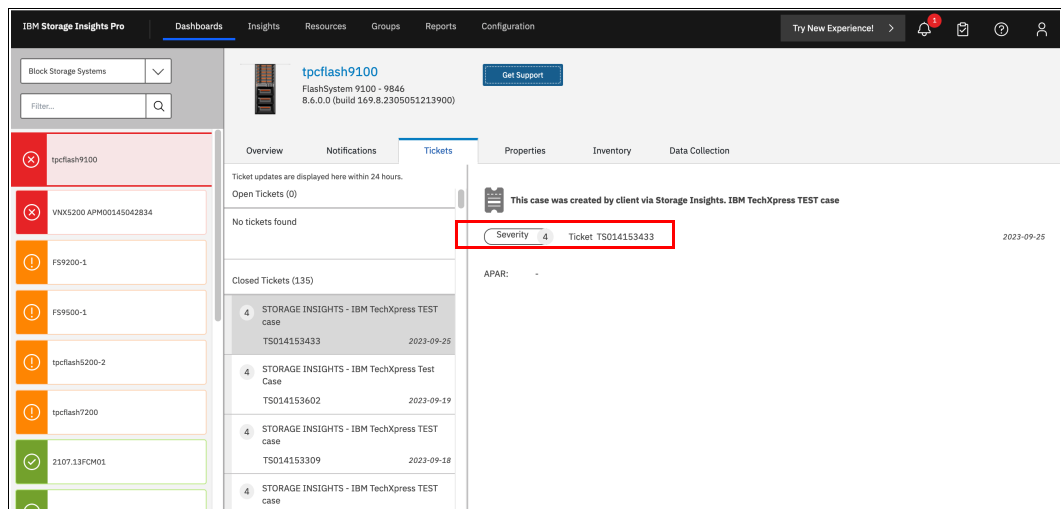


Figure 5-18 View tickets

- To quickly add logs to a ticket without having to browse to the system GUI or use IBM ECURep, click **Get Support** and **Add Log Package to Ticket**. A window opens that guides you through the process, as shown in Figure 5-19. After you enter the support ticket number, you can select which type of log package you want and add a note to the ticket with the logs.

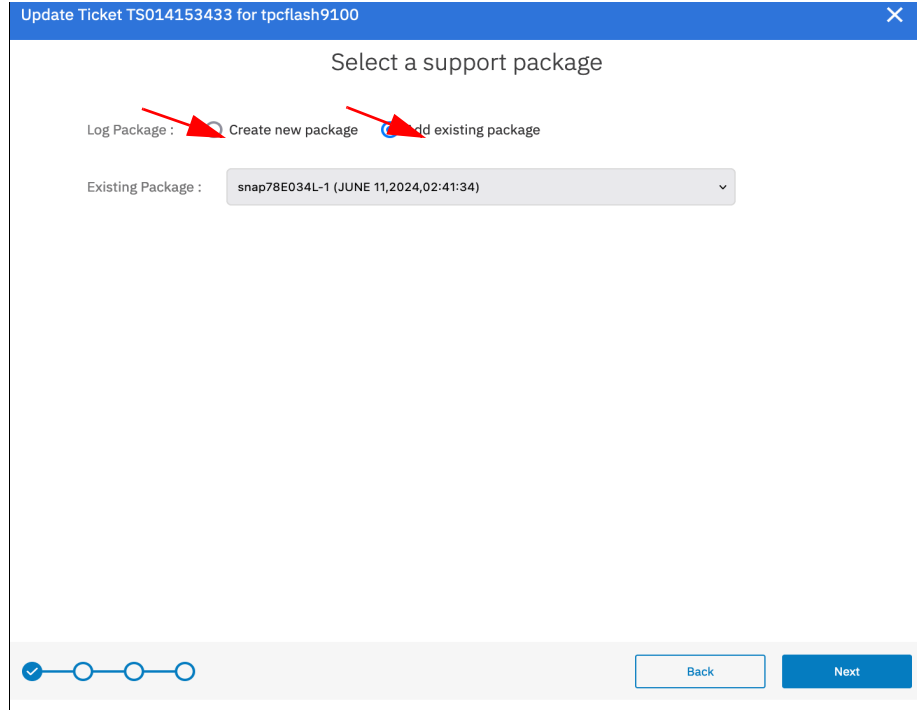


Figure 5-19 Adding a log package to the ticket

- Review the ticket Figure 5-20 to confirm what is going to be uploaded.

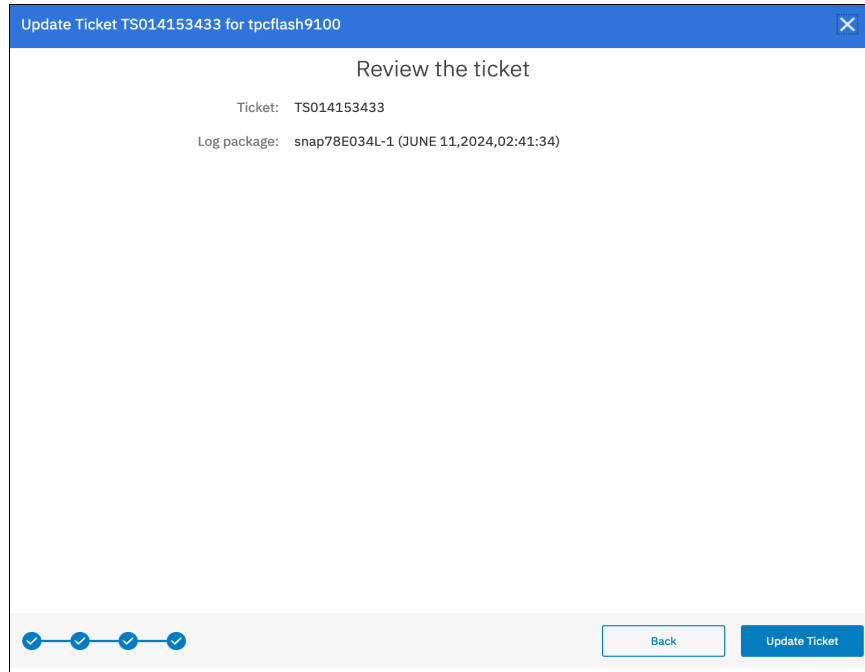


Figure 5-20 Confirming the log upload

4. After clicking **Update Ticket**, a confirmation opens, as shown in Figure 5-21. You can exit the wizard. IBM Storage Insights runs in the background to gather the logs and upload them to the ticket.

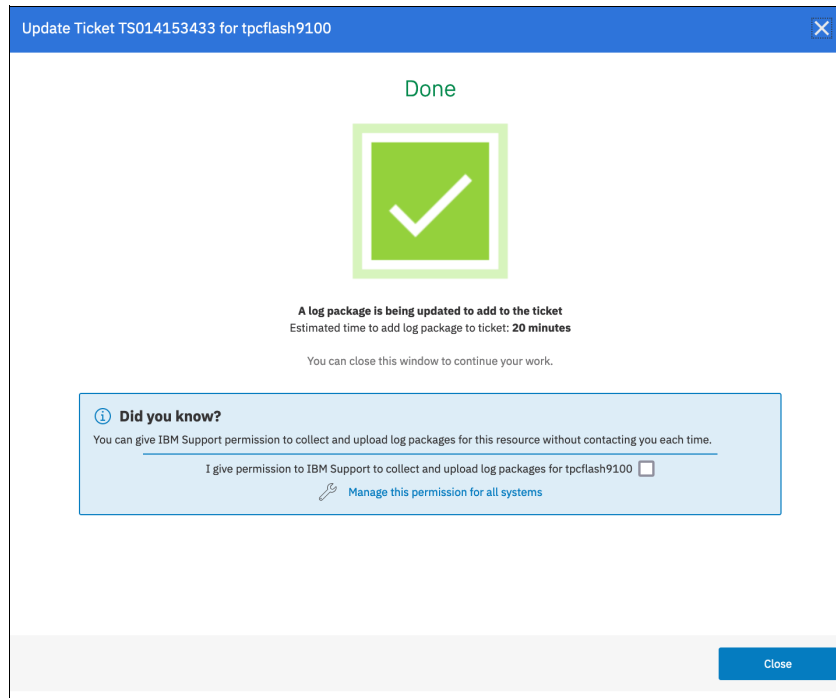


Figure 5-21 Log upload completed and processing

## 5.2.6 Enhancements to IBM Storage Insights Pro

The following section describes some enhancements to IBM Storage Insights Pro in Version 8.7.

### Ransomware threat detection

A key part of monitoring your system includes the detection of potential ransomware attacks. Starting with IBM Storage Virtualize 8.6.0, the ransomware threat detection mechanism was introduced as a Virtualize level only. To ensure that you have the latest storage metadata for detecting those types of attacks, compression and cyber resiliency statistics for volumes are collected every 5 minutes. With these statistics, IBM Storage Insights builds a historical model of a storage system and uses its built-in intelligence and formulas to identify when and where ransomware attacks might be occurring.

IBM Storage Virtualize software 8.7.0 and FlashCore modules (FCMs) with firmware 4.1 include the following enhancements to ransomware threat detection:

- ▶ IBM FCMs collect and analyze detailed ransomware statistics from every I/O with no performance impact.
- ▶ IBM Storage Virtualize runs an AI engine on every FlashSystem that is fed Machine Language (ML) models developed by IBM Research® trained on real-world ransomware.
- ▶ The AI engine learns what's normal for the system and detects threats by using data from the FCMs.
- ▶ IBM Storage Insights Pro collects threat information from connected FlashSystems. Alerts trigger SIEM/SOAR software to initiate a response.

- ▶ Statistics are fed back to IBM to improve ML models.

For more information about the IBM ransomware threat detection solutions, including those mentioned in this book, see [Ransomware protection solutions](#).

Also, for more information about how to mark volume snapshots as compromised after ransomware threat detection, see [Boost Your Defense with IBM Storage Insights](#).

### IBM Storage Virtualize 8.7.0 including Flash Grid

When you configure a Flash Grid, you can configure a maximum of 8 FlashSystem or SVC systems to be managed as one, which includes nondisruptive workload mobility between members of the grid.

IBM Storage Insights Pro works with the Flash Grid and provides an overview of your grid with grouping of your systems and the ability to nondisruptively move workloads, also called storage partitions, between systems in the grid. The goal is to provide a seamless integration and interaction between the on-premises and cloud-based management portals. Figure 5-22 shows the integration of IBM Storage Insights Pro with the IBM Storage Virtualize software GUI and the linkage to the IBM FlashSystems and IBM SAN Volume Controllers it monitors.

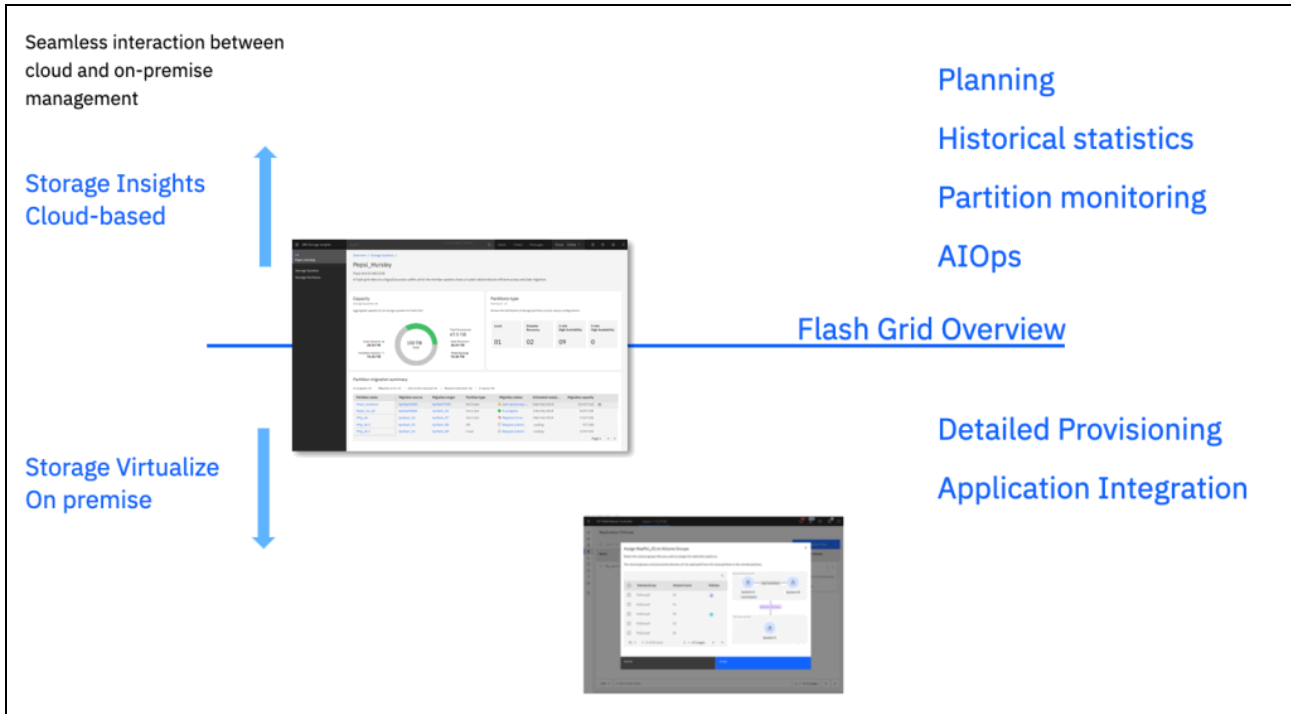


Figure 5-22 IBM Storage insights Pro and IBM Flash Grid integration







# Storage Virtualize troubleshooting and diagnostics

This chapter provides information to troubleshoot common problems that can occur in an IBM Storage Virtualize 8.7.0 environment. It describes situations that are related to IBM SAN Volume Controller (SVC), IBM FlashSystems, the storage area network (SAN) environment, optional external storage subsystems, and hosts. It also explains how to collect the necessary problem determination data.

This chapter includes the following topics:

- ▶ 6.1, “Troubleshooting” on page 108
- ▶ 6.2, “Collecting diagnostic data” on page 115

## 6.1 Troubleshooting

When you troubleshoot, follow a systematic approach to solve a problem. The goal of troubleshooting or problem determination is to understand why something does not work as expected and create a resolution to resolve this. Therefore, an important step is to make a proper problem description as accurate as possible. Then you need to collect the support data from all involved components of the environment for analysis. This might include a *snap* from the IBM Storage Virtualize system, logs from SAN or network switches and host OS logs.

An effective problem report ideally describes these items:

- ▶ the expected behavior
- ▶ the actual behavior
- ▶ if possible, how to reproduce the behavior
- ▶ a precise timeline

The following questions help define the problem for effective troubleshooting:

- ▶ What are the symptoms of the problem?
  - What is reporting the problem?
  - Which error codes and messages were observed?
  - What is the business impact of the problem?
  - Where does the problem occur?
  - Which exact component is affected, the whole system or for instance certain hosts, IBM Storage Virtualize nodes
  - Is the environment and configuration supported?
- ▶ When does the problem occur?
  - How often does the problem happen?
  - Does the problem happen only at a certain time of day or night?
  - What kind of activities was ongoing at the time the problem was reported?
  - Did the problem happen after a change in the environment, such as a code upgrade or installing software or hardware?
- ▶ Under which conditions does the problem occur?
  - Does the problem always occur when the same task is being performed?
  - Does a certain sequence of events need to occur for the problem to surface?
  - Do any other applications fail at the same time?
- ▶ Can the problem be reproduced?
  - Can the problem be re-created, for example by running a single command, a set of commands, or a particular application?
  - Are multiple users or applications encountering the same type of problem?
  - Can the problem be reproduced on any other system?

**Note:** For effective troubleshooting, it is crucial to collect log files as close to the incident as possible and provide an accurate problem description with a timeline.

## 6.1.1 Storage Insights

As discussed in Chapter 5, “IBM Storage Insights and IBM Storage Insights Pro” on page 91, IBM Storage Insights is an important part of monitoring to help ensure continued availability of IBM Storage Virtualize systems.

When IBM Support is needed, IBM Storage Insights simplifies uploading logs, speeds resolution with online configuration data, and provides an overview of open tickets all in one place.

IBM strongly recommends that all customers install and use this no-charge, cloud-based IBM application because it provides a single dashboard that provides a clear view of all your IBM block storage.

For detailed information and examples, refer to Chapter 5, “IBM Storage Insights and IBM Storage Insights Pro” on page 91.

## 6.1.2 Using the GUI

The IBM Storage Virtualize GUI is a good entry point to start troubleshooting with. There are two essential icons at the top that are accessible from any GUI panel.

As shown in Figure 6-1, the first icon shows IBM Storage Virtualize events, such as an error or a warning, and the second icon shows suggested, running, or recently completed background tasks.

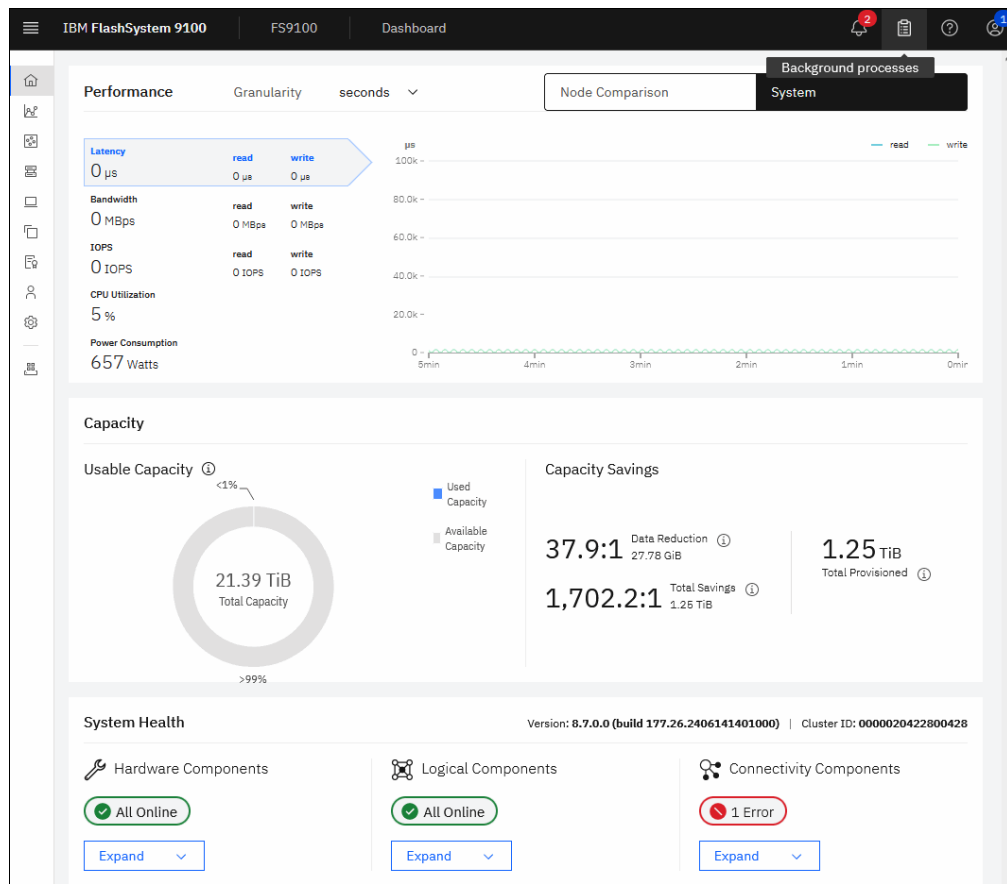


Figure 6-1 Events icon in the GUI

The GUI dashboard provides an at-a-glance view of the system's condition and notifies you of any circumstances that require immediate action. It contains sections for performance, capacity, and system health that provide an overall understanding of what is going on in the system.

The System Health section in the lower part of the dashboard provides information about the health status of hardware, logical, and connectivity components. If you click **Expand** in each of these categories, the status of the individual components is shown (see Figure 6-2). Clicking **More Details** takes you to the GUI panel related to that specific component, or shows more information about it.

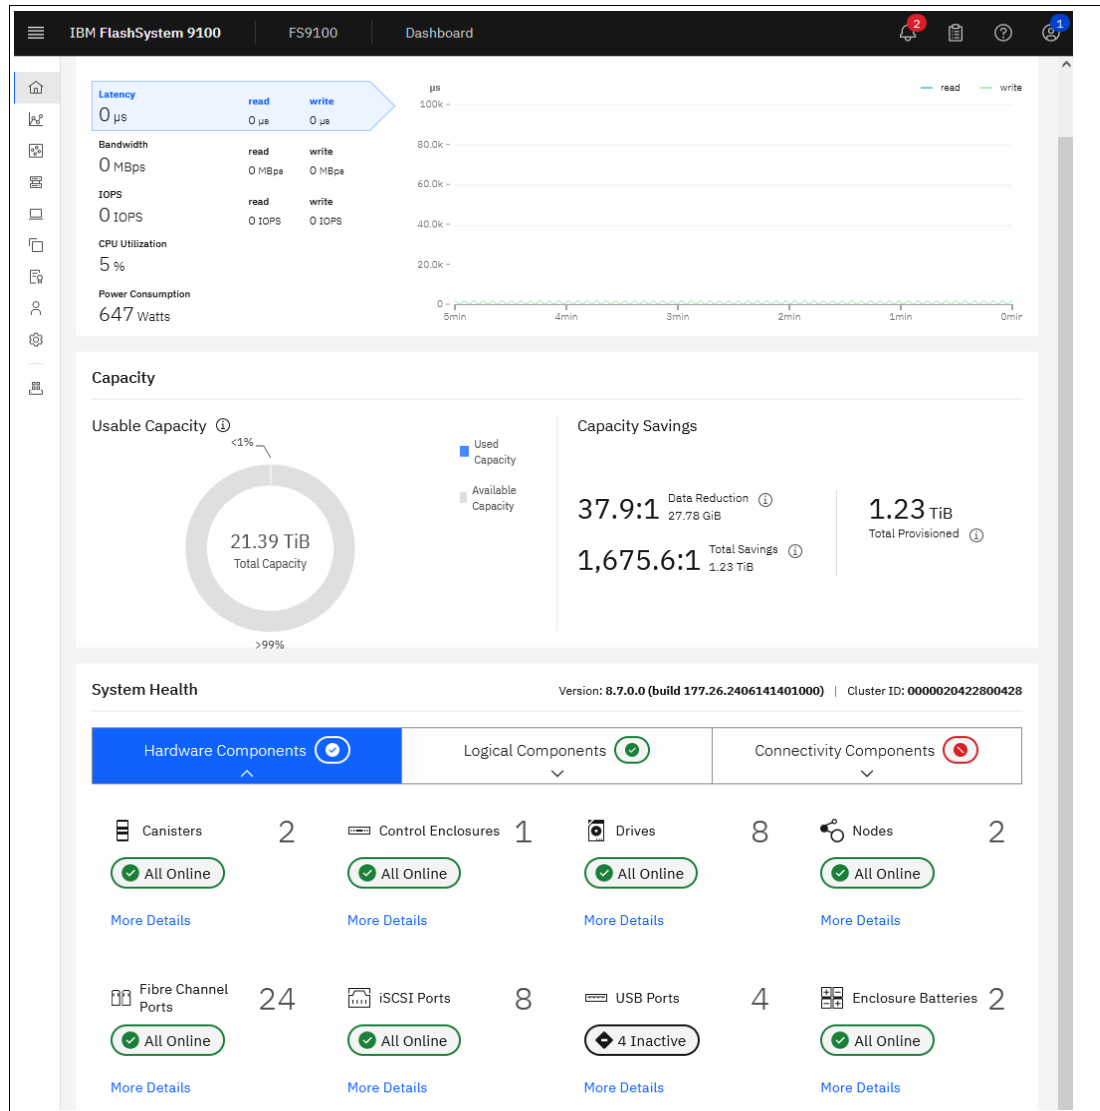


Figure 6-2 System Health expanded section in the dashboard

### 6.1.3 Recommended actions and fix procedure

A *Fix procedure*, sometimes also referred to as *Directed Maintenance Procedure*, assists you in fixing a problem without doing any harm. When one or multiple unfixed errors must be addressed, the management GUI provides the means to run the recommended fix procedure. Therefore, the first step in troubleshooting is to check the event log for Recommended Actions

in **Monitoring** → **Events**. The highest priority event, which is the event log entry with the lowest four-digit error code, is highlighted so that it is addressed first as shown in Figure 6-3.

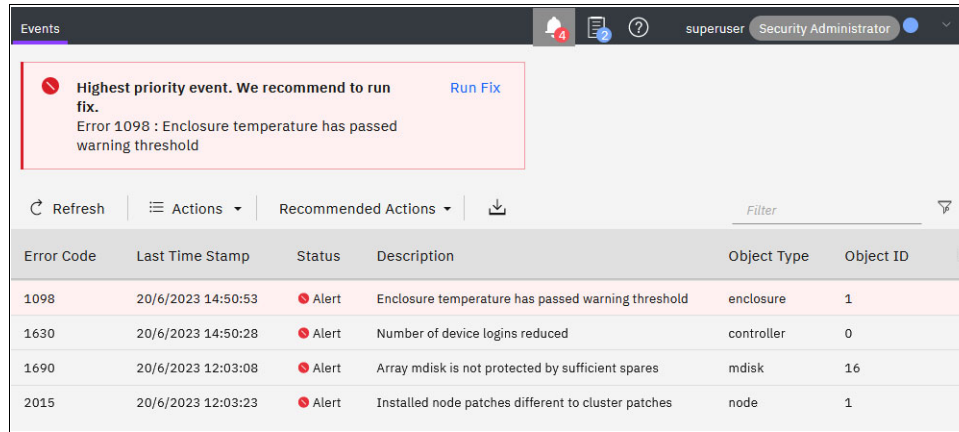


Figure 6-3 Recommended actions

Click **Run Fix** to start the Fix Procedure for this particular event. Fix Procedures help resolve a problem. In the background, a Fix Procedure analyzes the status of the system and its components and provides further information about the nature of the problem. This is to ensure that the actions taken do not lead to undesirable results, as, for instance, volumes becoming inaccessible to the hosts. The Fix Procedure then automatically performs the actions that are required to return the system to its optimal state. This can include checking for dependencies, resetting internal error counters and apply updated to the system configuration. Whenever user interaction is required, you are shown suggested actions to take and guided through the same. If the problem is fixed, the related error in the event log is eventually marked as fixed. Also, an associated alert in the GUI is cleared.

Error codes along with their detailed properties in the event log provide reference information when a service action is required. The four-digit *Error Code* is visible in the event log. They are accompanied by a six-digit *Event ID*, which provides additional details about this event.

Three-digit *Node Error Codes* are visible in the node status in the Service Assistant GUI. For more information about messages and codes, see [Messages and Codes](#).

## 6.1.4 Storage Virtualize failure recovery

An IBM Storage Virtualize system might encounter various kinds of failure recovery in certain conditions. These are known as Tier 1 (T1) through Tier 4 (T4) recovery.

- ▶ A T1 or Tier 1 recovery is a node warmstart (node assert) that is logged with error code 2030 in the event log.

A single node assert is a recovery condition that is deployed by the IBM Storage Virtualize software. A single node assert occurs when a single node attempts to run an invalid code path or detects a transient hardware problem.

A T1 recovery alias single-node warmstart is performed without suspending I/O. This task can be accomplished because the cluster is configured into redundant pairs of nodes or node canisters, and the clustering software helps ensure the deployment of a “replicated hardened state” across nodes. A single node can encounter an *assert condition*, perform a software restart recovery action (capturing first-time debug data), and return to the clustered system without the suspension of I/O.

After a warm restart, the assert condition is cleared and the node rejoins the cluster automatically. Typically, a single node assert restart takes 1–5 minutes. Host data I/O continues as the host OS multipath software redirects the I/O to the partner node of the same I/O group.

The event with error id 2030 is logged upon return of the cluster node to the system.

Right-click the **2030 event** and mark it as *Fixed* to prevent repeated alerts and notifications for the same event.

- ▶ A T2 or Tier 2 recovery is reported in the event log with error code 1001.

The cluster automatically initiated a warmstart to recover from an issue. This process (error code 1001) successfully restored the system and resumed I/O operations without data loss. However, temporary access interruption occurred. Host applications might need a restart, and it is recommended to examine the hosts' file systems afterward. Additionally, check the status of remote connections, replications, and Snapshot mappings.

After a T2 recovery all configuration commands are blocked until you re-enable them so that the unfixed event log entry with error code 1001 is marked as fixed. It is recommended that they are not reenabled until the recovery dumps and trace files from all nodes are collected and are reviewed by IBM Support to confirm that it is safe to do so.

The Service GUI is the preferred method for collecting logs of each node. Open a browser session to the Service GUI at [https://<cluster\\_ip>/service](https://<cluster_ip>/service). Select the **Collect Logs** pane from the left navigation bar, and then select the option to create a support package with the latest statesave.

- ▶ A Tier 3 or T3 Recovery is required when there is no more active cluster node and all nodes of the clustered system report node error 550 or 578. The *Recover System Procedure* recovers the system if the system state is lost from all cluster nodes.

The T3 Recovery procedure restores the system configuration to the state that it was in before the incident that caused this situation. Depending on the type of the IBM Storage Virtualize system and the configuration, this is achieved by retrieving the configuration and hardened system data. This data is stored on either a quorum Mdisk, quorum drive, or an IP quorum set up to store metadata. By using the information stored in the configuration backup `svc.config.backup.xml`, the system's configuration and state can be restored.

**Note:** Attempt to run the *Recover System Procedure* only after a complete and thorough investigation of the cause of the system failure. Try to resolve those issues by using other service procedures first.

Selecting **Monitoring** → **Events** shows information messages, warnings, and issues about the IBM Storage Virtualize system. Therefore, this area is a good place to check for problems in the system.

To display the most important events that must be fixed, use the **Recommended Actions** filter.

If an important issue must be fixed, look for the **Run Fix** button in the upper ribbon with an error message that indicates which event must be fixed as soon as possible. This fix procedure helps resolve problems. It analyzes the system, provides more information about the problem, suggests actions to take with the steps to follow, and finally checks to see whether the problem is resolved.

Always use the fix procedures to resolve errors that are reported by the system, such as system configuration problems or hardware failures.

**Note:** IBM Storage Virtualize systems detect and report error messages. However, events might be triggered by factors external to the system, for example back-end storage devices or the storage area network (SAN).

You can safely mark events as fixed. If the error persists or reoccurs, a new event is logged. To select multiple events in the table, press and hold the Ctrl key while clicking the events you want to fix.

Figure 6-4 shows **Monitoring** → **Events** window with Recommended Run Fix.

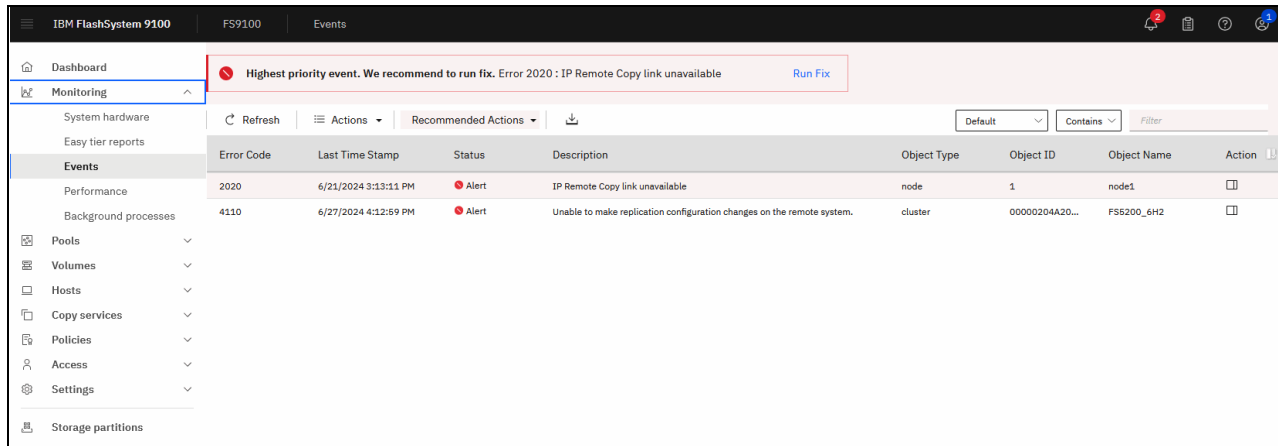


Figure 6-4 Monitoring → Events window

**Resolve alerts in a timely manner:** When an issue or a potential issue is reported, resolve it as quickly as possible to minimize its impact and to potentially avoid more serious problems with your system.

To obtain more information about any event, double-click or select an event in the table, and select **Actions** → **Properties**. You can also select **Run Fix Procedure** and properties by right-clicking an event.

The properties and details are displayed in a panel, as shown in Figure 6-5 on page 114. Sense Data is available in an embedded tab. You can review and click **Run Fix** to run the fix procedure.

The screenshot displays the IBM Storage Virtualize GUI. At the top, a notification banner reads: "Highest priority event. We recommend to run fix. Error 1042 : Managed Enclosure not present" with a "Run Fix" button. Below this is a table of events:

Error Code	Last Time Stamp	Status	Description
1042	3/10/2022 12:41:25 PM	Alert	Managed Enclosure not present
1042	3/10/2022 12:41:25 PM	Alert	Managed Enclosure not present
1042	3/10/2022 12:40:25 PM	Alert	Managed Enclosure not present
1042	3/10/2022 12:40:25 PM	Alert	Managed Enclosure not present
1401	3/29/2022 2:20:22 PM	Alert	Ethernet port failure
1401	3/29/2022 2:20:22 PM	Alert	Ethernet port failure
1620	3/10/2022 12:44:10 PM	Alert	A storage pool is offline
1620	3/10/2022 12:44:10 PM	Alert	A storage pool is offline
1620	3/10/2022 12:44:10 PM	Alert	A storage pool is offline
1690	3/10/2022 12:40:30 PM	Alert	Array mdisk is not protected by sufficient s
1690	3/10/2022 12:40:30 PM	Alert	Array mdisk is not protected by sufficient s
1691	3/8/2022 7:29:40 AM	Alert	Array mdisk is inconsistent
3124	3/29/2022 12:58:25 PM	Alert	No active quorum device found.

The detailed view for "Ethernet port failure" (Error Code: 1401) shows the following properties:

- First Time Stamp: 1648578022000
- Last Time Stamp: 1648578022000
- Fixed Time Stamp: 1
- Event Count: 1

**Properties:**

- Event ID: 071724
- Event ID Text: Fewer ethernet ports operational
- Sequence Number: 3195
- Object Type: node
- Object ID: 181
- Object Name: node2
- Secondary Object ID: 181
- Secondary Object Type: node2
- Copy ID: 181
- Reporting Node ID: 181
- Reporting Node Name: node2
- Root Sequence Number: 181
- Error Code: 1401
- Error Code Text: Ethernet port failure
- Dmp Family: IBM
- Status: alert
- Fixed: no
- Auto Fixed: no
- Notification Type: warning

Figure 6-5 Properties and Sense Data for an event

## 6.1.5 Using the command-line interface

Another option to investigate and resolve issues is to use the IBM Storage Virtualize command-line interface (CLI). Although the *fix procedures* automatically perform the necessary steps, it can sometimes be faster and more convenient to run these commands directly through the CLI. This is particularly the case when numerous events of the same kind need to be fixed, as this can be a strenuous task to click each individual event in the GUI.

Run the commands when you encounter the following issues:

- ▶ You experience a back-end storage or internode issue. For example:
  - Error code 1370: A managed disk (MDisk) error recovery procedure (ERP) occurred.
  - Error code 1630: The number of device logins was reduced.
  - Error code: 1230 or 1231 Login Excluded.
- ▶ You performed maintenance on the following items:
  - Back-end storage subsystems.
  - SAN devices like switches, cables, optical transceivers (SFPs).

**Important:** Run these commands when any type of change that is related to the communication between IBM Storage Virtualize systems and back-end storage subsystem occurs, such as back-end storage is configured or a SAN zoning change occurred. This process helps ensure that IBM Storage Virtualize recognizes the changes.

Common error recovery involves the following IBM Storage Virtualize CLI commands:

- ▶ **detectmdisk**  
Discovers changes in the SAN and back-end storage.
- ▶ **lscntroller** and **lsmdisk**



Provides the status of all controllers and MDisks. Pay attention to status values other than online, for instance *offline* or *degraded*.

► **lscontroller <controller\_id\_or\_name>**

Checks the controller that was causing the issue and verifies that all the worldwide port names (WWPNs) are listed as you expect. Also check whether the `path_counts` are distributed evenly across the WWPNs.

► **lsmdisk**

Determines whether all MDisks are online.

**Note:** When an issue is resolved by using the CLI, verify that the error disappears by selecting **Monitoring** → **Events**. If not, make sure to mark the error as fixed.

## 6.2 Collecting diagnostic data

Problem source identification and problem determination (PSI/PD) can be a challenging task in complex and heterogeneous IT environments. It is crucial to collect the right diagnostic data at the right time to enable the IBM Remote Support teams to assist you in resolving a problem. The following section outlines the steps so that you can collect diagnostic data to find and isolate problems in an IBM Storage Virtualize environment.

### 6.2.1 IBM Storage Virtualize systems data collection

When you encounter a problem with an IBM Storage Virtualize system and you need to open a case with IBM support, you are usually asked to provide a *support package* from the system. The support package is interchangeably referred to as a *Snap* or as *data collection*.

#### Checking for an automatically opened Call Home case

An IBM Storage Virtualize system that is configured for Call Home automatically reports events to IBM. A support case is automatically opened depending on the type of event. It is a good practice to first determine whether a case already exists.

To do so, check the View your cases section in the [IBM Let's Troubleshoot](#) portal.

Storage Virtualize systems that are configured to be monitored in Chapter 5, “IBM Storage Insights and IBM Storage Insights Pro” on page 91, show associated support cases there as well.

Alternatively, you can log in with your IBMid to [IBM Call Home Connect Cloud](#). Call Home Connect Cloud provides an enhanced live view of your assets, including the status of cases, warranties, maintenance contracts, service levels, and end of service information. Additionally, Call Home Connect Cloud offers links to other online tools such as IBM Storage Insights.

#### Determining which data to collect on IBM Storage Virtualize systems

The data that is needed for analysis depends on the type of problem to be analyzed. An IBM Storage Virtualize system stores different kinds of log files, message files, statistics, and traces. The following terms are commonly used in related publications:

- ▶ **Dump.** A node dump or full dump is collected whenever the software restarts for some reason. It is similar in nature to a core dump file, and it can be used by IBM Remote Technical Support and development teams to investigate a problem
- ▶ **Livedump.** A livedump is a binary data capture of the current state of the software. It causes only minimal impact to I/O operations. The contents of a livedump are similar to the contents of a dump with slightly less detailed information.
- ▶ **Statesave.** The term statesave is interchangeably used for either a dump or a livedump.

Four different types of *Snap* can be collected, *Snap Type 1* through *Snap Type 4*, colloquially often referred to as *Snap/1*, *Snap/2*, *Snap/3* or *Snap/4*. The Snap types vary in the amount of diagnostic information that is contained in the package:

- ▶ **Snap/1** includes Standard logs including performance stats. It is the fastest and smallest and contains no node dumps.
- ▶ **Snap/2** is the same as Snap/1 plus one existing statesave, which is the most recently created dump or livedump from the current config node. It is slightly slower than Snap/1 and can be large.
- ▶ **Snap/3** is the same as Snap/1 plus the most recent dump or livedump from each active member node in the clustered system.
- ▶ **Snap/4.** is the same as Snap/1 plus a fresh livedump from each active member node in the clustered, which is created upon triggering the data collection.

### **Statesaves, dumps, and livedumps**

When the Storage Virtualize software stack restarts unexpectedly, a dump file is created and written to a cluster node's boot drive. Similar to a core dump file, this information can be used by IBM Remote Technical Support and development teams to diagnose the cause of the software restart.

A livedump is a binary data capture of the current state of the software. It causes only minimal impact to I/O operations. Livedumps are preferred when the system is still operational and a detailed snapshot of the current state is needed. The contents of a livedump are similar to the contents of a dump with slightly less detailed information. Livedumps can be initiated manually or automatically based on certain events.

### **Determining which Snap to collect**

Two major factors determine which snap to collect for a specific support case:

- ▶ **Speed of collection.** Snap/1 is generated more rapidly, and it is much smaller.
- ▶ **Amount of data.** Collecting Snap/4 as soon as possible after a problem has occurred increases the likelihood that the livedump contains the data required to diagnose the problem.

Consider the following points:

- ▶ For issues that are related to interoperability with hosts or storage, collect Snap/4.
- ▶ For critical performance issues, collect Snap/1 and then collect Snap/4.
- ▶ For general performance issues, collect Snap/4.
- ▶ For issues that are related to replication (including 1920 errors), collect Snap/4 from both systems in the replication partnership.
- ▶ For issues that are related to DRPs, collect Snap/4.
- ▶ For 2030, 1196, or 1195 errors, collect Snap/3.
- ▶ For all other issues, collect Snap/4.

**Tip:** For urgent cases, start with collecting and uploading a Snap/1 followed by a Snap/4. This enables IBM Remote Support to more quickly begin an analysis while the more detailed Snap/4 is being collected and uploaded.

For more information about the required support package that is most suitable to diagnose different type of issues and their content, see [What data should you collect for a problem on IBM Storage Virtualize systems?](#)

**Recommendation:** After an issue is solved, it is a best practice to do some housekeeping and delete old dumps on each node by running the following command:

```
cleardumps -prefix /dumps <node_id | node_name>
```

## Support package collection and upload

The most commonly used method to collect and upload a support package is using the Storage Virtualize GUI. However, the use of Storage Insights is even more convenient because it automates uploading the collected package to IBM. This method is described in 5.2.5, “Managing existing support tickets by using IBM Storage Insights” on page 102.

By default, Storage Virtualize offers two options for automatic support package upload:

- ▶ Automatic upload by using the management interface  
You can configure Storage Virtualize to automatically collect and upload support packages to the IBM Support Center. This can be done through the GUI or CLI.
- ▶ Download and manual upload  
Alternatively, you can use Storage Virtualize to download the support package locally to your device. You can then manually upload it to the IBM Support Center if needed.

## Collecting data by using the GUI

To collect data by using the GUI, complete the following steps:

1. Select **Settings** → **Support** → **Support Package**. You can choose to download the support package to your workstation or upload it to the IBM Support Center. The latter option requires internet connectivity for the Storage Virtualize system, either directly or through a configured proxy server.
2. To automatically upload the support packages, click **Upload Support Package**.
3. Select **Create New Package and Upload**.
4. In the pop-up window, enter the IBM Support case number (TSxxxxxxx) and the type of support package to upload to the IBM Support Center. Press **Upload**. You can monitor the progress of the individual subtasks by clicking **View more details** as shown in Figure 6-6 on page 118.

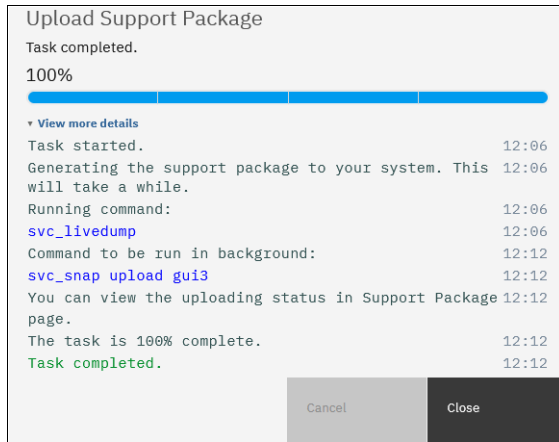


Figure 6-6 Upload Support Package details

## Collecting data by using the CLI

Log in to the CLI and run the command that matches the type of snap that is requested:

- ▶ Standard logs (Snap/1):
 

```
svc_snap gui1 upload pmr=TSxxxxxxxx
```
- ▶ Standard logs plus one existing statesave from current config node (Snap/2):
 

```
svc_snap gui2 upload pmr=TSxxxxxxxx
```
- ▶ Standard logs plus most recent statesave from each active cluster node (Snap/3):
 

```
svc_snap gui3 upload pmr=TSxxxxxxxx
```
- ▶ Standard logs plus new statesaves (Snap/4):
 

```
svc_livedump -nodes all -yes
svc_snap gui3 upload pmr=TSxxxxxxxx
```

To collect a Snap/4 using the CLI, a livedump of each active node must be generated by using the `svc_livedump` command. Then, the log files and newly generated dumps are uploaded by using the `svc_snap gui3` command, as shown in Example 6-1 on page 119. To verify whether the support package was successfully uploaded, use the `sainfo lscmdstatus` command (TSXXXXXX is the case number).

**Note:** The use of Service Assistant commands such as `sainfo` or `satask` requires superuser privileges.

### Example 6-1 The `svc_livedump` command

---

```
IBM_FlashSystem:FS9110:superuser>svc_livedump -nodes all -yes
Livedump - Fetching Node Configuration
Livedump - Checking for dependent vdisks
Livedump - Check Node status
Livedump - Preparing specified nodes - this may take some time...
Livedump - Prepare node 1
Livedump - Prepare node 2
Livedump - Trigger specified nodes
Livedump - Triggering livedump on node 1
Livedump - Triggering livedump on node 2
Livedump - Waiting for livedumps to complete dumping on nodes 1,2
Livedump - Waiting for livedumps to complete dumping on nodes 1
Livedump - Successfully captured livedumps on nodes 1,2

IBM_IBM FlashSystem:FLASHPF95:superuser>svc_snap gui3 upload pmr=TSxxxxxxxxx
Collecting data
Packaging files
Snap data collected in /dumps/snap.serial.YYMMDD.HHMMSS.tgz

IBM_FlashSystem:FS9110:superuser>sainfo lscmdstatus
last_command satask supportupload -pmr TSxxxxxxxxx -filename
/dumps/snap.serial.YYMMDD.HHMMSS.tgz
last_command_status CMMVC8044E Command completed successfully.
T3_status
T3_status_data
cpfiles_status Complete
cpfiles_status_data Copied 160 of 160
snap_status Complete
snap_filename /dumps/snap.serial.YYMMDD.HHMMSS.tgz
installcanistersoftware_status
supportupload_status Active
supportupload_status_data Uploaded 267.5 MiB of 550.2 MiB
supportupload_progress_percent 48
supportupload_throughput_KBps 639
supportupload_filename /dumps/snap.serial.YYMMDD.HHMMSS.tgz
```

---

If you do not want to automatically upload the snap to IBM, omit the `upload pmr=TSxxxxxxxxx` command option. When the snap creation is done, all collected files are packaged into a file that is compressed in gzip format that uses the following format:

```
/dumps/snap.<panel_id>.YYMMDD.hhmmss.tgz
```

The creation of the Snap archive takes a few minutes to complete. Depending on the size of the system and the configuration, it can take considerably longer, particularly if fresh livedumps are being created.

The generated file can be retrieved from the GUI by selecting **Settings** → **Support** → **Manual Upload Instructions** → **Download Support Package**, and then clicking **Download Existing Package**. Find the exact name of the snap that was generated by running the `svc_snap` command that was run earlier. Select that file, and click **Download**.

In certain circumstances it might be necessary to collect a livedump from an individual node of a clustered system at a certain point in time. This can be achieved through CLI commands `preplivedump` and `triggerlivedump`, followed by the targeted node's numeric id or name. The

livedump status of a node can be checked with the command `lslivedump`. When the status changes from *dumping* to *inactive*, the livedump file is ready to be copied off the system by using either the GUI or `scp` command.

*Example 6-2 preplivedump and lslivedump commands*

---

```
IBM_FlashSystem:FS9110:superuser>preplivedump 2
IBM_FlashSystem:FS9110:superuser>lslivedump 2
status
prepared
IBM_FlashSystem:FS9110:superuser>triggerlivedump 2
IBM_FlashSystem:FS9110:superuser>lslivedump 2
status
dumping
IBM_FlashSystem:FS9110:superuser>lslivedump 2
status
inactive
IBM_FlashSystem:FS9110:superuser>lsdumps 2
[...]
livedump.panel_id.YYMMDD.HHMMSS
[...]
```

---

## 6.2.2 Drive data collection: drivedumps

BM FlashCore Modules (FCMs) are a family of high-performance flash drives. The FCM design uses the NVMe protocol, a Peripheral Component Interconnect Express (PCIe) interface, and high-speed NAND memory (for example, Single-Level Cell) driven by a customizable Field-Programmable Gate Array (FPGA). FCMs deliver high throughput, inline compression, and consistent Input/Output Operations Per Second (IOPS) with predictable latency.

For deeper analysis in cases where drives or FCMs are involved, drivedumps are often useful. Drivedumps are particularly useful for troubleshooting issues with FCMs, as they capture the low-level state of the drive. Their data can help you understand problems with the drive, and they do not contain any data that applications write to the drive. In some situations, drivedumps are automatically triggered by the system.

To collect support data from a disk drive, run the `triggerdrivedump drive_id` command. The output is stored in a file in the `/dumps/drive` directory. This directory is located on one of the nodes that are connected to the drive.

Example 6-3 shows the usage of the `triggerdrivedump` command.

*Example 6-3 The triggerdrivedump command*

---

```
IBM_IBM FlashSystem:FS9110:superuser>triggerdrivedump 1
Drive dump on node id [5] successfully created
IBM_IBM FlashSystem:FS9110:superuser>

IBM_IBM FlashSystem:FS9110:superuser>lsdumps -prefix /dumps/drive
id filename
0 drivedump_7812345-1_1_220411_055205
IBM_IBM FlashSystem:FS9110:superuser>
```

---

Any snap that is taken after the trigger command contains the stored drivedumps. It is sufficient to provide Snap/1 for drivedumps.

## 6.2.3 Host multipath software

If a problem occurs that is related to host communication with an IBM Storage Virtualize system, collecting data from hosts and their multipath software is useful.

### Linux using device-mapper-multipath (dmmp)

To troubleshoot by using the multipath CLI, issue the `multipath -ll` command, which shows detailed information about the multipath devices.

Example 6-4 shows the output for the command `multipath -ll`, including the following information:

- ▶ Name of the mpath device (mpatha / mpathb).
- ▶ UUID of the mpath device.
- ▶ Discovered paths for each mpath device, including the name of the sd-device, the priority, and state information.

*Example 6-4 Output for the multipath -ll command*

---

```
root@myServer ~]# multipath -ll
mpatha (3600507680185801aa00000000000b79) dm-3 IBM      ,2145
size=100G features='1 queue_if_no_path' hwhandler='0' wp=rw
| -+- policy='service-time 0' prio=50 status=active
|   | - 16:0:0:2 sd1 8:176 active ready running
|   | ~- 18:0:0:2 sdm 8:192 active ready running
| ~-+- policy='service-time 0' prio=10 status=enabled
|   | - 16:0:1:2 sdg 8:96 active ready running
|   | ~- 18:0:1:2 sdt 65:48 active ready running
mpathb (3600507680185801aa00000000000b78) dm-4 IBM      ,2145
size=100G features='1 queue_if_no_path' hwhandler='0' wp=rw
| -+- policy='service-time 0' prio=50 status=active
|   | - 16:0:1:1 sde 8:64 active ready running
|   | ~- 18:0:1:1 sds 65:32 active ready running
| ~-+- policy='service-time 0' prio=10 status=enabled
|   | - 16:0:0:1 sdj 8:144 active ready running
|   | ~- 18:0:0:1 sdk 8:160 active ready running
```

---

Expand the command to `multipath -ll -v3` to print debug information.

You can also use the `multipathd` interactive console for troubleshooting. The `multipath -k` command opens an interactive interface to the `multipathd` daemon.

Entering this command opens an interactive multipath console. After this command is run, it is possible to enter `help` to get a list of available commands, which can be used within the interactive console. To exit the console, press `Ctrl-d`.

To display the current configuration, including the defaults, issue `show config` within the interactive console.

## AIX using multipath I/O

Table 6-1 shows some of the useful AIX `lspath` commands.

Table 6-1 Useful AIX `lspath` commands

Command	Result
<code>lspath</code>	Lists all paths for all hdisks with their status and parent FSCSI (Fibre Channel SCSI) device information.
<code>lspath -H -l hdisk1</code>	List all paths for the specified hdisk with its status and corresponding FSCSI device information. The output includes a column header.
<code>lspath -l hdisk1 -HF "name path_id parent connection path_status status"</code>	Lists more detailed information about the specified hdisk the parent FSCSI device and its path status.
<code>lspath -s disabled</code>	Lists all paths that have an operational status of disabled.
<code>lspath -s failed</code>	Lists all paths that have an operational status of failed.
<code>lspath -AHE -l hdisk0 -p vscsi0 -w "810000000000"</code>	Display attributes for a path and connection (-w) (-A is like <code>lsattr</code> for devices. If only one path exists to the parent device, the connection can be omitted by running: <code>lspath -AHE -l hdisk0 -p vscsi0</code> )
<code>lspath -l hdisk1 -a priority -F value -p fscsi0 -w 500507680d7e1264,0</code>	Lists the priority for a specific path.

Table 6-2 shows some of the useful AIX `lsmPIO` commands.

Table 6-2 Useful AIX `lsmPIO` commands

Command	Result
<code>lsmPIO</code>	Lists all disks and corresponding paths with state, parent, and connection information.
<code>lsmPIO -q</code>	Shows all disks with vendor ID, product ID, size, and volume name.
<code>lsmPIO -ql hdisk0</code>	Shows detailed disk information like: <ul style="list-style-type: none"> <li>▶ Vendor ID</li> <li>▶ Product ID</li> <li>▶ Capacity</li> <li>▶ Machine Type</li> <li>▶ Model Number</li> <li>▶ Host Group</li> <li>▶ Volume Name</li> <li>▶ Volume Serial Number</li> </ul>
<code>lsmPIO -Sl hdisk0   grep Path</code>	Shows path statistics.
<code>lsmPIO -ar</code>	Lists the parent adapter and remote port information (-a: adapter (local), and -r: remote port).
<code>lsmPIO -are</code>	Lists the parent adapter and remote port error statistics (-e: error).
<code>lsmPIO -z</code>	Lists all multipath I/O (MPIO) statistics.



## Windows using MPIO

Because IBM Storage Virtualize 8.3.0 is the last version that supports *Subsystem Device Driver Device Specific Module* (SDDDSM), you must use native Windows multipathing, which is provided by the installable feature MPIO.

You can manage the multipathing configuration by using the Windows GUI. It is also possible to use the CLI by using the tool `mpclaim.exe`, which is installed by default.

Table 6-3 shows some of the useful Windows `mpclaim.exe` commands.

Table 6-3 Useful Windows `mpclaim.exe` commands

Command	Result
<code>mpclaim.exe -e</code>	View the storage devices that are discovered by the system.
<code>mpclaim.exe -r -i -d _IBM_2145</code>	Manages FC devices with MPIO.
<code>mpclaim.exe -r -u -d _IBM_2145</code>	Removes MPIO management of FC devices.
<code>mpclaim.exe -r -i -d"MSFT2005iSCSIBusType_0x9 "</code>	Manages internet Small Computer Systems Interface (iSCSI) devices with MPIO.
<code>mpclaim.exe -r -u -d"MSFT2005iSCSIBusType_0x9 "</code>	Removes MPIO management of iSCSI devices.
<code>mpclaim.exe -r -i -a""</code>	Manages all storage devices with MPIO.
<code>mpclaim.exe -r -u -a""</code>	Removes MPIO management for all devices.
<code>mpclaim.exe -r</code>	View storage devices that are managed by Microsoft DSM.
<code>mpclaim.exe -L -M&lt;num&gt;</code>	Modifies the load-balancing policy.
<code>mpclaim.exe -s -d</code>	Checks the policy that your volumes are currently using.
<code>mpclaim.exe -s -d &lt;number&gt;</code>	Checks the policy for a specific disk.

Generic MPIO settings can be listed and modified by using Windows PowerShell cmdlets. Table 6-4 shows the PowerShell cmdlets, which can be used to list or modify generic Windows MPIO settings.

Table 6-4 Useful Windows PowerShell cmdlets

Command	Result
<code>Get-MSDSMSupportedHW</code>	The cmdlet lists hardware IDs in the Microsoft Device Specific Module (MSDSM) supported hardware list.
<code>Get-MPIOSetting</code>	The cmdlet gets Microsoft MPIO settings. The settings are as follows: <ul style="list-style-type: none"> <li>▶ PathVerificationState</li> <li>▶ PathVerificationPeriod</li> <li>▶ PDORemovePeriod</li> <li>▶ RetryCount</li> <li>▶ RetryInterval</li> <li>▶ UseCustomPathRecoveryTime</li> <li>▶ CustomPathRecoveryTime</li> <li>▶ DiskTimeoutValue</li> </ul>

Command	Result
Set-MPIOSetting	<p>The cmdlet changes Microsoft MPIO settings. The settings are as follows:</p> <ul style="list-style-type: none"> <li>▶ PathVerificationState</li> <li>▶ PathVerificationPeriod</li> <li>▶ PDORemovePeriod</li> <li>▶ RetryCount</li> <li>▶ RetryInterval</li> <li>▶ UseCustomPathRecoveryTime</li> <li>▶ CustomPathRecoveryTime</li> <li>▶ DiskTimeoutValue</li> </ul>

## VMware using VMware native multipathing

There are two methods that are used to obtain the multipath information from the VMware ESX host:

- ▶ ESXi CLI. Use the CLI to obtain the multipath information when performing troubleshooting procedures.
- ▶ vSphere Client and vSphere Web Client. Use this option when you are performing system maintenance.

### Command-line interface

To obtain logical unit number (LUN) multipathing information from the ESXi host CLI, complete the following steps:

1. Log in to the ESXi host console.
2. To get detailed information about the paths, run **esxcli storage core path list**.

Example 6-5 shows an example for the output of the **esxcli storage core path list** command.

#### *Example 6-5 Output of esxcli storage core path list command*

```
fc.5001438028d02923:5001438028d02922-fc.500507680100000a:500507680120000a-naa.6
00507680185801aa000000000000a68
UID:
fc.5001438028d02923:5001438028d02922-fc.500507680100000a:500507680120000a-naa.6
00507680185801aa000000000 000a68
Runtime Name: vmhba2:C0:T1:L54
Device: naa.600507680185801aa000000000000a68
Device Display Name: IBM Fibre Channel Disk
(naa.600507680185801aa000000000000a68)
Adapter: vmhba2
Channel: 0
Target: 1
LUN: 54
Plugin: NMP
State: active
Transport: fc
Adapter Identifier: fc.5001438028d02923:5001438028d02922
Target Identifier: fc.500507680100000a:500507680120000a
Adapter Transport Details: WWNN: 50:01:43:80:28:d0:29:23 WWPNN:
50:01:43:80:28:d0:29:22
Target Transport Details: WWNN: 50:05:07:68:01:00:00:0a WWPNN:
50:05:07:68:01:20:00:0a
```

Maximum I/O Size: 33553920

---

3. To list detailed information for all the corresponding paths for a specific device, run **esxcli storage core path list -d <naaID>**.

Example 6-6 shows the output for the specified device with the ID naa.600507680185801aa000000000000972, which is attached with eight paths to the ESXi server. The output is abridged for brevity.

*Example 6-6 Output of esxcli storage core path list -d <naaID>*

---

```
fc.5001438028d02923:5001438028d02922-fc.500507680100037e:500507680120037e-naa.600507680185801aa000000000000972
```

UID:

```
fc.5001438028d02923:5001438028d02922-fc.500507680100037e:500507680120037e-naa.600507680185801aa000000000000972
```

Runtime Name: vmhba2:C0:T3:L9

Device: naa.600507680185801aa000000000000972

Device Display Name: IBM Fibre Channel Disk (naa.600507680185801aa000000000000972)

Adapter: vmhba2

Channel: 0

Target: 3

LUN: 9

Plugin: NMP

State: active

Transport: fc

Adapter Identifier: fc.5001438028d02923:5001438028d02922

Target Identifier: fc.500507680100037e:500507680120037e

Adapter Transport Details: WWNN: 50:01:43:80:28:d0:29:23 WWPNN: 50:01:43:80:28:d0:29:22

Target Transport Details: WWNN: 50:05:07:68:01:00:03:7e WWPNN: 50:05:07:68:01:20:03:7e

Maximum I/O Size:

```
33553920fc.5001438028d02923:5001438028d02922-fc.500507680100037e:500507680130037e-naa.600507680185801aa000000000000972
```

UID:

```
fc.5001438028d02923:5001438028d02922-fc.500507680100037e:500507680130037e-naa.600507680185801aa000000000000972
```

Runtime Name: vmhba2:C0:T2:L9

Device: naa.600507680185801aa000000000000972

Device Display Name: IBM Fibre Channel Disk (naa.600507680185801aa000000000000972)

Adapter: vmhba2

Channel: 0

Target: 2

LUN: 9

Plugin: NMP

State: active

Transport: fc

Adapter Identifier: fc.5001438028d02923:5001438028d02922

Target Identifier: fc.500507680100037e:500507680130037e

Adapter Transport Details: WWNN: 50:01:43:80:28:d0:29:23 WWPNN: 50:01:43:80:28:d0:29:22

Target Transport Details: WWNN: 50:05:07:68:01:00:03:7e WWPNN: 50:05:07:68:01:30:03:7e

Maximum I/O Size:

```
33553920fc.5001438028d02921:5001438028d02920-fc.500507680100037e:500507680110037e-naa.600507680185801aa000000000000972
```

UID:

---

4. The command `esxcli storage nmp device list` lists the LUN multipathing information for all attached disks.

Example 6-7 shows the output for one of the attached disks. A portion of the output was omitted for brevity.

*Example 6-7 Output for esxcli storage nmp device list*

---

```
naa.600507680185801aa00000000000a68
  Device Display Name: IBM Fibre Channel Disk
(naa.600507680185801aa00000000000a68)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off; {TPG_id=1,TPG_state=AN0}{TPG_id=0,TPG_state=A0}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useAN0=0; lastPathIndex=1;
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba2:C0:T3:L54, vmhba1:C0:T2:L54, vmhba1:C0:T3:L54,
vmhba2:C0:T2:L54
  Is USB: false
```

---

### ***vSphere Client HTML5 and Web Client***

To obtain multipath settings for your storage in the HTML5 client, complete the following steps:

1. Select an ESXi host, and click the **Configure** tab.
2. Click **Storage Devices**.
3. Select the storage device that you want to verify.
4. Scroll down in the Properties tab and click **Edit multipathing...**

### ***vSphere Client (Thick Client for 6.x)***

To obtain multipath settings for your storage in vSphere Client, complete the following steps:

1. Select an ESXi host, and click the **Configuration** tab.
2. Click **Storage**.
3. Select a data store or mapped LUN.
4. Click **Properties**.
5. In the Properties dialog, select the extent, if necessary.
6. Select **Extent Device** → **Manage Paths** and obtain the paths from the Manage Path dialog.

A more thorough and in-depth discussion of VMware implementation and multipathing can be found in the following IBM Redbooks publication *IBM Storage Virtualize and VMware: Integrations, Implementation and Best Practices*, SG24-8549

## 6.2.4 More data collection

Data collection methods vary by storage platform, SAN switch, and operating system.

For an issue in a SAN environment when it is not clear where the problem is occurring, you might need to collect data from several devices in the SAN.

The following basic information must be collected for each type of device:

- ▶ Hosts:
  - Operating system: Version and level
  - Host Bus Adapter (HBA): Driver and firmware level
  - Multipathing driver level
- ▶ SAN switches:
  - Hardware model
  - Software version
- ▶ Storage subsystems:
  - Hardware model
  - Software version

For performance-related issues, it is helpful to have corresponding monitoring. IBM Storage Insights and IBM Spectrum Control are recommended. If required, you can export performance data from there for the related period taken from the historical data.



# Abbreviations and acronyms

<b>APIs</b>	application programming interfaces	<b>RAID</b>	redundant array of independent disks
<b>CA</b>	Certificate Authority	<b>RDMA</b>	Remote Direct Memory Access
<b>CG</b>	consistency groups	<b>RNICs</b>	RDMA-capable Ethernet NICs
<b>CLI</b>	command-line interface	<b>RoCE</b>	RDMA over Converged Ethernet
<b>CSM</b>	Copy Services Manager	<b>SAN</b>	storage area network
<b>CSR</b>	certificate signing request	<b>SDDDSM</b>	Subsystem Device Driver Device Specific Module
<b>DNS</b>	domain name server	<b>SME</b>	subject matter expert
<b>DR</b>	disaster recovery	<b>SMTP</b>	Simple Mail Transfer Protocol
<b>DRPs</b>	data reduction pools	<b>SSIC</b>	System Storage Interoperation Center
<b>EMEA</b>	Europe, Middle East, and Africa	<b>SVC</b>	SAN Volume Controller
<b>EOM</b>	End-of-Marketing	<b>T1</b>	Tier 1
<b>ERP</b>	error recovery procedure	<b>T4</b>	through Tier 4
<b>FC</b>	Fibre Channel	<b>TCP</b>	Transmission Control Protocol
<b>FCMs</b>	FlashCore Modules	<b>UDP</b>	user datagram protocol
<b>FCP</b>	Fibre Channel Protocol	<b>VASA</b>	vSphere APIs for Storage Awareness
<b>FCoE</b>	Fibre Channel over Ethernet	<b>VVOLs</b>	VMware vSphere virtual volumes
<b>FPGA</b>	Field-Programmable Gate Array	<b>WWNN</b>	worldwide node name
<b>FQDNs</b>	fully qualified domain names	<b>WWPN</b>	worldwide port name
<b>FRU</b>	FCM Field Replaceable Unit	<b>WWPNs</b>	worldwide port names
<b>HA</b>	high availability	<b>anyEKU</b>	Authentication. Additionally, Any Extended Key Usage
<b>HBA</b>	host bus adapter	<b>iSCSI</b>	internet Small Computer Systems Interface
<b>IOPS</b>	input/output operations per second	<b>vVols</b>	VMware Virtual Volumes
<b>ISCSI</b>	internet Small Computer Systems Interface		
<b>LED</b>	light-emitting diode		
<b>LU</b>	logical unit		
<b>LUN</b>	logical unit number		
<b>LUs</b>	logical units		
<b>MDisk</b>	managed disk		
<b>ML</b>	Machine Language		
<b>MPIO</b>	multipath I/O		
<b>MSDSM</b>	Microsoft Device Specific Module		
<b>Mdisks</b>	managed disks		
<b>NPIV</b>	N_Port ID Virtualization		
<b>NTP</b>	Network Time Protocol		
<b>NVMe</b>	Non-Volatile Memory Express		
<b>N_Port</b>	node ports		
<b>PCIe</b>	Peripheral Component Interconnect Express		
<b>PTF</b>	Program Temporary Fix		





# Related publications

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

## IBM Redbooks

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

- ▶ *Implementation Guide for IBM Storage FlashSystem and IBM SAN Volume Controller Updated for IBM Storage Virtualize Version 8.6*, SG24-8542
- ▶ *IBM Storage Virtualize and VMware: Integrations, Implementation and Best Practices*, SG24-8549
- ▶ *Ensuring Business Continuity: A Practical Guide to Policy-Based Replication and Policy-Based High Availability for IBM Storage Virtualize Systems*, SG24-8569
- ▶ *Introduction and Implementation of Data Reduction Pools and Deduplication*, SG24-8430
- ▶ *IBM Storage Insights Security Guide*, SC27-8774
- ▶ *Data Resiliency Designs: A Deep Dive into IBM Storage Safeguarded Snapshots*, REDP-5737

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

[ibm.com/redbooks](https://ibm.com/redbooks)

## Online resources

These websites are also relevant as further information sources:

- ▶ IBM Storage FlashSystem information:  
<https://www.ibm.com/flashsystem/>
- ▶ IBM SAN Volume Controller information:  
[https://www.ibm.com/products/san-volume-controller?mhsrc=ibmsearch\\_a&mhq=SAN%20Volume%20Controller](https://www.ibm.com/products/san-volume-controller?mhsrc=ibmsearch_a&mhq=SAN%20Volume%20Controller)
- ▶ IBM System Storage Interoperation Center (SSIC):  
<https://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

## Help from IBM

IBM Support and downloads

[ibm.com/support](https://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)

**Redbooks**

**Unleash the Power of Flash: Getting Started with IBM Storage Virtualize**

(0.2"spine)  
0.17"->0.473"  
90->249 pages







SG24-8561-00

ISBN 0738461776

Printed in U.S.A.

Get connected

