

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

Andy Harchen

Hartmut Lonzer

Jon Herd

Jonathan Wilkie

Vasfi Gucer



Storage





IBM Redbooks

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

July 2024

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

**First (July 2024)**

This edition applies to IBM Storage Virtualize Version 8.7.0.

This document was created or updated on July 16, 2024.



# 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 Benefits of IBM Storage Virtualize .....	2
1.2 IBM Storage Virtualize V8.7 supported product list .....	3
1.2.1 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 IBM Storage Virtualize V8.7.0 .....	6
1.4 Preparation and upgrading to IBM Storage Virtualize V8.7.0 .....	13
<b>Chapter 2. Initial configuration</b> .....	17
2.1 Prerequisites .....	18
2.2 System initialization .....	19
2.2.1 System initialization process .....	21
2.3 System setup .....	26
2.3.1 System Setup wizard .....	26
2.3.2 Configuring clustering by using Ethernet connections .....	51
2.3.3 Adding an enclosure in IBM FlashSystem .....	57
2.3.4 Adding a node or hot spare node in IBM SAN Volume Controller systems .....	59
2.3.5 Business continuity with policy-based high availability .....	63
2.3.6 Configuring quorum disks or applications .....	64
2.3.7 Configuring the local Fibre Channel port masking .....	67
2.3.8 Automatic configuration for IBM SAN Volume Controller back-end storage .....	69
<b>Chapter 3. Step-by-step configuration</b> .....	33
3.1 The Storage Virtualize GUI .....	34
3.1.1 Accessing the GUI .....	34
3.1.2 Brief introduction to the GUI .....	34
3.2 Network configuration .....	35
3.2.1 Management IP addresses .....	35
3.2.2 Service IP addresses .....	36
3.2.3 Additional Ethernet ports .....	36
3.2.4 Portsets .....	37
3.3 Pools and managed disks configuration .....	37
3.3.1 Provisioning policies .....	38
3.3.2 Types of pools .....	38

3.3.3	Ransomware threat detection . . . . .	38
3.3.4	Creating storage pools . . . . .	39
3.3.5	Creating RAID array managed disks in a storage pool. . . . .	40
3.3.6	Adding external managed disks into a storage pool . . . . .	41
3.3.7	Child pools . . . . .	42
3.4	Configuring volumes . . . . .	44
3.4.1	Creating volume groups . . . . .	44
3.4.2	Creating volumes . . . . .	44
3.4.3	Virtual volumes . . . . .	45
3.5	Configuring hosts . . . . .	45
3.5.1	Host attachment overview. . . . .	45
3.5.2	Fibre channel host connectivity. . . . .	46
3.5.3	Ethernet host connectivity. . . . .	47
3.5.4	Host objects . . . . .	48
3.5.5	Mapping volumes for host access. . . . .	50
3.6	Snapshots and replication. . . . .	50
3.6.1	Volume group snapshots . . . . .	50
3.6.2	Asynchronous policy-based replication. . . . .	52
	<b>Chapter 4. Verifying configuration and basic operations . . . . .</b>	<b>57</b>
4.1	Verifying the configuration. . . . .	58
4.1.1	System Health Dashboard . . . . .	58
4.1.2	Verifying network configuration. . . . .	58
4.1.3	Verifying storage configuration . . . . .	59
4.1.4	Verifying volume configuration . . . . .	59
4.1.5	Verifying host configuration. . . . .	60
4.2	Additional settings and basic operations. . . . .	61
4.2.1	Security settings . . . . .	61
4.2.2	Audit log . . . . .	63
4.2.3	Support settings . . . . .	63
4.2.4	Data migration. . . . .	64
4.2.5	SCSI unmap . . . . .	66
4.2.6	I/O throttling . . . . .	66
	<b>Chapter 5. IBM Storage Insights and IBM Storage Insights Pro. . . . .</b>	<b>69</b>
5.1	IBM Storage Insights overview . . . . .	70
5.1.1	IBM Storage Insights: Information and registration. . . . .	71
5.2	IBM Storage Insights monitoring. . . . .	71
5.2.1	Component health. . . . .	72
5.2.2	Capacity monitoring . . . . .	73
5.2.3	Performance monitoring . . . . .	74
5.2.4	Logging support tickets by using IBM Storage Insights . . . . .	76
5.2.5	Managing existing support tickets by using IBM Storage Insights . . . . .	82
5.2.6	Enhancements to IBM Storage Insights Pro. . . . .	85
	<b>Chapter 6. Storage Virtualize troubleshooting and diagnostics. . . . .</b>	<b>87</b>
6.1	Troubleshooting . . . . .	88
6.1.1	Storage Insights . . . . .	89
6.1.2	Using the GUI . . . . .	89
6.1.3	Recommended actions and fix procedure. . . . .	90
6.1.4	Storage Virtualize failure recovery . . . . .	91
6.1.5	Using the command-line interface. . . . .	94
6.2	Collecting diagnostic data . . . . .	95
6.2.1	IBM Storage Virtualize systems data collection . . . . .	95

6.2.2 Drive data collection: drivedumps . . . . .	100
6.2.3 Host multipath software . . . . .	101
6.2.4 More data collection . . . . .	107
6.3 Common problems and isolation techniques . . . . .	107
6.3.1 Interoperability . . . . .	108
6.3.2 Host problems . . . . .	108
6.3.3 Fibre Channel SAN and IP SAN problems . . . . .	113
6.3.4 Port issues and transceiver statistics . . . . .	115
6.3.5 Storage subsystem problems . . . . .	116
6.3.6 IP replication problems . . . . .	121
6.3.7 Short-distance partnership using RDMA . . . . .	122
6.3.8 Policy-based replication . . . . .	123
6.3.9 Data reduction pools . . . . .	123
6.3.10 Managing the physical capacity of overprovisioned storage controllers . . . . .	123
<b>Related publications . . . . .</b>	<b>125</b>
IBM Redbooks . . . . .	125
Online resources . . . . .	125
Help from IBM . . . . .	125



# 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 . . . . .	14
2-1 Technician port FlashSystem 9500 . . . . .	20
2-2 Technician port FlashSystem 7300 . . . . .	20
2-3 Technician port FlashSystem 5200 . . . . .	20
2-4 Technician port FlashSystem 5300 . . . . .	20
2-5 Technician port FlashSystem 5045 . . . . .	20
2-6 Technician port FlashSystem 5015 . . . . .	21
2-7 Technician port IBM SAN Volume Controller 2145-SV3 . . . . .	21
2-8 Technician port IBM SAN Volume Controller 2145-SV2 . . . . .	21
2-9 Logging in to Service Assistant by way of the technician port . . . . .	23
2-10 System Initialization: Canister detection . . . . .	23
2-11 System Initialization: Initialize the first enclosure . . . . .	24
2-12 System Initialization: Initialize the first IBM SAN Volume Controller node . . . . .	24
2-13 System Initialization: Enter Management IP . . . . .	25
2-14 System Initialization: Web-server restart timer counting down from 5 minutes. . . . .	25
2-15 System Initialization completed. . . . .	26
2-16 Logging in for the first time . . . . .	27
2-17 Initial Setup Window . . . . .	28
2-18 Setup Call Home . . . . .	29
2-19 Transmission Types for Call Home . . . . .	30
2-20 Setup Internal Proxy Server . . . . .	31
2-21 Connection Test to the Support Center . . . . .	32
2-22 System Location . . . . .	33
2-23 Summary page . . . . .	34
2-24 System Setup Welcome page. . . . .	35
2-25 Accept License Agreement. . . . .	36
2-26 Change password . . . . .	37
2-27 System Name . . . . .	38
2-28 License Functions . . . . .	39
2-29 DNS Server setup . . . . .	40
2-30 Date and Time . . . . .	41
2-31 Activate Encryption License . . . . .	42
2-32 Encryption licensed . . . . .	43
2-33 Change Call Home settings . . . . .	44
2-34 Setup Support Assistance. . . . .	45
2-35 System communicating with named IBM Support servers . . . . .	46
2-36 Remote support access settings. . . . .	47
2-37 Automatic Configuration for Virtualization. . . . .	48
2-38 Summary Page . . . . .	49
2-39 System Initialization . . . . .	50
2-40 Setup completed . . . . .	50
2-41 Dashboard . . . . .	51
2-42 No ISL connectivity . . . . .	52
2-43 Shared ISL connectivity . . . . .	53

2-44	Dedicated ISL connectivity . . . . .	54
2-45	Node IP address setup for Remote Direct Memory Access clustering . . . . .	55
2-46	Node IP addresses configured . . . . .	56
2-47	Setting the node discovery subnet . . . . .	57
2-48	Add Enclosure button . . . . .	58
2-49	Selecting the control enclosure to add . . . . .	58
2-50	Add Node button . . . . .	60
2-51	Adding a node . . . . .	60
2-52	IBM SAN Volume Controller is adding node to the cluster . . . . .	61
2-53	Node added . . . . .	61
2-54	Download IPv4 quorum button . . . . .	65
2-55	Download IP quorum application window . . . . .	65
2-56	IP quorum application that is deployed and connected . . . . .	66
2-57	Changing the quorum mode . . . . .	67
2-58	Applying a port mask by using a GUI . . . . .	68
2-59	Modify Connection dialog box . . . . .	68
2-60	Automatic Configuration wizard enablement . . . . .	70
2-61	Automatic configuration: Add Enclosure . . . . .	70
2-62	Defining a host cluster . . . . .	71
2-63	Hosts inside an IBM SAN Volume Controller host cluster . . . . .	72
2-64	Begin the automatic configuration process . . . . .	72
2-65	Automatic pool configuration . . . . .	73
2-66	Pools configuration . . . . .	73
2-67	Automatic configuration running commands . . . . .	74
2-68	Automatic configuration complete . . . . .	74
3-1	Welcome page with the dashboard . . . . .	34
3-2	Ethernet ports . . . . .	36
3-3	Add IP address . . . . .	37
3-4	Portsets . . . . .	37
3-5	Create Pool . . . . .	39
3-6	Create Pool panel . . . . .	39
3-7	Add Storage . . . . .	41
3-8	RAID array . . . . .	41
3-9	Child pools with different purposes . . . . .	42
3-10	Create Child Pool . . . . .	43
3-11	Add Host . . . . .	49
3-12	Snapshot Policies . . . . .	51
3-13	Suspend Policy . . . . .	52
3-14	Create Partnership . . . . .	53
3-15	Partnership Created . . . . .	54
3-16	Setup policy-based replication wizard . . . . .	54
4-1	System Health . . . . .	58
4-2	Portset Mappings . . . . .	58
4-3	Portset IP addresses . . . . .	59
4-4	Host with asymmetrical logins . . . . .	60
4-5	Audit log . . . . .	63
5-1	IBM Storage Insights System overview (classic view) . . . . .	71
5-2	IBM Storage Insights System overview (Carbon enhanced view) . . . . .	72
5-3	Component Health overview . . . . .	72
5-4	Ports in error . . . . .	73
5-5	Capacity area of the IBM Storage Insights system overview . . . . .	73
5-6	Capacity planning for one system . . . . .	74
5-7	System overview: Performance . . . . .	74

- 5-8 IBM Storage Insights: Performance view . . . . . 75
- 5-9 Filtered performance graph. . . . . 75
- 5-10 Performance List View . . . . . 76
- 5-11 Get Support (see highlighted area). . . . . 76
- 5-12 Get Support window . . . . . 77
- 5-13 Create Ticket wizard . . . . . 78
- 5-14 Add a note or attachment window. . . . . 79
- 5-15 Selecting a Severity Level window . . . . . 80
- 5-16 Review the ticket window . . . . . 81
- 5-17 Update ticket. . . . . 82
- 5-18 View tickets. . . . . 82
- 5-19 Adding a log package to the ticket . . . . . 83
- 5-20 Confirming the log upload. . . . . 84
- 5-21 Log upload completed and processing . . . . . 85
- 5-22 IBM Storage insights Pro and IBM Flash Grid integration. . . . . 86
- 6-1 Events icon in the GUI . . . . . 89
- 6-2 System Health expanded section in the dashboard . . . . . 90
- 6-3 Recommended actions . . . . . 91
- 6-4 Monitoring → Events window . . . . . 93
- 6-5 Properties and Sense Data for an event . . . . . 94
- 6-6 Upload Support Package details. . . . . 98





# Tables

1-1 IBM Storage Virtualize V8.7 supported product list . . . . .	4
1-2 IBM FlashSystem feature summary comparison chart . . . . .	4
6-1 Useful AIX lspath commands . . . . .	102
6-2 Useful AIX lsmpio commands . . . . .	102
6-3 Useful Windows mpclaim.exe commands . . . . .	103
6-4 Useful Windows PowerShell cmdlets . . . . .	103
6-5 Selected attributes of the lsportstats output . . . . .	115



# Examples

2-1	Reenabling the onboard Ethernet port 2 as the technician port	19
2-2	Listing node IPs currently not set	56
2-3	Running commands to change node IP	56
2-4	Changed node IP (output shortened for clarity)	56
2-5	Listing the I/O groups	59
2-6	Listing the candidate control enclosures	59
2-7	Adding a control enclosure	59
2-8	Adding an expansion enclosure	59
2-9	Listing I/O groups	62
2-10	Listing the candidate nodes	62
2-11	Adding a node as a spare	62
2-12	Adding a node to an I/O group	62
2-13	Single IO-group (two nodes) and one spare	62
2-14	Two IO-groups (four nodes) configured- no spare	63
2-15	Starting the IP quorum application on the Windows operating system	65
2-16	Viewing the local port mask	69
2-17	Setting a local port mask by running the chsystem command	69
4-1	Check if host_unmap is enabled	66
6-1	The svc_livedump command	99
6-2	preplivedump and lsplivedump commands	100
6-3	The triggerdrivedump command	100
6-4	Output for the multipath -ll command	101
6-5	Output of esxcli storage core path list command	104
6-6	Output of esxcli storage core path list -d <naaID>	105
6-7	Output for esxcli storage nmp device list	106
6-8	The lshost command	108
6-9	The lshost <host_id_or_name> command	109
6-10	The lsfabric -host <host_id_or_name> command	109
6-11	Incorrect WWPN zoning	114
6-12	Correct WWPN zoning	114
6-13	lsportstats command output	115
6-14	Issuing a lsmdisk command	117
6-15	Output of the svcinfo lscontroller command	119
6-16	Determining the ID for the MDisk	120



# 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®	Interconnect®
Easy Tier®	IBM FlashCore®	PowerHA®
FlashCopy®	IBM FlashSystem®	Redbooks®

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.

OpenShift, Red Hat, 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

Previously known as IBM Spectrum Virtualize, IBM Storage Virtualize is a cornerstone of the IBM Storage portfolio. It is a versatile storage solution enabling rapid deployment of block storage, including storage virtualization, for various workloads. This flexibility applies to on-premises, off-premises, or hybrid cloud environments.

IBM Storage FlashSystem® and SAN Volume Controller systems leverage IBM Storage Virtualize software to simplify infrastructure management. This software eliminates discrepancies in managing, functioning, and even supporting hybrid multicloud deployments.

Based on IBM Storage Virtualize Version 8.7, this IBM Redbooks® publication equips you with the practical knowledge to get started with these storage systems. We will explore essential considerations for implementation, guiding you through critical decisions like performance optimization and troubleshooting.

This book is intended for pre-sales and post-sales technical support personnel, as well as storage administrators.

## Authors

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



**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.



**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.



**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).



**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.



**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.

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](https://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](https://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 defines the concept of storage virtualization and provides an overview of its application in addressing the challenges of modern storage environment.

This chapter has the following sections:

- ▶ “IBM Storage Virtualize” on page 2
- ▶ “IBM Storage Virtualize V8.7 supported product list” on page 3
- ▶ “Changes and enhancements in IBM Storage Virtualize V8.7.0” on page 6
- ▶ “Preparation and upgrading to IBM Storage Virtualize V8.7.0” on page 13

## 1.1 IBM Storage Virtualize

IBM Storage Virtualize (previously known as IBM Spectrum Virtualize) is a key member of the IBM Storage portfolio. It is a highly flexible storage solution that enables rapid deployment of block storage including storage virtualization, for new and traditional workloads, on-premises, off-premises, or a combination of both.

**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 has not changed.

IBM FlashSystem Storage Systems and IBM SAN Volume Controllers are built with award-winning IBM Storage Virtualize software that simplifies infrastructure and eliminates the differences in management, function, and even hybrid multicloud support.

IBM Storage Virtualize is an offering that is available for the IBM SAN Volume Controller and IBM FlashSystem family of storage solutions. It provides a way to manage and protect huge volumes of data from mobile and social applications, enable rapid and flexible cloud services deployments, and deliver the performance and scalability that is needed to gain insights from the latest analytics technologies.

### 1.1.1 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 while enabling 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 easy 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 on page 3 shows the current IBM FlashSystem and IBM SAN Volume Controller Family.








<h2 style="margin: 0;">FlashSystem Family</h2>  <p style="text-align: center;"><b>FlashSystem 5015 &amp; 5045</b></p> <p>SAS based entry to the world of IBM Storage Virtualize and IBM Storage FlashSystem.</p>	<p style="text-align: center; color: #0070C0;">Released May 2024</p>  <p style="text-align: center; color: #0070C0;"><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 style="text-align: center;"><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 style="text-align: center;"><b>FlashSystem 9500</b></p> <p>Extreme performance for large, mixed and consolidated workloads</p>
<h2 style="color: #0070C0; margin: 0;">Powered by IBM Storage Virtualize</h2>			
 <p style="text-align: center;"><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 style="text-align: center;"><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 style="text-align: center;"><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 functionality is planned for a future release.

## 1.2 IBM Storage Virtualize V8.7 supported product list

Table 1-1 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
FS5000 (FS5045)	4680	3P2, 3P4	Current Product
FS5000 (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.

## 1.2.1 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 all models that can run Storage Virtualize V8.7.

Table 1-2 IBM FlashSystem feature summary comparison chart

	SVC	5015	5045	5300	7300	9500
<b>Machine Type</b>	2145 2147	4680	4680	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)

	SVC	5015	5045	5300	7300	9500
<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 128GB 386GB 764GB SV3 512GB 1TB 1.5TB	32GB 64GB	32GB 64GB	64GB 256GB 512GB	256GB 756GB 1.5TB	1TB 2TB 3TB
<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>	64 Gb/s FC or NVMe/FC (SV3) 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 1M SV3 2.5M	140k	400k	700k	1M	2.5M
<b>Max bandwidth (256K read miss)</b>	SA2 45GB/s SV3 100GB/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 <ul style="list-style-type: none"> <li>• Height: 8.7 cm (3.4 in.)</li> <li>• Width 44.6 cm (17.6 in)</li> <li>• Depth 82.6 cm (32.5 in)</li> </ul>	Control enclosure <ul style="list-style-type: none"> <li>• Height: 8.7 cm (3.4 in.)</li> <li>• Width: 48.3 cm (19.0 in.)</li> <li>• Depth: 55.6 cm (21.9 in.)</li> </ul>	Control enclosure <ul style="list-style-type: none"> <li>• Height: 8.7 cm (3.4 in.)</li> <li>• Width: 48.3 cm (19.0 in.)</li> <li>• Depth: 55.6 cm (21.9 in.)</li> </ul>	Control enclosure <ul style="list-style-type: none"> <li>• Height: 4.3 cm (1.7 in.)</li> <li>• Width: 44.6 cm (17.5 in.)</li> <li>• Depth: 77 cm (30.3 in.)</li> </ul>	Control enclosure <ul style="list-style-type: none"> <li>• Height: 8.8 cm (3.5 in.)</li> <li>• Width: 48.3 cm (19.0 in.)</li> <li>• Depth: 85 cm (33.5 in.)</li> </ul>	Control enclosure <ul style="list-style-type: none"> <li>• Height: 17.43 cm (6.8 in.)</li> <li>• Width: 44.6 cm (17.6 in.)</li> <li>• Depth: 82.6 cm (32.6 in.)</li> </ul>
<b>Weight</b>	25 kg (55 lb) to 30 kg (65 lb) depending on configuration	Fully configured (12 drives): 28.3Kg Fully configured (24 drives): 27.3Kg	Fully configured (12 drives): 28.3Kg Fully configured (24 drives): 27.3Kg	Fully configured (12 drives): 19.5Kg	Fully configured (24 drives): 46.6Kg	Fully configured (48 drives): 70.5Kg

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

This section describes the latest changes and enhancements in IBM Storage Virtualize V8.7.0.

**Important:** *IBM Storage Virtualize V8.7.0 will be the final release supporting all Remote Copy based features, including HyperSwap®, Metro Mirror, Global Mirror and Global Mirror with Change Volumes.* When planning new deployments, consider using policy-based replication and policy-based high availability to avoid causing an unnecessary migration.

### 1.3.1 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.

For more information, see [IBM System Storage Interoperation Center \(SSIC\)](#).

#### 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 used on each volume can significantly 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 may 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 what file systems are in which volumes in some recovery scenarios.

IBM Storage Virtualize V8.7.0 provides this new file level awareness for ransomware detection as follows:

1. Every 12 hours, the file system is automatically updated for each volume.
  - Can also updated by **analyzevdisk** or **analyzevdiskbysystem** CLI commands.
2. Background reads are sent to a volume.
3. Open source libraries used to determine file system.
4. Output is displayed in *lsvdiskanalysis's* file\_system field.
  - 15 character max for field.
  - Can display multiple file systems.
5. File system used by inferencing engine to improve ransomware detection.

### GUI updates

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

#### ***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.

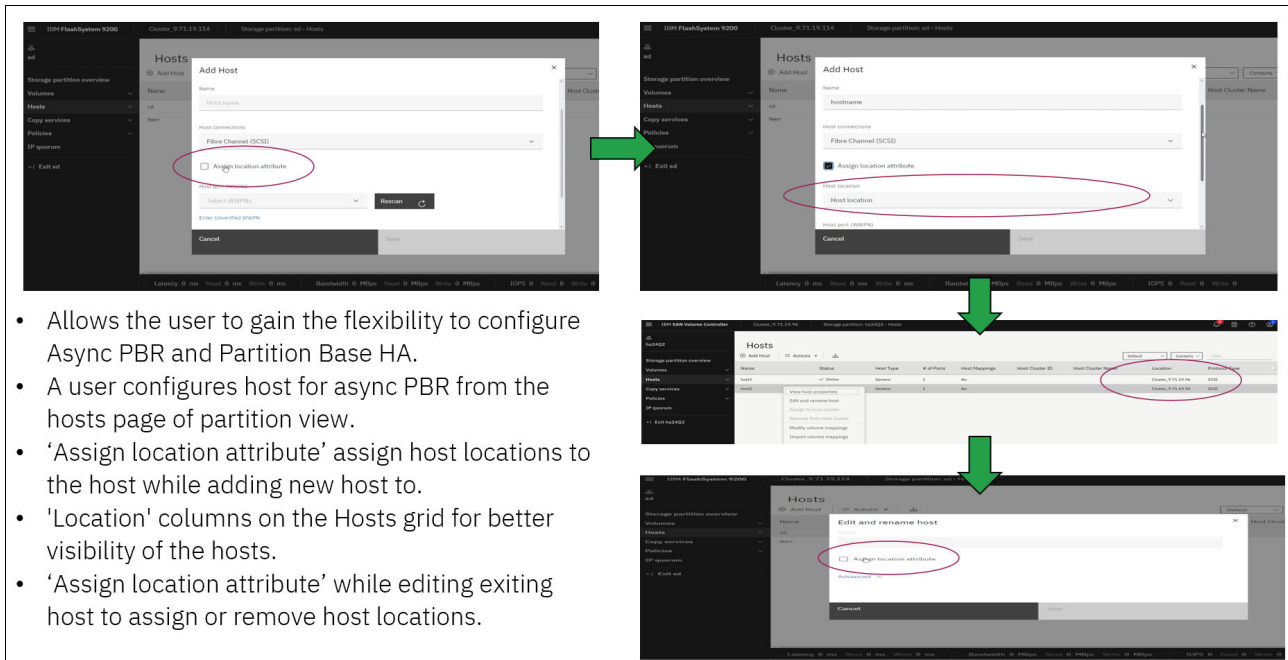


Figure 1-2 Async policy-based replication and partition base HA user experience improvements

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

Figure 1-3 on page 8 shows the volume group tile and assigning ownership groups to volume group enhancement.

- A user will be able to look on their dashboard logical components view and see a tile for volume groups and navigate to volume groups page.
- Ownership groups for restricting the volume groups access to users.
- If the user –
  - unrestricted user, the grid will show all volume groups
  - restricted user, the grid will only show the volume groups that have the same ownership group as that restricted user.
  - unrestricted user can assign ownership group while creating the volume group.

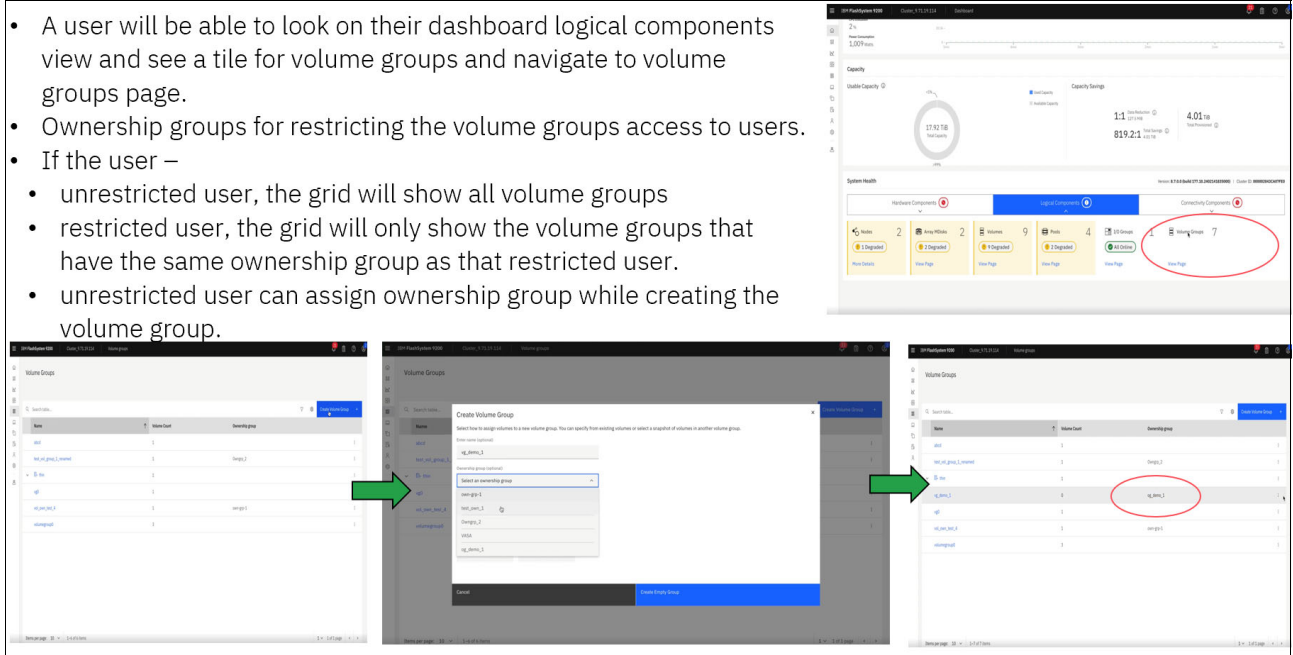


Figure 1-3 Volume group tile and assigning ownership groups to volume groups

### Performance panel carbonization

Storage Virtualize introduces GUI performance panel modernization 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.
- Easy to compare the IOPS, bandwidth and latency through 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.

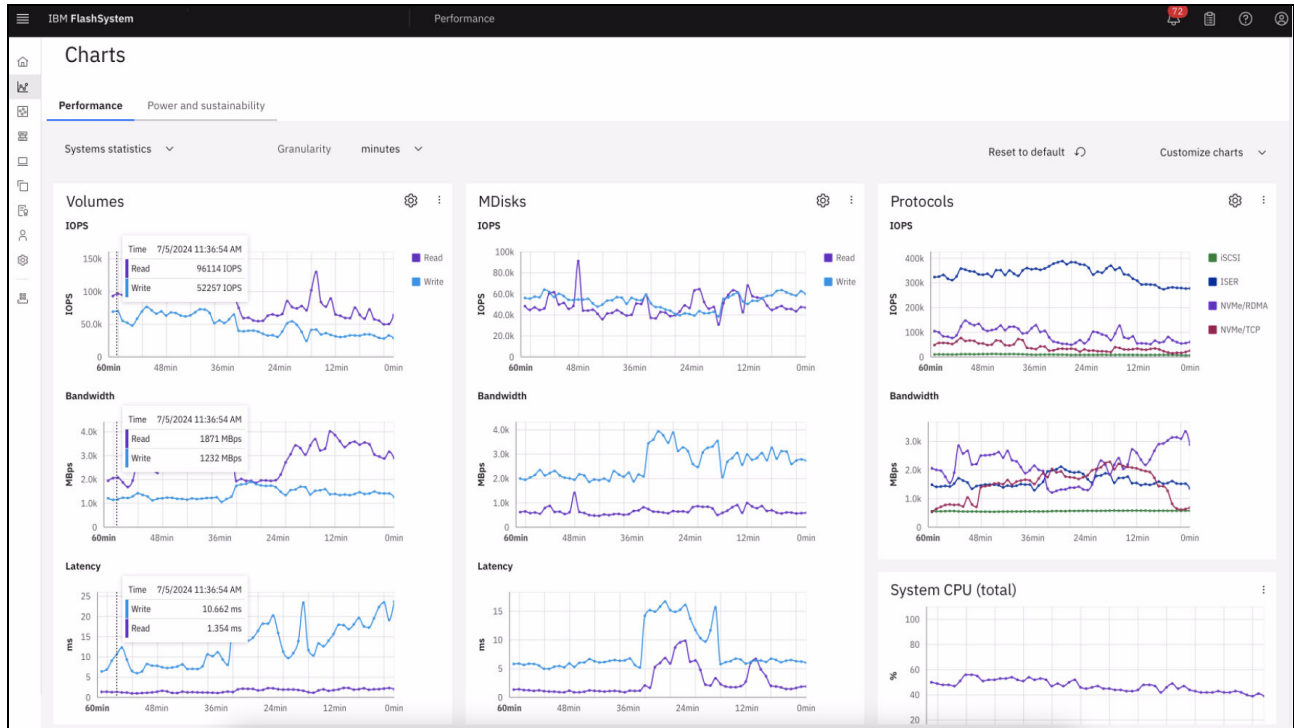


Figure 1-4 GUI performance panel

## Flash Grids and storage partition migration

A *Flash Grid* is a collection of single I/O group FlashSystem or SVC systems that *looks and feels* like a single storage solution, providing single pane of glass management and non-disruptive 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.
- ▶ **Non-linear object limits:** Volume, snapshot, and host counts are limited per system, not per I/O group, hindering scalability.
- ▶ **Feature restrictions:** Several advanced features, like 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.*

Flash Grid brings a novel approach to storage management, and Storage Virtualize 8.7.0 marks the first phase of its implementation. Here are the key features available in this release:

- ▶ **CLI-driven Flash Grid management:** Initial configuration and management are primarily done through the command-line interface (CLI), which uses AI-assisted storage partition migration utilizing 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 significant 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, designated as the Flash Grid owner, manages the membership of other systems within the grid.
- ▶ **Tiering for optimized performance:** Flash Grid empowers you to combine different hardware within a single grid, enabling you to tier storage based on performance and capacity needs.
- ▶ **Independent software updates:** Each member system in a Flash Grid can receive software updates independently, offering greater flexibility and reduced downtime for maintenance.

**Note:** It is planned to make the Flash Grid implementation, monitoring and management as part of the GUI in a future release. There will also be tighter integration with IBM Storage Insights to give AI capable operation to storage partitions migration across systems in the Flash Grid.

## Automatic Patch Updating and Automatic Drive Downloading

Storage Virtualize V8.7 simplifies 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 lightweight update to a function or service, which can be installed on a user's system. A patch install never requires a node reboot or reset.*

**Note:** A patch install may 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 and patches are stored on [IBM FixCentral](#). IBM Cloud® Call Home is used to access patches. Newer versions of IBM Storage Virtualize code can *include* older patches released in previous Storage Virtualize versions.

When developers identify a need, they create patches (such as bug fixes, security updates and so forth) to address issues in Storage Virtualize. These patches, along with existing ones, are published on IBM FixCentral, a centralized repository for software updates.

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

- ▶ The enhanced Patching framework empowers you to *schedule automatic patch application* for your Storage Virtualize systems. This eliminates the need for complex full PTF (Program Temporary Fix) or concurrent code upgrades, saving you valuable time and effort.
- ▶ Benefits users whose systems have patches that might need be frequently updated.

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

- ▶ Users can set up their systems in the knowledge that updating of vital patches will happen seamlessly, in the background.
- ▶ Automatic Patch Updating can be set up on a user's system, either using the GUI, or via CLI commands.

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

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

### ***Automatic Drive Download***

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

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

The Automatic Drive Download process is as follows:

- ▶ Building on the patching infrastructure, FCM drive firmwares are now built as patches.
- ▶ You can use the `applysoftware <firmware_patch>` CLI command to copy the firmware patch to all cluster nodes, streamlining the update process.
- ▶ When a new drive is added to the cluster, it automatically checks for and applies the latest available firmware, ensuring your storage remains up-to-date.
- ▶ Firmware patches are located on IBM FixCentral, alongside other software updates for Storage Virtualize.
- ▶ Newer versions of Storage Virtualize may include older firmware patches, eliminating the need to search for and apply them individually.

The benefits of Automatic Drive Download are:

- ▶ Automatic Drive Download allows you to easily maintain a standard firmware level across the system.
- ▶ You can be confident that any Field Replaceable Unit (FRU) replacements or drives added to your new array will be compatible with your system's firmware.
- ▶ Enabling automatic drive downloads ensures your array's firmware stays up-to-date. This not only optimizes performance but also unlocks access to exciting new features like Ransomware Threat Detection.

### ***Example scenario***

Let us consider the following scenario:

1. User has a FCM4 array using 4.1.4 firmware version. A drive fails and requires replacement.
2. A FRU arrives with version 4.0.4. The user performs the Dynamic Drive Pool (DMP) operation and replaces the failed drive with the replacement FRU.
3. As the drive attempts to rejoin the array, Automatic Drive Download functionality seamlessly takes over, verifies that version 4.1.4 is available and upgrades the drive.

## 64G Fibre Channel updates

Introducing support for dual port 64G Fibre Channel (FC) adapters on IBM FlashSystem 7300 and IBM FlashSystem 5300 systems. This will enable customers to use 64G FC adapters on low and mid end FlashSystem platform.

**Note:** At the time of writing this implementation is expected to be available in 2H 2024.

## 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.

Some of these changes are:

- ▶ Ability to configure system IP address on any port and with VLAN.
- ▶ Requirement of more routable data IP addresses:
  - Have 4 routable data IP addresses per port per node.
- ▶ System defined default management port set for system IP addresses.
- ▶ Limiting 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.

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

*IBM Storage Virtualize V8.7.0 is the final version that supports Remote Copy.*

**Note:** Remote Copy will be supported long-term on V8.7.0 for as long as 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 will not have replication capabilities if upgraded beyond V8.7.0.

- ▶ Global Mirror and Global Mirror with Change Volumes are replaced by policy-based replication.

- A migration procedure is available in the [product documentation](#).
- ▶ HyperSwap is replaced by policy-based HA.
- ▶ Migration is using storage partition migration.

**Note:** Refer to the IBM Redbooks *Implementation Guide for IBM Storage FlashSystem and IBM SAN Volume Controller Updated for IBM Storage Virtualize Version 8.6*, SG24-8542 for further details on Remote Copy and its related features.

### 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 [Storage Insights Overview](#).

**For more information on what is new with Storage Virtualize V8.7:** The following blog post [IBM Storage Virtualize 8.7.0 including Flash Grid](#) (*Barry Whyte and Andrew Martin : IBM Storage blog*) provides a good overview of the new features in Version 8.7.

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

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

Firstly, you need to confirm that your current hardware is able to support IBM Storage Virtualize V8.7.0.

Figure 1-5 on page 14 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 will be IBM Storage Virtualize 8.7.0.

IBM SAN Volume Controller and IBM FlashSystem 5015, 5035, 5200, 7200, 7300, 9100, 9200 & 9500										
Upgrade Support Matrix										
FROM / TO	8.5.0	8.5.1	8.5.2	8.5.3	8.5.4	8.6.0	8.6.1	8.6.2	8.6.3	8.7.0
8.5.0	n/a	Y	Y	Y	Y	Y	N	N	N	Y
8.5.1	N	n/a	Y	Y	Y	Y	N	N	N	N
8.5.2	N	N	n/a	Y	Y	Y	N	N	N	N
8.5.3	N	N	N	n/a	Y	Y	N	N	N	N
8.5.4	N	N	N	N	n/a	Y	N	N	N	N
8.6.0	N	N	N	N	N	n/a	Y	Y	Y	Y
8.6.1	N	N	N	N	N	N	n/a	Y	Y	Y
8.6.2	N	N	N	N	N	N	N	n/a	Y	Y
8.6.3	N	N	N	N	N	N	N	N	n/a	Y
8.7.0	N	N	N	N	N	N	N	N	N	n/a

IBM FlashSystem 5045					
Upgrade Support Matrix					
FROM / TO	8.6.0	8.6.1	8.6.2	8.6.3	8.7.0
8.6.0	n/a	Y	Y	Y	Y
8.6.1	N	n/a	Y	Y	Y
8.6.2	N	N	n/a	Y	Y
8.6.3	N	N	N	n/a	Y
8.7.0	N	N	N	N	n/a

IBM FlashSystem 5300		
Upgrade Support Matrix		
FROM / TO	8.6.3	8.7.0
8.6.3	n/a	Y
8.7.0	N	n/a

Figure 1-5 IBM Storage Virtualize upgrade support matrix

Take time to look at the above matrix and confirm your IBM Storage Virtualize hardware can upgrade to the IBM Storage Virtualize 8.7.0 level.

Secondly, you will need to check that any outstanding issues or errors are corrected or fixed prior to the upgrade. If required, place a service call with IBM for assistance with getting these resolved.

Also, be aware that there are certain restrictions with features at IBM Storage Virtualize 8.7.0 that might be applicable to your configuration. You need to take these into account before upgrading to IBM Storage Virtualize 8.7.0. See here for the list:

[V8.7.0.x Configuration Limits for IBM FlashSystem and SAN Volume Controller.](#)

Finally, you will need to schedule time to do the upgrade to the new IBM Storage Virtualize V8.7.0 software. This upgrade is concurrent, but it is recommended that it is done at a time of low I/O activity as the upgrade works in parallel to the normal running so at low I/O times the upgrade will happen sooner.

Also, checks need to be made for the multipathing driver considerations to ensure no host connectivity is lost during the nodes going offline as part of the code upgrade, Metro Mirror and Global Mirror relationships might also need to be reviewed and possibly revised as the code upgrade may affect the performance of these.

For upgrade process including the pre checks and code download can be found here in the IBM Documentation section: [Updating System Software.](#)



It is also important to check that the drive firmware on your IBM FlashSystem is also kept up to the latest level. This process is done separately to the IBM Storage Virtualize software updating.

Details of the drive firmware upgrade process can be found here: [Updating Drive Firmware](#).





# Initial configuration

This chapter describes the initial configuration of the following actual 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, before volume creation and provisioning.

This chapter includes the following topics:

- ▶ “Prerequisites” on page 18
- ▶ “System initialization” on page 19
- ▶ “System setup” on page 26

## 2.1 Prerequisites

Before initializing and setting up the system, ensure that the following prerequisites are met:

- ▶ The physical components fulfill all the requirements and are correctly installed, including:
  - 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.
  - The expansion enclosures (if available) 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 [Web browser requirements to access the management GUI](#).
- ▶ The required information for remote management of the system, including:
  - 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 [IBM Documentation](#).
  - Information that is used by a system when performing Call Home functions, such as:
    - 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 only 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:

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

After the IBM SSR completes their portion of the setup, see 2.3, “System setup” on page 26 to continue the setup process.

## 2.2 System initialization

This section provides step-by-step instructions about how to create the system cluster. Demonstration videos

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](#)

To start the initialization procedure, connect a workstation directly to the *technician port*. The technician port is a dedicated 1 Gb Ethernet (GbE) port located on the rear of each storage node canister. On all platforms except IBM FlashSystem 5015, it can only be used 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.

## 2.2.1 System initialization process

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

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, only initialize the first enclosure or node. The remaining enclosures or nodes can be added to the cluster later 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 completes, 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 could 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.



- The window that is used to log in to Service Assistant opens (Figure 2-9). This window is the first step of initializing the system. Enter the default superuser password (passw0rd) and click **Log in**.

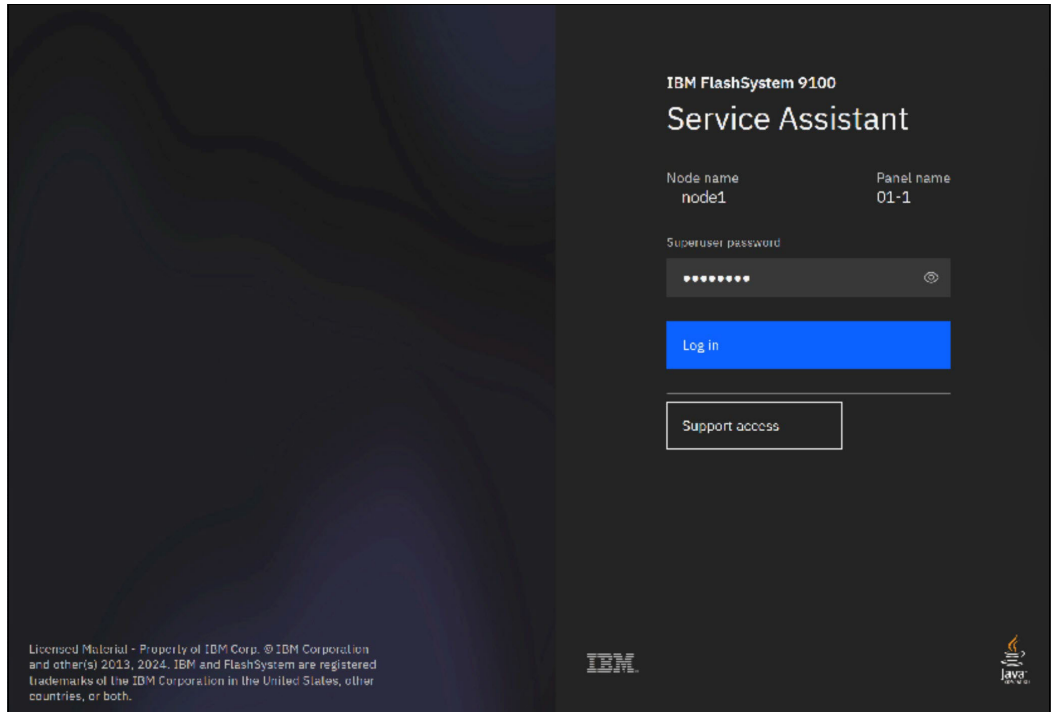


Figure 2-9 Logging in to Service Assistant by way of the technician port

- 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).

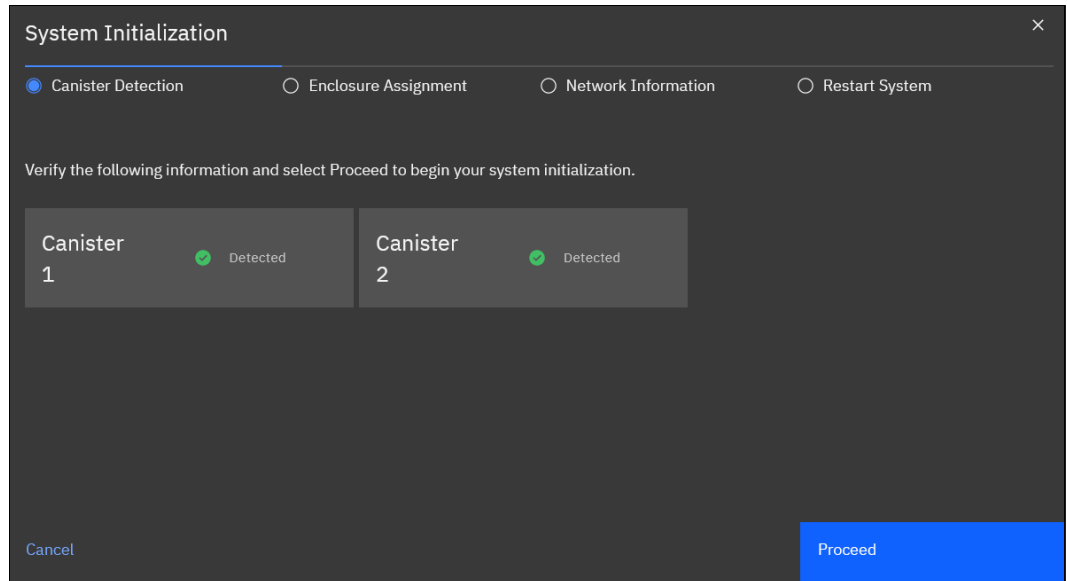


Figure 2-10 System Initialization: Canister detection

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**.

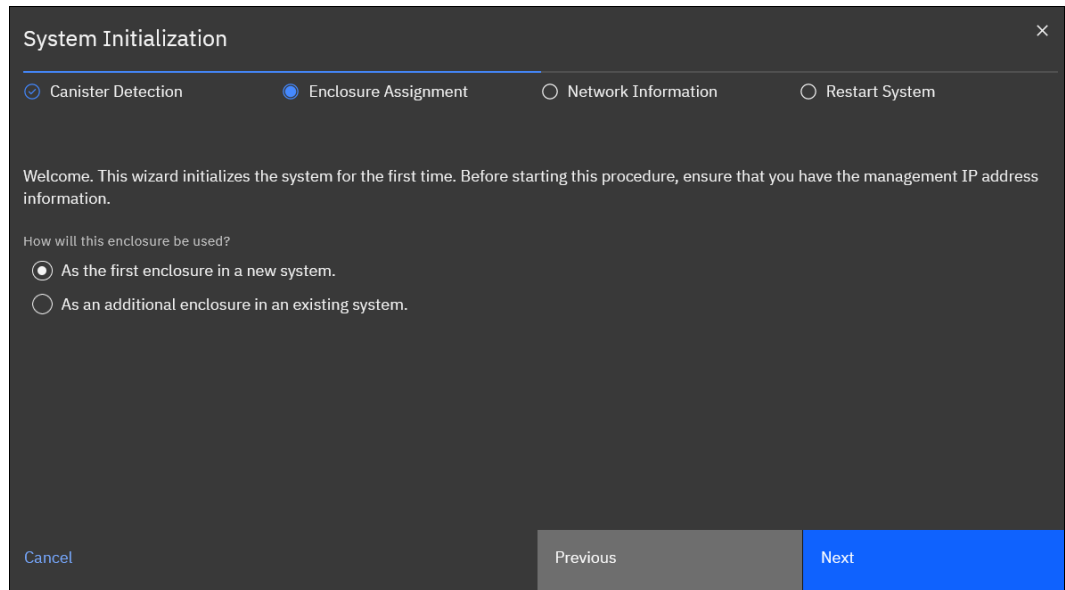


Figure 2-11 System Initialization: Initialize the first enclosure

If you select **As an additional enclosure in an existing system**, you are prompted to disconnect from the technician port and use the GUI of the system to which the new nodes are to be added.

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

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

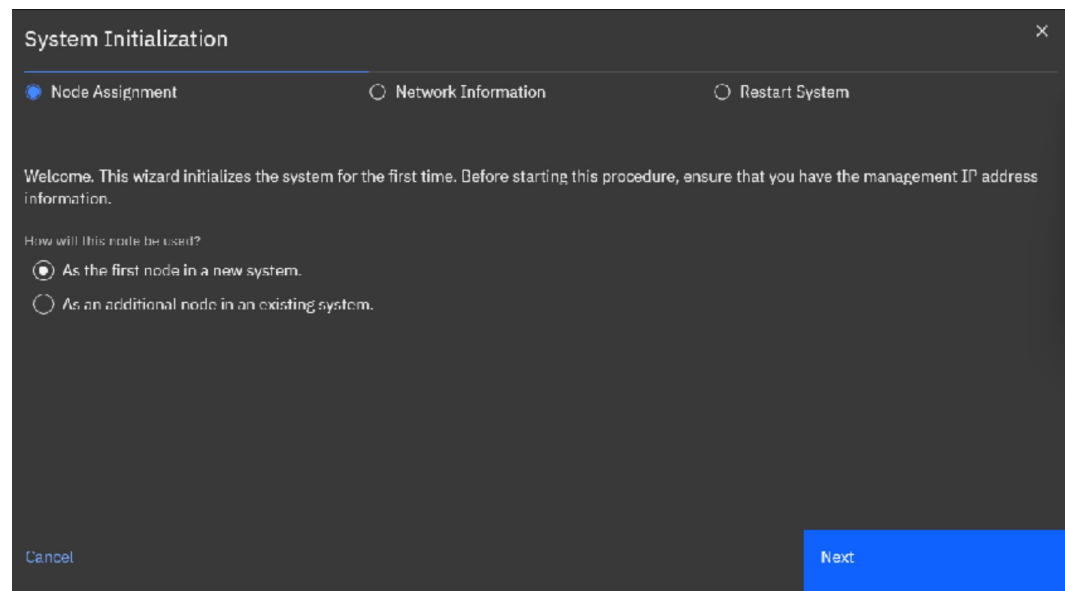


Figure 2-12 System Initialization: Initialize the first IBM SAN Volume Controller node

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**.

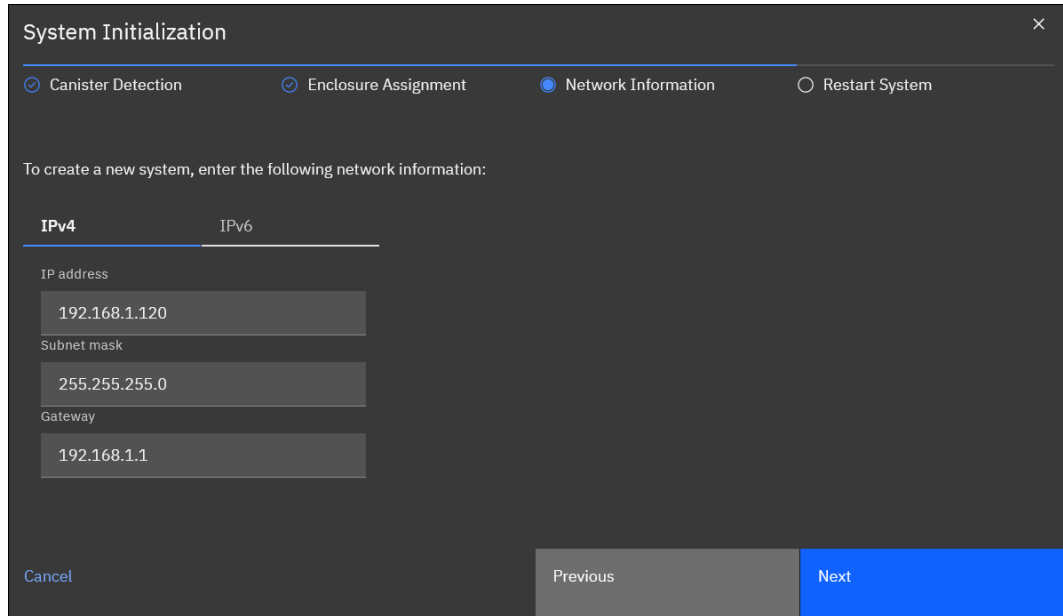


Figure 2-13 System Initialization: Enter Management IP

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.

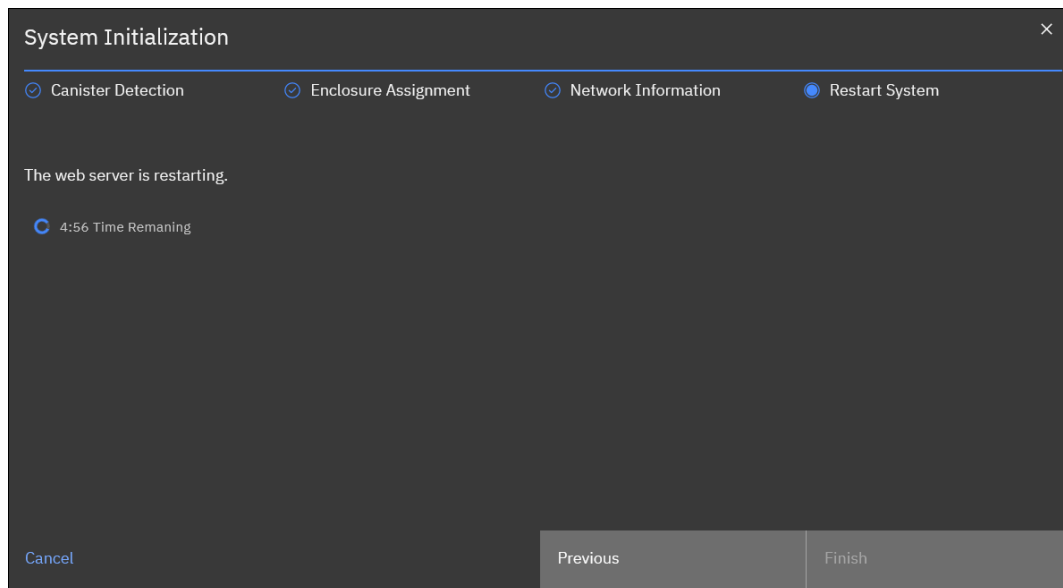


Figure 2-14 System Initialization: Web-server restart timer counting down from 5 minutes

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

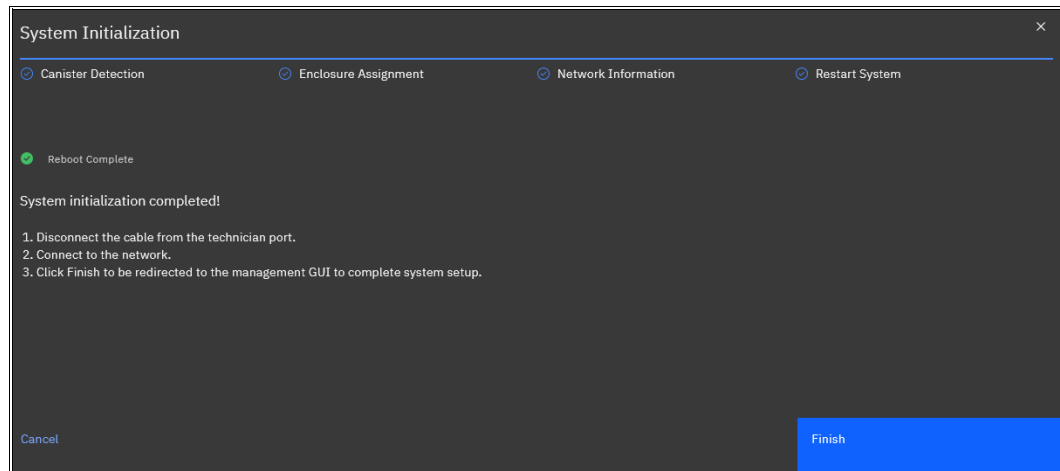


Figure 2-15 System Initialization completed

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 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. Click **Sign in**.

**Note:** The default password for the *superuser* account is `passw0rd` (*with the number zero, not the uppercase 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.

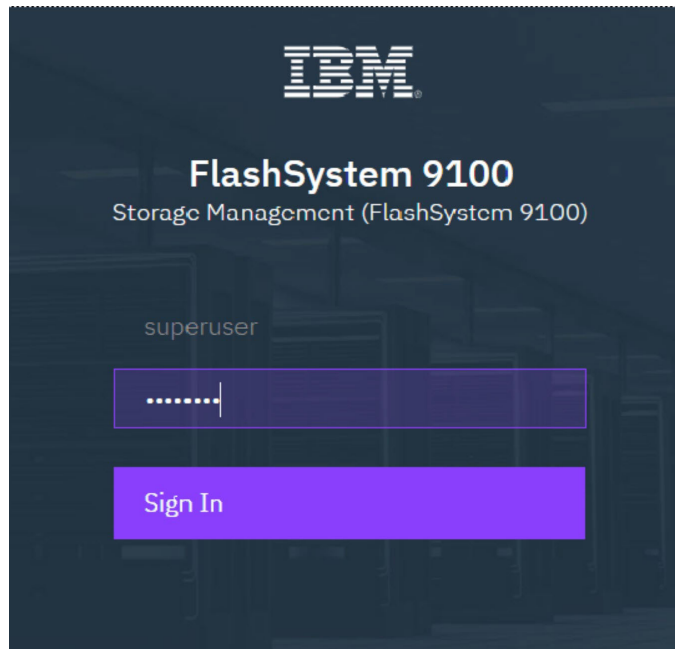


Figure 2-16 Logging in for the first time

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

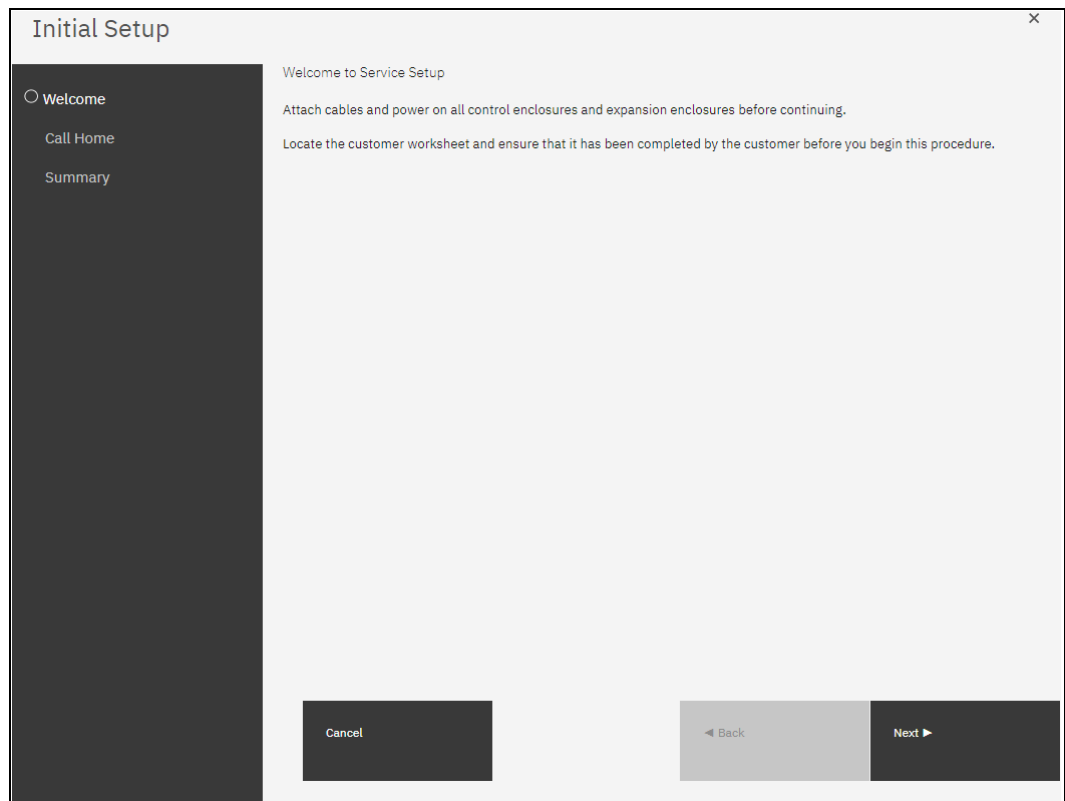


Figure 2-17 Initial Setup Window

3. The next step in the initial setup is the setup of Call Home. See Figure 2-18 on page 29. Call Home provides multiple benefits. It enables automatically creating tickets at IBM if errors occur, which improves the speed and efficiency by which calls are handled. Call Home also enables Storage Insights and Remote Support.

**Note:** Refer to “Call Home” on page 63 for more information on Call Home.




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. If you click **Next** you will be asked for the Transmission Type for Call Home.
5. System Setup prompts to select which transmission types to be used for Call Home. See Figure 2-19 on page 30.

**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 our case we choose **Send using Cloud services**. Click **Apply and Next** to setup the Internal Proxy Server. See Figure 2-20 on page 31. Enter the requested information.



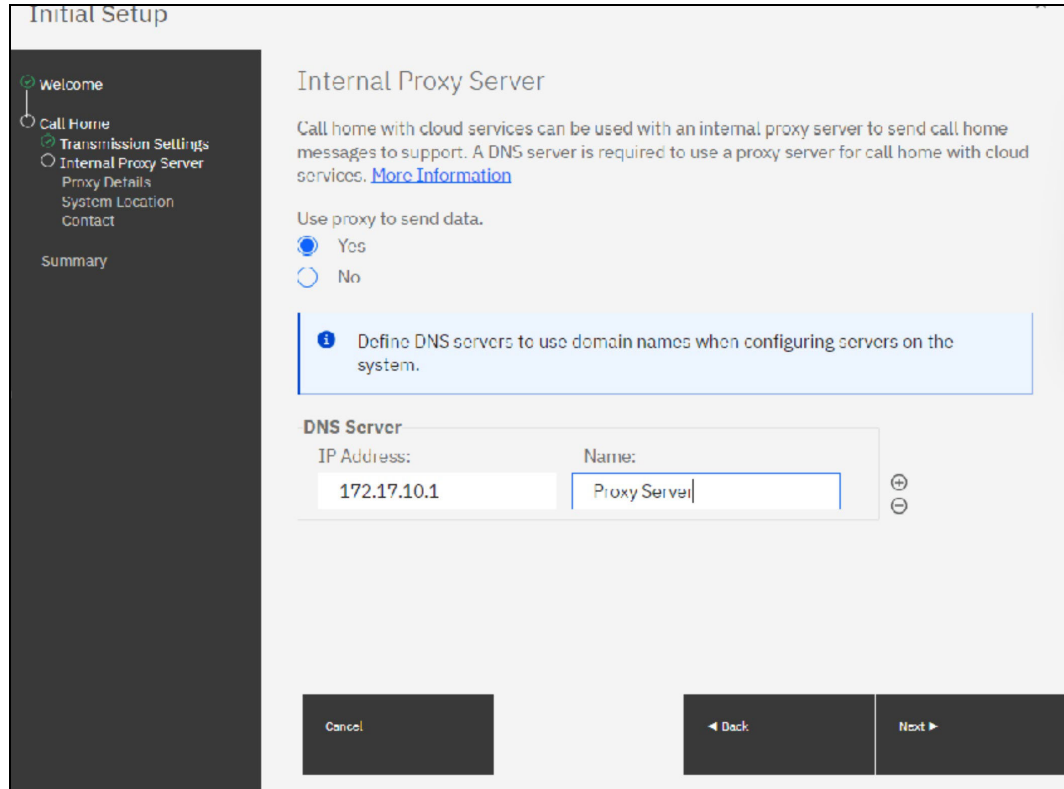


Figure 2-20 Setup Internal Proxy Server

7. After you setup the Proxy Server the system checks the connection to the Support Center. See Figure 2-21 on page 32.

**Initial Setup**

- ✓ Welcome
- Call Home
  - ✓ Transmission Settings
  - ✓ Internal Proxy Server
  - System Location
  - Contact
- Summary

**Connection to the support center was successful!**

**System Location**

Service parts should be shipped to the same physical location as the system.

Company name:

System address:

City:

State or province:

Postal code:

Country or region:

Machine location:

Figure 2-21 Connection Test to the Support Center

8. Enter all required information for the System Location. After clicking **Next** you should fill in the Contact information. Figure 2-22 on page 33 shows the panel for the contact information. Use the Company contact information to comply with privacy regulations. IBM may use the contact data if you allow it. To complete the registration click **Apply and Next**.

The screenshot shows the 'Initial Setup' window with a sidebar on the left containing the following items: Welcome (checked), Call Home (unchecked), Transmission Settings (checked), Internal Proxy Server (checked), System Location (checked), Contact (unchecked), and Summary. The main area is titled 'Contact' and contains the following text: 'Enter the contact information that support center can use to contact the customer to resolve system errors.' Below this is a blue callout box with an information icon and the text: 'Enter business-to-business contact information. To comply with privacy regulations, personal contact information for individuals with your organization is not recommended.' The form fields are: Name: lonzer; Email: lonzcr@dc.ibm.com; Phone (primary): +49-1234567; Phone (alternate): +49-1234560. Below the fields is a toggle switch for 'TDM may use my contact data to keep me informed of Storage related products, services and offerings.' which is currently set to 'Off'. At the bottom are three buttons: 'Cancel', 'Back', and 'Apply and Next'.

Figure 2-22 System Location

9. The System provides you the Summary information. If all is correct click **Finish**. See Figure 2-23 on page 34.

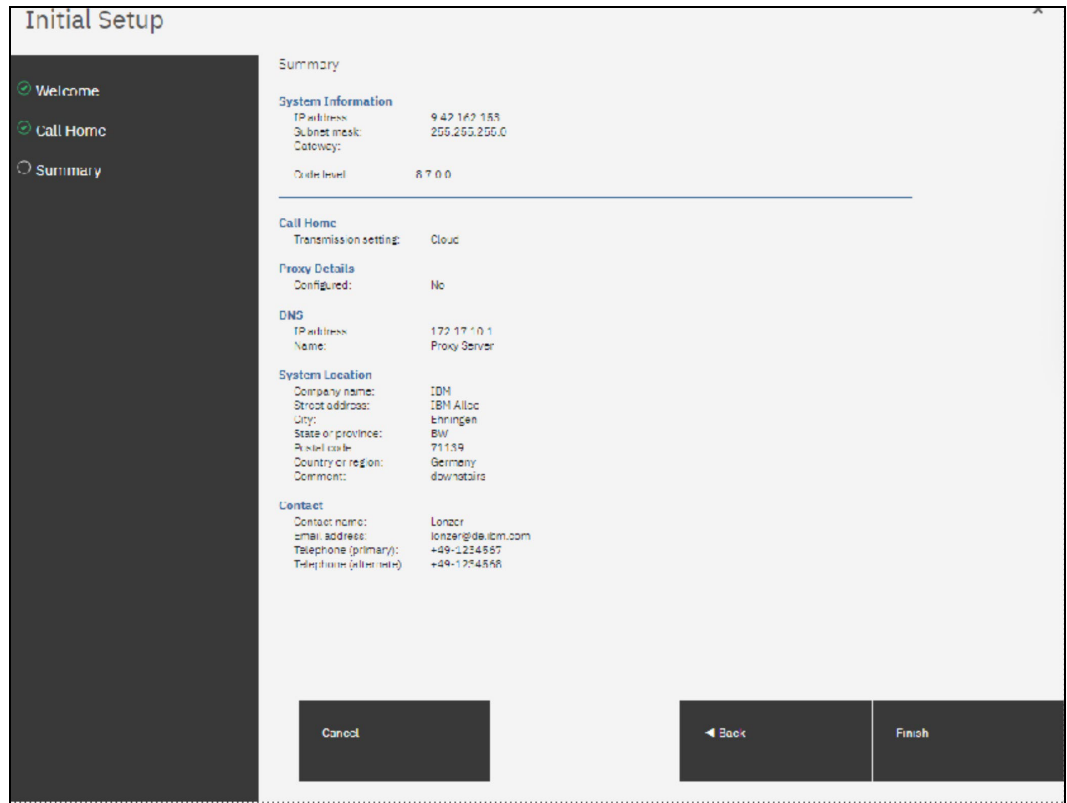


Figure 2-23 Summary page

10. The system saves the entered information and asks you to log in again. After login, you will be guided to the System Setup page. See Figure 2-24 on page 35.

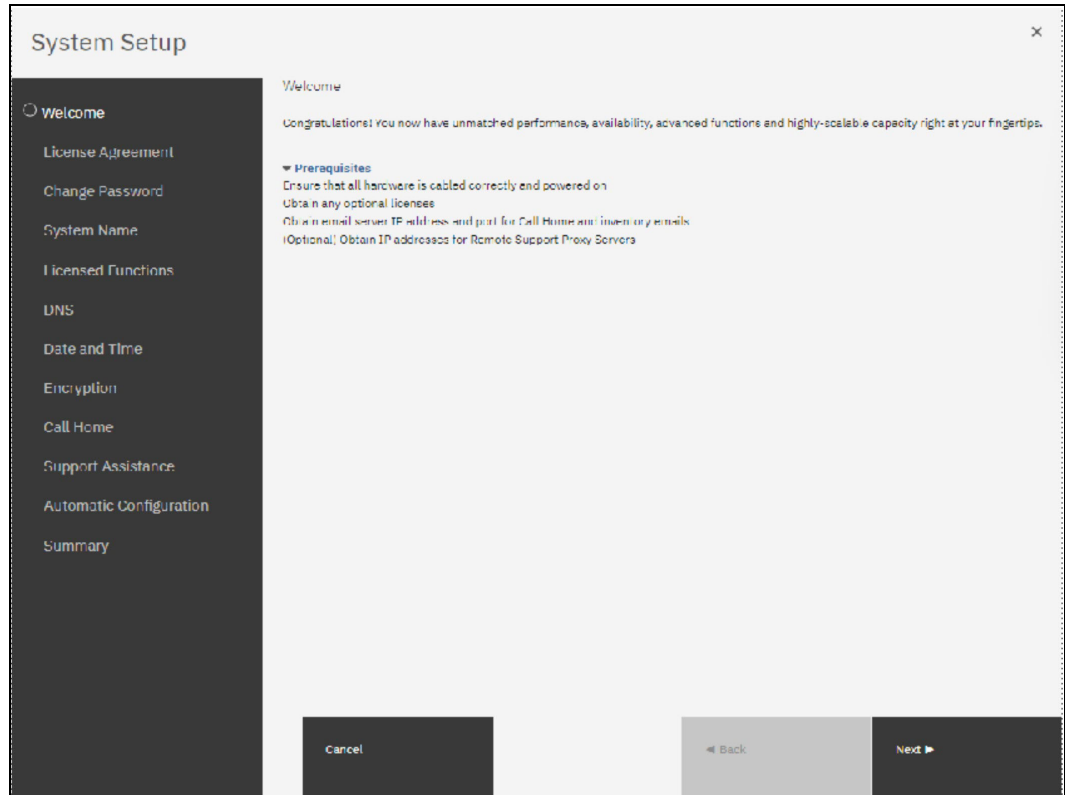


Figure 2-24 System Setup Welcome page

11. Clicking **Next** brings you to the License agreement page. Carefully read the license agreement. Select **I agree with the terms in the license agreement** if you want to continue the setup, otherwise the system will stop the setup. See Figure 2-25 on page 36.

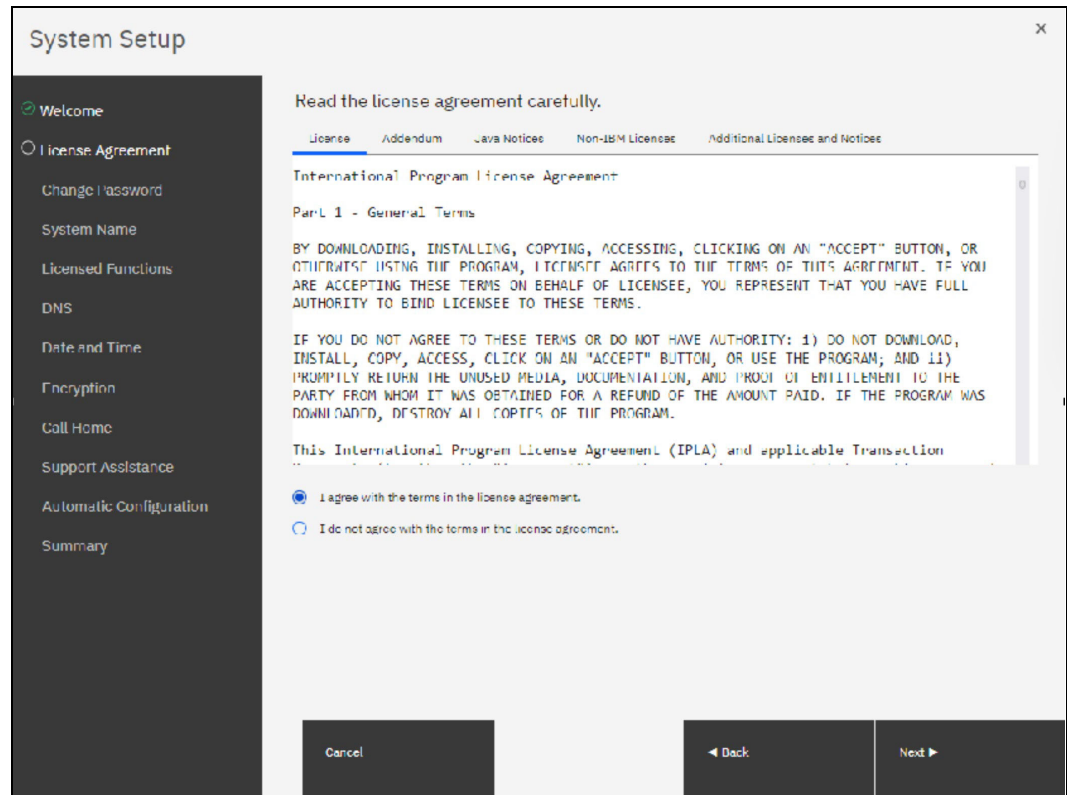


Figure 2-25 Accept License Agreement

- You are asked to change the password shown in Figure 2-26 on page 37. 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

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

Next step is to provide a System Name as shown in Figure 2-27 on page 38. Avoid the use of an underscore ( `_` ) in a system name. Although permitted here, it 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 to 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 by using the management GUI by clicking **Monitoring** → **System Hardware** and selecting **System Actions** → **Rename System**.

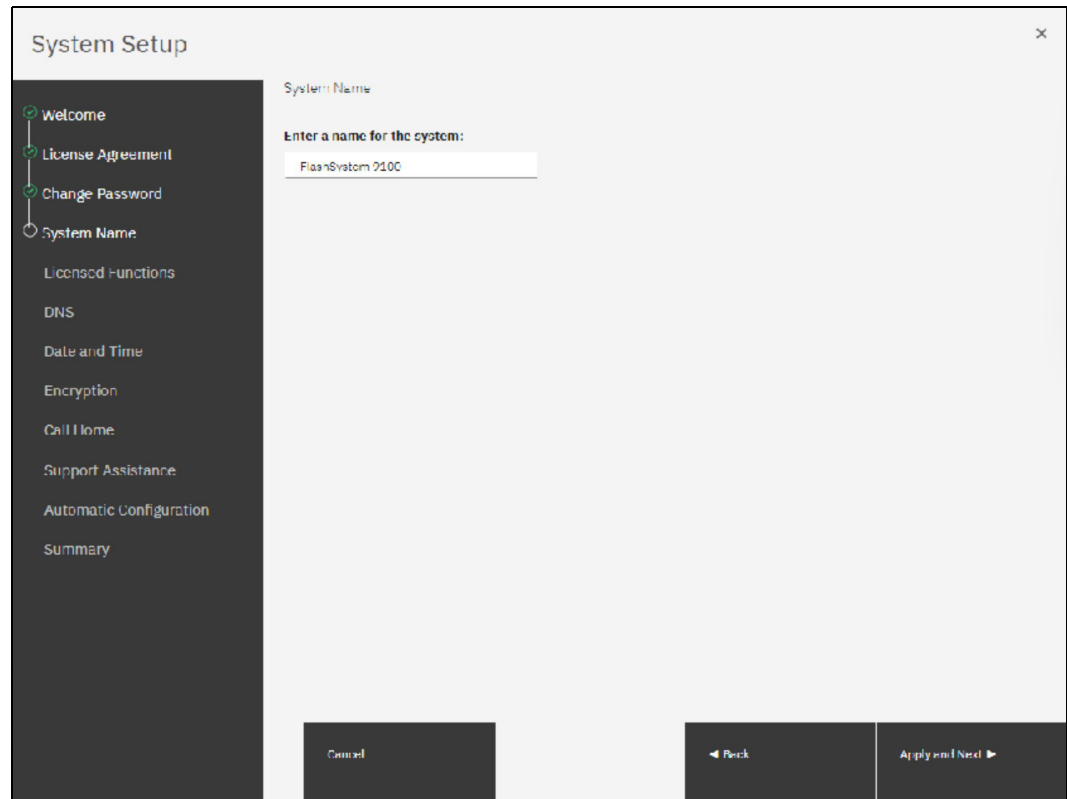


Figure 2-27 System Name

14. Click **Apply and Next**.

15. After you provide the System Name you will be asked to enter the required, additional licenses for each function. Figure 2-28 on page 39 shows you an example.






Figure 2-28 License Functions

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.

16. When done, click **Apply and Next**.

17. DNS can be configured on the system, as shown in Figure 2-29 on page 40. DNS helps the system to resolve the names of the computer resources that are in the external network if they are not indicated by an IP address.

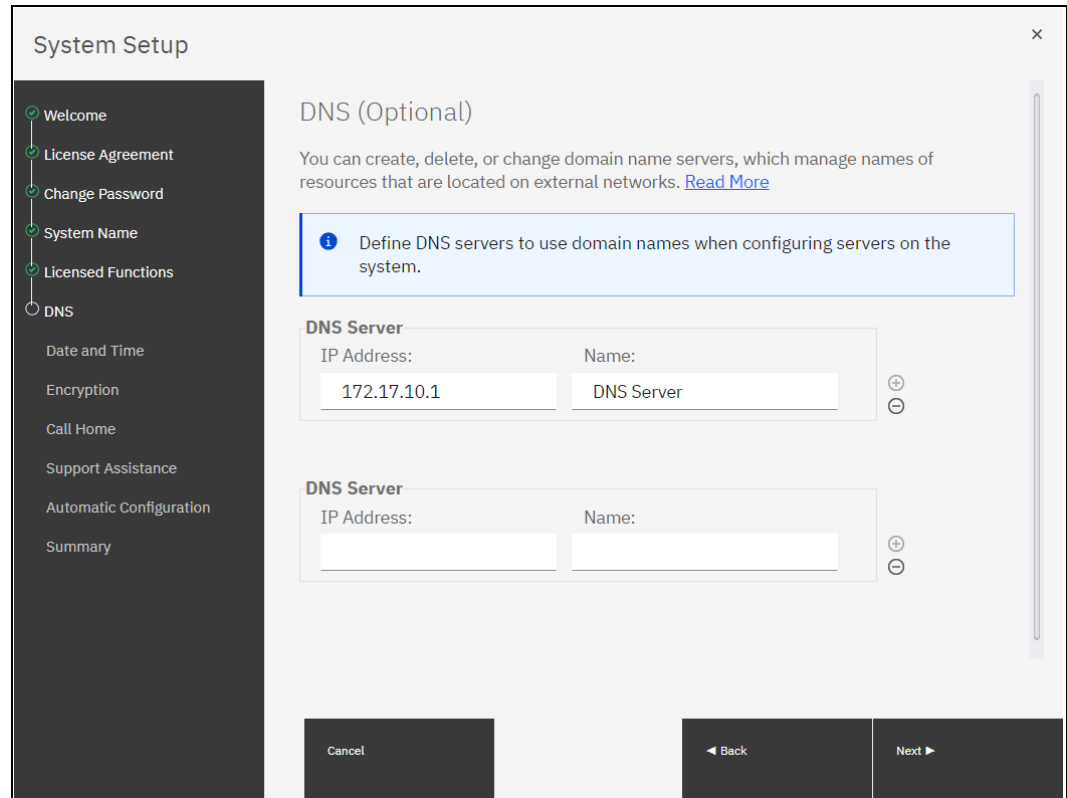


Figure 2-29 DNS Server setup

18. Following the DNS Setup you are asked to set Date and Time. Enter the date and time settings. In the example that is shown in Figure 2-30 on page 41, the date and time are set by using manually settings. Generally, 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.

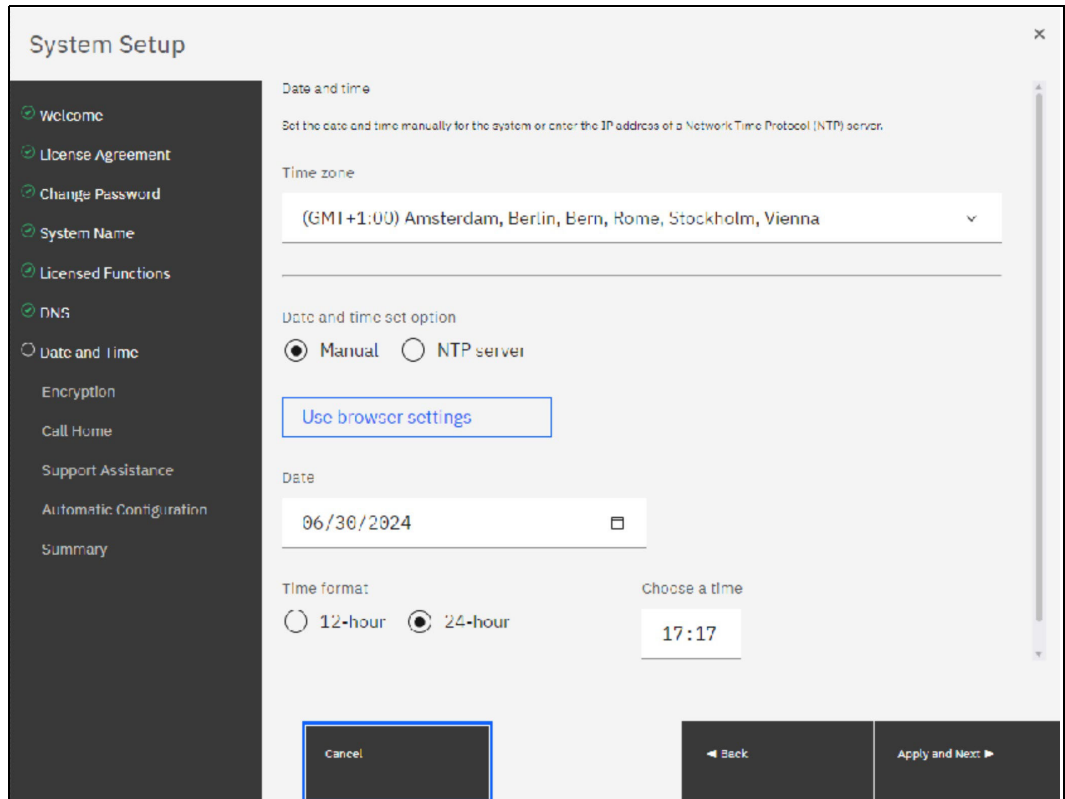


Figure 2-30 Date and Time

19. When done, click **Apply and Next**.

20. If you have purchased an Encryption License, activate it here. See Figure 2-31 on page 42.




Figure 2-31 Activate Encryption License

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.

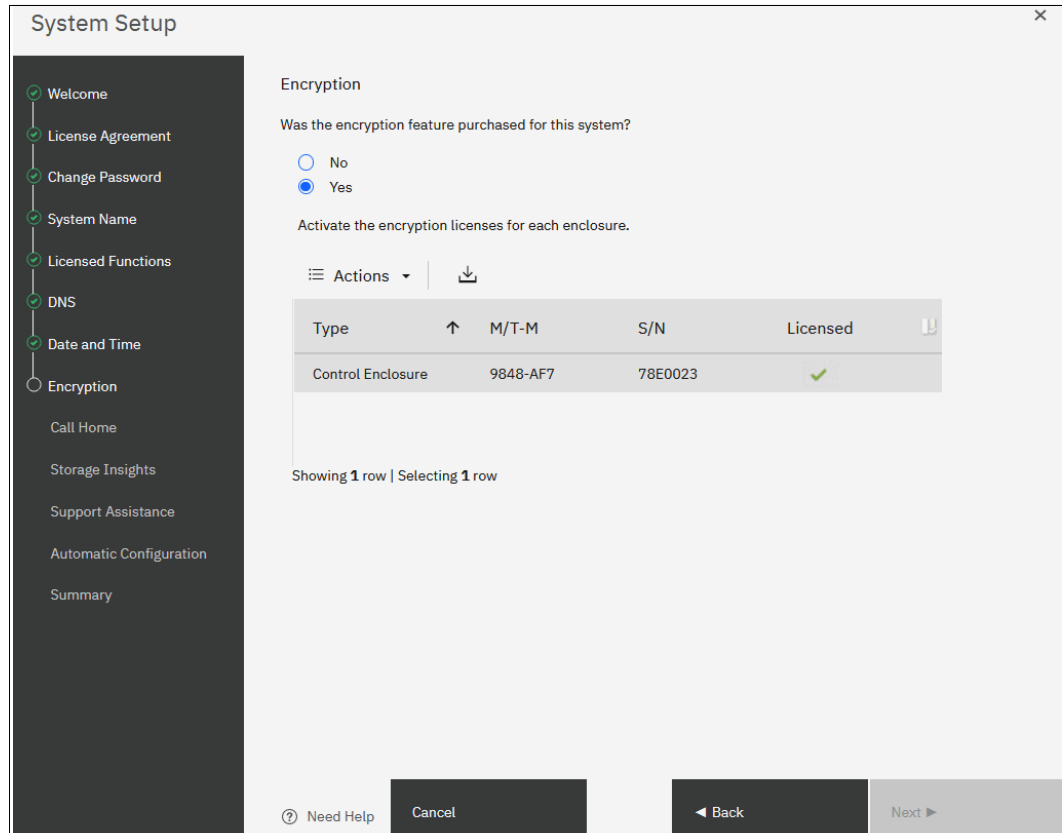


Figure 2-32 Encryption licensed

21. After the encryption license is activated, you see a green check mark for each enclosure, as shown in Figure 2-32 on page 43. After all the control enclosures show that encryption is licensed, click **Next**.
22. If you wish to modify your previously entered Call Home settings, you can do so here. See Figure 2-33 on page 44.

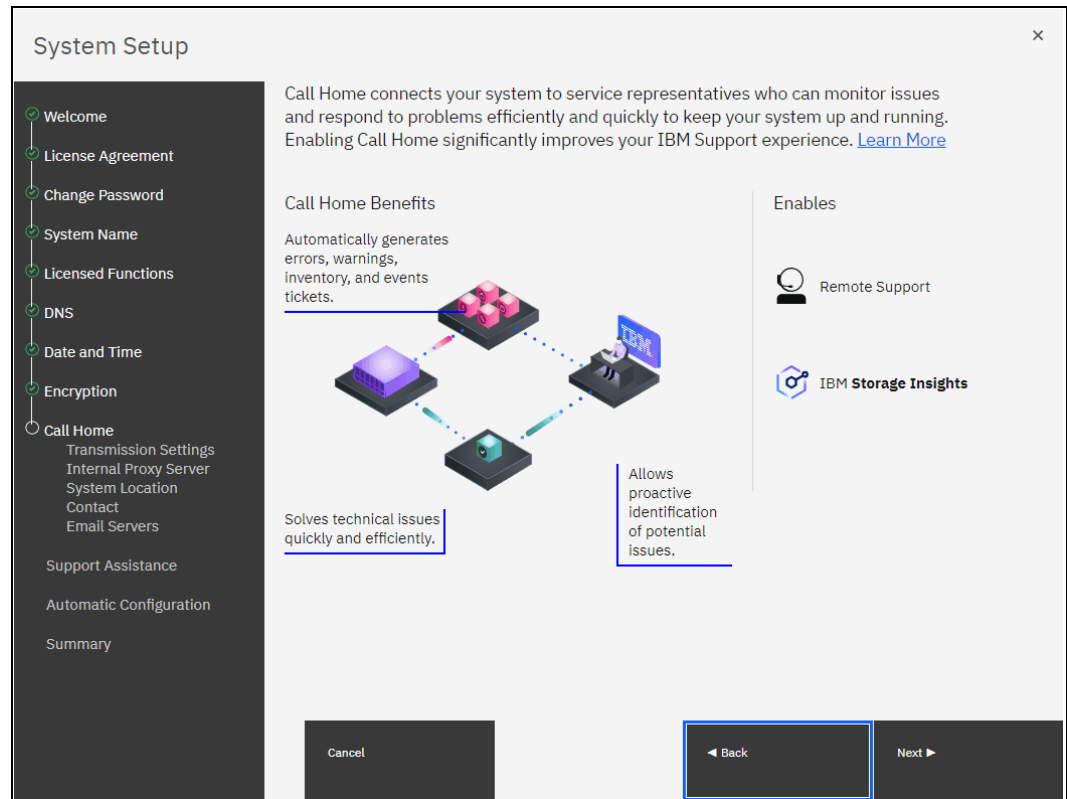


Figure 2-33 Change Call Home settings

23. If you want to use the **Support Assistance** offering from IBM you can choose your options here (Figure 2-34 on page 45). 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 64 for more information.

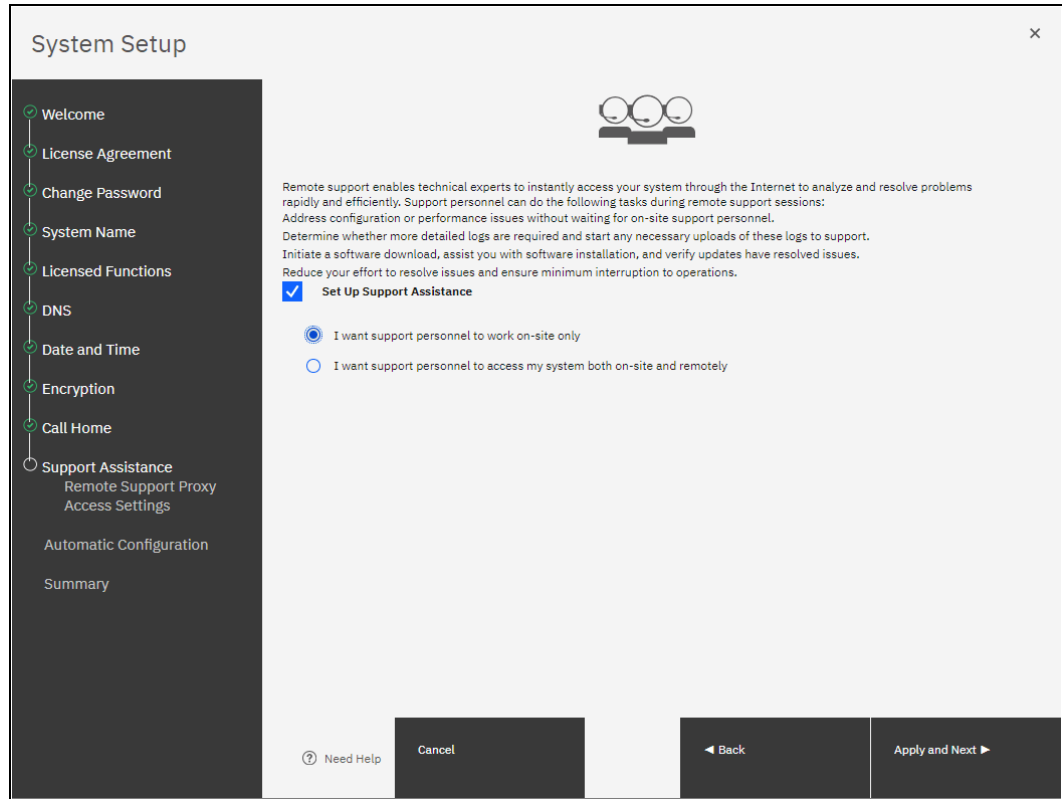


Figure 2-34 Setup Support Assistance

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 46. Click **Apply and Next**.

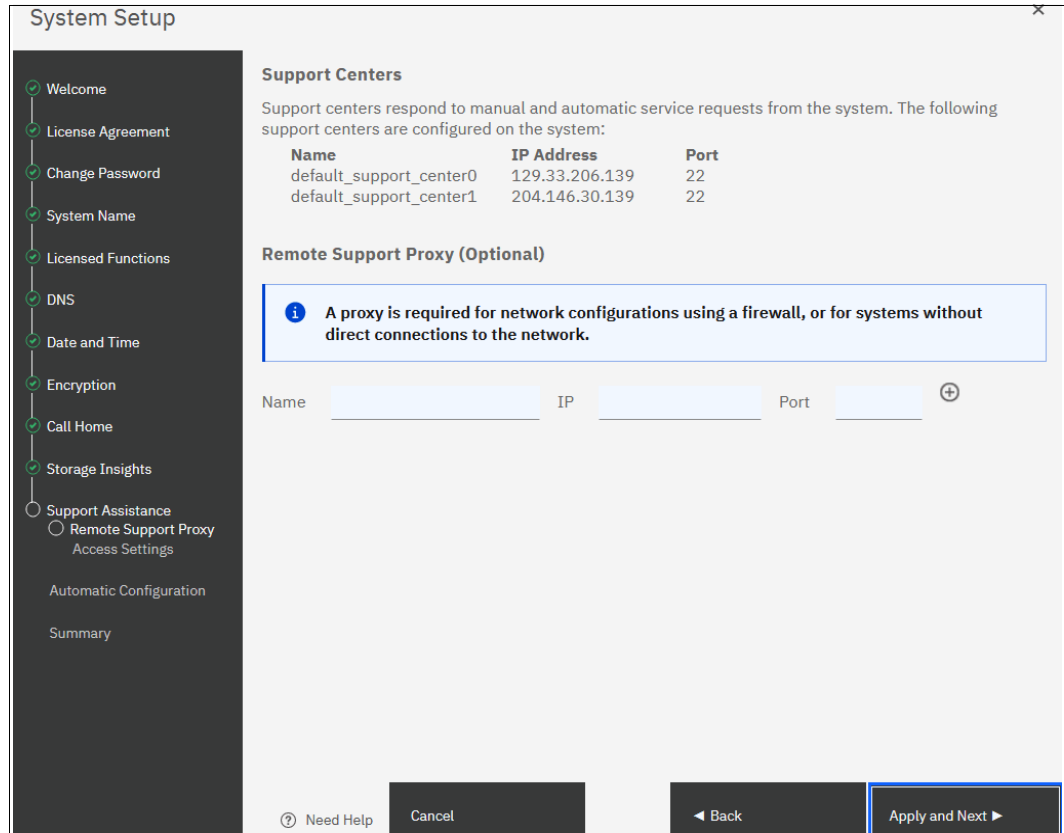


Figure 2-35 System communicating with named IBM Support servers

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 47.



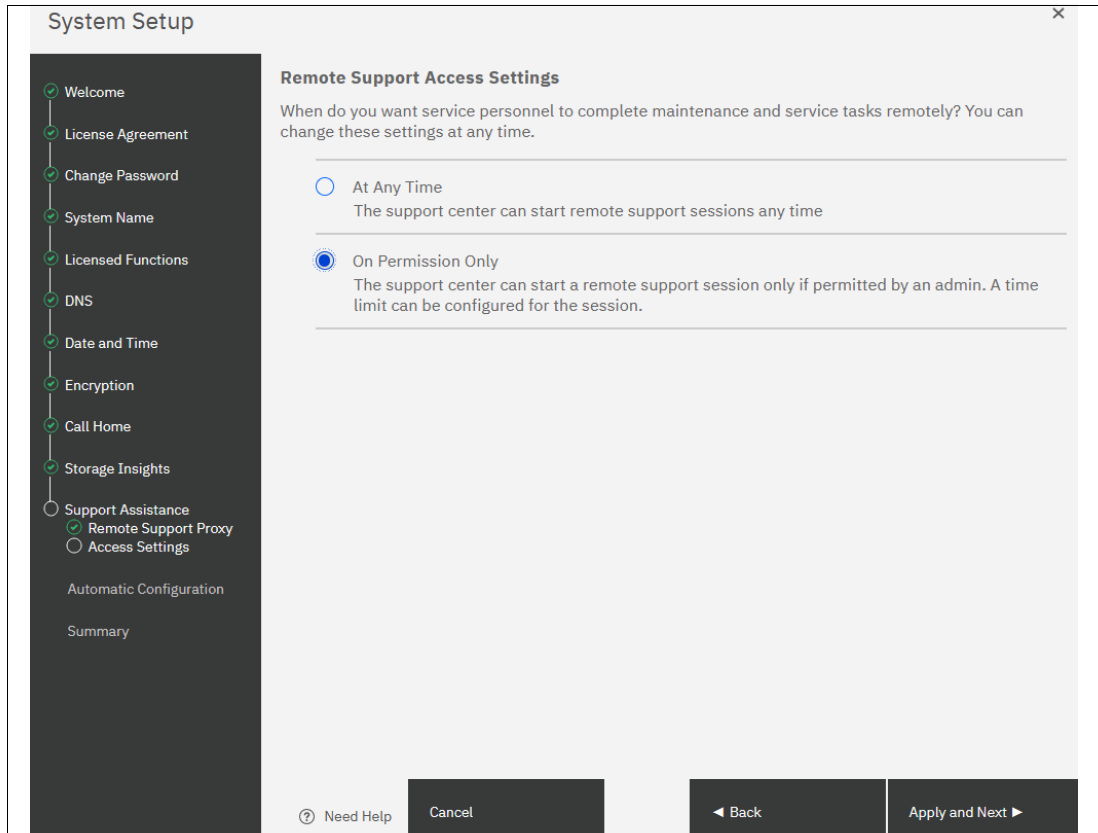


Figure 2-36 Remote support access settings

24. Click **Apply and Next**.

25. If you install your system below an IBM San Volume Controller, you can use **Automatic Configuration for Virtualization**. System Setup offers this option (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 on page 48. Click **Next** to continue.

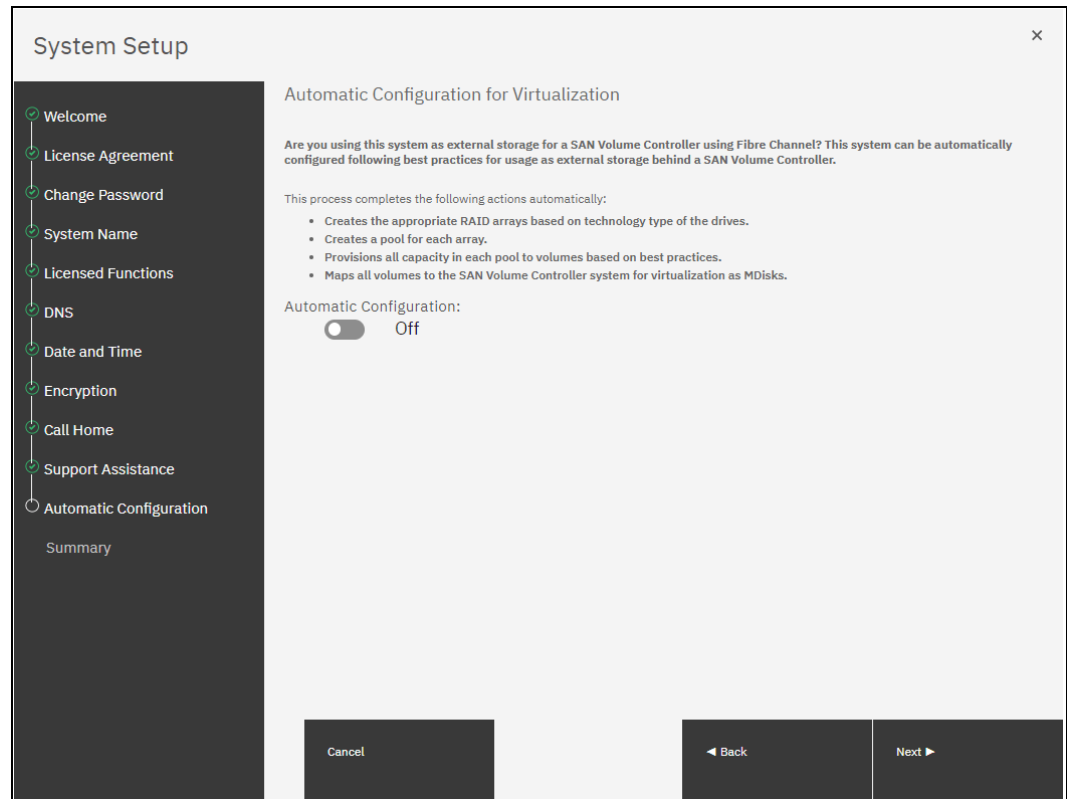


Figure 2-37 Automatic Configuration for Virtualization

For more information about how to enable Automatic configuration for IBM SAN Volume Controller on a running system after the System Setup wizard, see 2.3.8, “Automatic configuration for IBM SAN Volume Controller back-end storage” on page 69.

26. 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 49.

**System Setup**

Summary

**System Information**

System name:	FlashSystem 9100	Date:	Jun 30, 2024
Code level:	8.7.0.0	Time:	5:44:35 PM
Automatic Configuration:	Not enabled	Time zone:	Europe/Paris

---

**Licensed Functions**

External Virtualization:	100	SCU
FlashCopy:	20	TIB
Remote Mirroring:	20	TIB

---

**Call Home**

Transmission setting:	Cloud
-----------------------	-------

**Proxy Details**

Configured:	No
-------------	----

**DNS**

IP address:	172.17.10.1
Name:	DNS Server

**System Location**

Company name:	IBM
Street address:	IBM Allee
City:	Ehningen
State or province:	BW
Postal code:	71139
Country or region:	Germany
Comment:	downstairs

**Contact**

Contact name:	Lonzer
Email address:	lonzer@de.ibm.com
Telephone (primary):	+49-1234567

Buttons: Cancel, Back, Finish

Figure 2-38 Summary Page

When the system setup wizard completes, your IBM FlashSystem consists only of 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.

In the case of 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 2.3.3, “Adding an enclosure in IBM FlashSystem” on page 57.

For more information about how to add a node or hot spare node, see 2.3.4, “Adding a node or hot spare node in IBM SAN Volume Controller systems” on page 59.

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.

27. With pressing **Finish** you confirm your choices. See Figure 2-39.

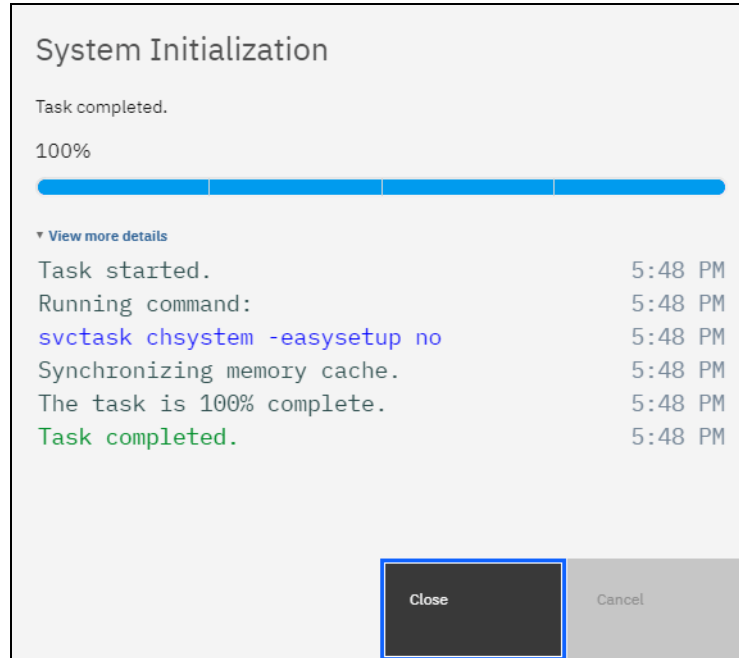


Figure 2-39 System Initialization

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

28. After Setup is completed you will be guided to the Management GUI. See Figure 2-40 on page 50.

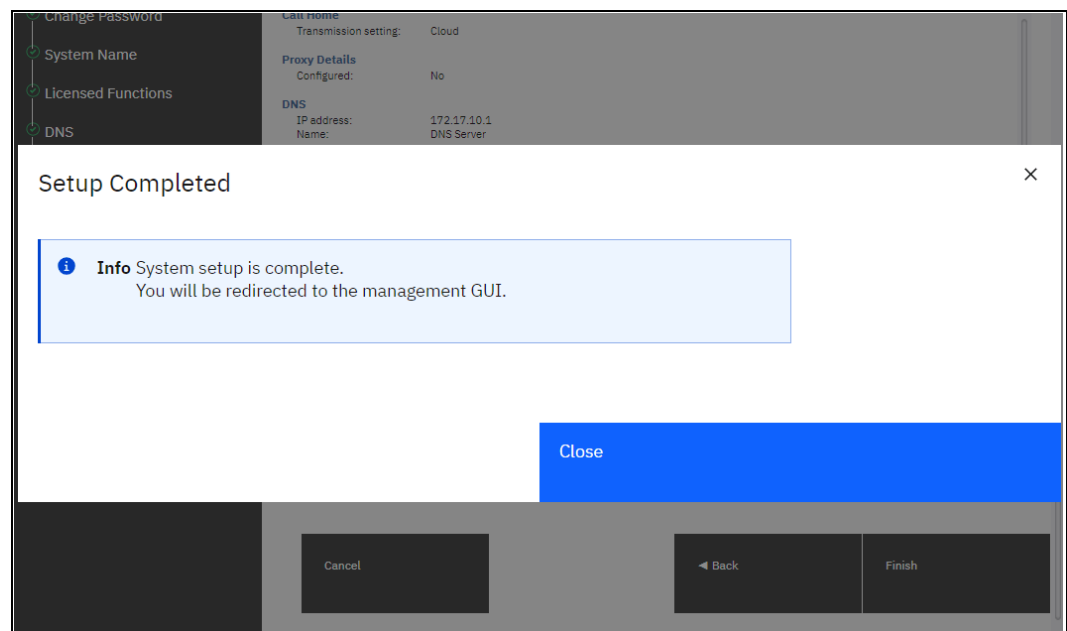


Figure 2-40 Setup completed

29. Clicking **Close and Finish** takes you to the Dashboard (see Figure 2-41 on page 51).




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 Configuring clustering by using Ethernet connections

The system supports high-performance node-to-node connections using Ethernet protocols with Remote Direct Memory Access (RDMA) technology, like RDMA over Converged Ethernet (RoCE) or iWARP. To enable this, each node needs an RDMA-capable adapter, and dedicated RDMA-capable Ethernet ports must be configured solely for node-to-node communication.

RDMA technologies, like RoCE and iWARP, allow RDMA adapters to directly transfer data between nodes, bypassing the CPU and caches for faster communication. This offers significant performance improvements over traditional iSCSI connections.

Up to four I/O groups can participate in either an IBM policy-based HA cluster, a standard topology cluster, or an enhanced stretched cluster (available only on IBM SAN Volume Controller). This section details the configuration steps specifically for systems designed for IP-based RDMA node-to-node traffic. For Fibre Channel (FC) SAN clustering, no special system configuration is required.

#### Planning clustered systems over RDMA and TCP-based Ethernet

Before implementing, select one of the following inter-Switch link (ISL) configurations based on the need to support systems:

##### **No ISL**

In this network configuration, no inter-switch links are used between the switches across two sites. With no inter-switch links between systems, the recommended operational range of this configuration is limited to 300 meters. This configuration is best suitable for small-scale enterprises for achieving high availability in a cost-effective manner. A graphical representation of No ISL network connections is shown in Figure 2-42 on page 52

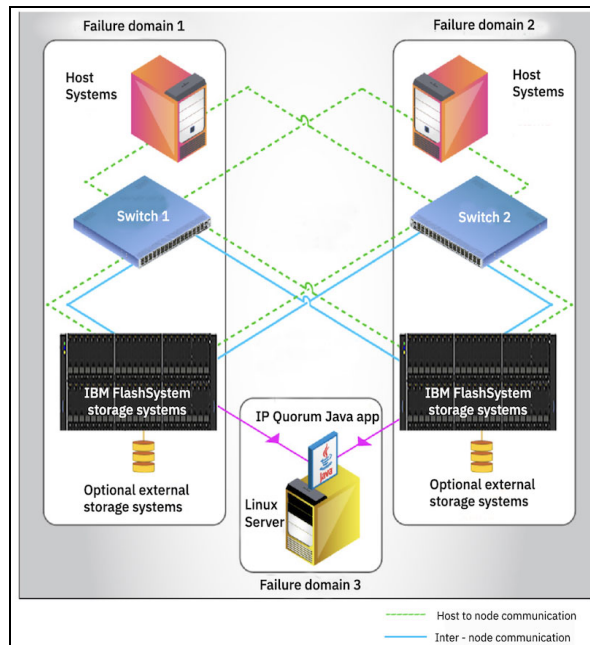


Figure 2-42 No ISL connectivity

### Shared ISL

This network configuration uses shared inter-switch links, which carry inter-node and host-node traffic with other local area network (LAN) traffic. At times, the mixed traffic can cause congestion at ISL. However, configure priority flow control (PFC) for inter-node and host-node traffic to avoid congestion and system instability due to multiple traffics at ISL. Also, ensure that the bandwidth of ISL is adequate to support all the network traffic flowing through it. This configuration is used for the policy-based HA systems that involve multiple traffics from various subnet and VLANs in the network. A graphical representation of Shared ISL network connections is shown in Figure 2-43

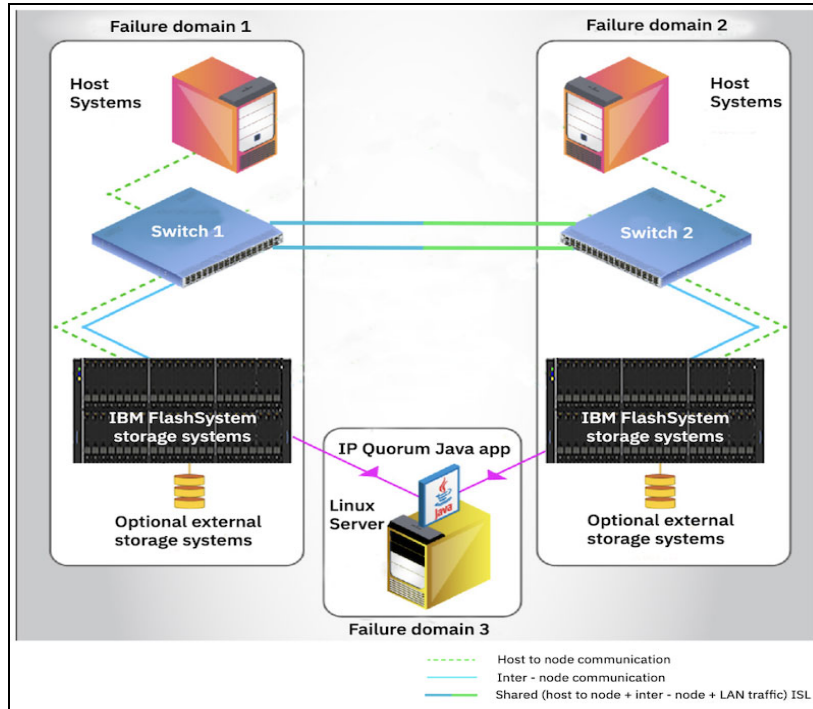


Figure 2-43 Shared ISL connectivity

**Note:** PFC (Priority Flow Control) is not supported when you use NVMe over TCP.

### Dedicated ISL

This configuration uses dedicated inter-switch links, which carry only inter-node and host-node traffic for the policy-based HA systems. In this configuration, host-node and inter-node traffic are isolated, so ISL bandwidth is not affected. A dedicated ISL configuration can prevent decreased system performance and slower host response times. However, setting up this configuration incurs extra cost. A graphical representation of dedicated ISL network connections is shown in Figure 2-44 on page 54

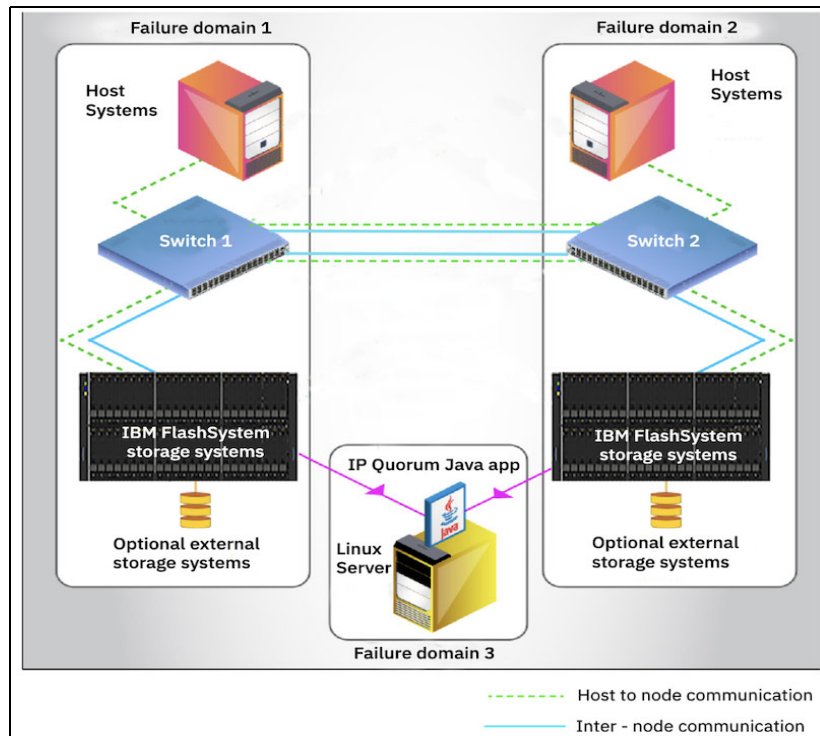


Figure 2-44 Dedicated ISL connectivity

## Prerequisites

Before RDMA clustering is configured, ensure that the following prerequisites are met:

- ▶ 25 gigabits per second (Gbps) RDMA-capable Ethernet cards are installed in each node.
- ▶ RDMA-capable adapters in all nodes use the same technology, such as RDMA over Converged Ethernet (RoCE) or internet Wide Area RDMA Protocol (iWARP).
- ▶ RDMA-capable adapters are installed in the same slots across all the nodes of the system.
- ▶ Ethernet cables between each node are connected correctly.
- ▶ The network configuration does not contain more than two hops in the fabric of switches. The router must *not* be placed between nodes that use RDMA-capable Ethernet ports for node-to-node communication.
- ▶ The negotiated speeds on the local and remote adapters are the same.
- ▶ The local and remote port (RPORT) virtual local area network (VLAN) identifiers are the same. All the ports that are used for node-to-node communication must be assigned to one VLAN ID, and ports that are used for host attachment must have a different VLAN ID.

If you plan to use VLAN to create this separation, you must configure VLAN support on all the Ethernet switches in your network before you define the RDMA-capable Ethernet ports on nodes in the system. On each switch in your network, set the VLAN to Trunk mode and specify the VLAN ID for the RDMA-ports that is to be in the same VLAN.

- ▶ A minimum of two dedicated RDMA-capable Ethernet ports are required for node-to-node communications to ensure best performance and reliability. These ports must be configured for inter-node traffic only and must not be used for host attachment, virtualization of Ethernet-attached external storage, or IP replication traffic.
- ▶ A maximum of four RDMA-capable Ethernet ports per node are allowed for node-to-node communications.



## Configuring node port IP addresses

To enable RDMA clustering, IP addresses must be configured on each port of each node that is used for node-to-node communication. Complete the following steps:

1. Connect to a Service Assistant of a node by browsing to `https://<node_service_IP>/service`. Then, select a node and click **Change Node IP** (see Figure 2-45).

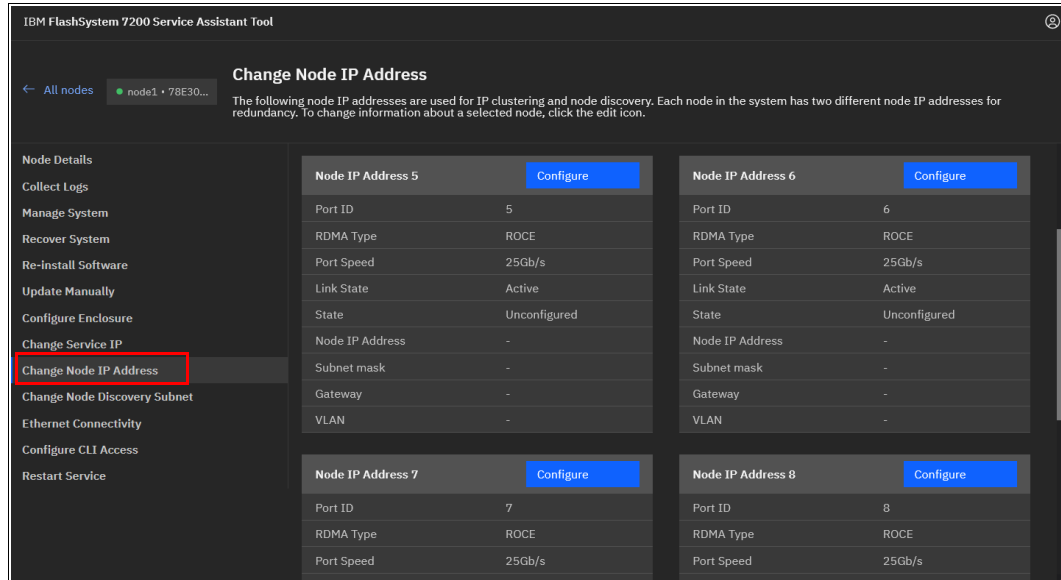


Figure 2-45 Node IP address setup for Remote Direct Memory Access clustering

2. Hover over a tile with a port and click **Configure** to set the IP address, netmask, gateway address, and VLAN ID for a port. The IP address for each port must be unique and cannot be used anywhere else on the system. The VLAN ID for ports that are used for node-to-node traffic must be the same on all nodes.
3. When the required information is entered, click **Save** and verify that the operation finished successfully, as shown in Figure 2-46. Repeat this step for all ports that you intend to use for node-to-node traffic, with a minimum of two and a maximum of four ports per node.

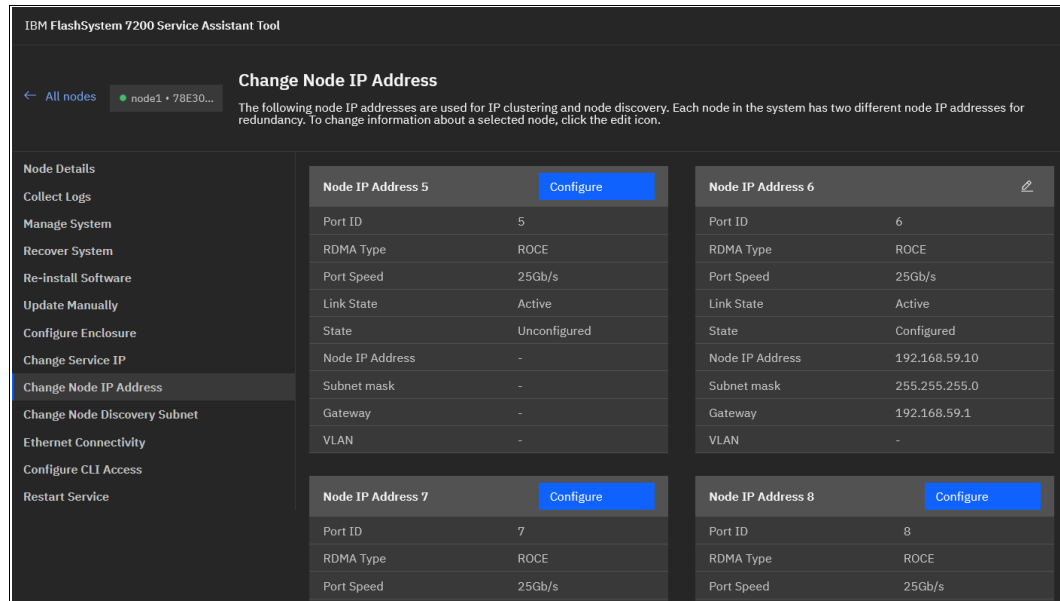


Figure 2-46 Node IP addresses configured

- To list the node IP configuration by using the CLI, run the **sainfo lsnodeip** command, as shown in Example 2-2.

*Example 2-2 Listing node IPs currently not set*

```
ITS0-FlashSystem:superuser>sainfo lsnodeip
port_id  rdma_type port_speed vlan link_state state      node_IP_address gateway subnet_mask
1                inactive unconfigured
2                inactive unconfigured
3                inactive unconfigured
4                inactive unconfigured
5          RoCE    25Gb/s    active unconfigured
6          RoCE    25Gb/s    active unconfigured
7          RoCE    25Gb/s    active unconfigured
8          RoCE    25Gb/s    active unconfigured
9          RoCE    25Gb/s    active unconfigured
10         RoCE    25Gb/s    active unconfigured
```

- Run the **satask chnodeip** commands to change node IP by using CLI, as shown in Example 2-3.

*Example 2-3 Running commands to change node IP*

```
superuser>satask chnodeip -ip 10.0.99.12 -gw 10.0.99.20 -mask 255.255.255.0 -port_id 5
superuser>satask chnodeip -ip 192.168.59.11 -gw 192.168.2.120 -mask 255.255.255.0 -port_id 6
```

- To list the changed node IP, run the **sainfo lsnodeip** again, as shown in Example 2-4.

*Example 2-4 Changed node IP (output shortened for clarity)*

```
ITS0-FlashSystem:superuser>sainfo lsnodeip
port_id  rdma_type port_speed vlan link_state state      node_IP_address
1                inactive unconfigured
2                inactive unconfigured
3                inactive unconfigured
4                inactive unconfigured
5          RoCE    25Gb/s    active  configured  10.0.99.12
6          RoCE    25Gb/s    active  configured  192.168.59.11
7          RoCE    25Gb/s    active  unconfigured
8          RoCE    25Gb/s    active  unconfigured
```

9	RoCE	25Gb/s	active	unconfigured
10	RoCE	25Gb/s	active	unconfigured

- Some environments might not include a stretched layer 2 subnet. In such scenarios, a layer 3 network (such as in standard topologies or long-distance RDMA node-to-node policy-based HA configurations) is applicable. To support the layer 3 Ethernet network, use the unicast discovery method for RDMA node-to-node communication. This method relies on unicast-based fabric discovery rather than multi cast discovery.
- To configure unicast discovery, use the `satask addnodediscoverysubnet`, `satask rmnodediscoverysubnet`, or `sainfo lsnodediscoverysubnet` commands. For more information, see [Command-line interface](#).

You can also configure discovery subnets by using the Service Assistant interface menu option **Change Node Discovery Subnet**, as shown in Figure 2-47.

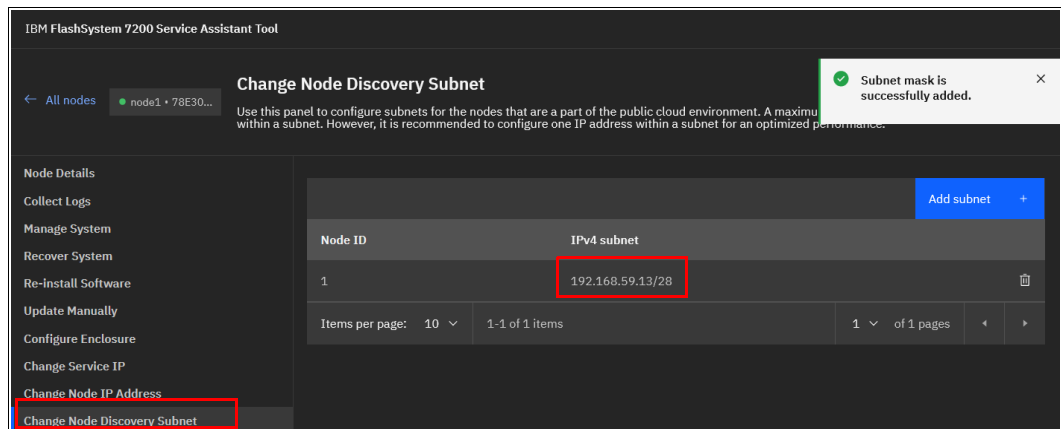


Figure 2-47 Setting the node discovery subnet

After the IP addresses are configured on all nodes in a system and the nodes to be partner nodes, from the Service Assistant GUI, navigate to **Ethernet Connectivity** to view which nodes are visible to the system.

Alternatively, run the `sainfo lsnodeipconnectivity` CLI command to verify that the partner nodes are visible on the IP network.

- When all the nodes that are joined to the cluster are connected, add the enclosure to the cluster.

### 2.3.3 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 performed by using the system GUI, 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.

To add an enclosure to the system, 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 on the system next to System Actions, as shown in Figure 2-48.



Figure 2-48 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 2-49. 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**.

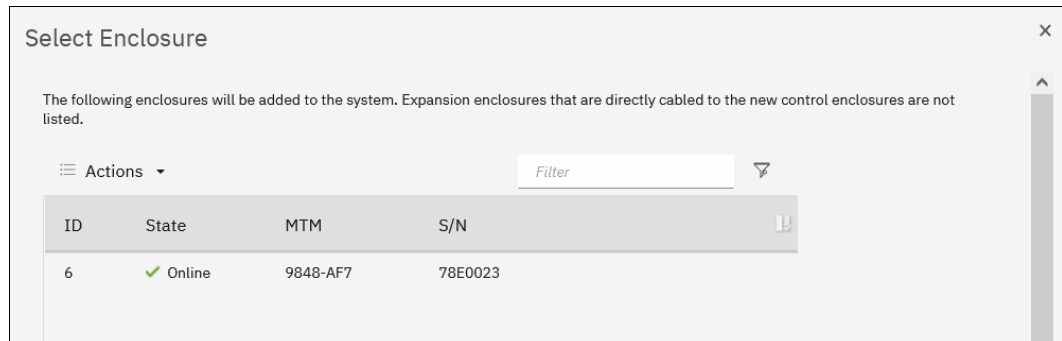


Figure 2-49 Selecting the control enclosure to add

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.
5. 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](#).
  - a. 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 2-5 shows that only `io_grp0` has nodes. Therefore, a new control enclosure can be added to `io_grp1`.

*Example 2-5 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
```

---

- b. To list control enclosures that are available to add, run the **lscontrolenclosurecandidate** command, as shown in Example 2-6. To list the expansion enclosures, run the **lsenclosure** command. Expansions that have the **managed** parameter set to no can be added.

*Example 2-6 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
```

---

- c. Add a control enclosure by running the **addcontrolenclosure** command, as shown in Example 2-7. The command triggers only the process, which starts in the background and can take 30 minutes or more.

*Example 2-7 Adding a control enclosure*

---

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

---

- d. To add an expansion enclosure, change its `managed` status to yes by running the **chenclosure** command, as shown in Example 2-8.

*Example 2-8 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
```

---

## 2.3.4 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:

1. 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 2-50.

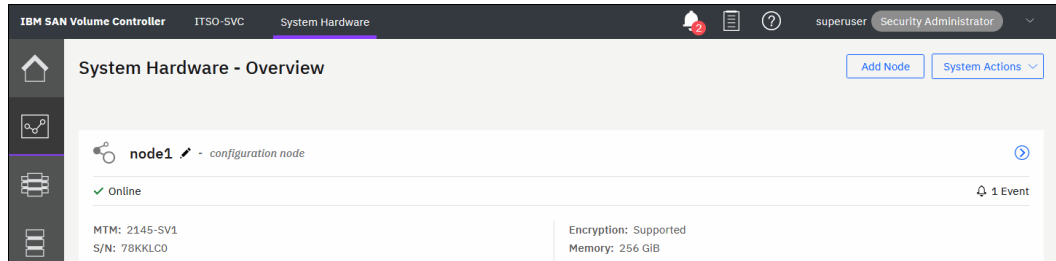


Figure 2-50 Add Node button

**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 2-51. To light the Identify light-emitting diode (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**.

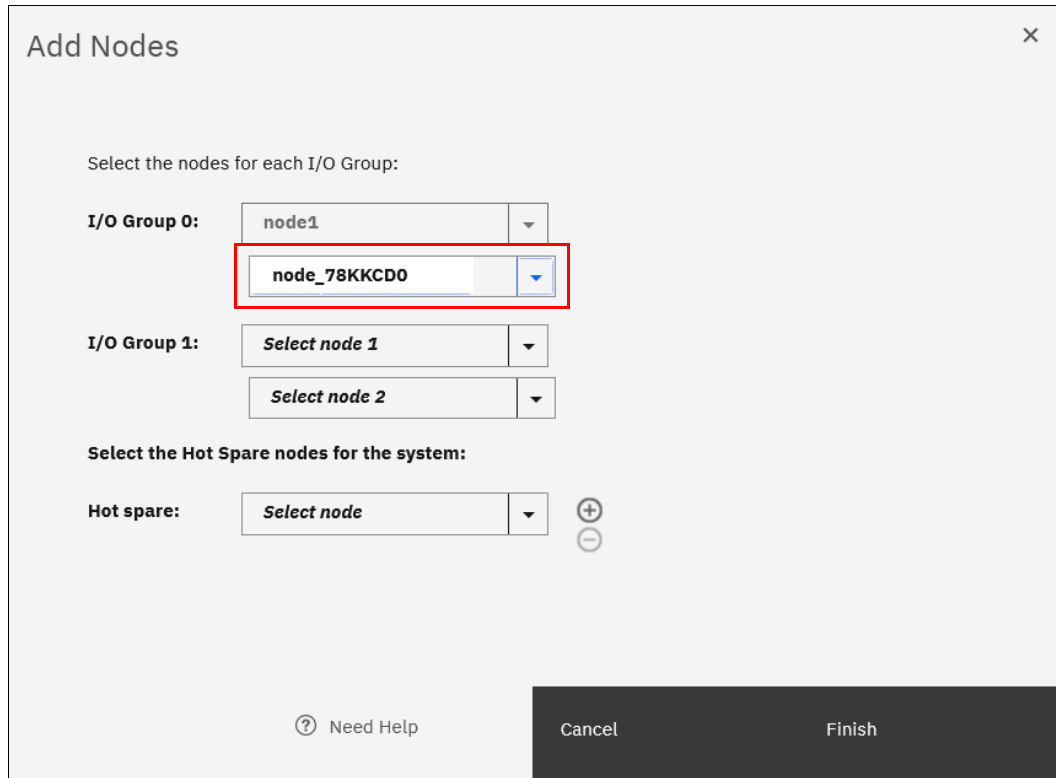


Figure 2-51 Adding a node

The Monitoring → Systems Hardware window now changes and shows that the node is added, as shown in Figure 2-52.

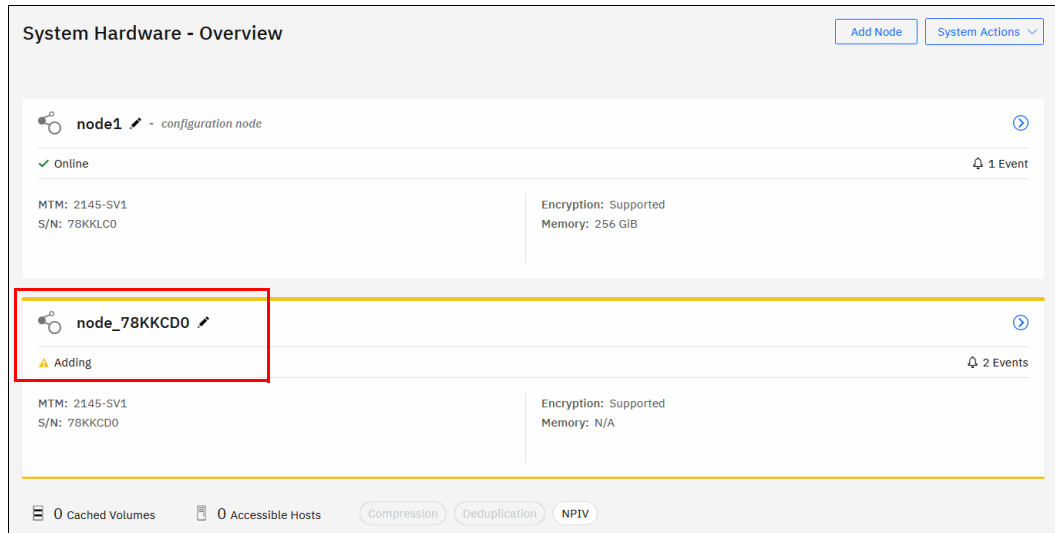


Figure 2-52 IBM SAN Volume Controller is adding node to the cluster

The node is added, as shown in Figure 2-53.

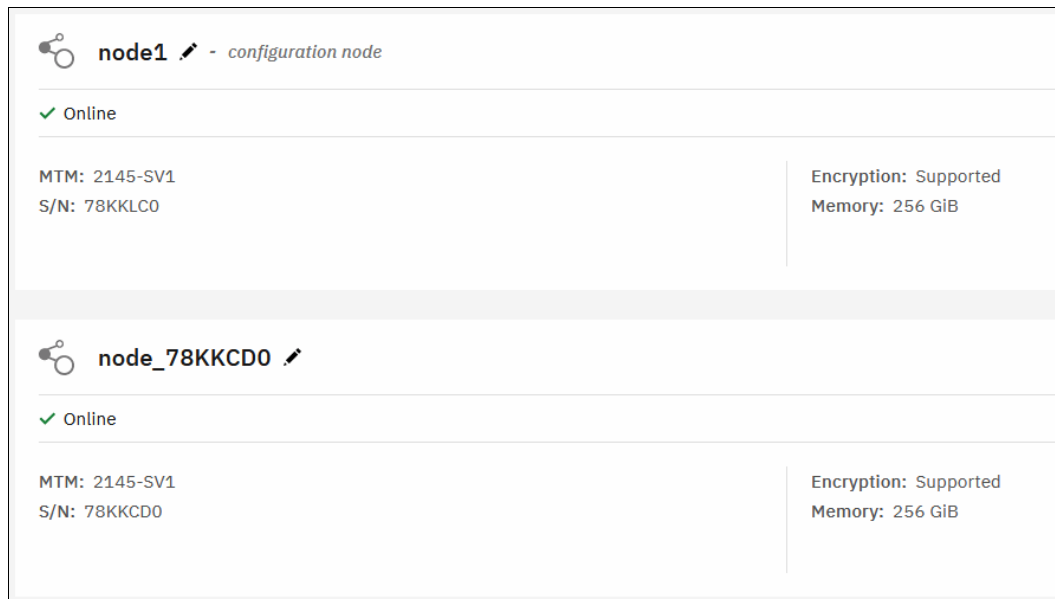


Figure 2-53 Node added

**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 completes.

To perform the same procedure by using a CLI, complete the following steps. For more information about the detailed syntax for each command, see this [IBM Documentation web page](#).

1. When adding nodes, check for unpopulated I/O groups by running **lsiogrp**. Each complete I/O group has two nodes. Example 2-9 shows that only `io_grp0` has nodes; therefore, a new control enclosure can be added to `io_grp1`.

*Example 2-9 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 **lscnodecandidate** command, as shown in Example 2-10.

*Example 2-10 Listing the candidate nodes*

---

```
BM_2145:ITS0-SVC:superuser>lscnodecandidate
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 78KKCHO          500507680C00D98A SV1      78KKCHO
2145-SV1      702D-D5FE-76AA-4034
```

---

3. Add a node by running the **addnode** command, as shown in Example 2-11. The command triggers only the process, which starts in the background and can take 30 minutes or more.

*Example 2-11 Adding a node as a spare*

---

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

---

Example 2-12 shows same command, but used to add a node to an I/O group `io_grp1`.

*Example 2-12 Adding a node to an I/O group*

---

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

---

4. Check the nodes in the system by using CLI. As shown in Example 2-13, 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 2-13 Single IO-group (two nodes) and one spare*

---

```
IBM_2145:ITS0-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_78KKLC0          500507680C00D990 online 0          io_grp0      yes
SV1      iqn.1986-03.com.ibm:2145.itso-svc.node178kklc0          78KKLC0
2 node2_78KKCDO          500507680C00D982 online 0          io_grp0      no
SV1      iqn.1986-03.com.ibm:2145.itso-svc.node278kkcd0          78KKCDO
```

---



```
3 spare1                    500507680C00D98F spare          no
SV1                        78KKLD0
```

---

A two-IO-group system with no spare is shown in Example 2-14.

*Example 2-14 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_78KKLC0  500507680C00D990 online 0          io_grp0    yes
SV1      iqn.1986-03.com.ibm:2145.itso-svc.node178kklc0 78KKLC0
2 node2_78KKCD0  500507680C00D982 online 0          io_grp0    no
SV1      iqn.1986-03.com.ibm:2145.itso-svc.node278kkcd0 78KKCD0
3 node3_78KKCHO  500507680C00D98A online 1          io_grp1    no
SV1      iqn.1986-03.com.ibm:2145.itso-svc.node378kkch0 78KKCHO
4 node4_78KKLD0  500507680C00D98F online 1          io_grp1    no
SV1      iqn.1986-03.com.ibm:2145.itso-svc.node478kkld0 78KKLD0
```

---

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**.

### 2.3.5 Business continuity with policy-based high availability

Business continuity ensures an organization can deliver services even during disruptions. While 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, while 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, ensuring business continuity. While not covered here, Storage Virtualize offers additional features like snapshots and Safeguarded snapshots to protect against data corruption or cyberattacks.

**Note:** Policy-based high availability is *not* supported by the IBM FlashSystem 5015. You need a FlashSystem model which supports clustering.

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.

## 2.3.6 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 former HyperSwap and Enhanced Stretched Cluster) topology systems, configure an active quorum device on a third, independent site. Because of the costs 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” will continue operating and processing I/O requests. The “loser” site will be 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 2-54. If you use IPv6 for management IP addresses, the Download IPv6 Application button is available and the IPv4 option is disabled. In our example, we select **Download IPv4 Application**.

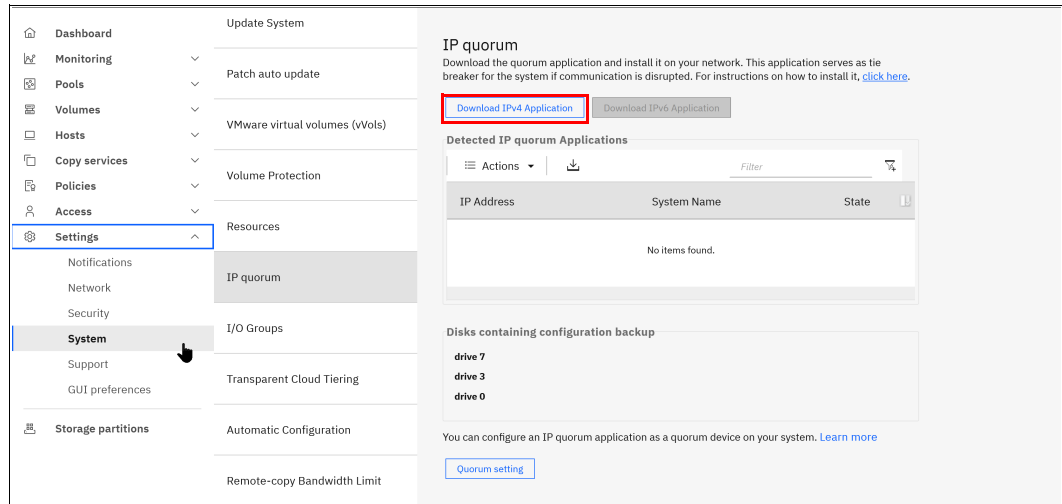


Figure 2-54 Download IPv4 quorum button

2. Click **Download...** and a window opens, as shown in Figure 2-55. 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 2-55 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 2-15.

*Example 2-15 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 2-56.




Figure 2-56 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, you can configure the following tie-breaker scenarios (a tie occurs when half of the nodes that were a member of the system are present):

- ▶ 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 Preferred, 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 is destroyed, 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 2-57.

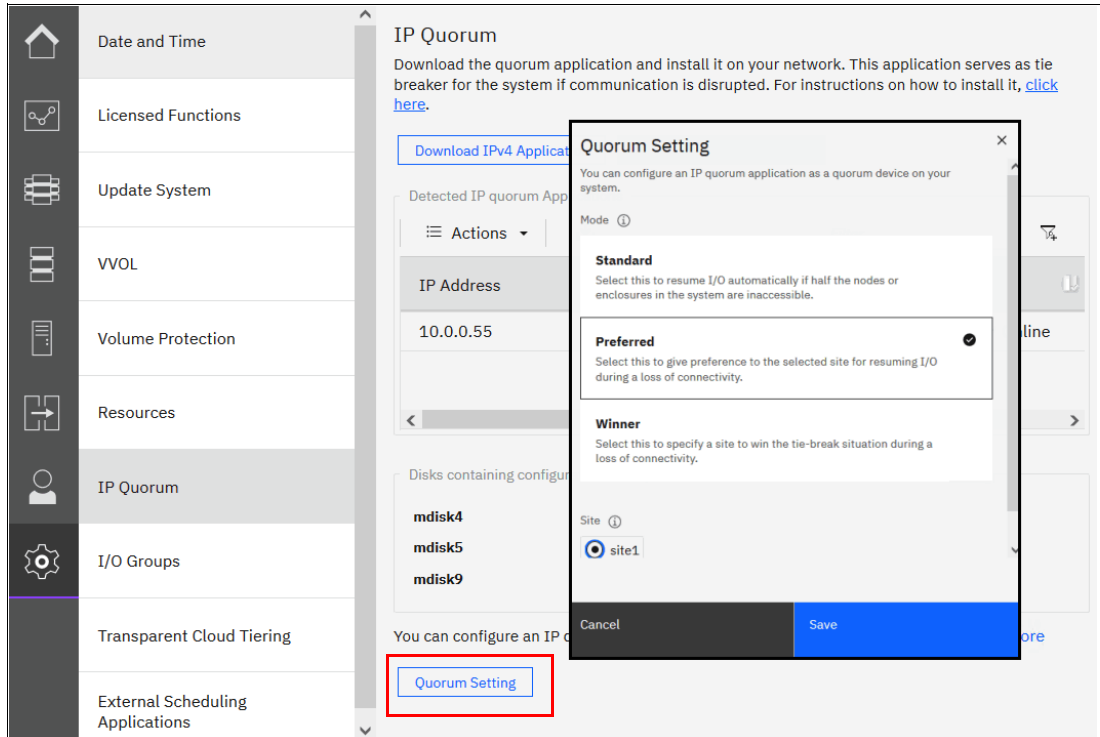


Figure 2-57 Changing the quorum mode

### 2.3.7 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 2-58.

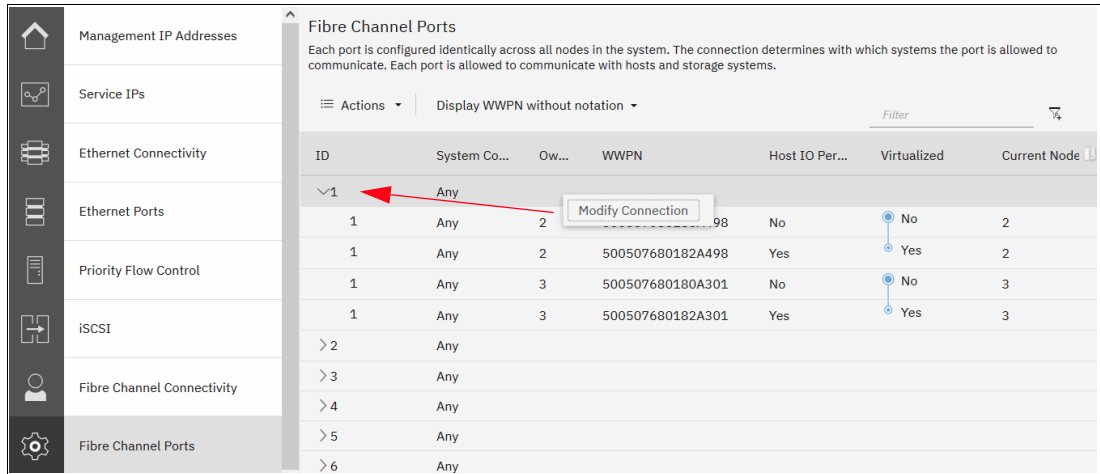


Figure 2-58 Applying a port mask by using a GUI

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 (see Figure 2-59 on page 68), 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.

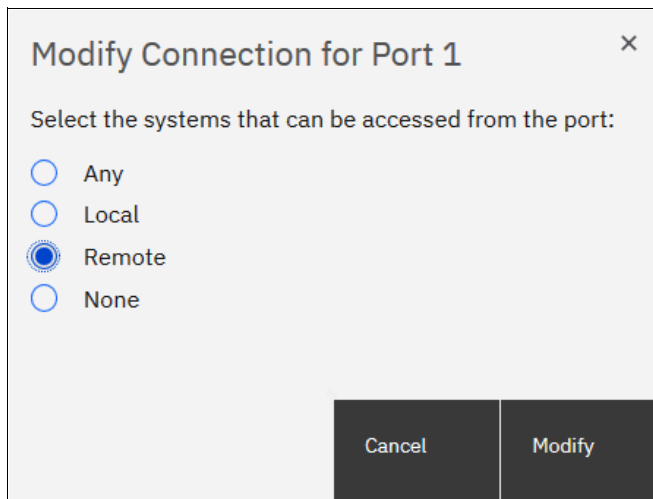


Figure 2-59 Modify Connection dialog box

The following types of traffic are allowed for each choice:

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



Complete the following steps:

1. Click **Settings** → **System** → **Automatic Configuration**. Then, select Automatic Configuration **ON** and click **Save**, as shown in Figure 2-60.




Figure 2-60 Automatic Configuration wizard enablement

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

Click **Add Enclosure** to start the adding process, or click **Skip** to move to the next step (see Figure 2-61).

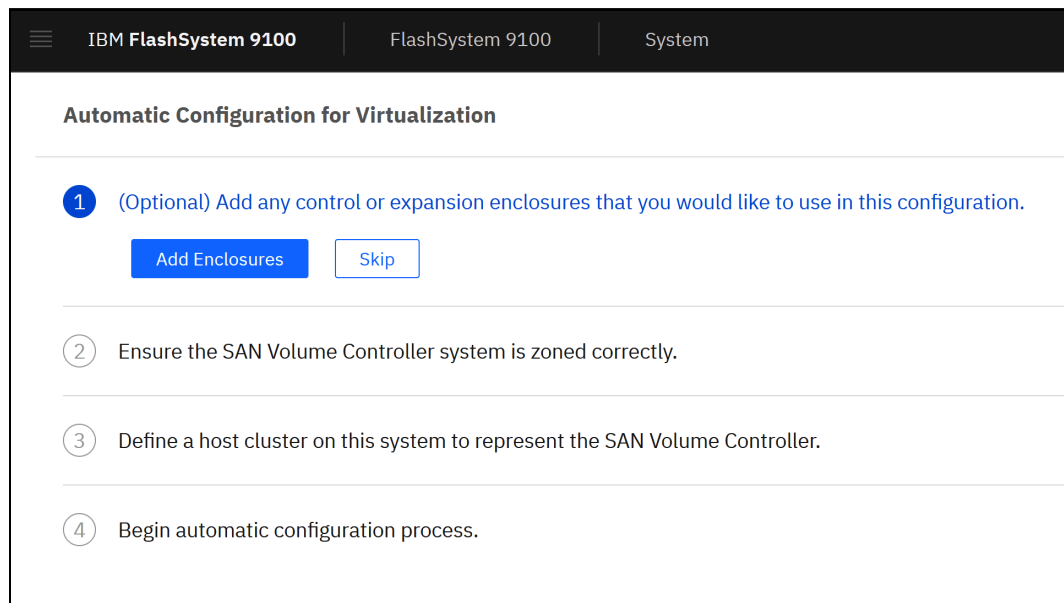


Figure 2-61 Automatic configuration: Add Enclosure

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



3. 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.
4. 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 2-62 shows that because the system detected an IBM SAN Volume Controller cluster with dual I/O groups, four WWNNs are selected.




Figure 2-62 Defining a host cluster

5. 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 2-63 on page 72. For convenience, name the host objects to match the IBM SAN Volume Controller node names or serial numbers.

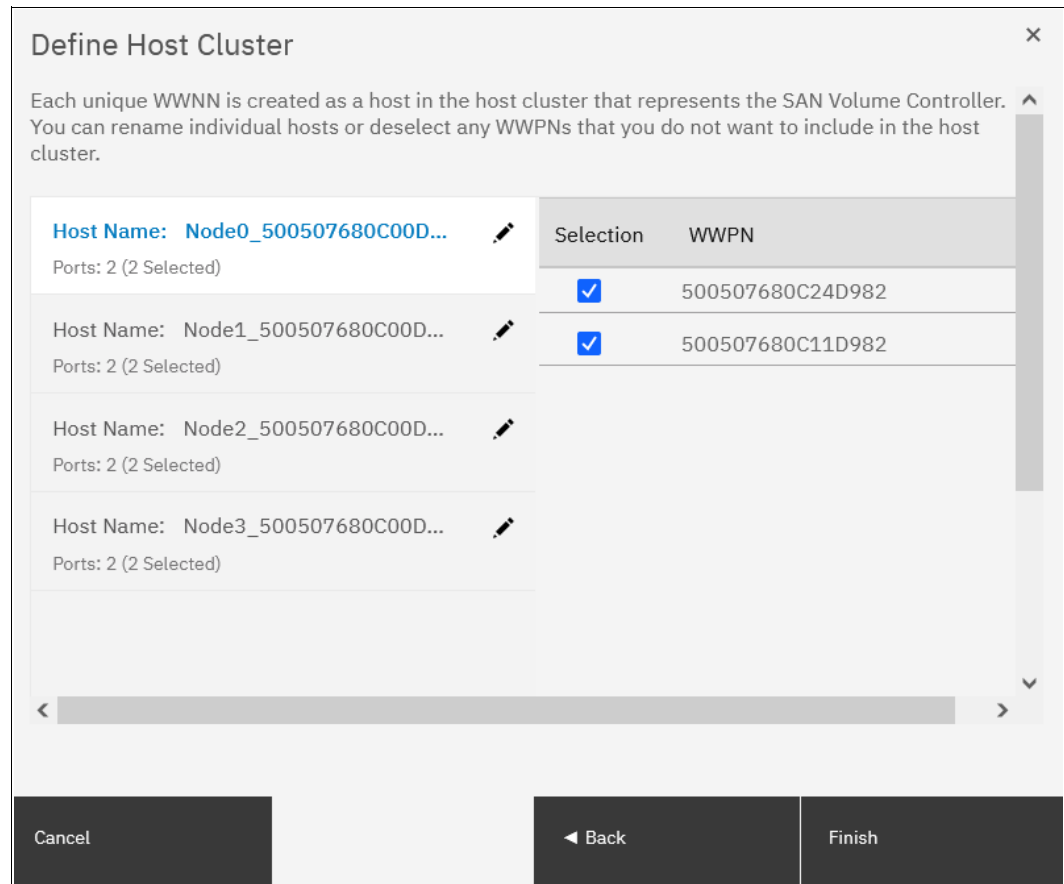


Figure 2-63 Hosts inside an IBM SAN Volume Controller host cluster

6. Click **Automatic Configuration** and check the list of internal resources that are used, as shown in Figure 2-64.

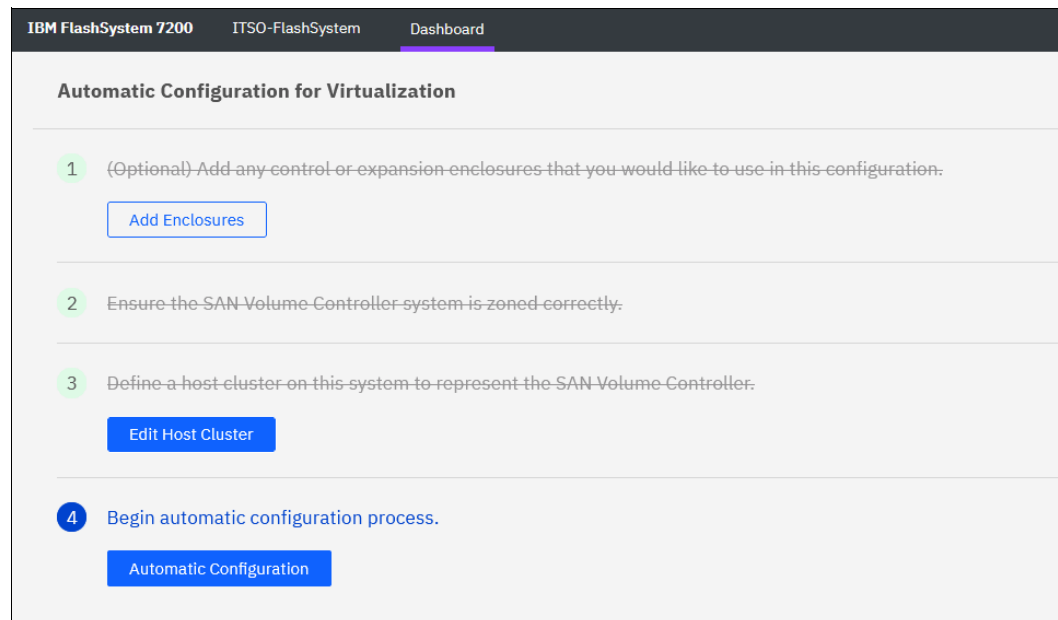


Figure 2-64 Begin the automatic configuration process

7. If the system uses compressed drives (FCM drives), you are prompted to enter your expected compression ratio (or total capacity that is to be provisioned to IBM SAN Volume Controller), as shown in Figure 2-65. If IBM SAN Volume Controller uses encryption or writes data that is not compressible, set the ratio to 1:1 and then, click **Next**.

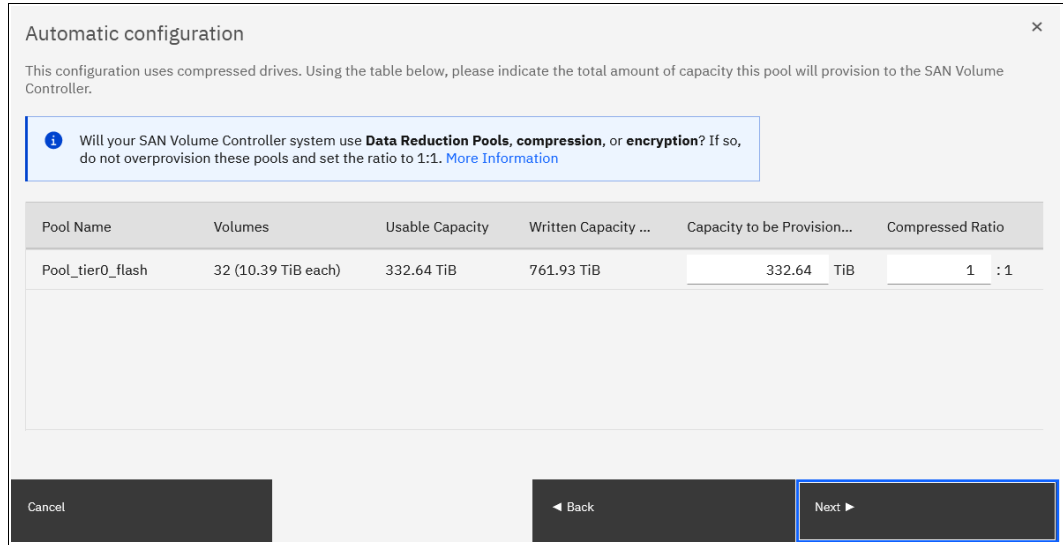


Figure 2-65 Automatic pool configuration

8. Review the pool (or pools) configuration, as shown in Figure 2-66, and click **Proceed** to trigger the commands that applies it.

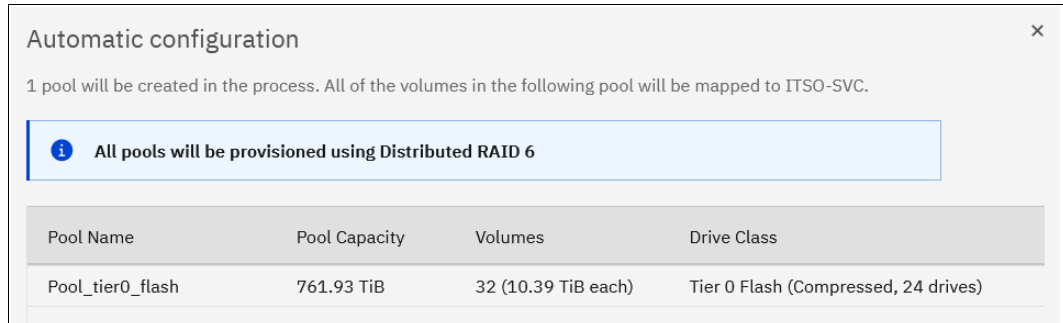


Figure 2-66 Pools configuration

- 9. When the Automatic Configuration for Virtualization wizard completes, you see the window that is shown in Figure 2-67. After clicking **Close**, you can proceed to the IBM SAN Volume Controller GUI and configure a new provisioned storage.

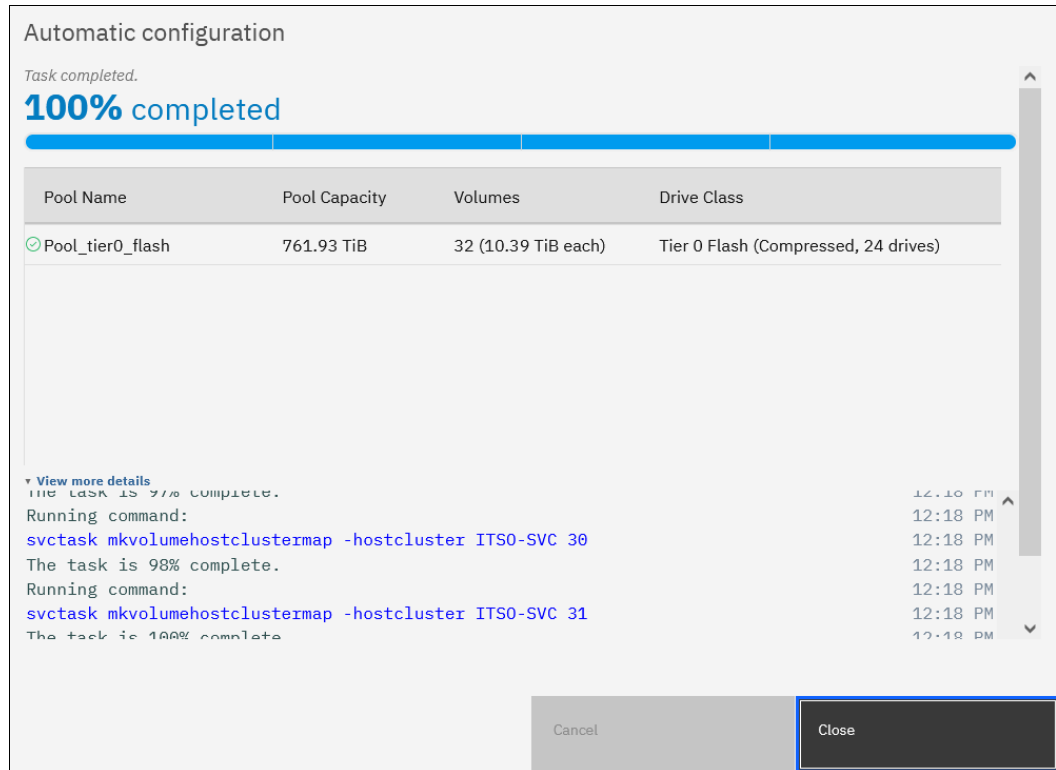


Figure 2-67 Automatic configuration running commands

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

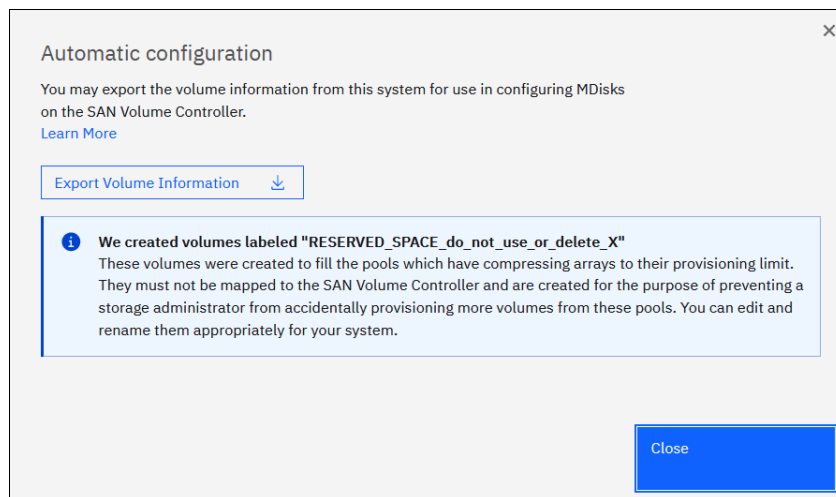


Figure 2-68 Automatic configuration complete



## Step-by-step configuration

This chapter describes the steps used to quickly implement a Storage Virtualize system utilizing the most commonly used functions. Where applicable it will provide links to further reading for advanced functions.

In this chapter we discuss the Storage Virtualize GUI, steps needed for network configuration, creating pools and assigning storage, configuring hosts, basic snapshot and asynchronous replication configuration.

This chapter has the following chapters:

- ▶ “The Storage Virtualize GUI” on page 34
- ▶ “Network configuration” on page 35
- ▶ “Pools and managed disks configuration” on page 37
- ▶ “Configuring volumes” on page 44
- ▶ “Configuring hosts” on page 45
- ▶ “Snapshots and replication” on page 50

### 3.1 The Storage Virtualize GUI

Throughout this chapter we will be using the Storage Virtualize GUI. The 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, be aware that the last action that is entered from the GUI is the action that takes effect if two users change the same object simultaneously.*

#### 3.1.1 Accessing the GUI

To access the IBM GUI, enter the IP address that was set 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 intended to be used for initial configuration and servicing the system only. For more information on user accounts see [Users documentation](#).

#### 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.

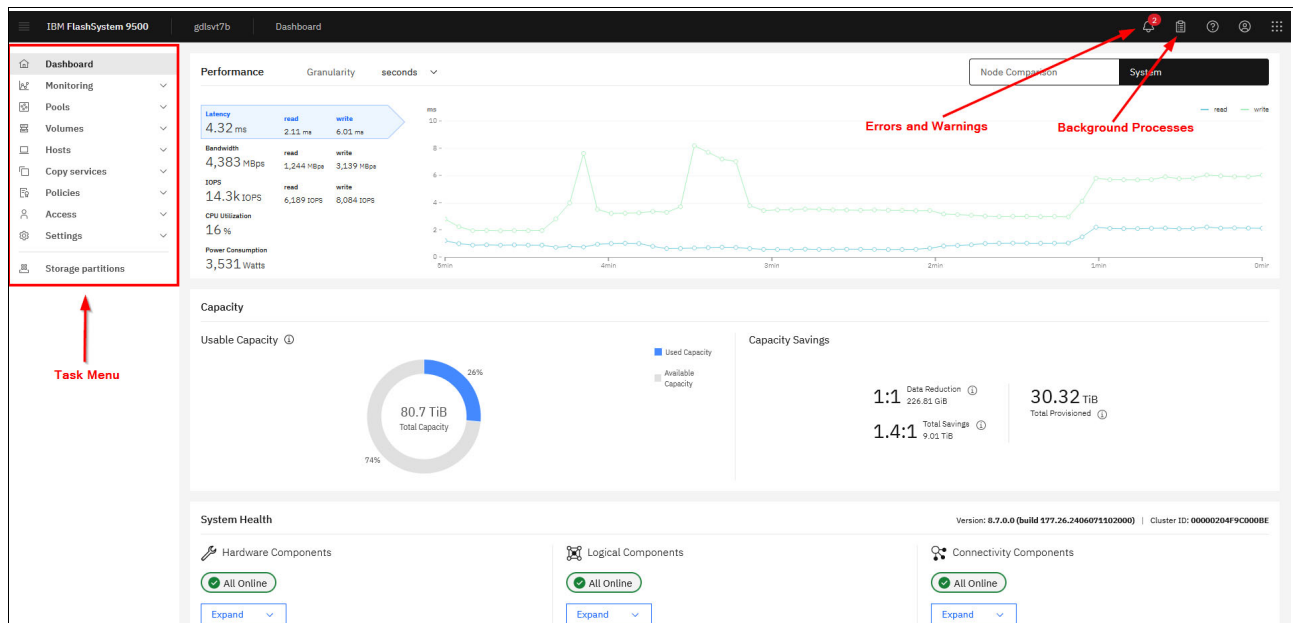


Figure 3-1 Welcome page with the dashboard

## Task menu

The IBM Storage Virtualize GUI task menu is always available on the left side 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 utilization. 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 utilization 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, logical, 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 that you are interested in is managed.

## 3.2 Network configuration

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

It also allows you to 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 if desired. Starting in 8.7.0 the management IPs no longer are required to use ports 1 and 2. The two management IPs can be configured on any ports. 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 will respond to requests to the management IP.

**Note:** The system will always use the management IP on the lowest numbered port for outbound communication, for example, Cloud Call Home, e-mail 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. These IPs should be changed 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, access IP quorum, or for remote support assistance.

### 3.2.3 Additional Ethernet ports

The Ethernet ports menu is for configuring ports to be used for host attachment, replication, and virtualizing back-end storage via iSCSI.

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

The screenshot shows the 'Ethernet Ports' configuration page in the IBM FlashSystem management interface. The page title is 'Ethernet Ports' and it includes a subtitle: 'The Ethernet ports can be used for iSCSI, ISER (SCSI), NVMe/RDMA and NVMe/TCP connections (if available), host attachment and remote copy.' The main content is a table with the following columns: Link State, Speed, Host Attach, Storage, and Replication. An 'Actions' menu is open over the table, showing options: Manage IP Addresses, Modify Remote Copy, Modify Host Attachment Support, Modify Storage Ports, and Modify Maximum Transmission Unit. The table data is as follows:

		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
node1	2	Inactive		No	No	No
node2	3	Active	25Gb/s	Yes	Yes	Yes
node1	3	Active	25Gb/s	Yes	Yes	Yes
node1	4	Active	25Gb/s	Yes	Yes	Yes
node2	4	Active	25Gb/s	Yes	Yes	Yes
node1	5	Active	25Gb/s	Yes	Yes	Yes
node2	5	Active	25Gb/s	Yes	Yes	Yes
node1	6	Active	25Gb/s	Yes	Yes	Yes
node2	6	Active	25Gb/s	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.





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 further details on portsets see [Portsets documentation](#).

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

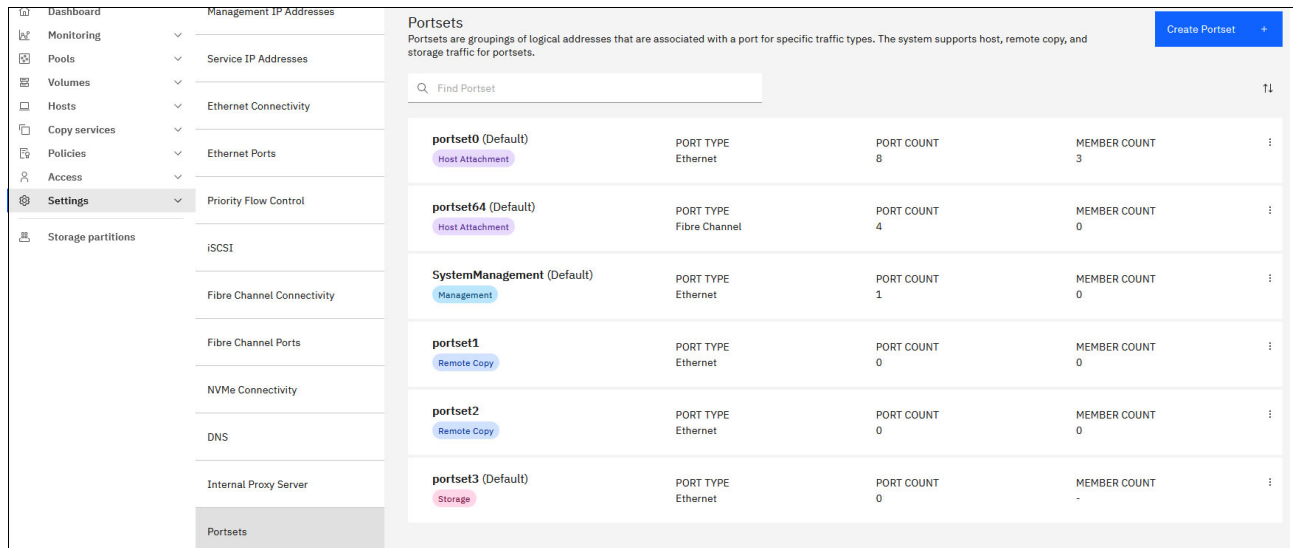


Figure 3-4 Portsets

## 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 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 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 themselves 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 utilizing a provisioning policy for pools. If a provisioning policy is assigned, any volumes created from the pool are provisioned based on the capacity savings method defined in the policy. A policy can be created and assigned or unassigned to a pool at any time, but it effects only volumes created while the policy was active To create a provisioning policy select **Policies** → **Provisioning policies** → **Create policy**

For more information see [Provisioning policy documentation](#).

**Note:** Provisioning policy will 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, 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 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 amount of usable capacity that is required 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 documentation](#).

### 3.3.3 Ransomware threat detection

Ransomware threat detection is automatically enabled provided 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 128GB 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.

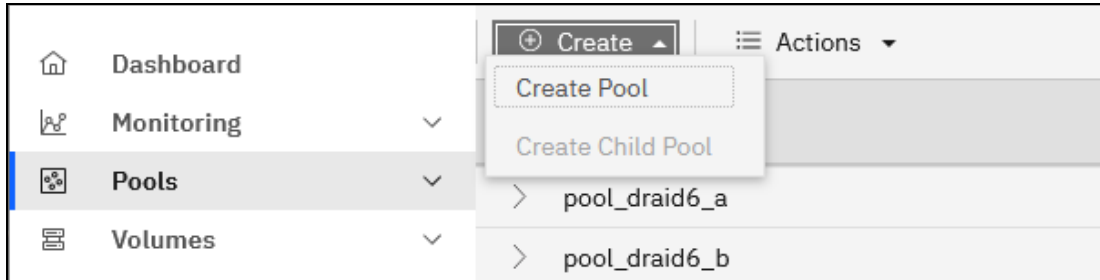


Figure 3-5 Create Pool

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

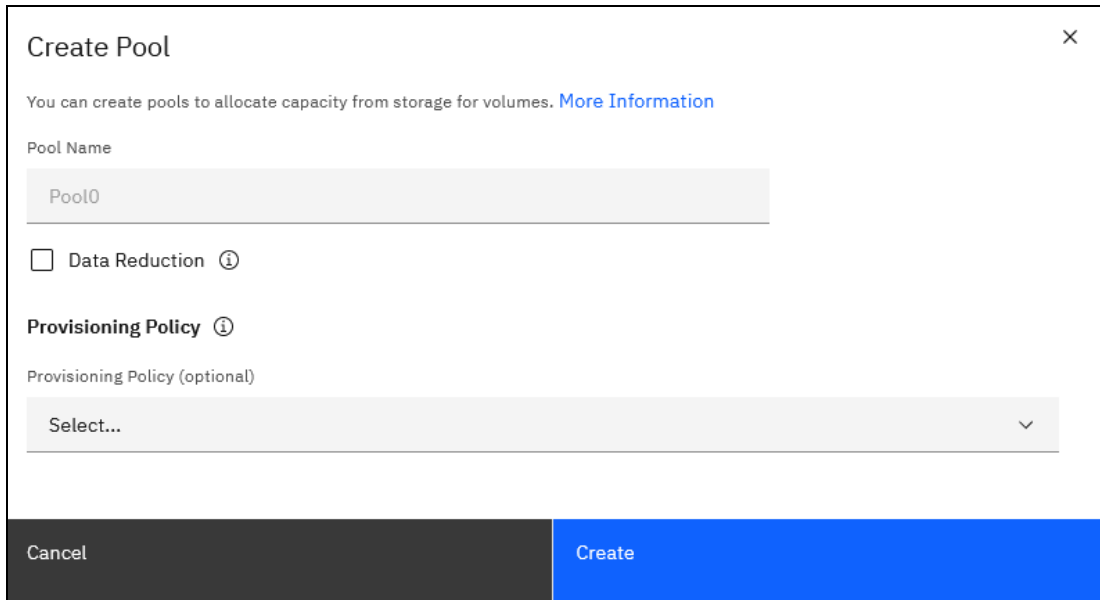


Figure 3-6 Create Pool panel

2. Select the **Data reduction** option if you wish 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.

Refer to the [Configuration Limits for IBM FlashSystem and SAN Volume Controller](#)

**Note:** 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 41.

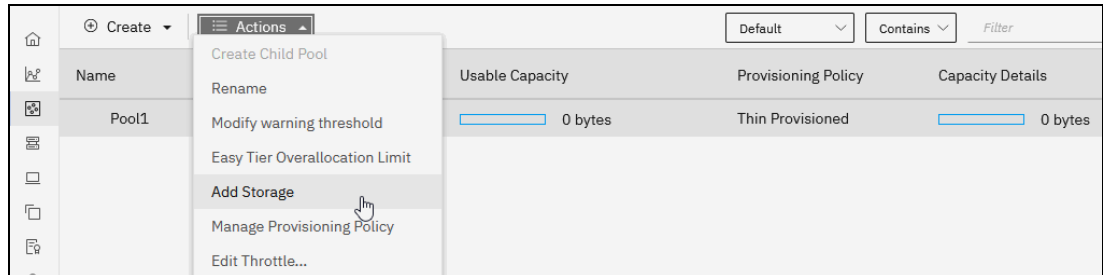


Figure 3-7 Add Storage

This will open a new panel with a suggested RAID array configuration based on the installed drives. See Figure 3-8 on page 41

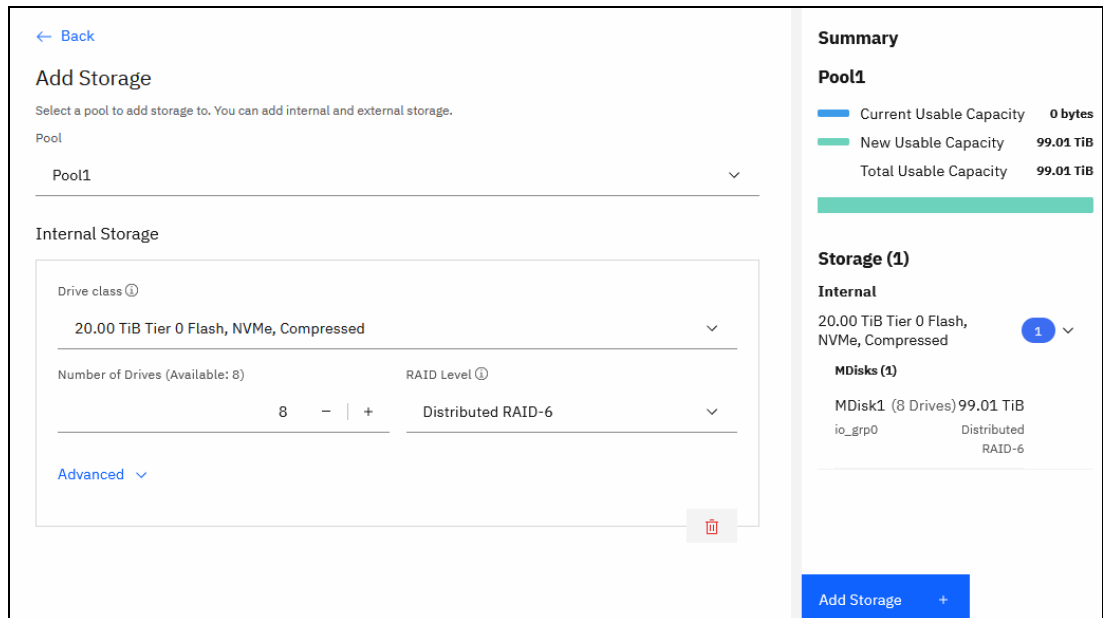


Figure 3-8 RAID array

### 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 makes your storage environment more flexible, cost-effective and easy 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 via 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](#).

Once external LUs are discovered by the IBM Storage Virtualize system they will be 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 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.

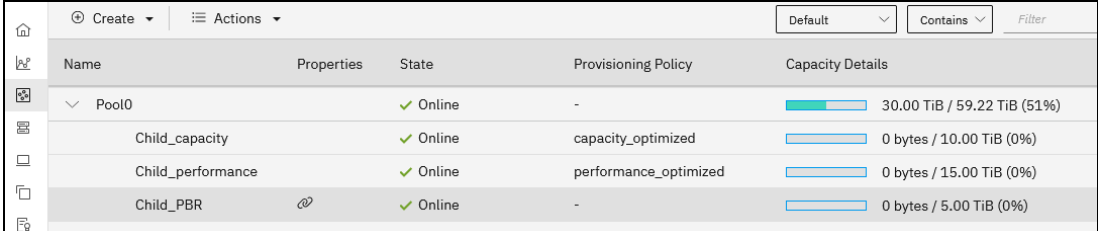
The storage tier setting can also be changed after the MDisk is assigned to the pool. For more information see [Easy tier documentation](#).

## 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 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 non-disruptively. 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 non-encrypted 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 non-encrypted 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 panel opens. See Figure 3-10.

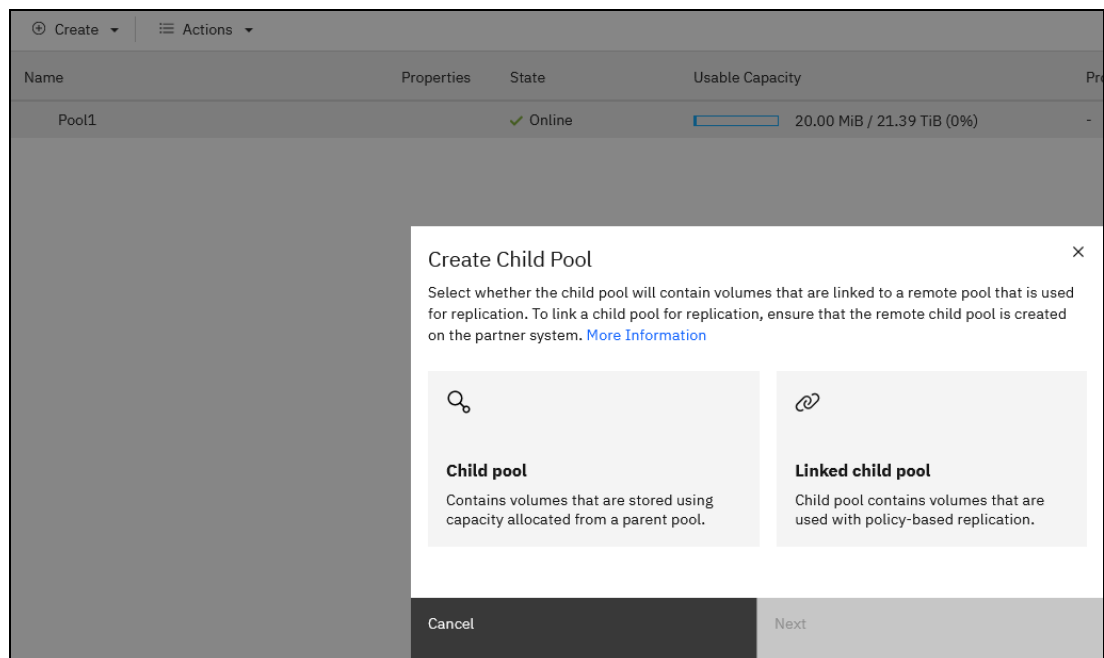


Figure 3-10 Create Child Pool

## 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 [Volume documentation](#).

### 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 together within a volume group. This ensures that the volumes within the group are managed and treated according to the same service level guidelines.
- ▶ Some applications or data sets may 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 together 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 may 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 on using policy-based replication, configuring it and assigning a replication policy to the volume group before creating volumes within the group will allow you to 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 will define the volume properties. However, if the pool has a pre-assigned provisioning policy, the capacity savings option will be locked and reflect the policy's settings.

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



The newly created volumes will 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. To increase the format rate for a volume (if desired), right-click the volume and select **Modify Mirror Sync Rate**. Then, choose the desired rate.

*Use caution to not over utilize the systems resources by formatting too many volumes at too high rate.* If experiencing a system performance problem after increasing the mirror sync rate, it can be decreased in the same manner.

**Note:** If you are using a system with IBM FlashCore® Modules the data written to the system will be 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), which 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 (also known as 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 documentation](#).

## 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 IBM Storage Virtualize storage system through various supported interconnect protocols.

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

### 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 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 documentation](#).

**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

### N\_Port ID Virtualization

IBM Storage Virtualize systems utilize 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 give each physically connected system port the ability to create four more NPIV ports.

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.

To view the virtual WWPNs to be used in system to host select **Settings** → **Network** → **Fibre Channel Ports**. Expand the twisty for each port. Columns will 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 will have two ports zoned to each IO group creating 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, with which you can insert target ports and multiple initiator ports in a single zone for easy of management but act the same 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 by using port 4791 between source and destination IPs; 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 IPs 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 first 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 panel will display. See Figure 3-11.

Figure 3-11 Add Host

**Tip:** The Host port drop down will show FCP initiator WWPNs currently logged into the system. If an expected WWPN is missing, check switch zoning and rescan the storage from the hosts. Some operating systems may log out if no LUNs are mapped to the host. If that is the case select **Enter Unverified WWPN** and enter the host WWPNs manually.

### 3.5.5 Mapping volumes for host access

In order for the host or host cluster to access the volumes they must be mapped to the host or host cluster. To perform the mappings, select the **volumes 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:** In most cases all volumes within a volume group would be mapped to the same host or host cluster and the mapping can easily be done within the volume group view.

## 3.6 Snapshots and replication

In this section we discuss configuration of 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 exact same time. This was achieved through the use of 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 FlashCopy® operations. It achieves this by offering a more straightforward setup process and separating the snapshot and clone features. With 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. In order 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 documentation](#).

#### 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 via the GUI, on a schedule 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 tailored to their specific needs.

- ▶ The snapshot policies 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. Additionally, 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 only be deleted before their expiration time by a security administrator.

**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 may shift forward and backward with daylight savings time changes.

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

Snapshot Policies						
Create policies and assign them to volume groups to automate the creation and retention of your snapshots.						
<input type="text" value="Search table..."/> <span style="float: right;"> </span>						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	Assign Policy
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.

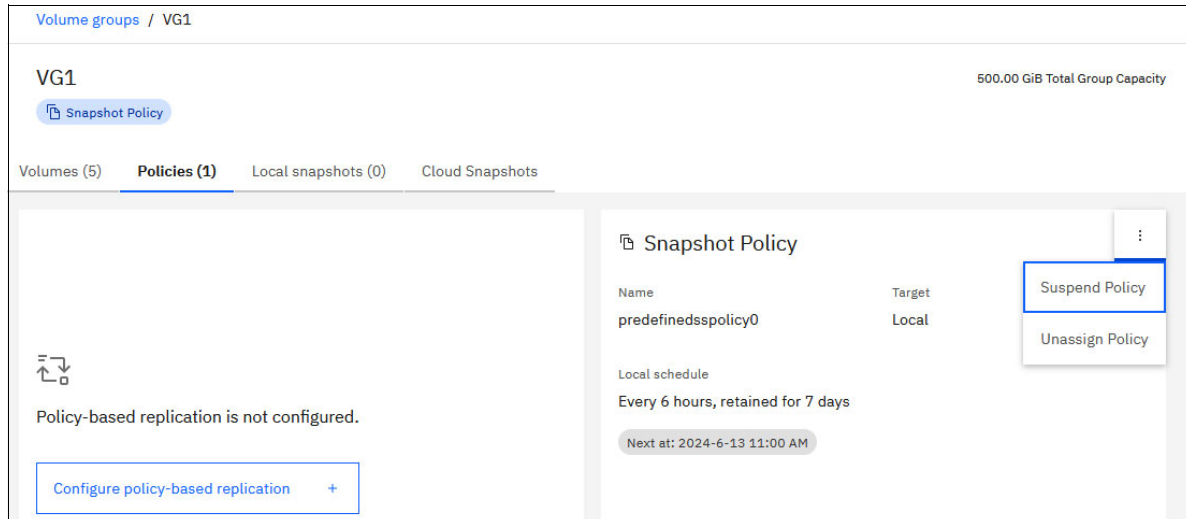


Figure 3-13 Suspend Policy

### 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 ensures mutual consistency between all volumes in the volume group.

To quickly configure policy based replication between two systems, both will require at least one IP address created and assigned to a replication portset along with at least one pool with storage created. Multiple IPs 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** the select **Test Connection**. If the partner meets requirements for policy based replication the Use policy-based replication selection can be checked. Fill out the rest of the panel and select **Create**. Repeat these steps on the partner system. See Figure 3-14 on page 53.



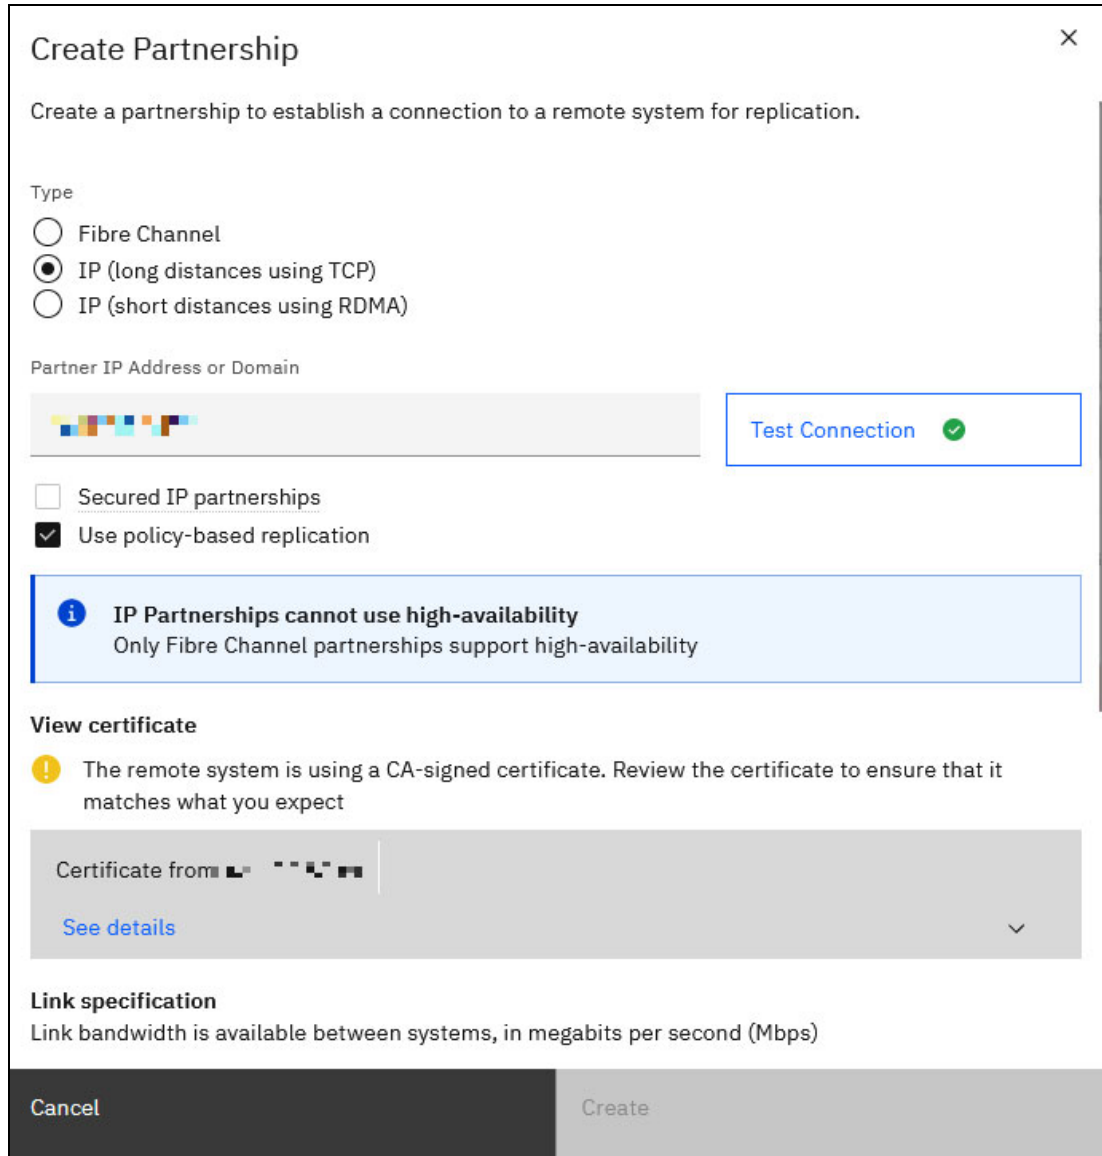


Figure 3-14 Create Partnership

2. Once the partnership shows a green dot and configured select **Setup policy-based replication**. See Figure 3-15 on page 54.

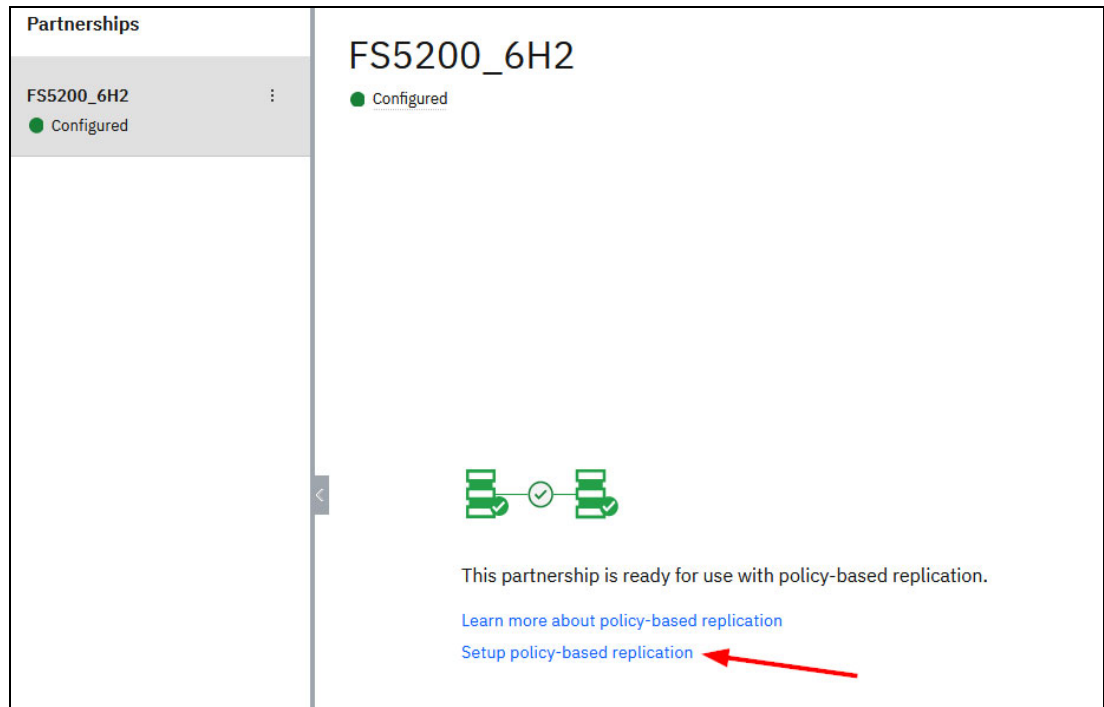


Figure 3-15 Partnership Created

3. Select **Setup policy-based replication** to be guided through the rest of the configuration. See Figure 3-16.

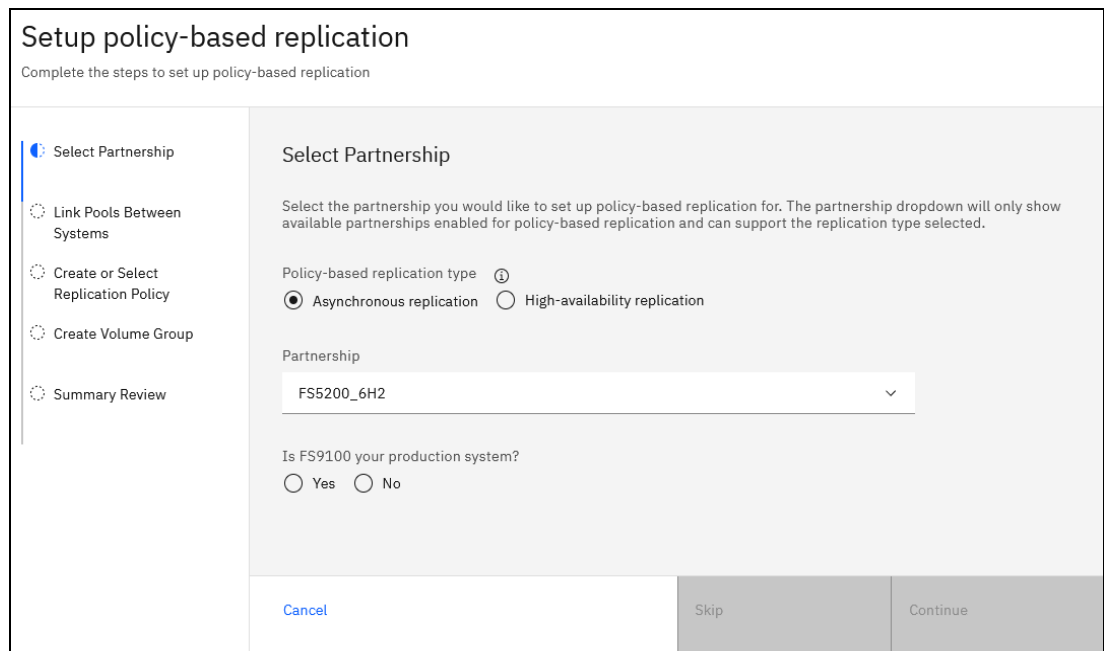


Figure 3-16 Setup policy-based replication wizard

For more information see IBM Documentation [Asynchronous disaster recovery replication](#).

Also see [Policy-based replication Redpaper](#).

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



## 4



# Verifying configuration and basic operations

This chapter provides steps to verify the configuration along with tips to resolve common implementation problems. It also provides additional information about functions that should be considered when implementing a new system.

In this chapter we discuss the system health dashboard, verifying configuration of objects configured in Chapter 3, “Step-by-step configuration” on page 33, system security, getting support from IBM, and data migration.

This chapter has the following sections:

- ▶ “Verifying the configuration” on page 58
- ▶ “Additional settings and basic operations” on page 61

## 4.1 Verifying the configuration

In this section we discuss some commons tasks to verify the configuration.

### 4.1.1 System Health Dashboard

The system dashboard provides a way to quickly assess the overall condition of the system and 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 quickly access relevant information. See Figure 4-1.

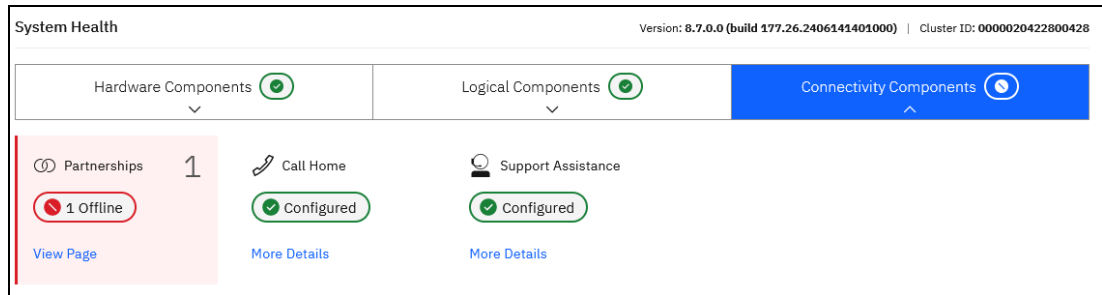


Figure 4-1 System Health

### 4.1.2 Verifying network configuration

The network IPs and WWPNs can be verified by selecting **Settings** → **Network** → **Portsets**. Selecting any of the portsets will allow viewing details of the ports and IP addresses assigned to that portset. See Figure 4-2.



Figure 4-2 Portset Mappings

Selecting **IP Addresses** or **Partnerships** gives further details. See Figure 4-3 on page 59

The screenshot shows a web interface for 'portset1'. It has a 'Portset Details' tab and a 'Portset Mappings' tab. Under 'Portset Mappings', there is a 'Back' button and a search icon. Below is a table with two rows of data. Each row has a visual representation of a portset on the left, followed by columns for PORT, I/O GROUP, NODE ID, SPEED, and VLAN.

PORT	I/O GROUP	NODE ID	SPEED	VLAN
1	io_grp0	2	1Gb/s	-
1	io_grp0	1	1Gb/s	-

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, check **Monitoring** → **Events** for any related events.

There is a column for Usable Capacity and a column for Written Capacity Limit. If utilizing over provisioned storage (for example, FCMS) the values will be different. The usable capacity represents the physical capacity available after the data is reduced via compression and deduplication. The Written Capacity Limit is the effective capacity of data 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.

The columns displayed can be modified by right clicking on the column titles bar. There are some useful capacity related columns that can be displayed.

If a volume is thin provisioned in a standard pool adding the columns Real Capacity and Used Capacity will 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 will display 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, thus providing 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** → **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 moving data to a Storage Virtualize system. To download the Comprestimator that can be installed on a server follow this [link](#).

### 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. e.g. both host WWPNs are logged into node1 while one host WWPN is logged into 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 into 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

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

**Tip:** If using Broadcom fibre channel switches the **fcping** command run via the switch CLI can be quickly used to verify zoning and WWPN connectivity.



## 4.2 Additional settings and basic operations

The following sections discuss additional settings that an administrator should be aware of when implementing a new systems 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 utilized settings are described below.

For more information see [Security documentation](#).

Also see [Storage Virtualize Security Feature Checklist](#).

#### Remote authentication

Remote authentication allows users to authenticate to the system 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 simplify user management and access to enforce password policies more efficiently, and to separate user management from storage management.

A remote user is authenticated on a remote LDAP server. A remote user does not need to be added 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 in which a user belongs are viewable by those users.

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

## 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 simplifies the certificate renewal process and prevents 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.

**Note:** An externally signed certificate must meet the following requirements:

- ▶ **X.509v3 Key Usage Extensions:** Include Digital Signature.
- ▶ **X.509v3 Extended Key Usage Extensions:** Include TLS Web Server Authentication and TLS Web Client Authentication. Additionally, 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 documentation](#).

## 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 ended up being 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.

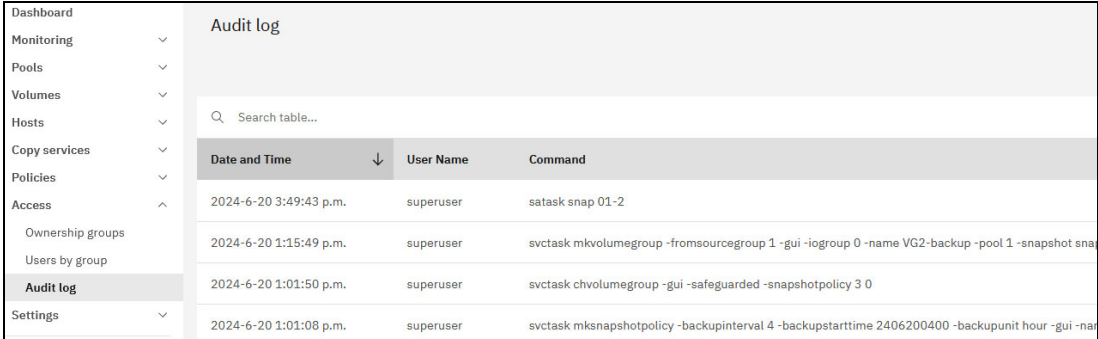
The audit log tracks action commands that are issued through an SSH session, management GUI, or Remote Support Assistance. It provides the following entries:

- ▶ 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.



Dashboard		Audit log	
Monitoring	▼	Search table...	
Pools	▼	<b>Date and Time</b>	<b>User Name</b>
Volumes	▼	2024-6-20 3:49:43 p.m.	superuser
Hosts	▼	2024-6-20 1:15:49 p.m.	superuser
Copy services	▼	2024-6-20 1:01:50 p.m.	superuser
Policies	▼	2024-6-20 1:01:08 p.m.	superuser
Access	▲		
Ownership groups	▼		
Users by group	▼		
<b>Audit log</b>			
Settings	▼		

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 embedded in all IBM Storage Virtualize storage products. By enabling call home, the health and functionality of your system is constantly monitored by IBM. Should a software or hardware error occur, 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 quickly develop an action plan for problem resolution.

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

- ▶ Cloud Services uses HTTPS to connect directly to IBM from the management IP assigned to the lowest physical port id over port 443
- ▶ Email services requires 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 following white paper. [IBM Storage Virtualize Products Call Home and Remote Support Overview](#).

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

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

### Support package

If you encounter a problem and contact the IBM Support Center, you will be asked to provide a support package (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 will advise 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](#).

## 4.2.4 Data migration

Data migration is an important part of an implementation; therefore, you must prepare a detailed data migration plan. You might need to migrate your data for one of the following reasons:

- ▶ Redistribute a workload within a clustered system across back-end storage subsystems.
- ▶ Move a workload:
  - On to newly installed storage
  - Off old or failing storage ahead of decommissioning it
  - To rebalance a changed load pattern
- ▶ Migrate data from:

- An older disk subsystem
- One disk subsystem to another subsystem

Because multiple data migration methods are available, choose the method that best fits your environment, operating system platform, type of data, and the application's service-level agreement (SLA).

Data migration methods can be divided into three classes based on:

- ▶ The host operating system; for example, by using the system's logical volume manager (LVM) or VMWare vMotion.
- ▶ Specialized data migration software.
- ▶ The system data migration features.

With data migration, apply the following guidelines:

- ▶ Choose the data migration method that best fits your operating system platform, type of data, and SLA.
- ▶ Choose where you want to place your data after migration in terms of the storage tier, pools, and back-end storage.
- ▶ Check whether enough free space is available in the target storage pool.
- ▶ To minimize downtime during the migration, plan all the required changes, including zoning, host definition, and volume mappings.
- ▶ Prepare a detailed operation plan so that you do not overlook anything at data migration time. Have the plan peer-reviewed and formally accepted by a suitable technical design authority within your organization (especially for a large or critical data migration).
- ▶ Perform and verify a backup before you start any data migration.
- ▶ You might want to use the system as a data mover to migrate data from a non-virtualized storage subsystem to another non-virtualized storage subsystem. In this case, you might have to add checks that relate to the specific storage subsystem that you want to migrate.

Be careful when you use slower disk subsystems for the secondary volumes for high-performance primary volumes because the system's cache might not buffer all of the writes. Flushing cache writes to slower back-end storage might affect the performance of your hosts.

- ▶ Consider storage performance. The migration workload might be much higher than expected during normal operations of the system. If application data is on the system to which you are migrating, the application performance might suffer if the system is overloaded. Consider the use of host or volume level throttles when performing migration on a production environment.

Non-Disruptive System Migration can be used to migrate data from one IBM Storage Virtualize system to another non-disruptively. You can create this specific remote-copy relationship that copies data from source volumes on a system that you are decommissioning to auxiliary volumes that are on another system. The nondestructive system migration is a remote-copy relationship type that is dedicated to volume migration between systems.

For further details see [Data migration documentation](#).

**Note:** On 8.7.0 systems where legacy remote-copy functions have been deprecated, `remote_copy` compatibility mode will need to be enabled to use Non-Disruptive System Migration.

## 4.2.5 SCSI unmap

File deletion in modern file systems is realized by updating file system metadata and marking the physical storage space that is used by the removed file as unused. The data of the removed file is not overwritten, which improves file system performance by reducing the number of I/O operations on physical storage that is required to perform file deletion.

However, this approach affects the management of the real capacity of volumes with enabled capacity savings. File system deletion frees space at the file system level, but physical data blocks that are allocated by the storage for the file still take up the real capacity of a volume.

To address this issue, file systems added support for the **SCSI UNMAP** (sometimes referred to as **TRIM**) command, which can be run after file deletion. It informs the storage system that physical blocks that are used by the removed file must be marked as no longer in use so that they can be freed. Modern operating systems run SCSI UNMAP commands to only storage that advertises support for this feature.

The SCSI UNMAP command on IBM Storage Virtualize systems, which enables hosts to notify the storage controller of capacity that is no longer required and can be reused or deallocated, which might improve capacity savings. This is particularly important with over provisioned storage. e.g. Systems with DRPs or FCMs.

To check if `host_unmap` is enabled the `lssystem` command can be used. The state can be changed with the `chsystem` command. Always leave `backend_unmap` enabled unless directed to change by IBM support. See Example 4-1.

*Example 4-1 Check if host\_unmap is enabled*

---

```
IBM_FlashSystem:FS9100:superuser>lssystem | grep unmap
host_unmap off
backend_unmap on
IBM_FlashSystem:FS9100:superuser>chsystem -hostunmap on
IBM_FlashSystem:FS9100:superuser>lssystem | grep unmap
host_unmap on
backend_unmap on
IBM_FlashSystem:FS9100:superuser>
```

---

Enabling `host_unmap` on the storage advertises support when a host queries the storage. In most cases the host operating system will need to rescan its disks or reboot before it will send SCSI UNMAP commands to the storage.

For further details on SCSI Unmap see [SCSI Unmap support in Spectrum Virtualize systems](#).

## 4.2.6 I/O throttling

Throttles are a mechanism to control the amount of resources that are used when the system is processing I/Os on supported objects. The system supports throttles on hosts, host clusters, volumes, copy offload operations, and storage pools. If a throttle limit is defined, the system either processes the I/O for that object, or delays the processing of the I/O to free resources for more critical I/O operations. A throttle is a useful mechanism for implementing Quality of service. e.g. Volumes used by a development are put in a child pool that has a throttle applied to prevent workload from impacting production.

The limit can be set in terms of the number of IOPS or bandwidth measured in megabytes per second (MBps), gigabytes per second (GBps), or terabytes per second (TBps). By default, I/O

throttling is disabled, but each volume can have up to two throttles that are defined: one for bandwidth and one for IOPS.

When deciding between the use of IOPS or bandwidth as the I/O governing throttle, consider the disk access profile of the application that is the primary volume user. Database applications generally issue large amounts of I/O operations, but transfer a relatively small amount of data. In this case, setting an I/O governing throttle that is based on bandwidth might not achieve much. A throttle that is based on IOPS is better suited for this use case.

Conversely, a video streaming or editing application issues a small amount of I/O, but transfers large amounts of data. Therefore, it is better to use a bandwidth throttle for the volume in this case.

An I/O governing rate of 0 does *not* mean that zero IOPS or bandwidth can be achieved for this volume. It means that no throttle is set for this volume.

I/O throttles can be applied to volumes, pools, or hosts. To apply an throttle select the **resource to throttle** → **Actions** → **Edit Throttle**.

For further details see [Throttles Documentation](#).







# IBM Storage Insights and IBM Storage Insights Pro

Managing storage systems can be complex. You need to keep an eye on performance, capacity, and overall health to ensure 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:

- ▶ “IBM Storage Insights overview” on page 70
- ▶ “IBM Storage Insights monitoring” on page 71

## 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 standalone and 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, which 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, refer to this page in the IBM Documentation [here](#).

### 5.1.1 IBM Storage Insights: Information and registration

For more information about IBM Storage Insights, see the following resources:

- ▶ [IBM Storage Insights Fact Sheet](#).
- ▶ [IBM Storage Insights Security Guide, SC27-8774](#).
- ▶ [This IBM Documentation web page](#).
- ▶ [Product registration](#) (used to sign up and register for this no-charge service).

## 5.2 IBM Storage Insights monitoring

With 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 [What's New in 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 panel, the classic version and the new Carbon enhanced version.

Figure 5-1 on page 71 shows the classic version view of the IBM Storage Insights dashboard.

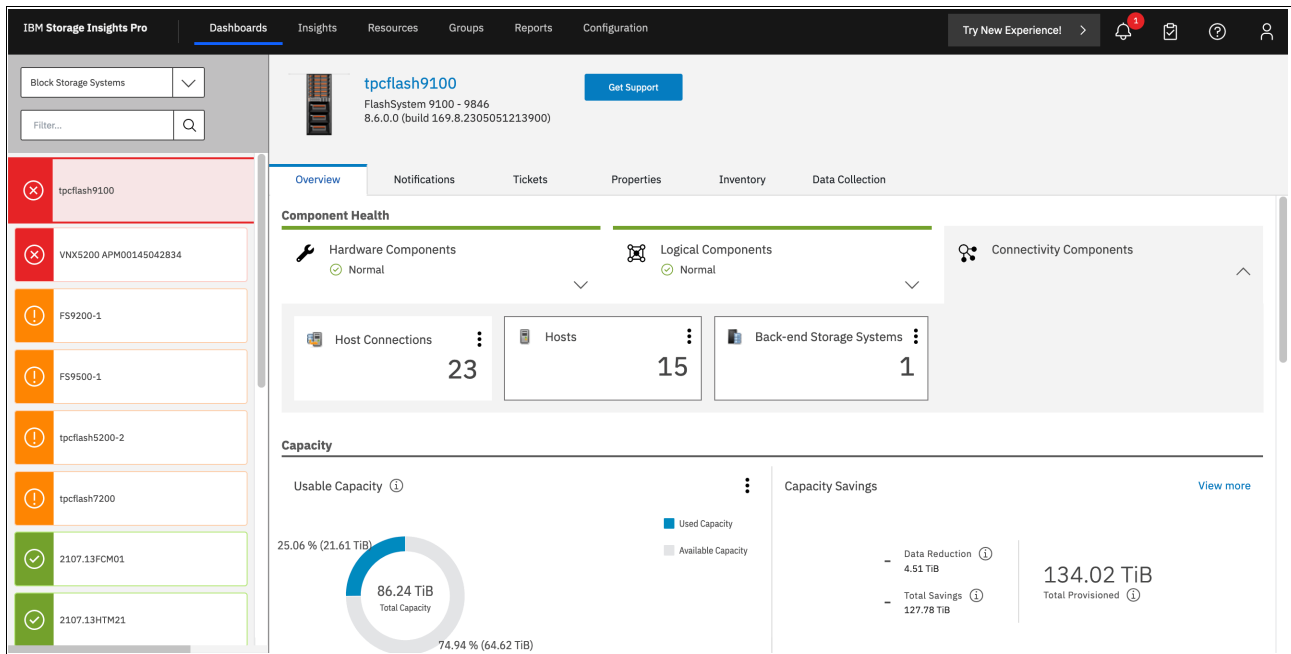


Figure 5-1 IBM Storage Insights System overview (classic view)

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

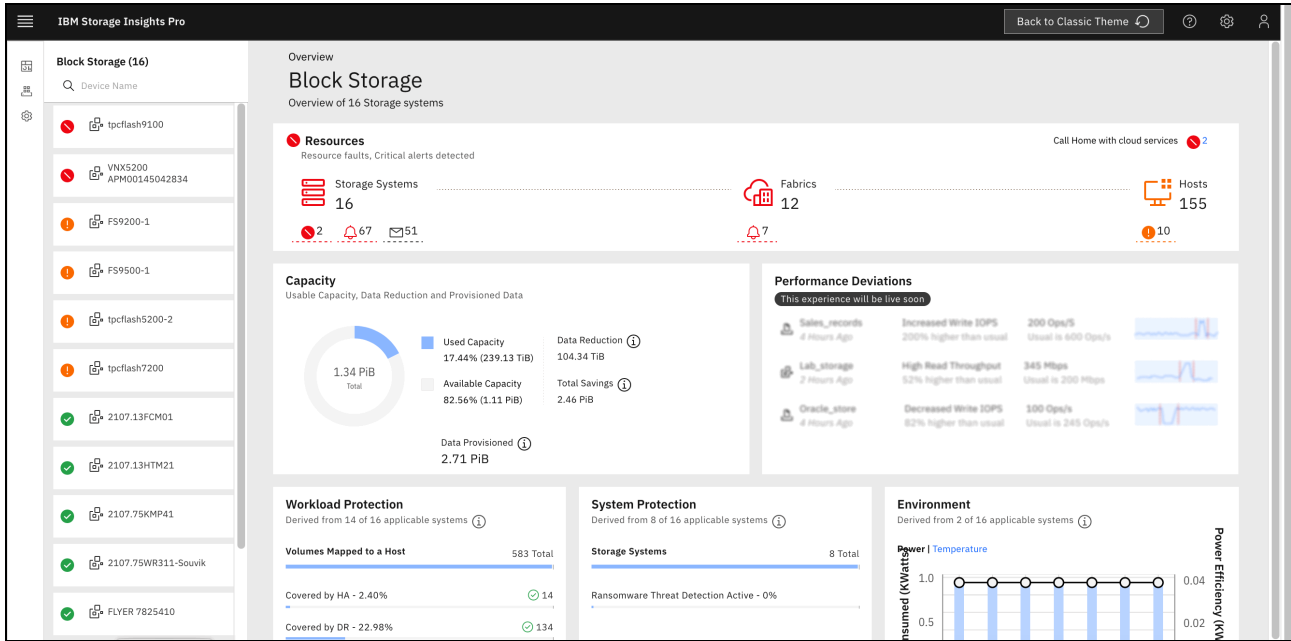


Figure 5-2 IBM Storage Insights System overview (Carbon enhanced view)

For the purposes of the next few examples we will 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, errors are shown here, as shown in Figure 5-3.

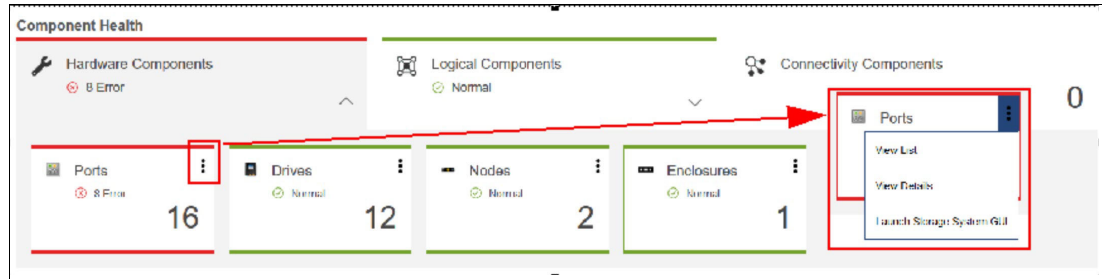


Figure 5-3 Component Health overview

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.

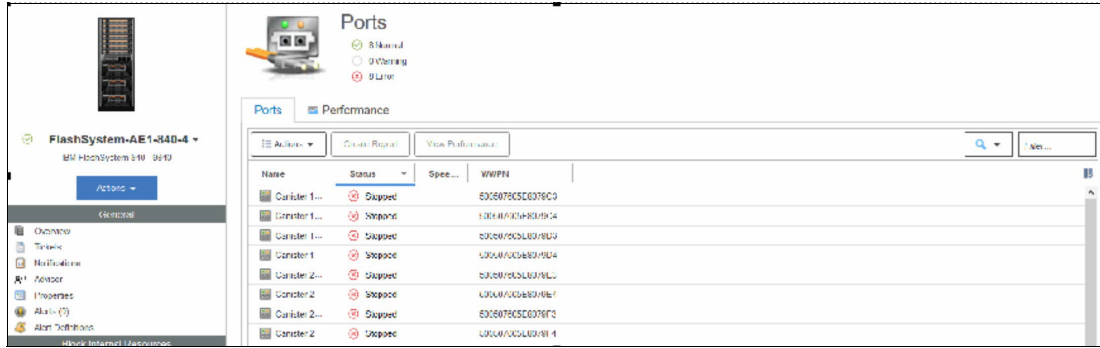


Figure 5-4 Ports in error

From here, it is now obvious which components have the problem and exactly what is wrong with them, so now you can log 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.

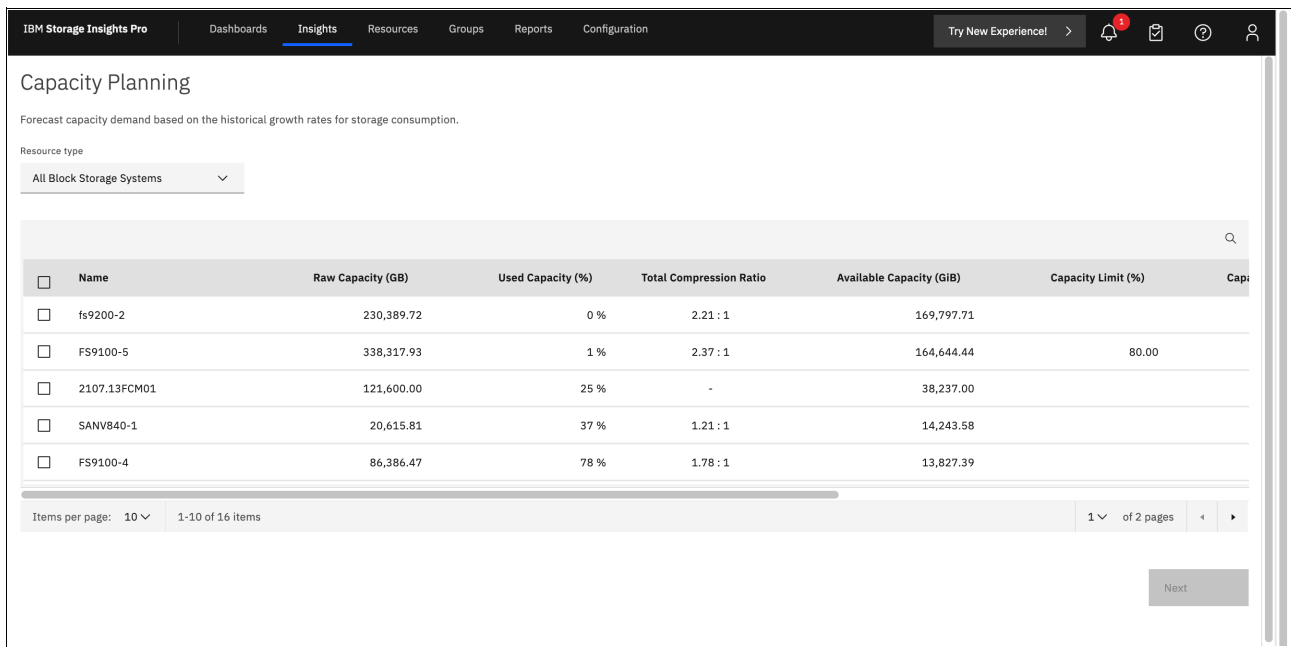


Figure 5-5 Capacity area of the IBM Storage Insights system overview

In the Capacity view, the user can click on 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. 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 73.

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

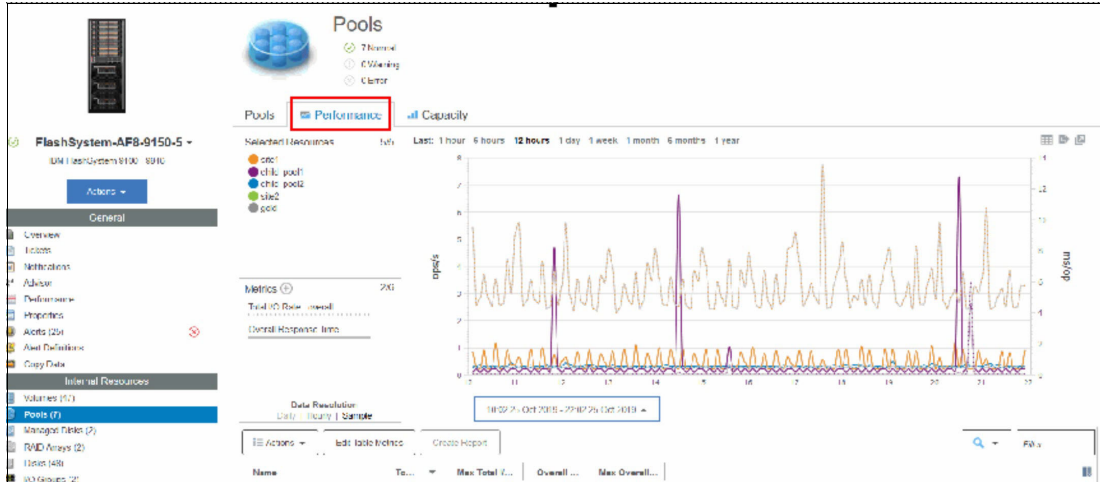


Figure 5-8 IBM Storage Insights: Performance view

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.

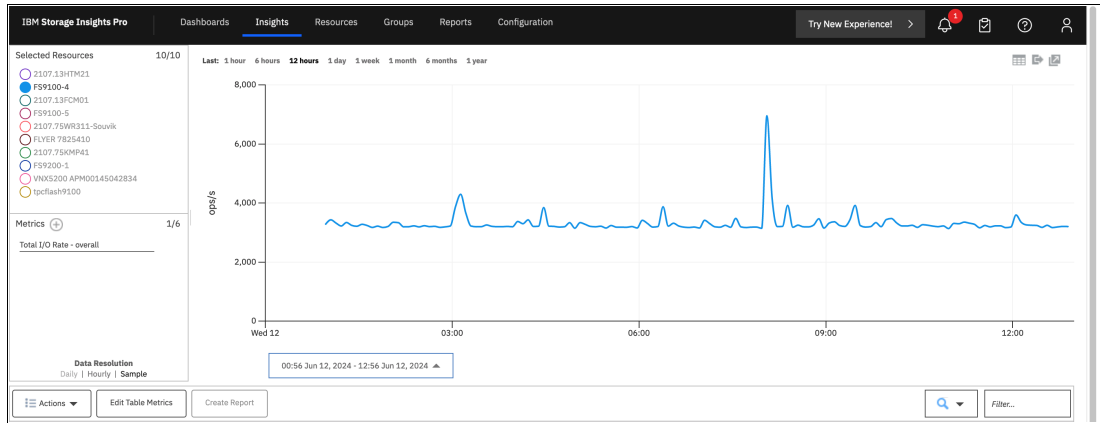


Figure 5-9 Filtered performance graph

Scrolling down the graph, the Performance List view is visible, as shown in Figure 5-10. 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.

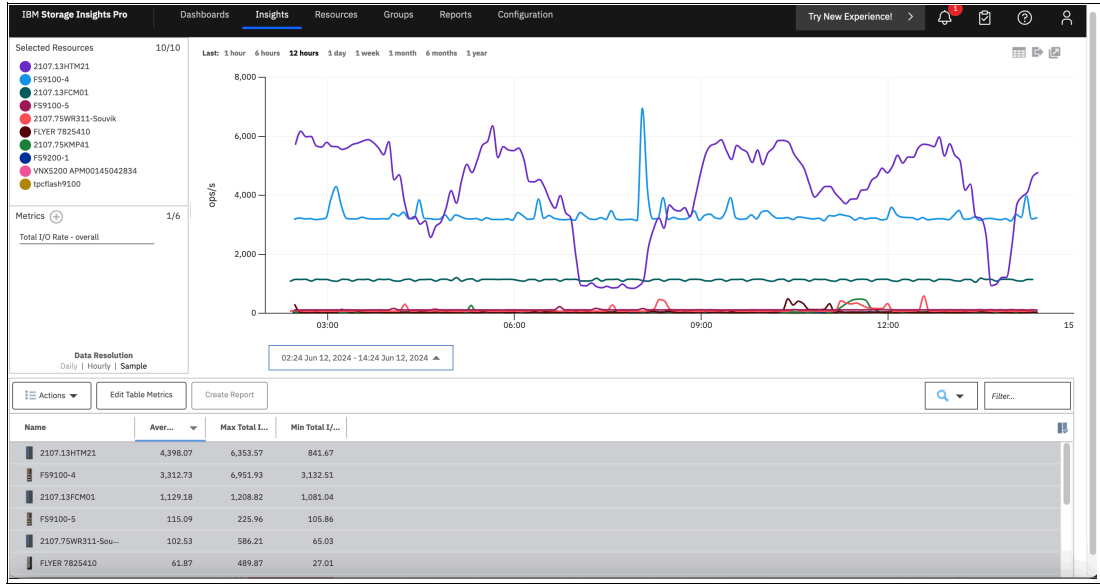


Figure 5-10 Performance List View

### 5.2.4 Logging support tickets by using IBM Storage Insights

With IBM Storage Insights, you can log existing support tickets that greatly 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.

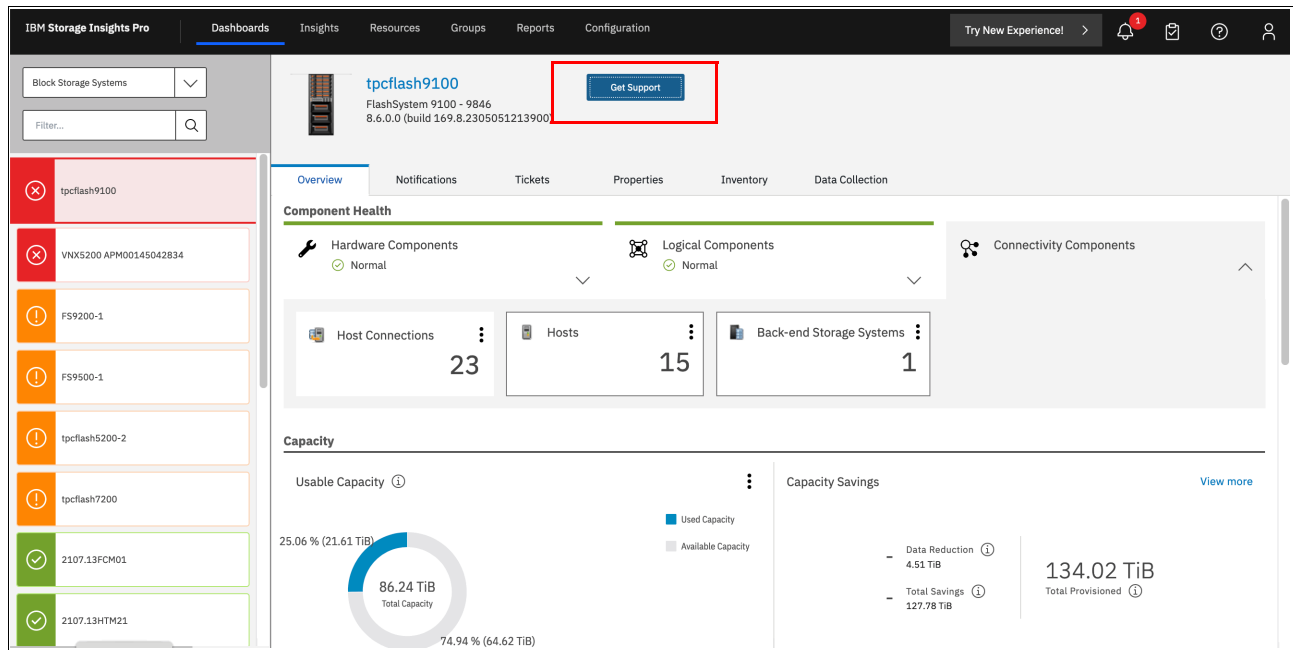


Figure 5-11 Get Support (see highlighted area)

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






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**.

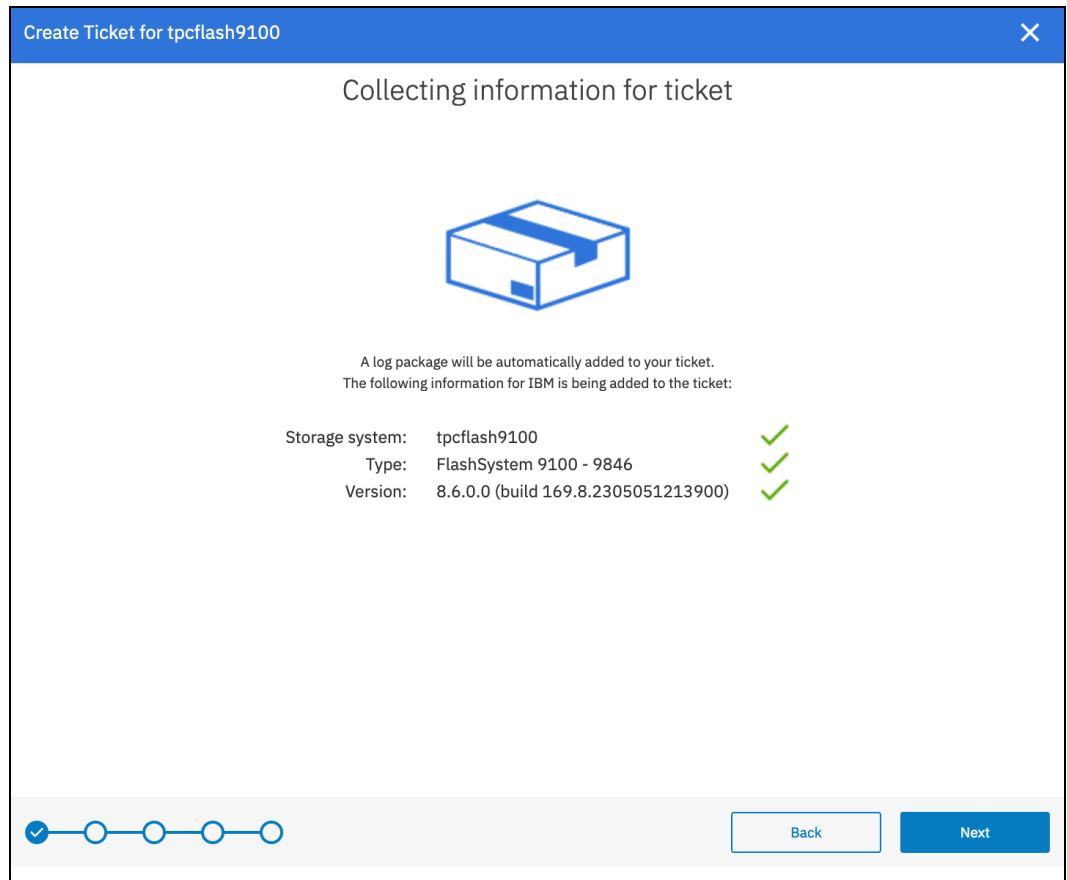


Figure 5-13 Create Ticket wizard

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. Once done, select **Next**.

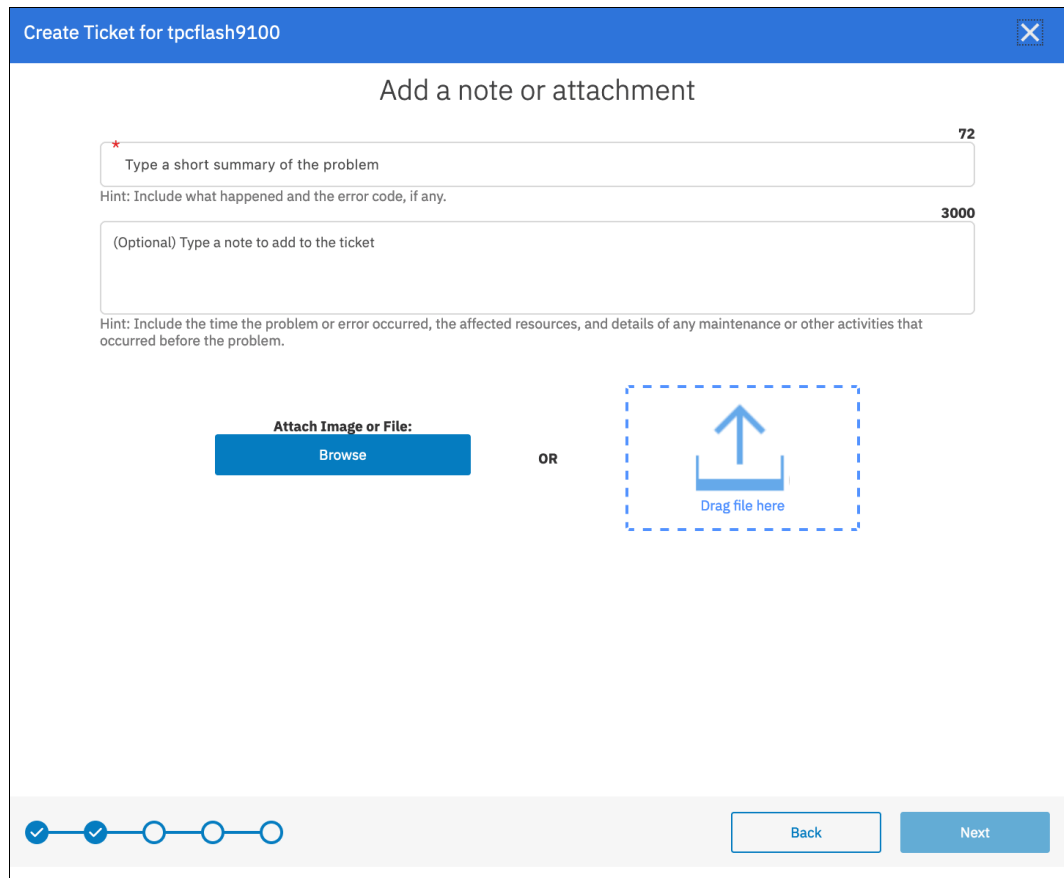


Figure 5-14 Add a note or attachment window

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

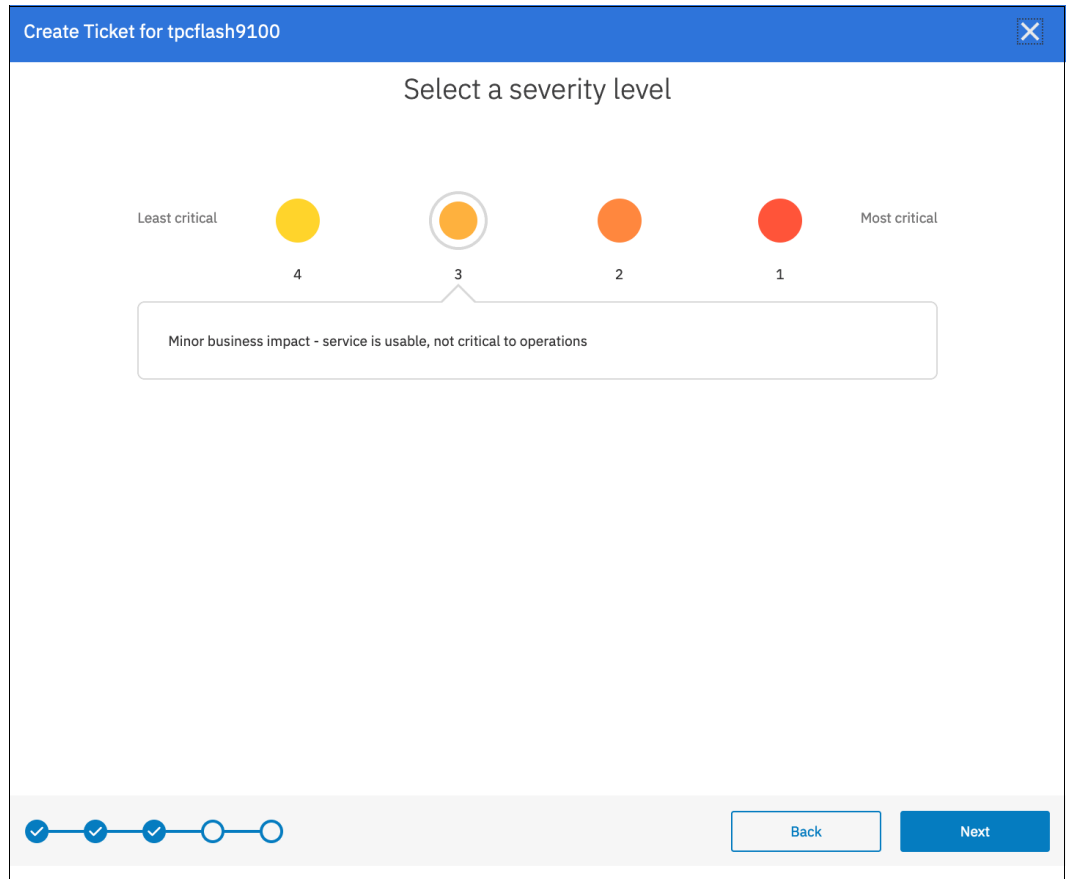


Figure 5-15 Selecting a Severity Level window

5. Choose whether this is a hardware or a software problem. Select the relevant option (for this example, the offline ports are likely caused by a physical layer hardware problem). Once done, click **Next**.

- Review the details of the ticket that will 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 should be attached to the ticket. For more information about the types of snap, see Figure 5-16. Click **Create Ticket**.

Create Ticket for tpcflash9100

### Review the 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

- Once done, select **Create Ticket**. A confirmation window opens, as shown in Figure 5-17 on page 82, and IBM Storage Insights automatically uploads the snap to the ticket when it is collected. Click **Close**.

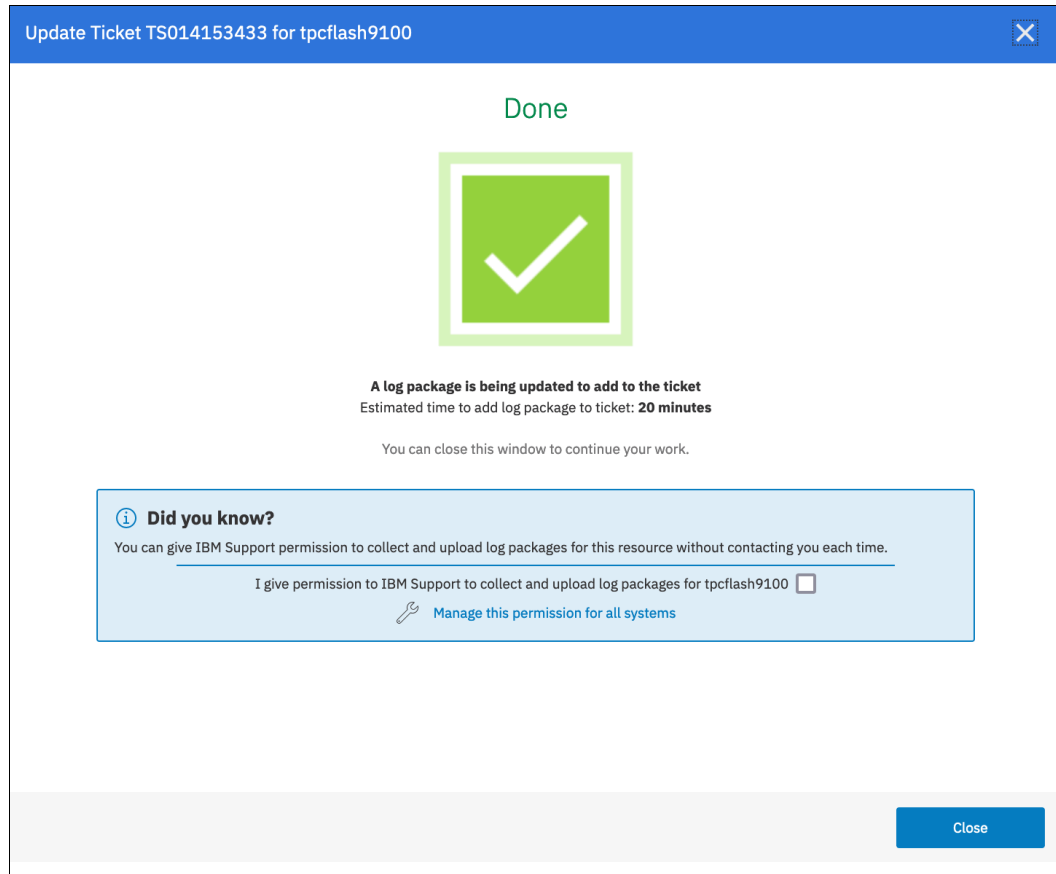


Figure 5-17 Update ticket

### 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.

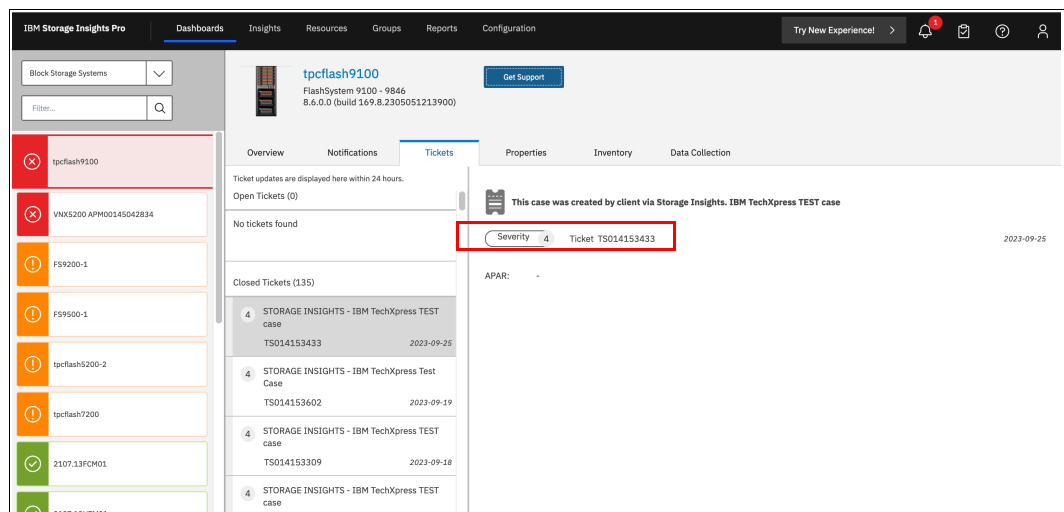


Figure 5-18 View tickets

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**.

2. 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 on page 83. After entering the support ticket number, you can select which type of log package you want and add a note to the ticket with the logs.

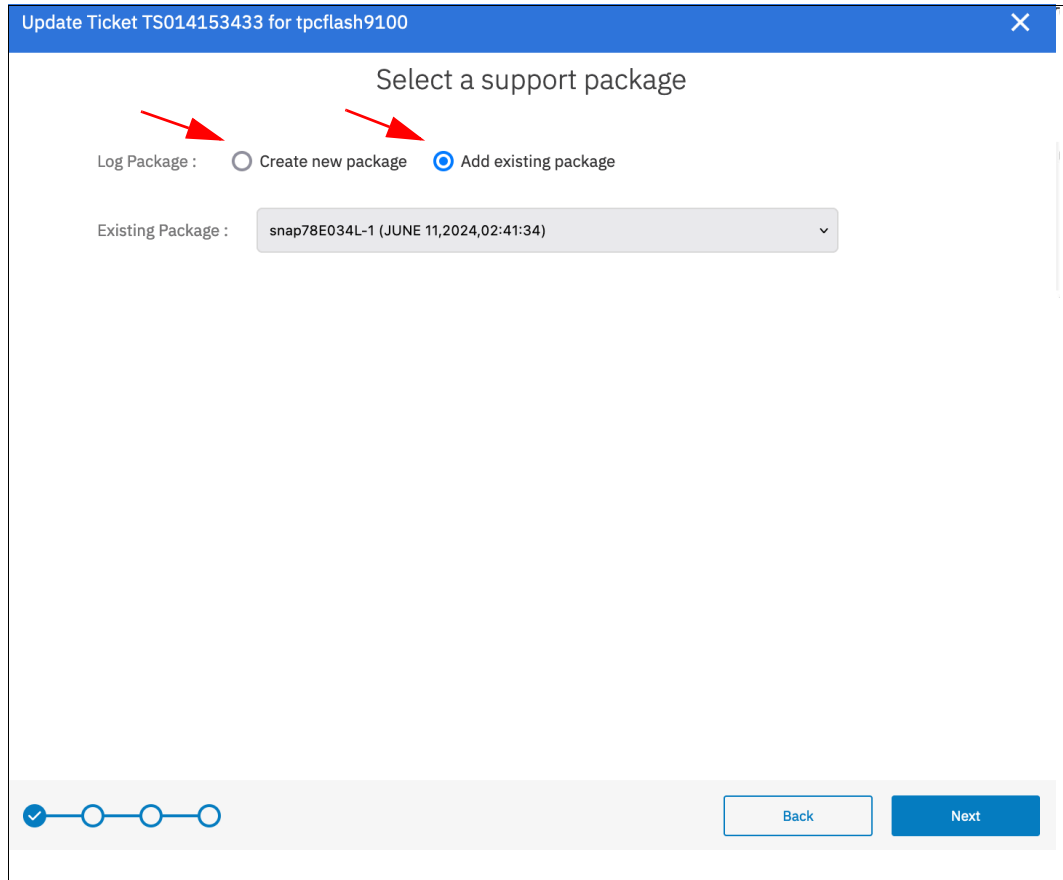


Figure 5-19 Adding a log package to the ticket

3. The review screen shown in Figure 5-20 allows you 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 on page 85. You can exit the wizard. IBM Storage Insights runs in the background to gather the logs and upload them to the ticket.



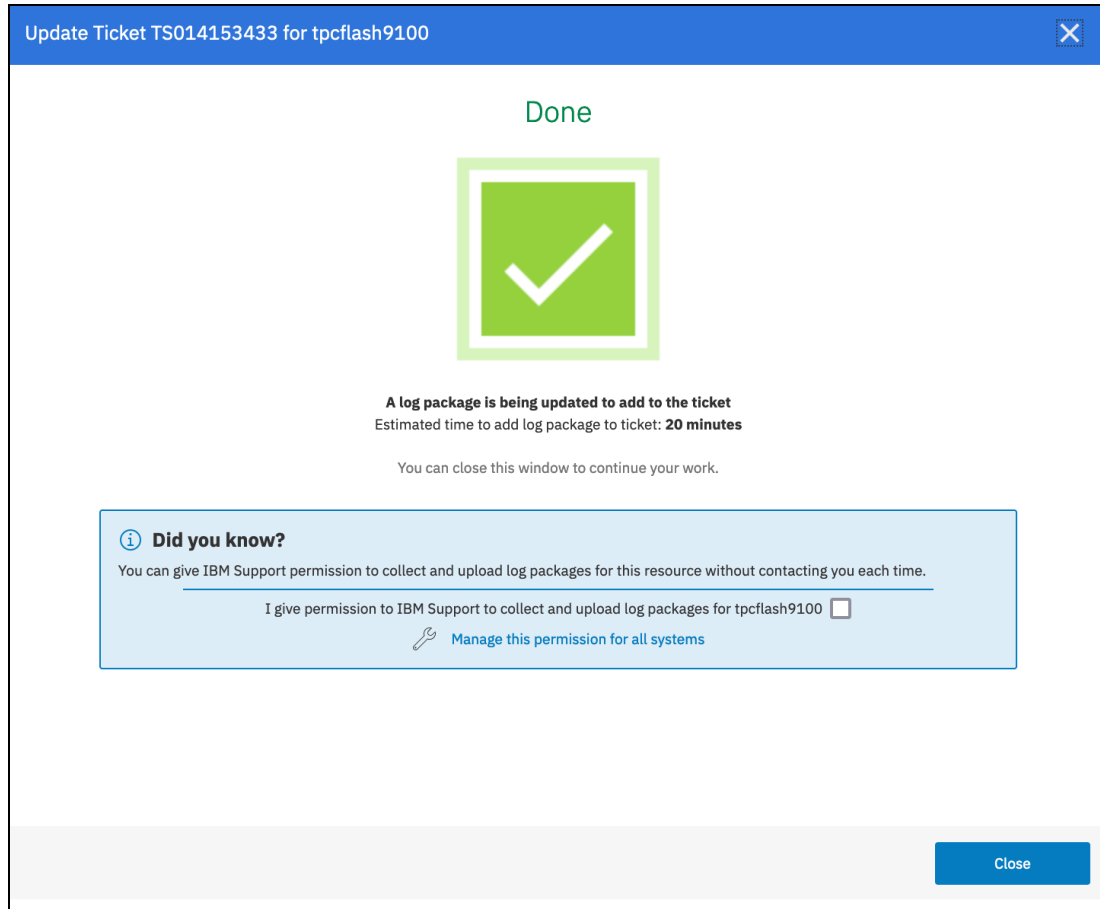


Figure 5-21 Log upload completed and processing

## 5.2.6 Enhancements to IBM Storage Insights Pro

In this section we discuss 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 860 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.

With IBM Storage Virtualize software 870 and FCM's with FCM firmware 4.1, the ransomware threat detection is further improved as follows:

- ▶ IBM FlashCore modules 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 using data from FCM.

- ▶ 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 on the IBM ransomware threat detection solutions including those mentioned in this book hsee this [link](#).

Also, refer to this [blog post](#) on how to mark volume snapshots those are created after ransomware threat detection as compromised.

### IBM Storage Virtualize 8.7.0 including Flash Grid

Flash Grid – scale out up to 8 FlashSystem or SVC systems to be managed as one, including non-disruptive workload mobility between members of the grid.

IBM Storage Insights Pro works in conjunction with the new Flash Grid and today provides an overview of your grid with grouping of your systems and the ability to non-disruptively move workloads (storage partitions) between systems in the grid. The end goal is to provide a seamless integration and interaction between the on-premise and cloud based management portals. Figure 5-22 on page 86 shows the integration on 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.

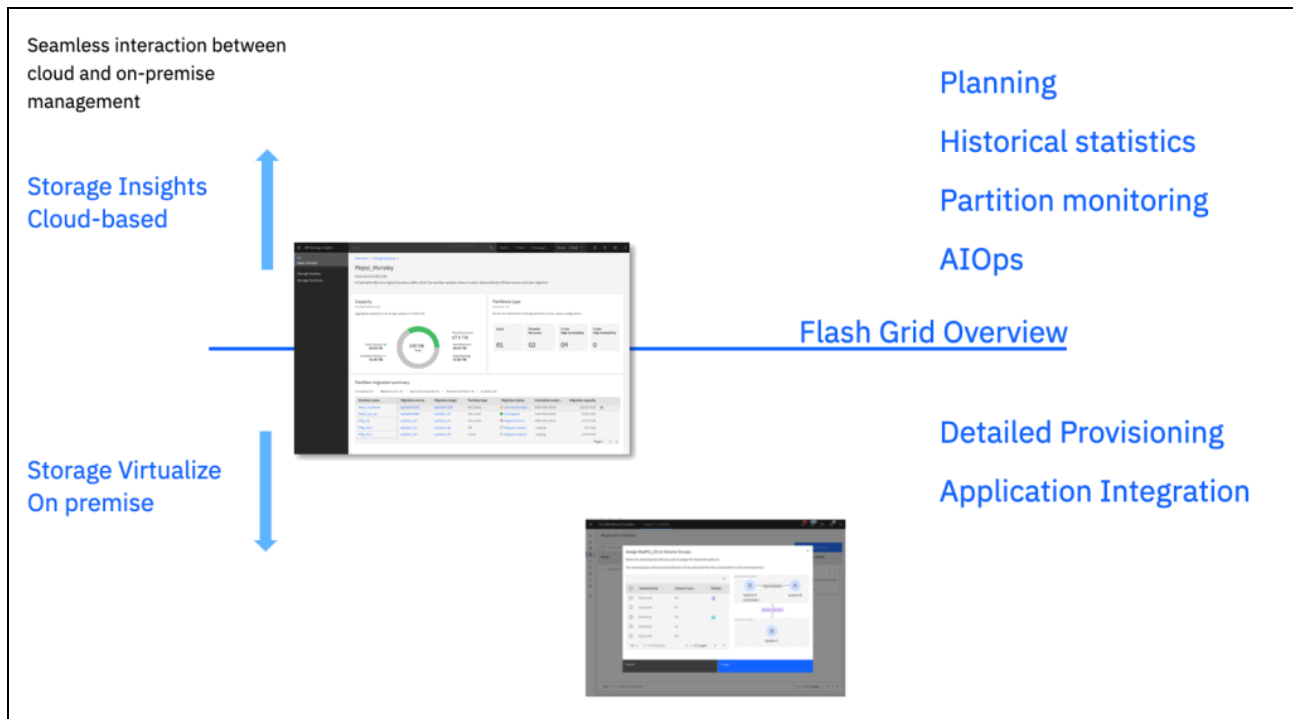


Figure 5-22 IBM Storage insights Pro and IBM Flash Grid integration



# 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 necessary problem determination data.

This chapter includes the following topics:

- ▶ “Troubleshooting” on page 88
- ▶ “Collecting diagnostic data” on page 95
- ▶ “Common problems and isolation techniques” on page 107

## 6.1 Troubleshooting

Troubleshooting should 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. An important step therefore is to make a proper problem description, which should be as accurately 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 recreated, 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 69, IBM Storage Insights is an important part of monitoring and ensuring 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 install and examples, refer to Chapter 5, “IBM Storage Insights and IBM Storage Insights Pro” on page 69.

### 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.

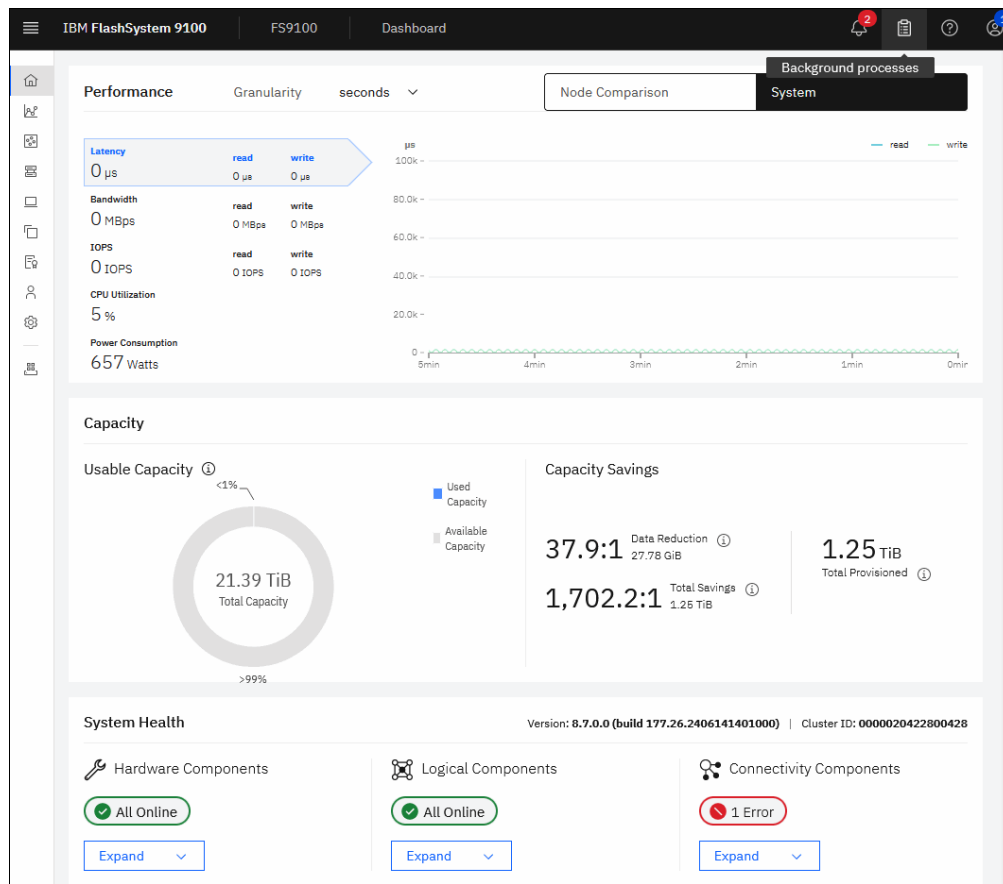


Figure 6-1 Events icon in the GUI

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 bottom 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 on **More Details** takes you to the GUI panel related to that specific component, or shows more information about it.

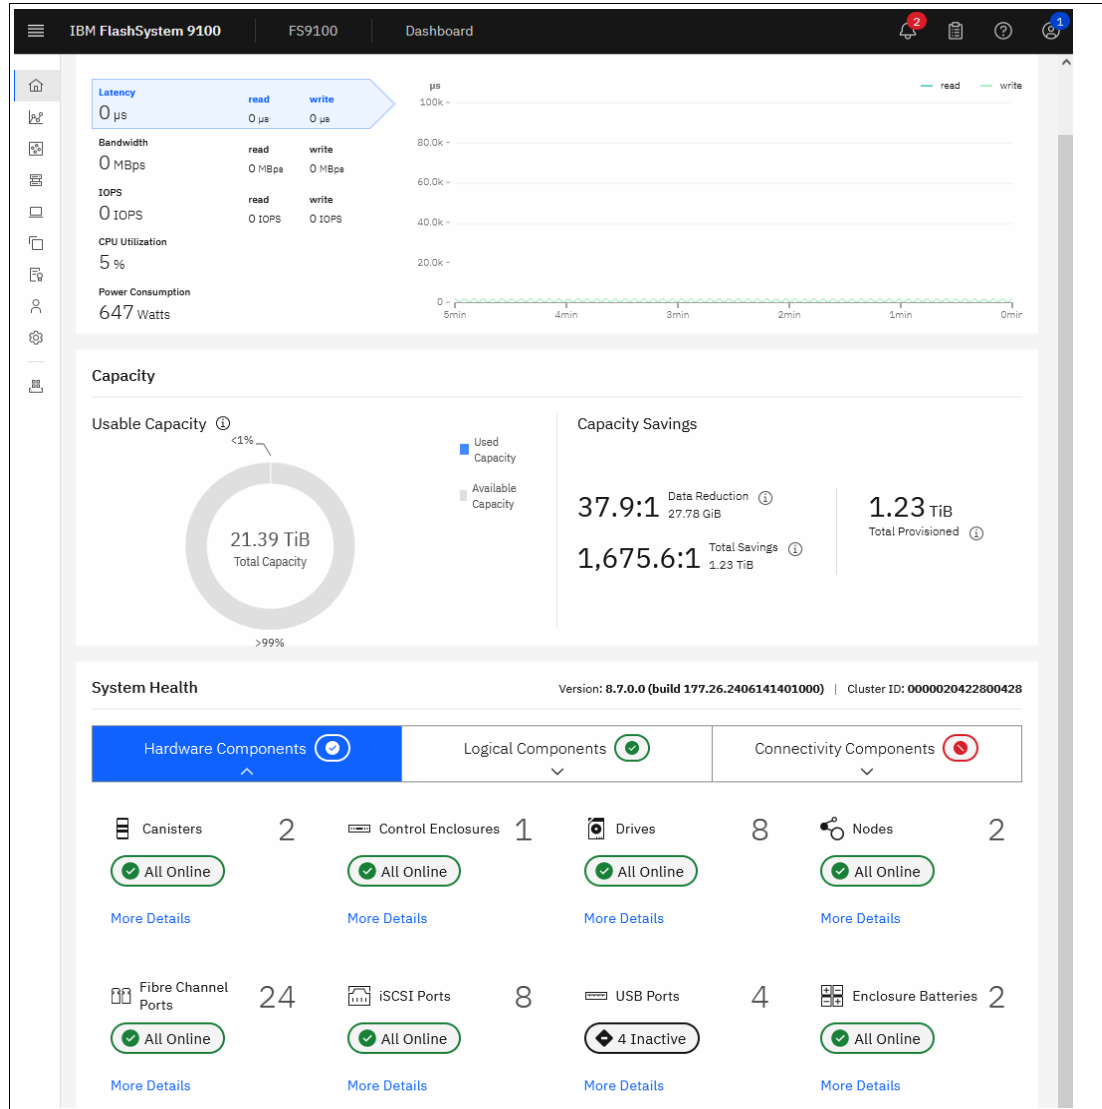
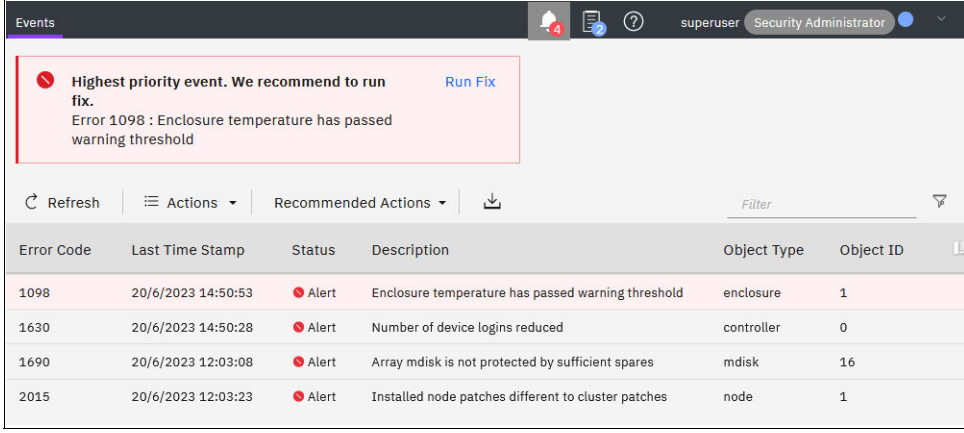


Figure 6-2 System Health expanded section in the dashboard

### 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. Whenever one or multiple unfixed errors need to 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, that is the event

log entry with the lowest four-digit error code, will be highlighted to be addressed first as shown in Figure 6-3.



Error Code	Last Time Stamp	Status	Description	Object Type	Object ID
1098	20/6/2023 14:50:53	Alert	Enclosure temperature has passed warning threshold	enclosure	1
1630	20/6/2023 14:50:28	Alert	Number of device logins reduced	controller	0
1690	20/6/2023 12:03:08	Alert	Array mdisk is not protected by sufficient spares	mdisk	16
2015	20/6/2023 12:03:23	Alert	Installed node patches different to cluster patches	node	1

Figure 6-3 Recommended actions

Click on **Run Fix** to launch the Fix Procedure for this particular event. Fix Procedures help resolving 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 required to return the system to its optimal state. This may include checking for dependencies, resetting internal error counters and apply updated to the system configuration. Whenever user interaction is required, you will be shown suggested actions to take and guided through the same. If the problem can be fixed, the related error in the event log eventually will be marked as fixed. Also, an associated alert in the GUI will be 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 - Node warmstart (node assert) will be 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, 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 ensures 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.

On 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 will be 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 may need a restart, and it is recommended to perform a sanity check on 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 being marked as fixed. It is recommended that they are not re-enabled until the recovery dumps and trace files from all nodes have been collected and were 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 and/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 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 ingesting 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. In combination with the information stored in the configuration backup `svc.config.backup.xml` the system's configuration and state will 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 may have been 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 will be 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 on page 93 shows **Monitoring** → **Events** window with Recommended Run Fix.

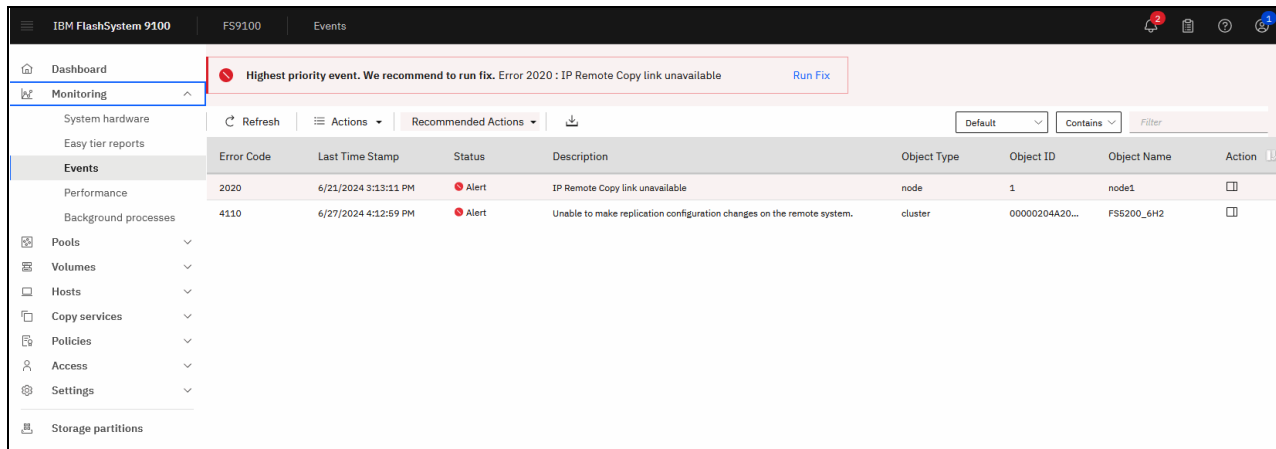


Figure 6-4 Monitoring → Events window

**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 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 pane, as shown in Figure 6-5. 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:

Properties	Sense Data:
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	
Secondary Object Type	
Copy ID	
Reporting Node ID	181
Reporting Node Name	node2
Root Sequence Number	
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 may be sometimes 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) has 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 as SFPs.

**Important:** Execute 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 ensures 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.

► **lscontroller** 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 if 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) may be a challenging task in complex and heterogeneous IT environments. It is crucial hence 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 to enable you to 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 most likely will be asked to provide a *support package* from the system. The support package may interchangeably be referred to as *Snap* or *data collection*.

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

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

To do so, check the My Cases section in the [IBM My Support](#) portal.

Storage Virtualize systems configured to be monitored in Chapter 5, “IBM Storage Insights and IBM Storage Insights Pro” on page 69, will show associated support cases there as well.

Alternatively, you may 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. For example, IBM Storage Insights.

#### What data to collect on IBM Storage Virtualize systems

The data that is needed for analysis depends on the kind 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 s 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:** Standard logs including performance stats.
  - fastest, smallest, no node dumps.
- ▶ **Snap/2:** Same as Snap/1 plus one existing statesave, the most recently created dump or livedump from the current config node.
  - slightly slower than snap/1, large.
- ▶ **Snap/3:** Same as Snap/1 plus the most recent dump or livedump from each active member node in the clustered system.
- ▶ **Snap/4:** Same as Snap/1 plus fresh livedump from each active member node in the clustered, which are 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.

### Which Snap to collect

Two major factors play a part in deciding which snap to collect for a specific support case:

- ▶ **Speed of collection**
  - A Snap/1 is generated more rapidly, and it is much smaller.
- ▶ **Amount of data**
  - Collecting a 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 quicker commence the 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, using Storage Insights is even more convenient, as 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 82.

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

- ▶ **Automatic upload via 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, Storage Virtualize allows you 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 (TSxxxxxxxx) and the type of support package to upload to the IBM Support Center. Press **Upload**. You can monitor the progress of the individual sub-tasks by clicking **View more details** as shown in Figure 6-6.




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. 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 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 completes, all collected files are packaged into a gzip-compressed tarball 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 may 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 command `lslivedump`. Once the status has changed from *dumping* to *inactive*, the livedump file is ready to be copied off the system 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 utilizes the NVMe protocol, a Peripheral Component Interconnect® Express (PCIe) interface, and high-speed NAND memory (e.g., Single-Level Cell) driven by a customizable Field-Programmable Gate Array (FPGA) to deliver high throughput, inline compression, and consistent IOPS (Input/Output Operations Per Second) 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 Type 1: Standard logs 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 multipathd 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 running this command, 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 whose operational status is disabled.
<code>lspath -s failed</code>	Lists all path whose operational status is 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>	Shows 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 -q1 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 -S1 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.

Besides managing the multipathing configuration by using the Windows GUI, it is 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 may 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

---

- 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 was omitted 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:

---

- 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. All other 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:

- Select an ESXi host, and click the **Configure** tab.
- Click **Storage Devices**.
- Select the storage device that you want to verify.
- 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:

- Select an ESXi host, and click the **Configuration** tab.
- Click **Storage**.
- Select a data store or mapped LUN.
- Click **Properties**.
- In the Properties dialog, select the extent, if necessary.
- 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.

## 6.3 Common problems and isolation techniques

Managing complex SAN environments with hundreds of disks, controllers, and switches can be challenging. Administrators need to effectively configure, monitor, and troubleshoot these systems.

Storage Virtualize simplifies this process with robust error logging and notification features. The system tracks internal events, identifies issues within the SAN, and helps isolate problems with attached hosts. These features empower administrators to quickly pinpoint issues and take corrective actions.

Furthermore, Storage Virtualize offers additional guidance:

- ▶ It suggests remedial actions in many cases.
- ▶ It verifies if a problem persists after troubleshooting steps.

Another valuable feature is the node's ability to track connected devices. This helps identify communication issues with hosts or back-end storage.

While Storage Virtualize provides excellent insight into its own health, external events like SAN zoning problems or host failures require troubleshooting outside the system.

For instance, a misconfiguration in SAN zoning could prevent the IBM Storage Virtualize cluster from functioning properly. This occurs because cluster nodes rely on FC SAN fabrics for communication.

To troubleshoot from the Storage Virtualize system's perspective, focus on these areas:

- ▶ The attached hosts. For more information, see 6.3.2, “Host problems” on page 108.
- ▶ The SAN. For more information, see 6.3.3, “Fibre Channel SAN and IP SAN problems” on page 113.

- ▶ The attached storage subsystem. For more information, see 6.3.5, “Storage subsystem problems” on page 116.
- ▶ The local FC port masking and portsets. For more information for Portmask see the options `-partnerfcportmask` `-localfcportmask` on IBM Documentation for the [lchsystem command](#).

### 6.3.1 Interoperability

When you experience events in an IBM Storage Virtualize environment, as an initial step, ensure that all components that comprise the storage infrastructure are interoperable, which applies to hosts, host OS, Host Bus Adapter (HBA), driver, firmware, SAN devices, and back-end devices. In a Storage Virtualize environment, the product support matrix is the main source for this information. For the latest IBM Storage Virtualize systems support matrix, see [IBM System Storage Interoperation Center \(SSIC\)](#).

It is crucial, to maintain up to date HBA firmware and device driver levels. This equally applies to multipath software and host OS patch levels. Failing to do so may lead to connectivity or interoperability issues, for instance host logins fail to reestablish after SAN maintenance activities. At worst case, this may lead to access loss.

### 6.3.2 Host problems

From host perspective, you can experience various situations that range from performance degradation to inaccessible disks. The first step in troubleshooting such issues is to check whether any potential interoperability issues exist.

After interoperability is verified, check the configuration of the host on the Storage Virtualize system’s side. The Hosts window in the GUI or the following CLI commands can be used to start a verification of potential host-related issues:

- ▶ `lshost`

**Note:** Depending on the connection type of the host (FC, FC direct attach, iSCSI, or NVMe, SAS) the output slightly differs in detail from each other.

This command displays the status of all configured host objects.

- Status *online*: the host ports are online in both nodes of an I/O group.
- Status *offline*: the host ports are offline in both nodes of an I/O group.
- Status *inactive*: the host has volumes that are mapped to it, but all its ports did not receive Small Computer System Interface (SCSI) commands in the last 5 minutes.
- Status *degraded*: there are no redundant logins to all nodes of an I/O group.

Example 6-8 shows the `lshost` command output.

Example 6-8 The `lshost` command

---

```
IBM_IBM FlashSystem:FLASHPFE95:superuser>lshost
0 Win2K8 2 4 degraded
1 ESX_67_B 2 4 online
2 ESX_67_A 2 1 offline
3 Server127 2 1 degraded
```

---





Based on this list, the host administrator must check and correct any issues found.

Hosts with a higher queue depth can potentially overload shared storage ports. Therefore, it is a best practice that you verify that the total of the queue depth of all hosts that are sharing a single target FC port is limited to 2048. If any of the hosts have a queue depth of more than 128, that depth must be reviewed because queue-full conditions can lead to I/O errors and extended error recoveries.

For more information about queue depths, see the following IBM Documentation web pages:

- ▶ [FC hosts](#)
- ▶ [iSCSI hosts](#)
- ▶ [iSER host](#)

Apart from hardware-related situations, problems can exist in such areas as the operating system or the software that is used on the host. These problems normally are handled by the host administrator or the service provider of the host system. However, the multipathing driver that is installed on the host and its features can help to determine possible issues.

For example, a volume path issue is reported, which means that a specific HBA on the server side cannot reach all the nodes in the I/O group to which the volumes are associated.

Faulty paths can be caused by hardware and software problems, such as the following examples:

- ▶ Hardware:
  - A faulty small form-factor pluggable transceiver (SFP) on the host or SAN switch.
  - Faulty fiber optic cables, for example damaged cables by exceeding the minimum permissible bend radius.
  - A faulty HBA.
  - Faulty SAN switch.
  - Contaminated SFP or cable connectors.
  - Patch panels
- ▶ Software or configuration:
  - Incorrect zoning, portset, or portmask.
  - Incorrect host-to-VDisk mapping.
  - Outdated HBA firmware or driver.
  - A back-level multipathing configuration or driver.

Based on field experience, it is a best practice that you complete the following hardware checks first:

- ▶ Whether connection error indicators are lit on the host, SAN switch or the Storage Virtualize system.
- ▶ Whether all the parts are seated correctly. For example, cables are securely plugged in to the SFPs and the SFPs are plugged all the way into the switch port sockets.
- ▶ Ensure that fiber optic cables are not damaged. If possible, swap a suspicious cable with a known good cables.

**Note:** When replacing or relocating fibre channel cables, always clean their pluggable connectors using proper cleaning tools. This even applies to brand new cables taken from sealed bags.

After the hardware check, continue to check the following aspects of the software setup:

- ▶ Whether the HBA driver level and firmware level are at the preferred and supported levels.
- ▶ Verify your SAN zoning configuration.
- ▶ The general SAN switch status and health for all switches in the fabric.
- ▶ The multipathing driver, and make sure that it is at the preferred configuration and supported level.
- ▶ Link layer errors reported by the host or the SAN switch may indicate so far undiscovered cable or SFP issues.

## **iSCSI or iSCSI Extensions for Remote Direct Memory Access configuration and performance issues**

This section describes the internet Small Computer Systems Interface (iSCSI) and iSCSI Extensions for Remote Direct Memory Access (RDMA) (iSER) configuration and performance issues.

For more information about the configuration, see the related [Configuration details for using RDMA-capable Ethernet ports](#) at the IBM Documentation web page.

### ***Link issues***

If the Ethernet port link does not come online, check whether the SFP or cables and the port support auto-negotiation with the switch. This issue is especially true for SFPs, which support 25 Gb and higher port speeds because a mismatch might exist in Forward Error Correction (FEC) that might prevent a port to auto-negotiate.

Another potential source of problems are 4X-splitter-cables and Direct Attach Copper (DAC) cables.

Longer cables are not only exposed to more noise or interference (high Bit Error Ratio (BER)); therefore, they require more powerful error correction codes.

Two IEEE 802.3 FEC specifications are important. For an auto-negotiation issue, verify whether a compatibility issue exists with SFPs at both end points:

- ▶ Clause 74: Fire Code (FC-FEC) or BASE-R (BR-FEC) (16.4 dB loss specification).
- ▶ Clause 91: Reed-Solomon (RS-FEC) (22.4 dB loss specification).

Use the `lshostiplogin` command to list the login session type, such as associated host object, login counts login protocol, and other details, for hosts that are identified by their iSCSI Qualified Name (IQN). The output is provided for ports, which logged in to Ethernet ports that are configured with IP addresses. The output shows, among other things, the protocol that is used.

The output in the protocol field indicates the connection protocol that is used by the configured IP host IQN to establish a login session that is referred by the login field. This value can be one of the following values:

- ▶ iSCSI
- ▶ iSER

### ***Priority flow control***

Priority flow control (PFC) is an Ethernet protocol that supports the ability to assign priorities to different types of traffic within the network. On most Data Center Bridging Capability Exchange (DCBX) protocol supported switches, verify whether Link Layer Discovery Protocol

(LLDP) is enabled. The presence of a virtual local area network (VLAN) is a prerequisite for the configuration of PFC. It is recommended to set the priority tag 0 - 7.

A DCBX-enabled switch and a storage adapter exchange parameters that describe traffic classes and PFC capabilities.

In Storage Virtualize systems, Ethernet traffic is divided into the following classes of service based on the feature use case:

- ▶ Host attachment (iSCSI or iSER)
- ▶ Back-end storage (iSCSI)
- ▶ Node-to-node communication (Remote Direct Memory Access (RDMA) clustering)

If challenges occur as the PFC is configured, verify the following attributes to determine the issue:

- ▶ Configure the IP address or VLAN by using `mkip`.
- ▶ Configure the class of service (COS) by using `chsystemethernet`.
- ▶ Ensure that the priority tag is enabled on the switch.
- ▶ Ensure that the `!sportip` output is as follows:
 

```
dcbx_state, pfc_enabled_tags
```
- ▶ The Enhanced Transmission Selection (ETS) setting is recommended if a port is shared.

For more information about problem solving, see [Resolving a problem with PFC settings](#).

### ***Standard network connectivity check***

Verify that the required TCP/UDP ports are allowed in the network firewall. The following ports can be used for various host attachments:

- ▶ Software iSCSI requires TCP port 3260.
- ▶ iSER or RDMA over Converged Ethernet (RoCE) host requires TCP port 3260.
- ▶ iSER or iWRAP host requires TCP port 860.

A comprehensive list of TCP/IP address and port requirements is available at IBM Documentation [IP address allocation and usage](#).

Verify that the IP addresses are reachable and the TCP ports are open.

**Note:** iSER host attachment is not supported on IBM FlashSystem 9500, IBM FlashSystem 7300, and IBM SAN Volume Controller SV3. It is, however, supported on other Storage Virtualize products.

### ***iSCSI performance analysis issues and tuning***

Follow the detailed steps for [iSCSI performance analysis and tuning](#).

**Tip:** If the host platform does not provide a mechanism to disable TCPDelayedAck, verify whether a smaller “Max I/O Transfer Size” with more concurrency (queue depth > 16) improves overall latency and bandwidth usage for the specific host workload. In most Linux distributions, this Max I/O Transfer Size is controlled by the `max_sectors_kb` parameter with a suggested transfer size of 32 KiB.

In addition, review network switch diagnostic data to evaluate potential issues as packet drop or packet retransmission. It is advisable to enable flow control or PFC to enhance the

reliability of the network delivery system to avoid packet loss, aiming to enhance the overall performance.

### 6.3.3 Fibre Channel SAN and IP SAN problems

Adding Storage Virtualize to your SAN is straightforward, but requires following basic rules to avoid access or performance issues.

Key SAN zoning configurations include:

- ▶ Host zones
- ▶ Storage zones for back-end storage
- ▶ A dedicated Storage Virtualize zone for cluster communication

Consider the following for maintaining performance:

- ▶ Dedicate FC ports for specific purposes: intra-cluster communication, replication, and host/storage traffic.
- ▶ Suboptimal FC port masking (`local_fc_port_mask` & `partner_fc_port_mask`) can lead to performance issues. Ensure proper zoning and masking to allow necessary traffic.

**Note:** Inter-node zoning restricts logins between FC ports, while `local_fc_port_mask` additionally controls how those logins are used. Further details on this can be found in [Fibre Channel port masking](#).

Some situations can cause issues in the SAN fabric and SAN switches. Problems can be related to a hardware fault or to a software problem on the switch. The following hardware defects are normally the easiest problems to find:

- ▶ Switch power, fan, or cooling units
- ▶ Installed SFP modules
- ▶ Fiber optic cables

Software failures often require data collection and IBM Support involvement.

Before escalating, check switch firmware and software release notes for known issues and updates. SAN connectivity issues frequently stem from zoning errors. Examples include incorrect WWPNs in zones or omitted host ports.

SAN zoning therefore should be done after thorough planning, using a unified naming schema for aliases, zones et cetera. It is equally beneficial for both the user and any supporting function, if the SAN switches follow a clear naming structure as well as a coordinated domain id assignment. While it is not a problem from a technical point of view to reuse switch domain ids across SAN fabrics, it unnecessarily complicates troubleshooting, as the switch nPort IDs (nPID) will show up multiple times in diagnostic data and the output of CLI commands as for instance `lspportfc`, `lstargetportfc` and `lsfabric`.

Existing SAN environment often have developed organically over time with no or little documentation. Document your SAN environment. Clear documentation and graphical layouts can significantly speed up problem resolution.

On Storage Virtualize systems, a part the worldwide port names (WWPN) is derived from the worldwide node name (WWNN) of the node canister in which the adapter is installed. The WWNN is part of each node's Vital Product Data (VPD), it impacts the WWPN's last four digits. The ports' WWPN also are derived from the PCIe slot the adapter is installed in and its port id. For more information, see [Worldwide node and port names](#).

So, the WWPNs for the different ports of the same node differ in the 6th and 5th last digit. For example:

```
50:05:07:68:10:13:37:dc
50:05:07:68:10:14:37:dc
50:05:07:68:10:24:37:dc
```

The WWPNs for ports on different nodes differ in the last 4 digits. For example, here are the WWPNs for port 3 and 4 on each node of a IBM FlashSystem:

```
50:05:07:68:10:13:37:dc
50:05:07:68:10:14:37:dc

50:05:07:68:10:13:37:e5
50:05:07:68:10:14:37:e5
```

As shown in Example 6-11, two ports that belong to the same Storage Virtualize node are zoned to a host FC port. Therefore, the result is that the host port will not log in to both nodes of that I/O group and the multipathing driver will not see redundant paths:

*Example 6-11 Incorrect WWPN zoning*

---

```
zone: z_Win2k19_FS9110_iogrp0
      50:05:07:68:10:13:37:dc
      50:05:07:68:10:14:37:dc
      20:00:00:e0:8b:89:cc:c2
```

---

The correct zoning must look like the zoning that is shown in Example 6-12.

*Example 6-12 Correct WWPN zoning*

---

```
zone: z_Win2k19_FS9110_iogrp0
      50:05:07:68:10:14:37:e5
      50:05:07:68:10:14:37:dc
      20:00:00:e0:8b:89:cc:c2
```

---

The following IBM FlashSystem error codes are related to the SAN environment:

- ▶ Error 1060: Fibre Channel ports are not operational.
- ▶ Error 1220: A remote port is excluded.

A bottleneck is another common issue that is related to SAN switches. The bottleneck can be present in a port where a host, storage subsystem, or Storage Virtualize device is connected, or in Inter-Switch Link (ISL) ports. The bottleneck can occur in some cases, such as when a device that is connected to the fabric is slow to process received frames, or if a SAN switch port cannot transmit frames at a rate that is required by a device that is connected to the fabric.

These cases can slow down communication between devices in your SAN. To resolve this type of issue, see the SAN switch documentation to investigate and identify what is causing the bottleneck and how fix it.

If you cannot fix the issue with these actions, use the method that is described in 6.2, “Collecting diagnostic data” on page 95, collect the SAN switch debugging data, and then contact the vendor for assistance or open a case with the vendor.

### 6.3.4 Port issues and transceiver statistics

Ports and their connections are involved in many support cases. The `lspportstats` command supports the administrator in troubleshooting ports of any kind on the IBM Storage Virtualize systems side. The output of the command contains many different details like the port type, WWPN, IQN, send and receive statistics, and SFP details. For example, you can check the physical error counter for a port and other interesting values of an SFP.

Example 6-13 shows the output for an FC port.

*Example 6-13 lspportstats command output*

---

```
IBM_FlashSystem:FS9110:superuser>lspportstats -node node1
Nn_stats_78E003K-1_230623_151121
<port id="1"
type="FC"
type_id="1"
wwpn="0x5005076810110214"
fc_wwpn="0x5005076810110214"
fcoe_wwpn=""
sas_wwn=""
iqn=""
hbt="80487" hbr="0" het="0" her="1465"
cbt="0" cbr="612" cet="260" cer="0"
lnbt="0" lnbr="0" lnet="955316" lner="955332"
rmbt="0" rmbr="0" rmet="0" rmer="0"
dtdt="242" dtdc="6" dtdm="956797"
dtdt2="242" dtdc2="6"
lf="14" lsy="21" lsi="0" pspe="0"
itw="54" icrc="0" bbcz="0"
tmp="46" tmpht="85"
txpwr="596" txpwr1t="126"
rxpwr="570" rxpwr1t="31"
hsr="0" hsw="0" har="0" haw="0"
/>
<port id="2"
type="FC"
type_id="2"

z
type_id="2"
[...]
```

---

Table 6-5 shows some of the most interesting attributes and their meanings.

*Table 6-5 Selected attributes of the lspportstats output*

Attribute	Information
Nn_stats_78F13MY-1_220413_152655	Data source stats file of the output.
lsy	Indicates the loss of sync error count.
itw	Invalid transmissionword error count.
icrc	Indicates the invalid cyclic redundancy check (CRC) error count.

Attribute	Information
txpwr	SFP TX power in microwatts.
rxpwr	SFP RX power in microwatts.

It is not possible to reset or clear the shown counter with a command at the moment. To examine the current trend of the values or whether they are increasing, a best practice is to compare two outputs of the command for differences. Allow some run time between the two iterations of the command.

For more information about the `lspportstats` command, see the [IBM Documentation for the lspportstats command](#).

### 6.3.5 Storage subsystem problems

Today, various heterogeneous storage subsystems are available. All these subsystems have different management tools, different setup strategies, and possible problem areas depending on the manufacturer. To support a stable environment, all subsystems must be correctly configured by following best practices and have no existing issues.

If you experience a storage-subsystem-related issue, check the following areas:

- ▶ Always check the [SSIC](#) to see whether the subsystem is supported.
- ▶ Storage subsystem configuration: Ensure that a valid configuration and best practices are applied to the subsystem.
- ▶ Storage subsystem controllers: Check the health and configurable settings on the controllers.
- ▶ Storage subsystem array: Check the state of the hardware, such as an FCM, solid-state drive (SSD), or disk drive module (DDM) failure or enclosure alerts.
- ▶ Storage volumes: Ensure that the LUN masking is correct.
- ▶ Host attachment ports: Check the status, configuration, and connectivity to storage SAN switches.
- ▶ Layout and size of redundant array of independent disks (RAID) arrays and LUNs: Performance and redundancy are contributing factors.

IBM Storage Virtualize has several CLI commands that you can use to check the status of the system and attached storage subsystems. Before you start a complete data collection or problem isolation on the SAN or subsystem level, first use the following commands and check the status from the IBM Storage Virtualize perspective:

- ▶ **lsccontroller <controller\_id\_or\_name>**

Checks that multiple WWPNs that match the back-end storage subsystem controller ports are available.

Checks that the `path_counts` are evenly distributed across each storage subsystem controller, or that they are distributed correctly based on the preferred controller. The total of all `path_counts` must add up to the number of MDisks multiplied by the number of IBM Storage Virtualize nodes.

- ▶ **lsmdisk**

Checks that all MDisks are online (not degraded or offline).



► **lsmdisk <MDisk\_id\_or\_name>**

Checks several of the MDisks from each storage subsystem controller. Are they online? Do they all have path\_count = number of back-end ports in the zone to IBM Storage Virtualize x number of nodes? An example of the output from this command is shown in Example 6-14. MDisk 0 is a local MDisk in an IBM FlashSystem, and MDisk 1 is provided by an external, virtualized storage subsystem.

*Example 6-14 Issuing a lsmdisk command*

---

```
IBM_IBM FlashSystem:FLASHPFE95:superuser>lsmdisk 0
id 0
name MDisk0
status online
mode array
MDisk_grp_id 0
MDisk_grp_name Pool0
capacity 198.2TB
quorum_index
block_size
controller_name
ctrl_type
ctrl_WWNN
controller_id
path_count
max_path_count
ctrl_LUN_#
UID
preferred_WWPN
active_WWPN
fast_write_state empty
raid_status online
raid_level raid6
redundancy 2
strip_size 256
spare_goal
spare_protection_min
balanced exact
tier tier0_flash
slow_write_priority latency
fabric_type
site_id
site_name
easy_tier_load
encrypt no
distributed yes
drive_class_id 0
drive_count 8
stripe_width 7
rebuild_areas_total 1
rebuild_areas_available 1
rebuild_areas_goal 1
dedupe no
preferred_iscsi_port_id
active_iscsi_port_id
replacement_date
over_provisioned yes
```

```

supports_unmap yes
provisioning_group_id 0
physical_capacity 85.87TB
physical_free_capacity 78.72TB
write_protected no
allocated_capacity 155.06TB
effective_used_capacity 16.58TB.

```

```

IBM_IBM FlashSystem:FLASHPFE95:superuser>lsmdisk 1
id 1
name flash9h01_itsosvcc11_0
status online
mode managed
MDisk_grp_id 1
MDisk_grp_name Pool1
capacity 51.6TB
quorum_index
block_size 512
controller_name itsoflash9h01
ctrl_type 6
ctrl_WWNN 500507605E852080
controller_id 1
path_count 16
max_path_count 16
ctrl_LUN_# 0000000000000000
UID 6005076441b53004400000000000000100000000000000000000000000000000
preferred_WWPN
active_WWPN many

```

NOTE: lines removed for brevity

---

Example 6-14 on page 117 shows that for MDisk 1 that the external storage controller has eight ports that are zoned to IBM Storage Virtualize systems, which has two nodes (8 x 2 = 16).

► **lsvdisk**

Checks that all volumes are online (not degraded or offline). If the volumes are degraded, are there stopped FlashCopy jobs present? Restart stopped FlashCopy jobs or seek IBM Storage Virtualize systems support guidance.

► **lsfabric**

Use this command with the various options, such as **-controller controllerid**. Also, check different parts of the IBM Storage Virtualize systems configuration to ensure that multiple paths are available from each IBM Storage Virtualize node port to an attached host or controller. Confirm that IBM Storage Virtualize systems node port WWPNs are also consistently connected to an external back-end storage.

### Determining the number of paths to an external storage subsystem

By using CLI commands, the total number of paths to an external storage subsystem can be determined. To determine the value of the available paths, use the following formulas:

```

Number of MDisks x Number of nodes per Cluster = Number of paths
MDisk_link_count x Number of nodes per Cluster = Sum of path_count

```

Example 6-15 shows how to obtain this information by using the **lscontroller <controllerid>** and **svcinfo lsnode** commands.

*Example 6-15 Output of the svcinfo lscontroller command*


---

```
IBM_IBM FlashSystem:FLASHPFE95:superuser>lscontroller 1
```

```
id 1
controller_name itsof9h01
WWNN 500507605E852080
MDisk_link_count 16
max_MDisk_link_count 16
degraded no
vendor_id IBM
product_id_low FlashSys
product_id_high tem-9840
product_revision 1430
ctrl_s/n 01106d4c0110-0000-0
allow_quorum yes
fabric_type fc
site_id
site_name
WWPN 500507605E8520B1
path_count 32
max_path_count 32
WWPN 500507605E8520A1
path_count 32
max_path_count 64
WWPN 500507605E852081
path_count 32
max_path_count 64
WWPN 500507605E852091
path_count 32
max_path_count 64
WWPN 500507605E8520B2
path_count 32
max_path_count 64
WWPN 500507605E8520A2
path_count 32
max_path_count 64
WWPN 500507605E852082
path_count 32
max_path_count 64
WWPN 500507605E852092
path_count 32
max_path_count 64
```

```
IBM_IBM FlashSystem:FLASHPFE95:superuser>svcinfo 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 500507681000000A online 0 io_grp0 no
AF8 iqn.1986-03.com.ibm:2145.flashpfe95.node1 01-2 1
2 F313150
```



Make sure that the zones are correctly configured, and that the zone set is activated. The zones that allow communication between the storage subsystem and the Storage Virtualize systems device must contain the WWPNs of the storage subsystem and WWPNs of the IBM Storage Virtualize system.

4. Collect all support data and contact IBM Support.

Collect the support data for the involved SAN, Storage Virtualize system, and external storage systems, as described in 6.2, “Collecting diagnostic data” on page 95.

### 6.3.6 IP replication problems

Two systems can be linked over native IP links that are connected directly or by Ethernet switches to perform remote copy functions. Remote copy over native IP provides a less expensive alternative to using FC configurations.

**Note:** IP replication that is configured over 25 Gbps ports does not use RDMA capabilities, and it does not provide a performance improvement compared to 10 Gbps ports. 100 Gbps ports do not support IP replication.

A system can be part of only two IP partnerships. IBM Storage Virtualize systems with pre-8.4.2.0 firmware are still limited to one IP partnership. Partnerships on low memory platform nodes share memory resources, which can lead to degraded performance.

Portsets replace the requirement for creating remote-copy groups for IP partnerships. Dedicated portsets can be created for remote copy traffic. The dedicated portsets provide a group of IP addresses for IP partnerships.

During updates of the software, any IP addresses that are assigned to remote-copy groups with an IP partnership are automatically moved to a corresponding portset. For example, if remote-copy group 1 is defined on the system before the update, IP addresses from that remote-copy group are mapped to portset 1 after the update. Similarly, IP addresses in remote-copy group 2 are mapped to portset 2.

The native IP replication feature uses the following TCP/IP ports for remote cluster path discovery and data transfer, therefore, these ports need to be open:

- ▶ IP partnership management IP communication: TCP port 3260.
- ▶ IP partnership data path connections: TCP port 3265.

If a connectivity issue exists between the cluster in the management communication path, the cluster reports error code 2021: Partner cluster IP address unreachable. However, when a connectivity issue exists in the data path, the cluster reports error code 2020: IP Remote Copy link unavailable.

For more information, see [Validating Ethernet partnerships configuration](#).

If the IP addresses are reachable and TCP ports are open, verify whether the end-to-end network supports a maximum transmission unit (MTU) of 1500 bytes without packet fragmentation. When an external host-based ping utility is used to validate end-to-end MTU support, use the “do not fragment” qualifier.

Fix the network path so that traffic can flow correctly. After the connection is made, the error auto-corrects.

The network quality of service largely influences the effective bandwidth usage of the dedicated link between the cluster. Bandwidth usage is inversely proportional to round-trip time (RTT) and the rate of packet drop or retransmission in the network.

**Note:** For standard block traffic, a packet drop or retransmission of 0.5% or more can lead to unacceptable usage of the available bandwidth.

Work with the network team to investigate over-subscription or other quality of service (QoS) issues of the link, with an objective of having the lowest possible (less than 0.1%) packet-drop percentage.

### 6.3.7 Short-distance partnership using RDMA

RDMA technology supports zero-copy networking, which makes it possible to read data directly from the main memory of one computer and write that data directly to the main memory of another computer. This technology bypasses CPU intervention during I/O, leading to lower latency and a faster rate of data transfer.

Refer to “*Chapter 8: Configuring FlashSystem and SVC partnerships over high-speed Ethernet*” in the IBM Redbooks *Ensuring Business Continuity: A Practical Guide to Policy-Based Replication and Policy-Based High Availability for IBM Storage Virtualize Systems*, SG24-8569.

For more information, see [Configuration details for using RDMA-capable Ethernet port for node-to-node communications](#).

#### Best practices to manage RDMA-capable Ethernet ports

The basic node tasks, such as adding a node or removing a node, are the same for both FC-based and RDMA-based connections between nodes. But, you might need to complete management actions on the RDMA-capable Ethernet ports before completing node-level management tasks.

Before completing managing tasks that are related to RDMA-capable Ethernet ports on a node, use the following best practices to manage these ports:

- ▶ If you already have a system that is configured to use RDMA-capable Ethernet ports, you must ensure that one redundant path is available before adding, removing, or updating settings for RDMA-capable Ethernet ports.
- ▶ Add, remove, or update settings on only one RDMA-capable Ethernet port at a time. Wait 15 seconds between these changes before updating other RDMA-capable Ethernet ports.
- ▶ If you are using a VLAN to create physical separation of networks, ensure that you follow these extra guidelines when completing management-related tasks:
  - VLAN IDs cannot be updated or added independently of other settings on a RDMA-capable Ethernet port, such as an IP address.
  - Before adding or updating VLAN ID information to RDMA-capable Ethernet ports, you must configure VLAN support on the all the Ethernet switches in your network. For example, on each switch, set VLAN to “Trunk” mode, and specify the VLAN ID for the RDMA-capable Ethernet ports that will be in the same VLAN.

#### Problem determination

It is a best practice to verify the setup against the [Configuration details for using RDMA-capable Ethernet port for node-to-node communications](#).

A first step is to review whether the node IP address is reachable and verify that the required TCP/UDP ports are accessible in both directions.

The following CLI command lists the port level connectivity information for node to node or clustering connectivity, and can be helpful to find the reason for connectivity error:

```
sainfo lsnodeipconnectivity
```

This command lists the port level connectivity information for node to node or clustering connectivity.

The [IBM Documentation for the lsnodeipconnectivity command](#) lists the different error\_data values with a description, and provides possible corrective actions.

### 6.3.8 Policy-based replication

Policy-based replication helps you to replicate data between systems with minimal management, significantly higher throughput, and reduced latency compared to the asynchronous remote copy function.

Refer to *Ensuring Business Continuity: Policy-Based Replication and Policy-Based High Availability for IBM Storage FlashSystem and IBM SAN Volume Controller*, SG24-8569 IBM Redbooks for more information.

#### **Checking the status and RPO**

The IBM Document [Checking volume group status and RPO](#) provides details on the different status and advise how to check in GUI and CLI.

### 6.3.9 Data reduction pools

Data reduction pools (DRPs) internally implement a Log Structured Array (LSA), which means that writes (including over-writes or updates) always allocate newer storage blocks. The older blocks (with invalid data) are marked for garbage collection later.

The garbage-collection process is designed to defer the work as much as possible because the more it is deferred, the higher the chance of having to move only a small amount of valid data from the block to make that block available to the free pool. However, when the pool reaches more than 85% of its allocated capacity, garbage collection must speed up to move valid data more aggressively to make space available sooner. This issue might lead to increased latency because of increased CPU usage and load on the back-end. Therefore, it is a best practice to manage storage provisioning to avoid such scenarios.

**Note:** If the usable capacity of a DRP exceeds more than 85%, I/O performance can be affected. The system needs 15% of usable capacity that is available in DRPs to ensure that capacity reclamation can be performed efficiently.

For more information about DRPs, see *Introduction and Implementation of Data Reduction Pools and Deduplication*, SG24-8430.

### 6.3.10 Managing the physical capacity of overprovisioned storage controllers

Drives and back-end controllers exist that include built-in hardware compression and other data reduction technologies that allow capacity to be provisioned over the available real physical capacity. Different data sets lead to different capacity savings, and some data, such

as encrypted data or compressed data, does not compress. When the physical capacity savings do not match the expected or provisioned capacity, the storage can run out of physical space, which leads to a write-protected drive or array.

To avoid running out of space on the system, the usable capacity must be monitored carefully by using the GUI of the Storage Virtualize system. The Storage Virtualize GUI is the only capacity dashboard that shows the physical capacity.

Storage Insights can monitor and report on any potential out-of-space conditions, and the Advisor function warns when the Storage Virtualize system almost at full capacity.

When the Storage Virtualize system pool reaches an out-of-space condition, the device drops into a read-only state. An assessment of the data compression ratio (CR) and the re-planned capacity estimation should be done to determine how much outstanding storage demand might exist. This extra capacity must be prepared and presented to the host so that recovery can begin.

The approaches that can be taken to reclaim space on the Storage Virtualize system in this scenario vary by the capabilities of the system, optional external back-end controllers, the system configuration, and planned capacity overhead needs.

In general, the following options are available:

- ▶ Add capacity to the Storage Virtualize system. Customers are encouraged to plan to add capacity to the system when needed.
- ▶ Reserve space in the Storage Virtualize system that makes it “seem” fuller than it really is, and that you can free up in an emergency situation. Storage Virtualize can create a volume that is not compressed, deduplicated, or thin-provisioned (a fully allocated volume). Create some of these volumes to reserve an amount of physical space, and give them a descriptive name (for example, “emergency buffer space”). If you are reaching the limits for physical capacity, you can delete one or more of these volumes to give yourself a temporary reprieve.

**Important:** Running out of space can be a serious situation. Recovery can be time-consuming. For this reason, it is imperative that suitable planning and monitoring be done to avoid reaching this condition.

See [Handling out of physical space conditions](#) for recovering from an out-of-space condition, including standard pools and DRPs.

**Note:** Power off all hosts accessing the pool to avoid host writes from impacting the success of the recovery plan.



# 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](https://ibm.com/services)

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(->Hide)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine:fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Draft Document for Review July 16, 2024 3:23 pm

8561 spine.fm 127



Redbooks

# Unleash the Power of Flash: Getting Started with

SG24-8561-00

ISBN



(1.5" spine)

1.5" <-> 1.998"

789 <-> 1051 pages



Redbooks

# Unleash the Power of Flash: Getting Started with

SG24-8561-00

ISBN



(1.0" spine)

0.875" <-> 1.498"

460 <-> 788 pages

Redbooks

## Unleash the Power of Flash: Getting Started with IBM Storage

SG24-8561-00

ISBN



(0.5" spine)

0.475" <-> 0.873"

250 <-> 459 pages

Redbooks

## Unleash the Power of Flash: Getting Started with IBM Storage Virtualize

(0.2" spine)

0.17" <-> 0.473"

90 <-> 249 pages

(0.1" spine)  
0.1" <-> 0.169"  
53 <-> 89 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(->Hide)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine:fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Draft Document for Review July 16, 2024 3:23 pm

8561 spine.fm 128



# Unleash the Power of Flash: Getting Started with

SG24-8561-00

ISBN

(2.5" spine)  
2.5" <-> mmm.n"  
1315 <-> mmm pages



# Unleash the Power of Flash: Getting Started with IBM Storage Virtualize 8.7 on IBM

SG24-8561-00

ISBN

(2.0" spine)  
2.0" <-> 2.498"  
1052 <-> 1314 pages







SG24-8561-00

ISBN

Printed in U.S.A.

Get connected

