

The Definitive Guide to IBM Storage FlashSystem 5300 Port Configuration

Bernd Albrecht

Hartmut Lonzer

Jon Herd

Sergei Kubin

Vasfi Gucer



Storage

Infrastructure Solutions



IBM Redbooks

**The Definitive Guide to IBM Storage FlashSystem 5300
Port Configuration**

October 2024

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

First Edition (October 2024)

This edition applies to IBM Storage Virtualize Version 8.7.

This document was created or updated on July 18, 2025.

© Copyright International Business Machines Corporation 2024. All rights reserved.

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

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	vii
Now you can become a published author, too!	ix
Comments welcome	ix
Stay connected to IBM Redbooks	ix
Chapter 1. IBM Storage FlashSystem 5300	1
1.1 The IBM FlashSystem 5300 introduction	2
1.2 IBM FlashSystem 5300 control enclosure overview	4
1.2.1 IBM FlashSystem 5300 hardware component overview	5
Chapter 2. Understanding IBM Storage FlashSystem 5300 ports	9
2.1 IBM Storage FlashSystem 5300 ports overview	10
2.2 Ports for management and monitoring	11
2.2.1 Management ports	11
2.2.2 Planning for management connections	13
2.3 Expansion enclosure ports	13
2.4 Fibre Channel SAN ports	13
2.4.1 FC ports and port attachment	14
2.4.2 FC ports planning and best practices	15
2.5 Ethernet SAN ports	16
2.5.1 Onboard and optional Ethernet SAN ports	16
2.5.2 Ethernet ports planning and best practices	17
Chapter 3. Step-by-step configuration guide	19
3.1 Setting up initial network access	20
3.1.1 Procedure	20
3.2 Management IP addresses	21
3.3 Service IP addresses	24
3.4 Port configuration (FC, iSCSI and so forth)	25
3.4.1 Portsets	25
3.4.2 Fibre channel ports (FC ports)	26
3.4.3 Ethernet ports	33
3.4.4 Configure protocols	38
3.5 Troubleshooting	40
Related publications	41
IBM Redbooks	41
Online resources	41
Help from IBM	41

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®
IBM FlashCore®

IBM FlashSystem®
Redbooks®

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

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

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

Preface

IBM® FlashSystem 5300 is an all-flash storage platform designed to empower businesses of all sizes. This compact powerhouse delivers exceptional performance and scalability, making it ideal for consolidating and protecting critical data.

This comprehensive IBM Redpaper explores the intricacies of IBM FlashSystem® 5300 port configuration, empowering IT professionals to optimize performance, enhance security, and ensure seamless integration within their existing infrastructure.


The target audience of this paper is storage administrators, system administrators and network specialists.

Authors

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



Bernd Albrecht is a Storage Advisory Partner Technical Specialist for DACH with over 32 years of technical sales experience at IBM. He brings deep expertise to his role. Previously, he spent 3 years as an OEM technical business development manager for Lenovo storage within IBM Germany. Based in Berlin, Bernd focuses on the IBM FlashSystem Family and IBM SAN Volume Controller products, with experience dating back to their launch in 2003. He is also a published IBM Redbooks® author.



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).



Sergei Kubin is a Senior Storage Support Engineer working in GBM Qatar. He holds an Electronics Engineer degree from Ural Federal University in Russia and has more than 15 years of experience in IT. In GBM, he provides support and guidance for customers using IBM and multi-vendor storage solutions. His expertise includes file and block storage, and storage area networks.



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
IBM USA

Lucy Harris, Evelyn Perez, Chris Bulmer
IBM UK

Vineet Sharma
IBM Dubai

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

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, 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>



IBM Storage FlashSystem 5300

IBM Storage FlashSystem 5300 stands out for its ability to deliver exceptional performance and scalability. This IBM Redpaper publication describes the IBM Storage FlashSystem (IBM FlashSystem) 5300 solution, which is a next-generation IBM FlashSystem control enclosure, and the new enhancements in the port configuration.

This chapter has the following sections:

- ▶ “The IBM FlashSystem 5300 introduction” on page 2
- ▶ “Placing your IBM Storage FlashSystem family storage system in your infrastructure” on page 2
- ▶ “IBM FlashSystem 5300 control enclosure overview” on page 4

1.1 The IBM FlashSystem 5300 introduction

This IBM Redpaper publication describes the IBM Storage FlashSystem (IBM FlashSystem) 5300 solution, which is a next-generation IBM FlashSystem control enclosure, and the new enhancements in the port configuration.

The IBM FlashSystem 5300 is an NVMe end-to-end platform that is targeted at the entry and midrange market and delivers the full capabilities of IBM FlashCore® technology.

The IBM FlashSystem 5300 also provides a rich set of software-defined storage (SDS) features that are delivered by IBM Storage Virtualize, including the following features:

- ▶ Data reduction and deduplication
- ▶ Dynamic tiering
- ▶ Thin provisioning
- ▶ Snapshots
- ▶ Cloning
- ▶ Replication
- ▶ Data copy services
- ▶ Transparent Cloud Tiering
- ▶ IBM HyperSwap® including 3-site replication for high availability (HA)
- ▶ Policy-based replication and policy-based HA (policy-based HA)
- ▶ Ransomware Threat Detection

Scale-out and scale-up configurations further enhance capacity and throughput for better availability.

The IBM FlashSystem 5300 is a high-performance storage solution that is based on a revolutionary 1U form factor. It consists of up to 12 NVMe Flash Devices in a 1U storage enclosure drawer with full redundant canister components and no single point of failure.

1. Placing your IBM Storage FlashSystem family storage system in your infrastructure

This section describes the possible use cases and where to place the IBM FlashSystem 5300 (or another system in the IBM FlashSystem family) in the client infrastructure. This solution addresses requirements and can be used to optimize and simplify an IT storage infrastructure.

Figure 1-1 on page 3 shows the current IBM FlashSystem and IBM SAN Volume Controller family.

FlashSystem Family  FlashSystem 5015 & 5045 SAS based entry to the world of IBM Storage Virtualize and IBM Storage FlashSystem.	Released May 2024  FlashSystem 5300 Next-gen award-winning NVMe efficiency leader. Densest 1U system unlocks the FlashCore Module value proposition along with significant data resilience.	 FlashSystem 7300 The 5th generation of the award-winning 7000 series, striking the right performance, scaling and value balance.	 FlashSystem 9500 Extreme performance for large, mixed and consolidated workloads
Powered by IBM Storage Virtualize			
 Storage Virtualize for Public Cloud Hybrid Cloud integration with all FlashSystem products allows seamless data migration from on premise to your chosen cloud provider.	 SAN Volume Controller (SVC) SV3 Storage Engine Celebrating 20 years in market, SVC supports more than 500 models of storage controller with seamless SVC and storage replacement.	 IBM FlashCore Modules The 4th generation of one of FlashSystem key value propositions with: • AI-enabled Ransomware Threat Detection embedded in hardware without any performance degradation • Data Reduction with hardware offload • Densest Tier-0 flash drive	

Figure 1-1 IBM Storage FlashSystem and IBM SAN Volume Controller 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.

The IBM FlashSystem 5300 can be used as a production data repository and a component of a DR solution because a primary system can send data in an efficient way into the hybrid multicloud infrastructure.

In particular, the IBM FlashSystem 5300 can meet the following customer requirements:

- ▶ First tier repository for production data.
 - ▶ Primary or target system for data replication or disaster recovery.
 - ▶ Provide HA services using policy-based HA.
 - ▶ Use Storage Virtualize capabilities to manage and virtualize older IBM or non-IBM storage and extend advanced Storage Virtualize functions (for example, data reduction) to the external capacity presented by the old storage.
- Old storage systems can be decommissioned, or their usage can be extended as an added pool of resources to the IBM FlashSystem 5300.
- ▶ Storage Virtualize in the IBM FlashSystem 5300 can provide the intelligent data migration tool from an outer storage to replace it or distribute application workload on more systems.
 - ▶ The IBM FlashSystem 5300 can use Transparent Cloud Tiering to move data into the cloud:
 - Use IBM Storage Virtualize for Public Cloud on Amazon AWS or other providers.
 - Use the Container Storage Interface (CSI) driver for Red Hat OpenShift Container Platform, which enables Cloud Pak foundation.
 - ▶ IBM software-defined storage (SDS) capabilities:

- Storage Scale support because the IBM FlashSystem 5300 can be used as a repository for several tiers in initial microprogram load (IML) processing.
- IBM Storage Protect as a cache or data repository.
- Copy Data Management.

Figure 2 shows an IBM FlashSystem 5300 as the main provider of advanced data services for on-premises and in a hybrid multicloud system.




Figure 2 IBM FlashSystem 5300 as the main provider of advanced data services

The client can expect the modern and advanced data services that are provided by a storage system to cover several scopes concurrently. The IBM FlashSystem products, which include the IBM FlashSystem 5300, all share this main characteristic.

Because all IBM FlashSystem products share the functions and software layer, it can be easier to select the suitable system that can match performance, capacity, and functional requirements.

Note: IBM FlashSystem 5300 provides proven 99.9999% availability, with an optional 100% guarantee when using IBM HyperSwap.

1.2 IBM FlashSystem 5300 control enclosure overview

The IBM FlashSystem 5300 system control enclosure manages your storage system, communicates with the host, and manages interfaces. In addition, it can house up to 12 NVMe-capable flash drives. These drives can be industry standard NVMe types or the exclusive IBM FlashCore Module (FCM) NVMe type and up to 12 optional Storage Class Memory (SCM) drives.

Figure 3 shows the IBM FlashSystem 5300 control enclosure front and rear view. In the front view, you can see the bezel removed and the 12 NVMe drive slots in two rows of six drives.



Figure 3 IBM FlashSystem 5300 control enclosure showing the front and rear view

1.2.1 IBM FlashSystem 5300 hardware component overview

The IBM FlashSystem 5300 control enclosure is a 1U rack-mounted NVMe flash memory enclosure that is based on IBM flash technology. It provides the primary management interface (GUI) and the host interface configuration. The IBM FlashSystem 5300 control enclosures support FC Protocol (FCP and FC-NVMe), iSCSI, NVMe over RDMA, and NVMe over TCP protocol interfaces.

Note: RoCE protocol is not supported on IBM FlashSystem 5300. iWARP is only supported as part of an IP Replication partnership between two Storage Virtualize products using RDMA.

Figure 4 shows the front view of the IBM FlashSystem 5300 control enclosure with the bezel removed. Also shown are six NVMe drives that are installed in upper slots 1–6 and six fillers in lower slots 7–12.




Figure 4 IBM FlashSystem 5300 control enclosure front view with bezel removed and drive slot locations

The IBM FlashSystem 5300 provides the following NVMe technologies:

- ▶ Supports unique world class IBM FCMs with inline compression.
- ▶ Supports industry-standard NVMe drives.
- ▶ Option for SCM type drives.
- ▶ NVMe-oF on FC.
- ▶ NVMe over RDMA or NVMe over TCP protocol.
- ▶ IBM FlashSystem 5300 system is offered as one model:
 - 4662 model 7H2: IBM FlashSystem 5300 NVMe control enclosure.

Figure 5 shows a top view of the IBM FlashSystem 5300 enclosure. Highlighted are the various components of the control enclosure and the two canisters.




Figure 5 Top view of the IBM FlashSystem 5300 enclosure

- ▶ Control enclosure:
 - Two canisters that are placed side by side.
 - 12 NVMe drive slots.
 - Six enclosure fan assemblies.
- ▶ Each canister contains the following components and quantities:
 - CPU (1)
 - DIMM Slots (4)
 - Battery (1)
 - Canister Fans (3)

- Power Supply PSU (1)
- PCIe adapters (0–2)
- PCIe riser cards (2)
- PCIe adapter blanking plates (0-2)

Note: The number of PCIe adapters is configurable at product ordering time and can be added or removed by a sales MES. MES (Miscellaneous Equipment Specification) refers to any server hardware modification, including adding, improving, removing, or a combination of these actions. The server's serial number remains unchanged.

Figure 6 shows the rear view of the IBM FlashSystem 5300 control enclosure. You can see the two canisters side by side, the interface cards, power supply units, and the various USB and Ethernet ports. All components are concurrently maintainable, except for the passive midplanes. All external connections are from the rear of the system.




Figure 6 Rear view of IBM FlashSystem 5300 control enclosure

Figure 7 shows a more detailed view of the rear of a canister.




Figure 7 IBM FlashSystem 5300 canister rear view

In Figure 7, you can see the RJ45 and USB ports in the canister. Also shown are the two new, on board, planer SFP ports on the left-hand side of the canister. These ports can be used for both external storage virtualization and host attachment. The PCIe adapter slots are shown with blanking plates in place to ensure and maintain the correct air flow for cooling through the canister. The IBM FlashSystem 5300 will allow 2 x PCIe adapters per canister, 4 x adapters per IBM FlashSystem 5300 enclosure.

For information on adapter support, see [IBM FlashSystem 5300 Node Canister Overview](#).

Additional notes:

All drives in the IBM FlashSystem 5300 are dual-port and hot-swappable.

The onboard compression cards are utilized only when data reduction volumes are configured. At present, this hardware offload is used for compression and SHA hashing for deduplication.



Understanding IBM Storage FlashSystem 5300 ports

The IBM Storage FlashSystem 5300 offers a robust set of ports that provide critical connections for data transfer, management, and system health. Understanding these ports and their functionalities is essential for configuring and maintaining optimal performance of your FlashSystem storage solution. This chapter delves into the various types of ports available on the FlashSystem 5300, explaining their purposes, functionalities, and configuration considerations.

This chapter has the following sections:

- ▶ “IBM Storage FlashSystem 5300 ports overview” on page 10
- ▶ “Ports for management and monitoring” on page 11
- ▶ “Expansion enclosure ports” on page 13
- ▶ “Fibre Channel SAN ports” on page 13
- ▶ “Ethernet SAN ports” on page 16

2.1 IBM Storage FlashSystem 5300 ports overview

By understanding the distinct functionalities and configuration considerations of your FlashSystem 5300 ports, you can optimize your storage solution for peak performance and ensure the continued accessibility and security of your critical data.

Figure 2-1 shows the node canister ports.




Figure 2-1 FlashSystem 5300 node canister ports

Important: Unlike previous offerings, the IBM FlashSystem 5300 assigns logical port numbers differently from physical port numbers. The management port is always port 1. This means physical port 1 becomes logical port 2, and physical port 2 becomes logical port 3.

A fixed set of ports is available on each node canister. Those ports are always present:

- ▶ 1x RJ45 dedicated management port.
- ▶ 1x RJ45 dedicated technician port.
- ▶ 1x USB Type A port for attaching encryption key media and service tasks.
- ▶ 2x Ethernet SFP ports for host I/O, and replication over Ethernet SAN.

Each canister has two slots for Host Interface Cards (HIC). Both nodes in the control enclosure must have the same set of cards installed. The following HIC cards can be added to each node canister to expand its connectivity:

- ▶ 2-port 12Gb SAS card for expansion enclosure attachment (one card per node only).
- ▶ 2-port 64Gb Fibre Channel card for host I/O, clustering and replication.
- ▶ 4-port 32Gb Fibre Channel card for host I/O, clustering and replication.
- ▶ 4-port 10Gb Ethernet card for host I/O, clustering and replication.

Note: The list above is valid at the moment of writing this publication and can be extended in future.

HIC cards can be installed into any of two node slots, with the exception of a SAS adapter which is supported in slot 2 only. Manufacturing starts populating slots with slot #1. Adapter types can be mixed within a single node.

Both nodes in a control enclosure must have the same HIC card configuration, which means that they must have identical cards in a same slot.

Possible adapter configurations are shown in Figure 2-2.

Node1								Node2							
Slot 1				Slot 2				Slot 1				Slot 2			
Empty (onboard Ethernet only)				Empty (onboard Ethernet only)				Empty (onboard Ethernet only)				Empty (onboard Ethernet only)			
Empty (onboard Ethernet only)				SAS Adapter				Empty (onboard Ethernet only)				SAS Adapter			
10GbE	10GbE	10GbE	10GbE	Empty				10GbE	10GbE	10GbE	10GbE	Empty			
32GB FC	32GB FC	32GB FC	32GB FC	Empty				32GB FC	32GB FC	32GB FC	32GB FC	Empty			
64GB FC				Empty				64GB FC				Empty			
10GbE	10GbE	10GbE	10GbE	SAS Adapter				10GbE	10GbE	10GbE	10GbE	SAS Adapter			
32GB FC	32GB FC	32GB FC	32GB FC	SAS Adapter				32GB FC	32GB FC	32GB FC	32GB FC	SAS Adapter			
64GB FC				SAS Adapter				64GB FC				SAS Adapter			
10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE
32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC	32GB FC
64GB FC				64GB FC				64GB FC				64GB FC			
64GB FC				32GB FC	32GB FC	32GB FC	32GB FC	64GB FC				32GB FC	32GB FC	32GB FC	32GB FC
10GbE	10GbE	10GbE	10GbE	64GB FC				10GbE	10GbE	10GbE	10GbE	64GB FC			
10GbE	10GbE	10GbE	10GbE	32GB FC	32GB FC	32GB FC	32GB FC	10GbE	10GbE	10GbE	10GbE	32GB FC	32GB FC	32GB FC	32GB FC

Figure 2-2 Possible adapter cards combinations

2.2 Ports for management and monitoring

System management interfaces, such as CLI and GUI, can be accessed through various methods.

2.2.1 Management ports

We first discuss the management ports.

Technician port

There is a technician port on each FlashSystem 5300 node canister. The technician port is an RJ45 1Gb Ethernet port, which can auto-negotiate down to 100Mbps and 10Mbps. It can be visually identified by blue stripes on both sides of the connector, and by a black gear symbol on the node’s faceplate. See Figure 2-1 on page 10.

The technician port is used for initial system setup and recovery tasks, such as resetting the superuser password. It requires a direct connection to a workstation (no LAN) and provides access to a dedicated management interface.

Important: The technician port should be disconnected during normal operation. It is solely for management purposes and does not provide data access.

The port is set up to assign an IP address to the workstation that is attached to it with DHCP. IP address on the interface is 192.168.0.1 and it cannot be changed.

For more information, see [Initializing the system with the technician port](#) and [Using technician port for node access](#).

Important: Restricting physical access to the storage system is essential for safeguarding the technician port due to its elevated management privileges.

USB port

A single type A USB port is found on each node. It serves two purposes:

- ▶ To be used as an emergency management interface. A very limited set of commands can be executed on the node by putting the command to `satask.txt` file on a USB pen drive and plugging that pen drive to the node's USB port. Tasks that can be performed with USB management interface are:
 - Check node hardware status.
 - Set or change the node's service IP address.
 - Reset the superuser password.

For more information, see [USB flash drive interface](#).

- ▶ To attach USB flash drives with the master encryption key for the system. USB flash drives are one of the supported methods of storing encryption keys. The node reads the master encryption key from the USB drive during the boot process. After this initial read, the USB drive is not accessed until the next system startup. For more information, see [Encryption with USB flash drives](#).

Tip: If required, USB ports can be disabled to comply with organizational security policies.

Primary management port

The dedicated management port is an RJ45 1Gb Ethernet port, which can auto-negotiate down to 100 Mbps and 10 Mbps. It is a primary interface for all system management and monitoring tasks. No host I/O access is possible on this port, ensuring physical separation of management and data networks.

The dedicated management port is identified by #3 on the FlashSystem 5300 node faceplate.

Important: On FlashSystem 5200 and most of the other platforms, the dedicated or shared primary management port is usually port #1. On FlashSystem 5300, hardware port number assignment is different, and the port has physical #3, while logically it is still recognized as port id 1.

The node's service IP address is accessible through the primary management port. If a node becomes the Configuration node within a cluster, the system's management IP address (or Cluster IP address) is also assigned to this port.

On a new system, only the default service IP is available on the dedicated management port. The default address is 192.168.70.121 on node 1 (in the left chassis slot) and 192.168.70.122 on node 2, in the right slot of the chassis.

The system management IP address is assigned during initial cluster setup and can be modified as needed.

Beyond providing access to the system's GUI, CLI, and REST API, the primary management port facilitates outbound communication for services like DNS and Call Home.

Optional secondary management port

The secondary management IP (or cluster IP #2) can be configured on the system's first host I/O port (physical port #1), which is assigned logical ID #2 within the Storage Virtualize configuration. Refer to Table 2-2 on page 17 for relationship between physical and logical port numbers.

2.2.2 Planning for management connections

Each system node requires at least one Ethernet connection for management and monitoring. A unique service IP address is necessary for each node, while a single cluster management IP (supporting both IPv4 and IPv6 concurrently) is shared across the cluster.

The cluster management IP must be unique from all service IPs. While it can reside on the same subnet as the service IPs, it is not required to do so.

An optional secondary management IP can be configured as a failover option. It is only activated if the primary management port becomes unavailable. Outbound connections will always originate from the primary management IP when active.

Depending on the set of features that are in use, set up management network firewall to pass the following traffic to and from the system:

- ▶ Management - from administrators hosts to the system's CLI (SSH) and GUI/REST API (HTTPS) interfaces.
- ▶ Monitoring - to Storage Insights data collector host (if standalone collector is used).
- ▶ Network services - from the system to NTP and DNS servers.
- ▶ Remote user authentication - from the system to LDAP server.
- ▶ Event notifications - from the system to SMTP, SNMP and syslog servers.
- ▶ Replication management - control plane of replication, traffic between management ports of Storage Virtualize systems in replication partnership.
- ▶ IP Quorum - to and from hosts running IP Quorum application.

In addition to that, some features require outside connectivity:

- ▶ Transparent Cloud Tiering - to cloud provider endpoint.
- ▶ Multi-Factor Authentication and Single Sign-On - to authentication service provider (can work through proxy).
- ▶ Call Home and Remote Assist - to IBM Support Center (can work through proxy).

For more information, see [Planning management connections](#).

2.3 Expansion enclosure ports

Expansion enclosures require the system to be equipped with SAS HIC. Only a single SAS card can be installed into each node, and it is supported only in card slot 2.

SAS card has four ports, but only ports 1 and 3 can be used for expansion attachment.

SAS card supports only expansion enclosure attachment. SAS host attachment is not supported.

For more information, see [Connecting expansion enclosures to control enclosure](#).

2.4 Fibre Channel SAN ports

Up to eight 32Gb Fibre Channel (FC) SAN ports can be installed into each FlashSystem 5300 node by using one or two 4-port Host Interface Cards. All FC ports support FCP-SCSI and FC-NVMe.

The system can be configured without any Fibre Channel (FC) ports, operating solely as an Ethernet-attached device.

2.4.1 FC ports and port attachment

Two types of optional FC Host Interface Cards are available:

- ▶ 4-port 32Gbps FC HIC
- ▶ 2-port 64Gbps FC HIC

Cards can be installed in either node slot, allowing for flexibility in card type selection within a node. However, both nodes must have identical card configurations.

A card can be ordered with 32Gbps or 64Gbps LC SFPs. An SFP is capable of auto-negotiating up to two steps down, which means that minimum supported attachment speeds are 8Gbps or 16Gbps respectively.

FC ports can be used for:

- ▶ Host I/O with FCP-SCSI and FC-NVMe.
- ▶ Clustering (communication between control enclosures of a single system).
- ▶ Replication (communication between different Storage Virtualize systems).

Ports can be attached to SAN switches, and also support direct attachment to hosts and another Storage Virtualize system for clustering or replication.

When using direct attachment between two control enclosures, consider the examples shown in Figure 2-3. Each node requires connectivity to both nodes in the opposing control enclosure. In clustering configurations, such as HyperSwap, it is required to have redundant connections - four links per node. In replication-based configurations, such as policy-based high availability, one connection to each remote node is sufficient. However, *two* are recommended for maximum throughput and performance.




Figure 2-3 Direct attached configurations for replication and clustering

FlashSystem 5300 exclusively operates in NPIV mode, allowing each physical FC port to utilize multiple WWPNs for SAN fabric registration. This mode is mandatory and cannot be altered.

In NPIV mode, virtual ports (WWPNs) can be migrated between equivalent physical ports on different nodes within the same I/O group. However, maintaining consistent SAN fabric connectivity is crucial. All equivalent physical ports in an I/O group must be connected to the same SAN fabric.

For example, ports #1 and #3 of both nodes can be connected to Fabric A, and ports #2 and #4 to Fabric B. For more information, see [Zoning requirements for N_Port ID virtualization](#).

Every physical port of FlashSystem 5300 registers in the SAN switch three WWPNs:

- ▶ Physical WWPN: allows external storage virtualization, replication, clustering traffic.
- ▶ FCP-SCSI host WWPN: allows host I/O with FCP-SCSI.
- ▶ FC-NVMe host WWPN: allows host I/O with FC-NVMe.

During an NPIV failover, host WWPNs migrate to the partner node while the physical WWPN remains static. This allows up to five WWPN logins per physical FC port.

WWPN is assigned according to adapter and port location. Figure 2-4 shows the WWPN numbering scheme.

PCI slot	Adapter port	Physical WWPN	NPIV WWPN for FCP-SCSI hosts	NPIV WWPN for FC-NVMe hosts
1	1	500507681211xxx	500507681215xxx	500507681219xxx
1	2	500507681212xxx	500507681216xxx	50050768121axxx
1	3	500507681213xxx	500507681217xxx	50050768121bxxx
1	4	500507681214xxx	500507681218xxx	50050768121cxxx
2	1	500507681221xxx	500507681225xxx	500507681229xxx
2	2	500507681222xxx	500507681226xxx	50050768122axxx
2	3	500507681223xxx	500507681227xxx	50050768122bxxx
2	4	500507681224xxx	500507681228xxx	50050768122cxxx

Figure 2-4 Adapter port number to WWPN relationship

2.4.2 FC ports planning and best practices

Consider the information below when planning FC SAN configuration.

- ▶ In HyperSwap or policy-based high availability configurations, use [FC portsets](#) or SAN zoning to dedicate ports for node-to-node traffic or replication.

All FC ports of the Storage Virtualize system can serve any type of traffic (node to host, node to external storage, node to another node), and all types can flow simultaneously. *However, for performance reasons, in configurations involving HyperSwap or policy-based replication, it is recommended to dedicate a group of ports to be used exclusively for this purpose and keep host I/O on the other ports.*

HyperSwap communication relies on inter-node (intra-cluster) links. Unlike replication, port sets cannot be dedicated for this traffic, and SAN zoning often proves insufficient for isolation. So, [FC port mask](#), the `localfcportmask` defines the ports that are used for inter-node communication within a storage system, must be applied on the system to limit the system from using any ports except dedicated for inter-node. This will optimize inter-node communication and prevent congestion.

Replication traffic demands dedicated inter-cluster links for optimal performance. To achieve this, port sets can be configured to exclusively handle replication data.

- ▶ To optimize performance and prevent potential bottlenecks, it's recommended to utilize all available FC ports, even if initial bandwidth calculations suggest fewer are sufficient. Distributing the workload across multiple ports enhances overall system performance and accommodates potential burst I/O.

It is important to use FC SAN zoning or [FC portsets](#) to limit a number of host paths down to a recommended number, which is 4 (with 8 paths maximum allowed).

- ▶ Use advanced zoning techniques, such as [Brocade Peer Zoning](#) or [Cisco Smart Zones](#) when designing host access zones. Peer zoning combines the benefits of single-initiator zoning with the benefits of a minimal administrative overhead. An example of a peer zone is:
 - Virtual WWPNs (FCP-SCSI or FC-NVMe) of a pair of equivalent ports (ports with the same location) of both nodes of a FlashSystem 5300 IO group, as principal members.
 - WWPNs of all the hosts that need to be attached, as non-principal members.

As non-principal members cannot communicate directly with each other, host HBAs are isolated and can only interact with principal members.

Note that for back-end storage virtualization, replication and clustering, traditional zones must be used.
- ▶ Exclude FC-NVMe virtual WWPNs from all zones unless the NVMe protocol is actively utilized.

2.5 Ethernet SAN ports

Each FlashSystem 5300 node has two onboard Ethernet ports for data transfers. Up to eight Ethernet ports per node can be added with two optional Host Interface Cards.

There is a significant difference in attachment and use options between onboard and optional ports.

2.5.1 Onboard and optional Ethernet SAN ports

Both 25 GbE onboard ports and 10 GbE optional ports are shipped without optics. Supported fiber and copper connection options are summarized in Table 2-1.

Table 2-1 Ethernet port media options

Card	SFP Module	Cable
On-Board 10/25GbE	10GbE/25GbE SFP28 Feature code ACHP	OM3 LC
	None	25G DAC (up to 3m)
On-Board 10/25GbE and 4-port 10GbE	1GbE/10GbE SFP+ Feature code ACHQ	OM3 LC
	1GbE/10GbE RJ45 Feature code ACJ2	Cat5/Cat6 UTP
	None	10G DAC / AOC

Both onboard 25GbE and optional 10GbE ports support host access with SCSI and NVMe protocols for host attachment and IP replication, however optional 10GbE ports are RDMA-capable, while 25GbE are not. This results in a wider list of possible use options for 10GbE ports.

- ▶ For onboard 10/25GbE ports, the following protocols and applications are supported:
 - Host I/O with iSCSI
 - Host I/O with NVMe/TCP

- Long-distance replication over IP link for policy-based replication (IP replication over TCP)
- ▶ For ports in an optional 4-port 10GbE card, in addition to the protocols supported for 25GbE, feature that require RDMA are available:
 - Node-to-Node communication with RDMA (Ethernet clustering or iSER clustering).
 - Short-distance partnerships using RDMA for synchronous links (High Speed Replication or IP replication with iWARP RDMA).

2.5.2 Ethernet ports planning and best practices

Consider the following when planning and implementing Ethernet SAN:

- ▶ Equip the system with 4-port 10GbE cards if you plan to build a HyperSwap cluster or policy-based HA system using Ethernet inter-site links. 10GbE card supports RDMA and is the only available solution to achieve that without using Fibre Channel cards. For more information, see [Short distance partnerships using RDMA](#).

Note: HyperSwap configuration requires an approved [SCORE](#) request. IBM recommends policy-based HA.

- ▶ For optimal performance and reliability, dedicate separate ports for IP-based clustering or replication traffic. Avoid combining host access and replication on the same port. This can be achieved using [Network Portsets](#).
- ▶ If in HyperSwap configuration it is not possible to dedicate inter-site links for intra-cluster traffic, configure [Priority Flow Control](#) (QoS) to make sure that system traffic is prioritized over other traffic types.
- ▶ For host access, it is recommended to separate host to storage (iSCSI or NVMe/TCP) traffic from other types of traffic in your LAN. This can be achieved by building a separate physical network, or using dedicated storage access ports on a host side, separating networks with VLANs and using QoS to prioritize storage traffic.
- ▶ Use recommendations given in the [iSCSI performance analysis tuning](#) article:
 - Utilize all available storage ports.
 - Verify that your network supports Jumbo frames end-to-end, and enable them by setting MTU to 9000 on ports designated for host access.
 - Disable delayed TCP ACK on the hosts.
- ▶ When working with the FlashSystem 5300, carefully distinguish between logical and physical port numbers. These are distinct identifiers within the system. Table 2-2 shows relationship between onboard ports physical and logical numbers.

Table 2-2 Physical and logical ports

Physical port number (left to right)	1	2	3
Logical port number	2	3	1
Port designation	Host IO IP Replication Secondary management	Host IO IP Replication	Dedicated (primary) management



Step-by-step configuration guide

In this chapter provide a detailed, step-by-step guide on configuring different types of FlashSystem 5300 ports:

This chapter has the following sections:

- ▶ “Setting up initial network access” on page 20
- ▶ “Management IP addresses” on page 21
- ▶ “Service IP addresses” on page 24
- ▶ “Port configuration (FC, iSCSI and so forth)” on page 25
- ▶ “Troubleshooting” on page 40

3.1 Setting up initial network access

Before you begin, review the [Planning worksheets](#).

Initial setup of an FS5300 is done on the T-port. These are 1Gb RJ45 ports.

3.1.1 Procedure

Perform the following steps:

1. Ensure the system is powered on.
2. Configure an Ethernet port on the personal computer to enable Dynamic Host Configuration Protocol (DHCP) configuration of its IP address and DNS settings. If you do not have DHCP, you must manually configure the personal computer. Specify the static IPv4 address 192.168.0.2, subnet mask 255.255.255.0, gateway 192.168.0.1, and DNS 192.168.0.1.
3. Locate the technician port on each node canister, as shown in Figure 3-1.




Figure 3-1 Where to find the technician port

4. Disconnect the personal computer from all networks. Connect an Ethernet cable between the port of the personal computer that is configured in step 2 and the technician port in the left canister (1) that is shown in Figure 3-1.
5. Once the personal computer is connected through Ethernet, open a supported web browser and navigate to <https://install>. If your network does not use DHCP for automatic IP assignment, use the static IP address 192.168.0.1 instead.
6. The browser is automatically directed to the initialization tool.
7. To set up the system's management IP address, follow the on-screen prompts provided by the initialization wizard.
8. Once the initialization process is complete, disconnect the cable between the personal computer and the technician port. You can then continue with the initial system setup as outlined in see [Completing the initial system setup \(customer task\)](#).

Alternatively, you can configure the system using the service address on the management port. The default IP addresses are 192.168.70.121 for the first node and 192.168.70.122 for the second node in an enclosure. This method offers two options:


- **Command Line:** Use SSH with the command **satask mkcluster** (for advanced users).
- **GUI:** Use the graphical user interface for a more user-friendly approach.

3.2 Management IP addresses

During system initialization, a single management IP address is assigned to be used through all the management ports of the cluster. This IP serves as the central point for managing the entire cluster, regardless of which node is currently designated as the configuration node. For proper functionality, ensure all management ports are connected to your network.

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.


Management IPs are included by default in the newly defined System Management portset (ID 72) for managing the system. See Figure 3-2.



portset0 (Default)	PORT TYPE	PORT COUNT	MEMBER COUNT
Host Attachment	Ethernet	0	0
portset64 (Default)	Fibre Channel	4	1
SystemManagement (Default)	Management Ethernet	1	0
portset1	Remote Copy	0	0
portset2	Remote Copy	0	0
portset3 (Default)	Storage	0	-
portset4	Remote Copy	0	0

Figure 3-2 Management portset after initialization

You can see that one port is currently defined (Figure 3-3). You can view details about this port or check for additional defined ports.



SystemManagement (Default)	PORT TYPE	PORT COUNT
Management Ethernet	Ethernet	1

Figure 3-3 Management portset after initialization - Details

Figure 3-4 shows the portset mappings.




Figure 3-4 Management portset after initialization - Portset mappings

Another new feature is the VLAN support. This functionality enables the creation of Virtual Local Area Networks (VLANs) within your existing network infrastructure. VLANs provide a method for segmenting your network into logical sub-divisions, thereby enhancing both security and network performance. VLAN configurations can be added and modified at your convenience to best suit your evolving network management needs. See Figure 1-5.




Figure 3-5 Modify management IP address

Starting with IBM Storage Virtualize Version 8.7.0, system administrators now have the option to configure a second management IP address for increased redundancy and manageability. Additionally, management IP addresses are no longer restricted to ports 1 and 2. This expanded flexibility allows for a more customized and efficient network configuration.

Only two management IPs can be configured, but you can now use any Ethernet port for the two allowed management IPs.

The system-defined default management port set restricts the number of configurable system IP addresses to two.

Figure 1-6 shows how to add a second management port.




Figure 3-6 Add management IP -1

Figure 1-7 shows how to select a port number for the second management IP.




Figure 3-7 Add management IP - select port

Other related improvements in IBM Storage Virtualize V 8.7:

Increased data IP flexibility: You can now configure up to 4 routable data IP addresses per port, per node. This provides more flexibility for network configuration and traffic management.

Faster failover: The configuration node failover time has been reduced by 10%. This means the system recovers from a failure of the config node more quickly, minimizing downtime.

Unified CLI commands: Common Command-Line Interface (CLI) commands (**mkip**, **rmip**, **lsip**) are now available for both data and management IP addresses. This simplifies managing IP addresses by providing a consistent interface for both types.

New command for system management IPs: There is a new command named **chip** for managing system and management IP addresses. For more information, see [chip](#).

For more information, see [Release Note for systems built with IBM Storage Virtualize](#).

Enabling DNS resolution is also recommended. This allows the system to translate hostnames into IP addresses, simplifying network operations and improving overall usability.

3.3 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.

This service IP address is always assigned to port 1, even if you change the management IP address to a different port. You need to change these default service IPs to addresses that are readily accessible on your network. This is necessary for remote management and service tasks using the Service Assistant Interface.

The service IP address allows access to the Service Assistant Interface, accessible through a web browser or SSH client. This interface provides functionalities for maintenance and service tasks on the system.

Important: While the service IPs are used to access the Service Assistant Interface, their importance goes beyond that. These IPs are also crucial for various system functions such as:

- ▶ **Key server access:** Communication with a key server for security purposes.
- ▶ **IP quorum:** Establishing a quorum for cluster management and data consistency.
- ▶ **Remote support assistance:** Enabling remote technicians to access the system for troubleshooting or maintenance.

Therefore, it is vital to configure the service IPs with addresses that are readily accessible on your network.

3.4 Port configuration (FC, iSCSI and so forth)

In this part we discuss the configuration of ports for host attachment, remote copy and clustering.

3.4.1 Portsets

Portsets are collections of logical port addresses grouped based on specific traffic types. This allows for efficient management and isolation of different network traffic flows.

Note: Fibre Channel (FC) port masking does not affect traffic between hosts and storage devices. It applies only to communication between nodes within a system and replication traffic between systems. FC port masking is deprecated after Storage Virtualize Version 8.5.

The system offers pre-defined Fibre Channel and Ethernet portsets for specific traffic types: host attachment, system management, remote copy, and back-end storage virtualization. For more information, see [Portsets documentation](#).

For more information, see [Planning for more than four fabric ports per node canister](#).

Note: A host definition is configured to access storage devices through a single Fibre Channel portset.

Figure 3-8 shows how to create or modify portsets.




Figure 3-8 Portsets

For a high volume of similar devices, consider creating dedicated portsets (Figure 3-9 on page 26). This allows for granular grouping based on functionalities (for example, cluster and server groups) to optimize network traffic flow and simplify management.




Figure 3-9 Create a portset

Choose a portset type: Fibre Channel or Ethernet. See Figure 3-10.




Figure 3-10 Select a portset type

Host attachment portsets you can specifying an ownership group (optional). This group defines user access and simplifies management for specific sets of hosts.

3.4.2 Fibre channel ports (FC ports)

FC ports are only available if you order FC adapter cards and the necessary SFP+ transceiver modules for those cards.

For more information, see [Fibre channel connectivity](#).

SAN fabrics with long-distance fiber connections

FC ports require separately purchased FC adapter cards and compatible SFP+ transceiver modules for connection to the storage network. The maximum distance these connections can reach depends on the type of FC cable you use. See Figure 3-11 on page 27 for the maximum distances achievable over FC connections using Short Wavelength (SW) SFP+ transceiver modules to determine suitable cabling for your desired configuration.

Formfactor	FC Speed	OM2	OM3	OM4/OM5
		(50µm)	(50µm)	(50µm)(b)
SFP+ (LC)	2 Gbps	300m	500m	
SFP+ (LC)	4 Gbps	150m	380m	400m
SFP+ (LC)	8 Gbps	50m	150m	150m
SFP+ (LC)	10 Gbps	82m	300m	550m
SFP+ (LC)	16 Gbps	35m	100m	125m
SFP+ (LC)	32 Gbps	20m	70m	100m
SFP+ (LC)	64 Gbps	20m	70m	100m
QSFP(MTP)	16 Gbps	N/A	66m	100m
QSFP(MTP)	32 Gbps	N/A	70m	100m

Figure 3-11 Maximum distances achievable over FC connections using Short Wavelength (SW) SFP+ transceiver modules

For longer distances, Long Wavelength (LW) SFP+ transceivers with mono mode cable are required. However, the maximum achievable speeds are limited by the cable type, distance, and the FC adapter's 40 Buffer Credits. Here is a summary of the limitations:

- ▶ 8 Gb: Up to 10km
- ▶ 16 Gb: Up to 5km
- ▶ 32 Gb: Up to 2.5km

Tip: Due to the use of small block sizes in cluster-to-cluster communication, it is recommended to limit the actual distance to around one-third of the maximum values mentioned above. This helps to ensure optimal performance.

FC switches are recommended for distances exceeding these limitations.

Determining FC card and speeds

The System Hardware - Overview allows you to identify the specific FC cards installed in your system. See Figure 3-12.




Figure 3-12 System Hardware - Overview - 32 Gbit/s FC card

This system allows you to view information about your connected C) ports, including their speed (for example, 32 Gbit/s or 64 Gbit/s) and port numbers.

A 4-port 32 Gbit/s FC adapter card provides four FC connections, each capable of transferring data at 32 gigabits per second. See Figure 3-12 on page 27.

The system also provides a 2-port 64 Gbit/s FC adapter card for higher bandwidth needs. See Figure 3-13.

Important: While a 64 Gbit/s FC adapter card option exists with four physical ports, only two of those ports are usable.




Figure 3-13 System Hardware - Overview - 64Gbit/s FC card

Selecting a dedicated single port displays information about its number and speed. See Figure 1-14.




Figure 3-14 Identify a single FC port

Usage of FC ports

On the **System Hardware - Overview** view you will find only the physical WWPN used for cluster communication, remote mirroring and external virtualization. You can find the FC port WWPN details by selecting **Settings** → **Fibre Channel Ports**. See Figure 3-15 on page 29.

Physical versus logical WWPN: A physical WWPN identifies the entire Fibre Channel (FC) adapter card. Logical port WWPNs identify specific functions within the card. These logical WWPNs can be taken over by another FC card on another node in case of a node failure.

ID	System Connection	Owning Node	WWPN	Host Count	Portset Count	Login Count	Host ID Permitted	Virtualized	Current Node	Protocol Type
~1	Any									
1	Any	1(Right)	50050768121180B1	0	0	5	No	No	1	SCSI
1	Any	1(Right)	50050768121180B1	1	1	1	Yes	Yes	1	SCSI
1	Any	1(Right)	50050768121180B1	1	1	0	Yes	Yes	1	NVMe
1	Any	2(Left)	50050768121180A4	1	1	0	Yes	Yes	2	NVMe
1	Any	2(Left)	50050768121180A4	0	0	5	No	No	2	SCSI
1	Any	2(Left)	50050768121180A4	1	1	1	Yes	Yes	2	SCSI
~2	Any									
2	Any	1(Right)	50050768121180B1	1	1	1	Yes	Yes	1	SCSI
2	Any	1(Right)	50050768121180B1	1	1	0	Yes	Yes	1	NVMe
2	Any	1(Right)	50050768121180B1	0	0	5	No	No	1	SCSI
2	Any	2(Left)	50050768121180A4	1	1	1	Yes	Yes	2	SCSI
2	Any	2(Left)	50050768121180A4	1	1	0	Yes	Yes	2	NVMe
2	Any	2(Left)	50050768121180A4	0	0	5	No	No	2	SCSI
~3	Any									
3	Any	1(Right)	50050768121180B1	1	1	0	Yes	Yes		SCSI
3	Any	1(Right)	50050768121180B1	1	1	0	Yes	Yes		NVMe
3	Any	1(Right)	50050768121180B1	0	0	0	No	No		SCSI
3	Any	2(Left)	50050768121180A4	1	1	0	Yes	Yes		SCSI
3	Any	2(Left)	50050768121180A4	1	1	0	Yes	Yes		NVMe
3	Any	2(Left)	50050768121180A4	0	0	0	No	No		SCSI
~4	Any									
4	Any	1(Right)	50050768121180B1	1	1	0	Yes	Yes		NVMe

Figure 3-15 Fibre Channel ports

On IBM FlashSystem storage units with IBM Storage Virtualize 8.7, NPIV is *always* on by default.

N-Port ID Virtualization:

The system supports N_Port ID Virtualization (NPIV) technology for Fibre Channel (FC) connections. NPIV is an industry standard that allows a single physical FC adapter to act as multiple virtual ports. Each virtual port can have its own unique World Wide Port Name (WWPN) and World Wide Node Name (WWNN) to register with the Storage Area Network (SAN) fabric.

NPIV benefits include:

- Improved host behavior during node restarts by delaying virtual port activation until I/O is needed.
- Masking path failures caused by offline nodes from host multipathing drivers.
- Efficient use of physical FC ports by allowing separate WWPNs for host I/O and clustering on the same port.

For successful NPIV configuration, see [N-Port ID Virtualization](#) for proper cabling and zoning procedures.

In some cases, FC fabric management or zoning scripts might require a different WWPN format for copy-and-paste operations. See Figure 3-16 on page 30.




Figure 3-16 FC ports change WWPN view

To determine the correct WWPN for use, refer to Table 3-1.

Table 3-1 WWPN usage

Host IO permitted	Virtualized	Protocol type	Usage
Yes	Yes	SCSI	FC SCSI Host connect
Yes	Yes	NVMe	NVMe/FC Host connect
No	No	SCSI	Cluster communication, remote mirroring, external virtualization

It is important to zone the correct WWPNs. Hosts cannot access storage through physical, non-virtualized WWPNs. When using virtualized WWPNs, the selection determines whether the host can access via FC SCSI or NVMe/FC, allowing control over the access protocol.

For direct FC connections, there is no selection for access protocols. To confirm direct operating system access through FC, you can check the [IBM System Storage Interoperation Center \(SSIC\)](#).

Important: Direct FC connections (without NPIV virtualization) do not support NVMe/FC, only FC SCSI connections are available.

Each host or external storage system does a full fabric login as part of the FC communication process. The FS5300 maintains the information of each device, which is registered and accessing FS5300 nodes, similar to the name space of a fiber switch.

You can check the FC SCSI host connections under **Settings** → **Fibre Channel Connectivity**. See Figure 3-17 on page 31.




Figure 3-17 Fibre Channel Connectivity

You can also check the state of the connection. Only FC SCSI connections will be shown.

You can check for active NVMe connections within your system management interface by navigating to **Settings** → **NVMe Connectivity**. See Figure 3-18.




Figure 3-18 NVMe FC Connectivity

It is possible to filter by hosts and nodes in the listings and also to export this listings to CSV files.

Modify FC ports

The system allows you to change the FC port being used. See Figure 3-19 on page 32.




Figure 3-19 Modify FC port actions

You can manage FC port assignments. This allows you to add or remove FC ports from portsets for better zoning control.

For advanced configurations, you can also change how hosts access an FC port. This might involve using NPIV to assign different virtual WWPNs or modifying security settings. See Figure 3-20.




Figure 3-20 FC Modify port connection

- ▶ **Any:** Selected Port can be used for any type of traffic like local cluster communication, remote replication or host communication (using NPIV WWPN for host).
- ▶ **Local:** Selected Port can be used for only local cluster communication or host communication (using NPIV WWPN for host).
- ▶ **Remote:** Selected Port can be used for only remote replication or host communication (using NPIV WWPN for host).
- ▶ **None:** Selected Port cannot be used for cluster communication or remote replication. Host communication is always allowed in any type (using NPIV WWPN for host).

Note: Host communication is always allowed irrespective of any type selected but only using NPIV WWPN of the port.

3.4.3 Ethernet ports

The Ethernet ports menu is used for configuring network connectivity. This includes assigning ports for host attachment, storage replication, and iSCSI back-end storage virtualization. See Figure 3-21.




Figure 3-21 Ethernet ports

The `lsportethernet` command displays information about the Ethernet ports on a system. The output of the command includes details about each Ethernet port, such as:

- **Status:** Whether the port is up, down, or experiencing any errors.
- **Speed:** The connection speed of the port (for example, 1Gbps, 10 Gbps).
- **Connected:** Indicates if there is a physical cable connection established on the port.
- **Possible usage:** This might provide clues about how the port is being used, such as "Host Attachment," "iSCSI," or "Replication. See Figure 3-28 on page 36.

If you plan to change the Maximum Transmission Unit (MTU) size of a port, it is only possible if no IP address is configured on that port and its reference port. As shown in Figure 3-22 on page 34, *both ports 2 need to be free of IP addresses in this scenario.*

Restriction: NVMe/TCP and clustering are only supported with an MTU size of 1500 bytes.

To change the MTU size of a port, you can either:

- Select the desired port and then go to **Actions** → **Modify Maximum Transmission Unit**.
- Right-click on the port and choose **Modify Maximum Transmission Unit** from the context menu.

See Figure 3-22 on page 34.




Figure 3-22 Modify MTU size

To configure an IP address on a port, select the port and go to **Actions** → **Manage IP addresses** or use the right mouse button.

On the next screen, select **Add IP address**. See Figure 3-23.




Figure 3-23 Add IP address

On the next screen, you can assign an IP address to the selected port. Click **Add IP Address** to open a configuration window. Here, you can specify the IP address, subnet mask, and other relevant settings. Additionally, you might have the option to add the port to a specific portset for further management. See Figure 3-24 on page 35.




Figure 3-24 Add IP Address menu

Select a portset. In this scenario we use Portset0, the default predefined Portset for Ethernet host attach.

Once the configuration is complete, the assigned IP address will be displayed for the port. You can refer to Figure 3-25 for illustration.




Figure 3-25 Assigned IP address is displayed

You can configure a port with up to four routable IP addresses. However, it is important to use separate VLANs for each IP address to avoid network conflicts.

To manage existing IP addresses on a port, click the **overflow menu (three dots)** located to the right of the IP address entry. This menu allows you to modify, duplicate, or delete the IP address configuration. Refer to Figure 3-26 on page 36 for illustration.




Figure 3-26 Multiple IP addresses

To check your configuration using CLI use the `lsip` command. See Figure 3-27.

```
IBM_FlashSystem:POK-5300:87redbooks>lsip
id node_id node_name port_id portset_id portset_name IP_address prefix vlan gateway owner_id owner_name
0 1 node1 1 72 SystemManagement 129.40.203.30 24 129.40.203.254
1 2 node2 2 0 portset0 192.168.2.2 24
2 1 node1 2 0 portset0 192.168.2.1 24
3 1 node1 3 1 portset1 192.168.3.1 24 10
4 2 node2 3 1 portset1 192.168.3.2 24 10
```

Figure 3-27 `lsip` command

You can use the `mkip` command to add additional IP addresses. With the `chip` command you can modify IP addresses and with the `rmip` command remove IP addresses.

More details are available by the `lsportip` command.

The `lsportethernet` command can be used to display information about Ethernet ports on your system. See Figure 3-28.

```
IBM_FlashSystem:POK-5300:87redbooks>lsportethernet
port_id node_id node_name MAC duplex speed link_state dcbb_state rdma_type adapter_location adapter_port_id host storage replication eth_clustering management
1 1 node1 98:be:94:86:03:31 Full 1Gb/s active unsupported TCP 0 3 no no no no yes
2 1 node1 98:be:94:86:03:2f inactive unsupported TCP 0 1 yes yes yes no yes
3 1 node1 98:be:94:86:03:30 inactive TCP 0 2 yes yes yes no no
1 2 node2 98:be:94:86:00:7a Full 1Gb/s active unsupported TCP 0 3 no no no no yes
2 2 node2 98:be:94:86:00:78 inactive TCP 0 1 yes yes yes no yes
3 2 node2 98:be:94:86:00:79 inactive TCP 0 2 yes yes yes no no
```

Figure 3-28 `lsportethernet`

The `lsportethernet` command can show you whether Data Center Bridging (DCBX) is supported on a port. DCBX can contribute to achieving lossless Ethernet, which is important for some applications.

The current configuration displays only "TCP" as the RDMA type on internal ports. This suggests that other RDMA types like iWARP (RDMA over Converged Ethernet) might not be supported with the existing hardware.

If you plan on using clustering or remote copy functionality that relies on RDMA, you will likely need to add a 4-port 10Gb/s Ethernet card that supports iWARP.

For more information, see [Ethernet connectivity](#).

Limiting access to ports

Sometimes, ports can be used for multiple purposes. For example, in Figure 3-27 on page 36, port 3 is being used with portset 1 for remote mirroring.

In such scenarios, you might want to restrict direct host access to these ports for security reasons. You can achieve this through the Ethernet ports menu. Select the desired port and go to **Actions** → **Modify Host Attachment Support**. Alternatively, right-click on the port and choose **Modify Host Attachment Support** from the context menu. See Figure 3-29.




Figure 3-29 Modify Host Attachment Support

In your current IO group configuration, any changes made to port X will be applied to all nodes within the group. View the current configuration details for these ports within the Ethernet Ports menu. See Figure 3-30.




Figure 3-30 Ethernet ports with no host access to port 3

Modifying storage ports and remote copy settings follow a similar approach. Use the same method to access configuration options, including those specific to remote copy functionality.

3.4.4 Configure protocols

In this section we discuss how to configure protocols.

Configure host access protocol

Define IP addresses and VLANs for network connectivity. Specify the communication protocol used by the host to interact with the storage system in the host definition. See Figure 3-31.




Figure 3-31 Add host

If you have defined multiple host portsets, select the intended portset for a host during configuration using the Advanced option. See Figure 3-32.




Figure 3-32 Define NVMe/TCP host with advanced options

To show and verify iSCSI connections, navigate to **Settings** → **Network** → **Ethernet Connectivity** menu. NVMe connections will not be seen there.

For NVMe connections like NVMe/TCP, use the **Settings** → **Network** → **NVMe Connectivity** menu.

Configure protocol for Remote Copy

The protocol used for Remote Copy is defined during the partnership definition process. See Figure 3-33.




Figure 3-33 Create a partnership

When configuring remote copy protocols over IP, you have two main options:

- ▶ **TCP:** This is a widely supported protocol that works with any standard Ethernet network. However, it may not offer the highest performance for remote copy operations.
- ▶ **RDMA (Remote Direct Memory Access):** This protocol can provide significantly faster data transfer speeds for remote copy compared to TCP.

Note: Remote copy with RDMA will require an Ethernet adapter with RDMA capability and a maximum Round Trip Time (RTT) ≤ 1 ms.

Remote copy with RDMA you will also need a corresponding Portset to define. See Figure 3-34.




Figure 3-34 Create portset for Remote Copy with RDMA

Clustering by using Ethernet connections

Note: Clustering will require an Ethernet adapter with RDMA capability.

Information for Configuring clustering by using Ethernet connections you will find in the "Configuring clustering by using Ethernet connections" in IBM Redbooks *Unleash the Power of Flash: Getting Started with IBM Storage Virtualize Version 8.7 on IBM Storage FlashSystem and IBM SAN Volume Controller*, SG24-8561.

3.5 Troubleshooting

For troubleshooting port configurations, see [Troubleshooting](#).

Also, refer to the “*Troubleshooting chapter*” in IBM Redbooks *Unleash the Power of Flash: Getting Started with IBM Storage Virtualize Version 8.7 on IBM Storage FlashSystem and IBM SAN Volume Controller*, SG24-8561.

For problem resolution, see [Resolving a problem](#).

Related publications

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

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.

- ▶ *Unleash the Power of Flash: Getting Started with IBM Storage Virtualize Version 8.7 on IBM Storage FlashSystem and IBM SAN Volume Controller*, SG24-8561
- ▶ *Ensuring Business Continuity: A Practical Guide to Policy-Based Replication and Policy-Based HA for IBM Storage Virtualize Systems*, SG24-8569

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

Online resources

These websites are also relevant as further information sources:

- ▶ [IBM Storage FlashSystem](#)
- ▶ [IBM SAN Volume Controller information](#)
- ▶ [IBM System Storage Interoperation Center \(SSIC\)](#)

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5734-00

ISBN 0738461784

Printed in U.S.A.

Get connected

