

IBM b-type Gen 8 Installation, Migration, and Best Practices Guide

Vasfi Gucer
Nezih Boyacioglu
Dave Burchwell
Harshit Chaudhary
Mark Detrick
Petar Dimitrov
Silviano Gaona
Kenneth Griffo
David Green
Tony McKenzie
Hartmut Lonzer
Vijay Shenoy
Paul Smith
Marcus Thordal
Tim Werts



Preface

Introduction

The IBM® b-type Gen 8 Fibre Channel directors and switches provide a resilient, scalable, and secure foundation for modern data centers, supporting virtualization, hybrid cloud, and next-generation Non-Volatile Memory Express (NVMe) and flash storage workloads. This IBM Redbooks® publication helps administrators design, implement, and migrate to an IBM b-type Gen 8 storage area network (SAN). It provides an overview of key hardware and software components and explains how to install, monitor, optimize, and troubleshoot SAN environments. By using this book, you can learn best practices for SAN design, along with techniques for managing, monitoring, and maintaining a high-performance storage infrastructure.

Purpose of this document

This document provides guidance for installing, migrating, and implementing best practices for IBM b-type Gen 8 Virtual Fabrics. It covers architecture, configuration, and operational considerations to help ensure reliable and efficient deployments.

Who should read this document

This publication is intended for:

- Storage administrators
- SAN architects
- System engineers
- IBM Business Partners

Readers are expected to have a basic understanding of Fibre Channel networking and IBM storage environments.

Section overview

This publication includes the following sections:

- Introduction to IBM b-type Gen 8
- IBM b-type Gen 8 product review
- Fabric Operating System (FOS) 10.x overview
- IBM SANnav Management Suite overview
- Fibre Channel SAN design and implementation
- Virtual Fabrics overview
- FICON design and implementation
- Unified Storage Fabric mode
- Migration use cases
- SAN security
- Brocade Fabric OS administrative tools and diagnostics
- IBM b-type Gen 8 SAN automation (REST)
- IBM Storage SAN design guidelines and best practices
- Appendix A: Optics and cables

Acknowledgments

Thanks to the following people for their contributions to the project.

Elias Luna, Phil Gerrard, and Andy Benbenek from IBM USA

Rahul Kulkarni and Om Dhumal from IBM India

Thanks to the authors of the previous edition of this book, *IBM b-type Gen 7 Installation, Migration, and Best Practices Guide*, SG24-8497: Brian Larsen, Kameswara Bhaskarabhatla, Matt Brennan, AJ Casamento, Jack Consoli, Mark Detrick, Johnny Fonseca, and Tim Jeka of Broadcom.

Feedback

This document reflects practical experience implementing these technologies in production environments. If you encounter issues not covered here, or if you have suggestions for improving the guidance, your feedback helps make this resource more valuable for the community.

[Contact IBM Redbooks.](#)

Authors

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

Vasfi Gucer

Vasfi Gucer leads projects for the IBM Redbooks team, using his 20+ years of experience in systems management, networking, and software. A prolific writer and global IBM instructor, his focus shifted to storage and cloud computing in the past 8 years. Vasfi holds multiple certifications, including IBM Certified Senior IT Specialist, PMP, ITIL V2 Manager, and ITIL V3 Expert.

Nezih Boyacioglu

Nezih Boyacioglu is a Storage Area Network (SAN) and enterprise storage specialist with more than 23 years of industry experience. He began his career working with IBM Spectrum® Protect (formerly IBM Tivoli® Storage Manager) and tape storage systems before specializing in IBM Storage Virtualize solutions, including IBM SAN Volume Controller, IBM Storage Virtualize, and IBM FlashSystem®. He also has extensive experience with IBM Power® servers and SAN infrastructure. Nezih currently works as an independent consultant and serves as Tech Solutions Director at Istanbul Pazarlama. He is a co-author of several IBM Redbooks® publications and holds multiple certifications in enterprise storage support and IBM Storage Virtualize software.

Dave Burchwell

Dave Burchwell is a Field Application Engineer (FAE) and SAN Specialist at Broadcom, based in San Francisco, California, where he supports the Brocade Fibre Channel products offered within IBM's b-type storage networking portfolio. His high-tech career began in 1977 in the U.S. Army, supporting missile batteries and command-and-control data centers. After joining ITT Courier/Servcom in 1980, he dedicated his career to deploying and supporting Storage Area Networks in raised-floor mainframe data centers. Dave joined Data Switch Corporation and served in numerous technical roles as the organization evolved through General Signal Networks, Inrange Technologies, CNT, McDATA, Brocade, and ultimately Broadcom. During this tenure, he was instrumental in achieving IBM FICON certification for the company's SAN directors. Today, he continues to leverage his decades of expertise to help OEMs and global customers deploy highly resilient FICON/FC SANs.

Harshit Chaudhary

Harshit Chaudhary is a Delivery Consultant at IBM Expert Labs with over 11 years of experience in enterprise data infrastructure and Hybrid Cloud solutions. He specializes in designing, architecting, and implementing high-performance storage environments using IBM Storage Virtualize, IBM Enterprise Storage systems, and SAN technologies. He partners with IBM clients and Business Partners to deliver complex storage solutions, business-critical data migrations, and resilient disaster recovery architectures. He holds a Bachelor of Technology (B.Tech.) in Information Technology and maintains an extensive portfolio of advanced industry certifications across hyperscale cloud platforms, enterprise storage and storage networking technologies.

Mark Detrick

Mark Detrick is a Principal Solutions Architect who has worked for Broadcom for 25 years. Involved with many Fortune 500 large-scale IBM projects worldwide, Mark plays a consultative role in storage network design and deployment. Mark is an expert in many data center technologies, including mainframe SAN design and extension; distributed systems large-scale SAN design, FC Routing, and extension; IP routing and Ethernet switching; and WAN technologies. Mark has over 30 years of experience in data center networking environments. He understands Broadcom's ASIC technology and is a consultative resource within Brocade and externally to customers, OEMs, and resellers. Mark applies his in-depth technical knowledge of fabrics, protocols, flow control, TCP/IP, and large-scale SAN/LAN architectures. Mark is adept at next-generation data center solutions, including cloud and the many types of virtualizations found in data centers.

Petar Dimitrov

Petar Dimitrov is a SAN Switch and SAN Central Subject Matter Expert (SME) for the IBM AG region. Based in Sofia, Bulgaria, he has more than 10 years of experience at IBM, supporting b-type (Brocade) and c-type (Cisco) Fibre Channel switches, as well as the IBM Storage Virtualize family of products.

Kenneth Griffo

Kenneth Griffo is the Worldwide Sales Lead for IBM's Fibre Channel Storage Networking business. Prior to that Ken was a Client Systems Manager for key clients of IBM. Throughout his 35+ year professional career, Ken has served as a field salesperson, sales manager, Storage Product Management Program Director, and worldwide sales leader, focused on growing IBM's storage business since 1996. Ken began his career with IBM back in 1987 as a professional hire from EDS Corporation. He has a B.S. in Mathematics from the University of Notre Dame and MBA from The University of Michigan.

Silviano Gaona

Silviano Gaona is a Field Application Engineer for Brocade, a Broadcom Company. Based in Saratoga Springs, Utah, Silviano has more than 30 years of experience as a Field Application Engineer, System Engineer, education developer and trainer, and Customer Support Engineer supporting a wide range of Storage Area Network solutions. Silviano has an educational background in Computer Information Systems.

David Green

David Green works with the IBM SAN Central team to troubleshoot performance and other problems on storage networks. He authored or contributed to several IBM Redbooks publications. He is a regular speaker at IBM Technical University.

Tony McKenzie

Tony McKenzie is a Senior Field Applications Engineer for Broadcom based in Holly Springs, North Carolina, focusing on Brocade Fibre Channel products and Emulex fibre channel host bus adapters. Tony has been in the storage industry for 21 years as a SAN Product Field Engineer, Systems Engineer, and Field Applications Engineer. Tony is also a Brocade Distinguished Architect.

Hartmut Lonzer

Hartmut Lonzer brings 49 years of technical and sales experience at IBM to his role as a Brand Technical Storage Specialist for the DACH region. Based at IBM's headquarters in Ehningen, Germany, he specializes in the IBM FlashSystem family and IBM SAN Volume Controller products, having worked with these solutions since their introduction. His expertise in IBM Fibre Channel SAN technologies dates back to the ESCON era, giving him decades of experience in enterprise storage networking.

Vijay Shenoy

Vijay Shenoy is a Storage Area Network (SAN) Subject Matter Expert (SME) and a member of IBM's SAN Central team, based in Research Triangle Park (RTP), North Carolina. With more than 20 years of experience in enterprise SAN storage technologies, Vijay has provided advanced technical and solution-level support for a wide range of IBM SAN hardware products. His expertise spans SAN architecture, implementation, troubleshooting, and customer support. Vijay holds a degree in Computer Science.

Paul Smith

Paul Smith is a Storage Consultant working in the IBM Technology Expert Labs team. He specializes in IBM FlashSystem, DS8K and providing design and implementations for B & C type SAN technologies. Regularly working with IBM clients to deliver Open and zSystems storage solutions and providing detailed knowledge of data migration techniques across different business critical architectures. Based in Sutton Coldfield, United Kingdom, Paul has more than 32 years of experience as a field engineer and remote technical support specialist, supporting a wide range of IBM hardware and software solutions. Paul's educational background is in electronics and internet technology.

Marcus Thordal

Marcus Thordal is a Product Manager for Brocade, a Broadcom company, based in San Jose, California. Leveraging more than 25 years of industry experience, his product management responsibilities include Brocade Fabric OS and Support Link. Marcus holds a degree in Chemical Engineering.

Tim Werts

Tim Werts is a Global Marketing Manager for Brocade, a Broadcom Company. Based in Columbia, South Carolina, USA, Tim has been in the storage industry since 1987 as a tape and disk qualification, SAN System Engineer, Fibre Channel Business Development, and now Technical Marketing. Tim holds a degree in Electrical Engineering.

IBM b-type Gen 8 platforms overview

As technology evolves, IT organizations are held to an even higher set of standards, driven by the rapid innovation of Enterprise AI, massive data growth, and highly consolidated workloads. The only way to accelerate IT delivery and keep pace with these ever-increasing demands is to transition from traditional rule-based automation to real-time, AI-

powered autonomy. The storage network must not only monitor behavior, but also understand, learn from, and dynamically adapt to infrastructure changes.

Brocade Gen 8 technology delivers a comprehensive set of capabilities that integrate embedded SAN AI technology directly into the fabric architecture. This 128 Gbps performance supports enterprise AI workloads and modern workloads while embedded SAN AI technology automates infrastructure management, helping streamline operations and improve efficiency.

This breakthrough technology transforms millions of fabric data points into real-time automated decisions, serving as the foundation for an advanced self-learning, self-optimizing, and self-healing autonomous SAN.

The combination of these features gives IT administrators unmatched operational visibility across hosts, virtual machines (VMs), and storage arrays, establishing a completely self-governing infrastructure that safeguards application uptime without human intervention.

Self-learning

Brocade Gen 8 technology modernizes SAN management by introducing AI-powered autonomy that continuously learns, adapts, and responds to changing environments. Rather than relying solely on manual threshold configurations, the fabric uses integrated telemetry and machine-learning algorithms to map and understand behavioral patterns across the entire storage ecosystem:

- **SAN Fabric Intelligence (SAN FI):** Natively provides a unified framework that simplifies ecosystem monitoring by consolidating server, VM, and storage data into customizable, pattern-based dashboards. By delivering summarized telemetry from MAPS, Flow Vision, port statistics, and congestion analytics, it introduces streamlined troubleshooting workflows that significantly reduce the mean time to resolution (MTTR). Furthermore, SAN FI enables end-to-end path enumeration, allowing administrators to easily distinguish between device and fabric latencies while gaining visibility into microburst congestion episodes.
- **Automated Fabric Analytics:** Processes millions of real-time telemetry data-points to automatically baseline normal fabric behavior. When abnormal deviations or early-stage performance anomalies emerge, the integrated SAN AI engine immediately isolates the root cause and flags the exact points of interest across switches, hosts, and storage arrays without requiring human correlation.
- **VM Insight / VMID+:** Enables direct virtual machine visibility and monitoring across the SAN fabric, removing any operational dependency for native VMID support on the downstream storage arrays.
- **Fabric Performance Impact (FPI) Monitoring:** Interoperates with the Monitoring and Alerting Policy Suite (MAPS) to deliver real-time visibility into link performance, automatically detecting and pinpointing exactly which devices are causing or are impacted by structural congestion.

Self-optimizing

By using the real-time actionable intelligence gathered from its self-learning capabilities, Brocade Gen 8 Fibre Channel SANs automatically execute dynamic performance tuning. This ensures that high-priority application flows remain completely unimpeded, even when running mixed AI, OLTP, backup, and batch workloads concurrently on the same infrastructure. Self-optimizing includes the following features:

- **Adaptive Traffic Optimizer:** It is a built-in Gen 8 innovation that autonomously eliminates congestion to guarantee application performance. Enabled by default without configuration, ATO segregates traffic into distinct Performance Groups based on protocol (separating NVMe from SCSI) and destination port speeds. Through continuous periodic checks, it dynamically provisions up to 23 Performance Groups across Virtual Channels, allocating bandwidth only to active traffic types (such as prioritizing 64G SCSI when it dominates). By using independent flow control and optimal buffer allocation, this self-optimizing engine prevents slower devices from bottlenecking high-speed workloads, lowering I/O response times without manual intervention.
- **Dynamic Path Selection:** This feature optimizes fabric utilization by distributing input traffic across multiple paths—such as ISLs, ICLs, or trunk groups—in ratios proportional to their available bandwidth. By balancing data loads based on path capacity, it mitigates potential network congestion. When fabric topology changes alter path availability, DPS dynamically and non-disruptively redistributes the traffic using an exchange-based routing policy.

Self-healing

When anomalies or potential infrastructure risks are discovered, Brocade Gen 8 executes immediate, real-time remediation to protect data delivery and application availability. It acts directly inline at the earliest warning signs, triggering trusted self-correction mechanisms faster than humanly possible:

- **Dual Forward Error Correction (FEC):** Operates natively at the physical layer to actively detect, capture, and correct optical link bit errors in real time. Running concurrently with ultra-low port switching latencies (580 ns), Dual FEC repairs data transmission errors at the microsecond level before a frame can be dropped or corrupted.
- **AI-Driven Inline Mitigation and Signaling:** Automatically acts on real-time telemetry to neutralize anomalies. Through Fabric Performance Impact Notifications (FPIN), the fabric instantly alerts connected end devices and host adapters when a path develops degradation, enabling immediate traffic adjustments.
- **Port Decommissioning and Multipathing Integration:** This feature works with the host's multipath I/O (MPIO) drivers via Fabric Performance Impact Notifications (FPIN). When a link degrades, the fabric sends an FPIN alerting the host to instantly decommission and bypass the unhealthy path. This ensures traffic is seamlessly rerouted over healthy physical paths before application availability is impacted.
- **Hardware Congestion Signaling:** Senses microsecond-level congestion backpressure at the center of the data path and instantly notifies physical end devices, allowing for the dynamic mitigation of fabric congestion impacts before they cascade.
- **Autonomous Error Recovery:** Works entirely behind the scenes to perform real-time error corrections, link resets, and hardware-level credit recoveries without manual intervention, significantly reducing the time administrators spend on active troubleshooting and minimizing operational risk.

Quantum-safe security: Hardening your SAN against evolving risk

As enterprises extend application infrastructures to support distributed enterprise AI architectures, attack surfaces expand geometrically. Furthermore, the rapid emergence of quantum computing threatens to render current cryptographic standards completely obsolete. Bad actors are actively executing "harvest now, decrypt later" strategies—intercepting and storing encrypted data today with the intent of decrypting it when quantum processors become widely viable.

Brocade Gen 8 mitigates this risk by embedding advanced security controls directly into the fabric architecture, effectively hardening storage environments around the clock:

- **Post-Quantum Cryptography (PQC):** Gen 8 introduces quantum-resistant encryption algorithms designed to withstand future quantum-enabled attacks. By pairing these cryptographic protocols with 256-bit in-flight encryption on switch-to-switch links, the fabric provides confidentiality for mission-critical storage traffic.
- **Principle of Least Privilege Architecture:** Fabric OS is hardened to strictly minimize internal and external attack surfaces. By using mandatory role-based access controls, administrative actions are isolated to narrowly scoped operational profiles, eliminating superfluous internal access points and preventing unauthorized fabric manipulation.
- **Hardware-Rooted Trust Anchors:** Gen 8 uses advanced anti-tampering technology and secure boot mechanisms to continuously validate firmware and hardware integrity. By validating cryptographic roots of trust, the system blocks the installation or execution of unverified, malicious software.

IBM b-type Gen 8 platforms and naming conventions

The following figure shows the IBM b-type Gen 8 platforms: the IBM SAN512B-8 8-slot director, the IBM SAN256B-8 4-slot director, and the IBM SAN56B-8 switch.




Figure: IBM b-type Gen 8 platforms

The IBM b-type Gen 8 and Brocade naming conventions for the Gen 8 platform offerings are listed in the following table.

IBM b-type Gen 8 name	IBM machine type and model	Brocade name
IBM Storage Networking SAN512B-8	8635-F88	X8-8
IBM Storage Networking SAN256B-8	8635-F84	X8-4
IBM Storage Networking SAN56B-8	8634-F56	G820

IBM b-type Gen 8 product review

This section provides a hardware review of each Gen 8 director, supported blades, and mid-range switch. New to the Gen 8 offering is the ability to upgrade from a Gen 7 Director to a Gen 8-supported director.

The last section reviews the supported transceiver modules for the Gen 8 platforms.

IBM b-type Gen 8 SAN director overview

The IBM b-type Gen 8 Director is a modular platform, purpose-built to power and secure large-scale storage environments. This director provides a stable, scalable, and high-performance foundation that supports growth, workload consolidation, and reliable operations, making it well-suited for mission-critical enterprise and AI workloads. With the integration of new technologies that are accelerating the delivery of data and services, the network must evolve to keep pace with innovations in storage and demands of modern applications to maximize the value of investments in next-generation data center infrastructure.

With up to 384 × 128G ports, the IBM b-type Gen 8 Director delivers exceptional bandwidth and throughput required to support growing numbers of devices, applications, and workloads without compromising performance. High bandwidth and ultra-low latency eliminate I/O bottlenecks, delivering maximum performance for high-transaction workloads and next-generation storage. This performance headroom enables organizations to consolidate workloads onto fewer systems. To further maximize utilization, up to 128 UltraScale ICL links allow the director to scale efficiently with fewer chassis by preserving device ports.

The IBM b-type Gen 8 Director offers flexible deployment options, multiprotocol support, and mixed-blade capability, enabling organizations to adapt their infrastructure and optimize resources to meet the evolving requirements of next-generation storage and servers.

As workloads become more demanding, cyber threats become more sophisticated, and the spectrum of threats expands, it is essential to fortify the storage network against hidden risks. The network must deliver optimal performance, operate autonomously, streamline management, and protect critical data.

The IBM b-type Gen 8 Director with Gen 8 technology combines 128G performance, quantum-safe security, and AI-powered autonomy to enable the most robust foundation for modern data center architecture. Integrated security featuring quantum-resistant encryption safeguards SAN fabrics against cybersecurity threats in the era of quantum computing. Embedded SAN AI technology modernizes SAN management, enabling the SAN to operate autonomously and efficiently.

By hardening the SAN against evolving security risks and enabling AI-powered autonomy to learn, adapt, and respond to changing environments, IBM b-type Gen 8 provides a secure, intelligent, and high-performance storage networking platform. This enables faster decision-making, improves operational efficiency, and maintains high levels of resiliency.

The IBM b-type Gen 8 Director is engineered to meet continuous data growth and demanding application requirements in large-scale Fibre Channel and IP storage environments. Its modular design delivers the capacity, throughput, and resiliency needed for all-flash and NVMe workloads, while enabling flexible, scalable SAN architectures.

IBM b-type Gen 8 SAN512B-8 (8635-F88) director overview

Designed to meet continuous data growth and critical application demands, the IBM b-type Gen 8 Director is purpose-built to power large-scale storage environments that require increased capacity, greater throughput, and higher levels of resiliency and operational efficiency. This modular building block enables organizations to build the highest-performing data center SAN fabrics that are required for all-flash and NVMe storage environments.

The IBM b-type Gen 8 Director modular design provides flexibility with two customizable chassis that can scale on-demand to accommodate more devices, applications, and workloads. Both chassis use IBM b-type Gen 8 UltraScale ICL technology to scale out modular SANs while preserving ports for device connectivity. This technology also enables a flexible SAN design that supports core-edge or mesh topologies.

The 14U IBM b-type Gen 8 SAN512B-8 Director is built for large enterprise networks and has eight horizontal blade slots to provide up to 384 128 Gbps line rate ports for device connectivity. Another 16 UltraScale InterChassis Link (ICL) connections provide the equivalent of 128 ports for chassis-to-chassis interconnect.

The 9U IBM b-type Gen 8-4 SAN256B-8 Director is built for mid-size networks and has four horizontal blade slots to provide up to 192 128 Gbps line rate ports for device connectivity. An additional eight UltraScale ICL connections provide the equivalent of 64 ports for chassis-to-chassis interconnect. Each blade slot within the IBM b-type Gen 8 chassis can be populated with an optional port or extension blade. For device connectivity, the following blade is available:

- The IBM b-type Gen 8 FC128-48 Fibre Channel port blade provides 48 front-end 128 Gbps Fibre Channel ports with backward-compatible support for 16, 32, 64, and 128 Gbps Fibre Channel connectivity. Ports 24 through 47 are Ethernet-capable Flexports that can operate either as Fibre Channel ports or as Ethernet ports supporting 10, 25, 50, and 100 GbE connectivity. When configured for Ethernet operation, these ports support Universal Storage Fabric (USF) services through an IPS logical switch. The FC128-48 port blade does not support Fibre Channel over Ethernet (FCoE). In the IBM b-type Gen 8 SAN512B-8 Director, up to eight slots (slots 1-4 and 9-12) can be populated with port blades or extension blades, enabling scalable Fibre Channel and extension services within a single chassis.

IBM b-type Gen 8 Directors build upon years of innovation and use the core technology of IBM b-type Gen 8 systems to consistently deliver five nines of availability in the world's most demanding data centers. Delivering nondisruptive

software upgrades, hot-pluggable components, and a no-single-point-of-failure design, the IBM b-type Gen 8 offers a highly resilient solution for today's enterprise-class storage environments.

Notes:

- The operational port speed is determined through negotiation between both ends of the link and is limited by the highest common speed supported and enabled on the connected ports and installed transceivers.
- The WWN cards have been replaced with a Chassis Identifier (CID) that contains the chassis WWNs. The CID contains redundant WWN EEPROMs (SEEPROMs). The Chassis Identifier (CID) is not a field-replaceable unit (FRU) and is installed during the assembly of the chassis.

This device includes the following key product features:

- Redundant and hot-swappable SFP+ and Gen 8 ICL OSFP (Octal Small Form-factor Pluggable) transceivers; port, control processor (CP), and core routing (CR) blades; power supply assemblies, and fan assemblies that enable a high-availability platform and allow nondisruptive software upgrades for mission-critical SAN applications.
- Support for the following features when using the FC128-48 blade:
 - Port blade with 48 front-end 128G Fibre Channel ports for attachment of storage devices, servers, and SAN switches. The ports connect through the director core blades to provide high-performance, low-latency fabric connectivity.
 - Supports 64G and 128G secure optics only. See this web page.
 - Each SFP+ FC port supports 16, 32, 64, and 128G FC speeds.
 - A 128G optical transceiver can auto-negotiate to 128G, 64G, or 32G.
 - A 64G optical transceiver can auto-negotiate to 64G, 32G, or 16G.
 - Each SFP+ Ethernet port supports 10, 25, 50, and 100G Ethernet speeds.
 - A 128G optical transceiver can be configured to 100GbE, 50GbE, or 25GbE.
 - A 64G optical transceiver can be configured to be 50GbE, 25GbE, or 10GbE.
 - Brocade Trunking (BT) technology groups up to eight ports to create high-performance 1024G ISL trunks between switches using 128G ports.
- Support for the following features when using the ICLX8-8 blade:
 - Up to eight ICL ports on each ICLX8-8 blade provide 64 OSFP links.
 - An OSFP optical transceiver can operate at Gen 8 or Gen 7 ICL port speeds.
- Universal ports that automatically operate as `E_Ports` or `F_Ports` as required and can be configured as `EX_Ports`, `M_Ports` (mirror ports), and FICON ports.
- ClearLink™ Diagnostic port (`D_Port`) functions on FC ports.
- Data compression and encryption capabilities through the port blades when ports are configured as ISLs.

Notes:

- Device control processors and management modules include batteries for RTC/NVRAM backup. Do not attempt to replace these batteries. Dispose of hardware components that contain these batteries as required by local ordinances and regulations.

- The CID is not a FRU and is installed during the assembly of the chassis.

Hardware components

The device has a modular and scalable mechanical construction that allows a wide range of flexibility in installation, fabric design, and maintenance. The device can be mounted with the cables facing the front or the rear of the equipment rack. It consists of the following hardware components:

- Up to eight slots for hot-swappable port blade assemblies, providing the following Fibre Channel port support based on the port blade:
 - FC128-48 blades provide up to 384 128G FC ports. FC128-48 blades can also support up to 192 100G Ethernet ports, which reduces the FC port count.
 - For a list of supported transceivers for these blades, see this web page.
- Two half-size slots for control processor (CP) blades:
 - A single active CP blade can control all the ports in the device.
 - The standby CP blade assumes control of the device if the active CP blade fails.
- Two slots on the nonport-side for core routing (CR) blades:
 - The CR blade interconnects all port blades.
 - CR blades now mount internally behind fan assemblies.
 - ICL connectivity via connected ICL port blades.
 - Both CR blades are active and can be hot-swapped.
- Two half-size slots for ICL blades:
 - Each ICLX8-8 blade has 64 front-end ports that are mapped to the eight Gen 8 ICL OSFP ports.
 - ICL ports support Gen 8 and Gen 7 ICL operating speeds.
 - ICL ports allow interconnection with neighboring director chassis.
 - Both ICL blades are active and can be hot-swapped.
 - For a list of supported transceivers for these blades, see this web page.
- Modular, hot-swappable field-replaceable units (FRUs):
 - Three fan assemblies, available with Non-Portside Intake (NPI) or Non-Portside Exhaust (NPE) airflow configurations to support different data center cooling requirements.
 - Up to six 3000W power supplies are located on the port side of the device, available with the NPI or NPE airflow.
 - Redundant primary power connections ensure high availability. Each power supply assembly has its own connector, so the number of primary power connections is three for optimum efficiency and redundancy.
 - Port blades use SFP+ optical transceivers.
 - ICL blades use Gen 8 ICL OSFP optical transceivers.
- One Chassis Identifier (CID) located on non-port side of the device behind fan assembly 1.
- A cable management comb. These combs are installed on the chassis below the blades for cable management.

Port-side view of the device

The following figure shows the port-side view of the IBM SAN512B-8 Director with installed blades. The following components are identified:

1. Port Blades (FC128-48)

2. ICL Blades (ICLX8-8)
3. Control Processor Blades (CPX8)
4. Rack ears with slot number labels
5. Power supply assemblies
6. 2AWG Panduit LCD2-14AF lug for building ground connection




Figure: Port-side view of IBM SAN512B-8 Director

Although not shown in the previous figure, the chassis label that contains the serial number, SKU, and WWN is located on the top front of the chassis.

Port-side slot numbering

The IBM SAN512B-8 Director contains eight full-width slots and four half-width slots, for a total of 12 slots. Facing the port side of the device, from the top, the full-width slots are numbered 1 through 4. The half-width slots are in the middle, numbered 5 (left side top half-width slot), 6 (left side bottom half-width slot), 7 (right side top half-width slot), and 8 (right side bottom half-width slot). The remaining full-width slots are numbered 9 through 12, counting from top to bottom of the chassis. Slots contain guide pins and connectors designed for specific blade types.

Only install the control processor (CP), ICL, and port blades into slot numbers as follows:

- Slots 1–4 and slots 9–12 are restricted to port blades.
- Slots 5–6 are restricted to ICLX8-8 blades.
- Slots 7 and 8 are restricted to CP blades. The blade installed in slot 7 is designated `CP0`, and the blade in slot 8 is designated `CP1` in CLI command and message output.

Note: Depending on the fan assemblies that are installed, airflow can be configured as Non-Portside Exhaust (NPE) or Non-Portside Intake (NPI). In an NPE configuration, airflow moves from the port side to the non-port side of the chassis. In an NPI configuration, airflow moves from the non-port side to the port side of the chassis.

Non-port side view of the device

The following figure shows the non-port-side view of the IBM SAN512B-8 Director with all fan and Core Routing blades installed. The following components are identified:

1. Fan assemblies
2. Core Routing blades (CR128-8)




Figure: Non-port side of the IBM SAN512B-8 Director (example configuration)

More information

All documentation for the IBM b-type Gen 8 SAN512B-8 (8635-F88) is available at [this web page](#).

IBM b-type Gen 8 SAN256B-8 (8635-F84) director overview

Note: The WWN cards have been replaced with a Chassis Identifier (CID) that contains the chassis WWNs. The CID contains redundant WWN EEPROMs (SEEPROMs). The Chassis Identifier (CID) is not a field-replaceable unit (FRU) and is installed during the assembly of the chassis.

This device includes the following key product features:

- Redundant and hot-swappable SFP+ and Gen 8 ICL OSFP transceivers; port, control processor (CP), and core routing (CR) blades; power supply assemblies, and fan assemblies that enable a high-availability platform and allow nondisruptive software upgrades for mission-critical SAN applications.

- Support for the following features when using the FC128-48 blade:
 - Port blade with 48 front-end 128G Fibre Channel ports for connecting storage arrays, servers, and SAN switches. All ports connect through the director core blades to provide high-performance fabric connectivity. Ports 24–47 are Flexports that can operate as either Fibre Channel ports or Ethernet ports supporting speeds up to 100GbE. When configured for Ethernet operation, these ports support Universal Storage Fabric (USF) services.
 - Supports 64G and 128G secure optics only. See this web page.
 - Each SFP+ FC port supports 16, 32, 64, and 128G FC speeds.
 - A 128G optical transceiver can auto-negotiate to 128G, 64G, or 32G.
 - A 64G optical transceiver can auto-negotiate to 64G, 32G, or 16G.
 - Each SFP+ Ethernet port supports 10, 25, 50, and 100G Ethernet speeds.
 - A 128G optical transceiver can be configured to 100GbE, 50GbE, or 25GbE.
 - A 64G optical transceiver can be configured to 50GbE, 25GbE, or 10GbE.
 - Brocade Trunking technology groups up to eight ports to create high-performance 1024G ISL trunks between switches using 128G ports.
- Support for the following features when using the ICLX8-4 blade:
 - Up to four ICL ports on each ICLX8-4 blade, providing 32 OSFP links.
 - An OSFP optical transceiver can operate at Gen 8 or Gen 7 ICL port speeds.
- Universal ports that can be configured as `E_Ports`, `F_Ports`, `EX_Ports`, `M_Ports` (mirror ports), and FICON ports.
- ClearLink Diagnostic port (`D_Port`) functions on FC ports.
- Data compression and post-quantum cryptography (PQC) encryption capabilities through the port blades when ports are configured as ISLs.

Note: The CID is not a FRU and is installed during the assembly of the chassis.

Hardware components

The device has a modular and scalable mechanical construction that allows a wide range of flexibility in installation, fabric design, and maintenance. The device can be mounted with the cables facing the front or the rear of the equipment rack. It consists of the following hardware components:

- Up to four slots for hot-swappable port blade assemblies, providing the following Fibre Channel port support based on the port blade:
 - FC128-48 blades provide up to 192 128G FC ports. FC128-48 port blades can also support up to 96 100G Ethernet ports, which reduces the FC port count.
 - For a list of supported transceivers for these blades, see this web page.
- Two half-size slots for control processor (CP) blades:
 - A single active CP blade can control all the ports in the device.
 - The standby CP blade assumes control of the device if the active CP blade fails.
- Two slots on the nonport-side for core routing (CR) blades:
 - The CR blade interconnects all port blades.
 - CR blades now mount internally behind fan assemblies.

- ICL connectivity via connected ICL port blades.
- Both CR blades are active and can be hot-swapped.
- Two half-size slots for ICL blades:
 - Each ICLX8-8 blade has 32 front-end ports that are mapped to the four Gen 8 ICL OSFP ports.
 - ICL ports support Gen 8 and Gen 7 ICL operating speeds.
 - ICL ports allow interconnection with neighboring director chassis.
 - Both ICL blades are active and can be hot-swapped.
 - For a list of supported transceivers for these blades, see this web page.
- Modular, hot-swappable field-replaceable units (FRUs):
 - Three fan assemblies, available with NPI or NPE airflow.
 - Up to four 3000W power supplies are located on the port side of the device, available with the NPI or NPE airflow.
 - Redundant primary power connections ensure high availability. Each power supply assembly has its own connector, so the number of primary power connections is two for optimum efficiency and redundancy.
 - Port blades use SFP+ optical transceivers.
 - ICL blades use Gen 8 ICL OSFP optical transceivers.
- One Chassis Identifier (CID) located on non-port side of the device behind fan assembly 1.
- A cable management comb. The comb is installed on the chassis along the sides of the blades for cable management.

Note: Device control processors and management modules include batteries for RTC/NVRAM backup. Do not attempt to replace these batteries. Dispose of hardware components that contain these batteries as required by local ordinances and regulations.

See the following documentation for additional specifications:

- The "Power Supply Specifications (per PSU)" in the [IBM b-type Gen 8 Director Technical Specifications](#) for maximum output power, input voltage, input line frequency, and other specifications for your power supply model.
- The "Power Supply Requirements" section in the [IBM b-type Gen 8 Director Technical Specifications](#) for the minimum number of power supplies required for operation and redundancy when different input voltages are used, such as low line and high line AC.
- The "Power Consumption" sections in the [IBM b-type Gen 8 Director Technical Specifications](#) for power output data and the minimum number of power supplies for supported input voltages.

Redundant primary power connections ensure high availability. Each power supply assembly has its own connector; therefore, the number of primary power connections is two for optimum efficiency and redundancy.

The device includes the following hardware requirements:

- Port blades use small form-factor pluggable (SFP+) optical transceivers.
- Core routing blades use OSFP optical transceivers. For a list of supported transceivers for these blades, see this web page.
- Chassis door, which must be installed to meet EMI compliance certification.
- Two vertical cable management finger assemblies. The comb is installed on the chassis along the sides of the blades for cable management.

For more information about the IBM b-type Gen 8 SAN256B-8 (8635-F84), see the [IBM Storage Networking SAN256B-8](#)

Note: When non-port-side exhaust (NPE) fan and power supply assemblies are installed, air flows through the side vent and exhausts to the non-port side. When non-port-side intake (NPI) fan and power supply assemblies are installed, air flows from the non-port-side and exhausts through the side vent.

When installing the chassis in a four-post rack with the airflow diversion rack mount kit, airflow is redirected to or from the port side of the chassis. Therefore, if NPE fan and power supply assemblies are installed, air flows into port-side air vents and exhausts to the non-port side. If NPI fan and power supply assemblies are installed, air flows into non-port-side air vents and exhausts through to the port-side.

Port-side view of the device

The following figure shows the port-side view of the IBM SAN256B-8 Director with installed blades.




Figure: Port-side view of the IBM SAN256B-8 Director (example configuration)

Although not shown in the previous figure, the chassis label that contains the serial number, SKU, and WWN is located on the top front of the chassis.

The following components are identified:

1. Port blades (FC128-48)
2. ICL blades (ICLX8-4)
3. Control Processor blades (CPX8)
4. Rack ears with slot number labels
5. Power supply assemblies
6. 2AWG Panduit LCD2-14AF lug for building ground connection

Port-side slot numbering

The IBM SAN256B-8 Director contains four full-width slots and four half-width slots, for a total of eight slots. Facing the port side of the device, from the top, the full-width slots are numbered 1 and 2. The half-width slots are in the middle, numbered 3 (left side top half-width slot), 4 (left side bottom half-width slot), 5 (right side top half-width slot), and 6 (right side bottom half-width slot). The remaining full-width slots are numbered 7 and 8 counting from top to bottom of the chassis.

Slots contain guide pins and connectors that are designed for specific blade types. Only install the control processor (CP), ICL, and port blades into slot numbers as follows:

- Slots 1–2 and slots 7–8 are restricted to port blades.
- Slots 3–4 are restricted to ICLX8-4 blades.
- Slots 5–6 are restricted to CP blades. The blade installed in slot 5 is designated as **CP0**, and the blade in slot 6 is designated as **CP1** in CLI command and message output.

Non-port-side view of the device

The following figure shows the non-port-side view of the IBM SAN256B-8 Director with all fan and power supply assemblies installed.




Figure: Non-port side of the IBM SAN256B-8 Director (example configuration)

1. Fan assemblies
2. Core routing blades (CR128-4)

Note: The core routing blades are installed behind the fan assemblies and appear as slot 9 and 10 in `slotshow` output.

Use the Chassis Airflow Diversion and Port Side Exhaust Kit to divert airflow so that air fully exhausts to the port side of the chassis or is fully drawn from the port side of the chassis while mounted in a four-post rack.

More information

For more information about the IBM b-type Gen 8 SAN256B-8 (8635-F84), see this web page.

IBM b-type Gen 8 SAN56B-8 (8634-F56) switch overview

Note: The IBM SAN56B-8 switch is secure-booted.

The IBM SAN56B-8 switch offers the following features and capabilities:

- Enterprise-class 56-port Gen 8 Fibre Channel switch supported with Fabric OS (FOS) `v10.0.0` and later that offers 56 × 128 Gbps SFP+ ports in a 1U form-factor.
- Each of the 56 SFP+ ports supports 16, 32, 64, and 128 Gbps Fibre Channel speeds.
- 128 Gbps auto-sensing Fibre Channel switch and router ports:
 - A 128 Gbps optical transceiver can auto-negotiate to 128 Gbps, 64 Gbps, or 32 Gbps.
 - A 64 Gbps optical transceiver can auto-negotiate to 64 Gbps, 32 Gbps, or 16 Gbps.
- Dynamic Ports on Demand (Dynamic-POD) scaling from a base configuration of 24 ports to 56 ports using four 8-port POD licenses.
- Universal ports self-configure as `E_Ports`, `F_Ports`, `N_Ports`, or `D_Ports`. `EX_Ports` can be activated on a per-port basis with the optional Integrated Routing license. A Diagnostic Port (`D_Port`) provides diagnostics, troubleshooting, and verification services for the physical media.
- In-flight 128 Gbps data compression and encryption provide efficient link utilization and security. The number of ports that can be enabled with compression and encryption is listed in the following table.

Note: The operational port speed is determined by the highest speed supported by the installed transceivers and permitted by the port configuration at both ends of the link.

Table: IBM SAN56B-8 encryption and compression: number of supported ports

Feature	SAN56B-8 8634-F56 (Switch Type 205)	FC128-48 (Blade ID 237)
Number of front-end port groups (ASICs)	1 Port group	1 Port group
Number of ports per ASIC/port group that can be enabled with encryption or compression at any speed	32 Ports	32 Ports
Total number of ports per switch/blade that can be enabled with encryption or compression (row 1 × row 2)	32 Ports	32 Ports

Note: All the ports are connected to a single ASIC.

- Support for NPI and NPE airflow configurations to accommodate different data center cooling requirements.
- Hardware-enabled input and output (I/O) latency statistics collection.
- Hardware-enabled VM support.
- IBM 128G SFP+ optical transceivers support short wavelength (SWL) optical media among the switch ports.
- IBM 64G SFP+ optical transceivers support any combination of short wavelength (SWL), long wavelength (LWL), and extended long wavelength (ELWL) optical media among the switch ports.
- High-performance C3338R processor with two cores operating at 1.5 GHz delivers high performance, scalability, and advanced Fabric Vision functionality.
- One 1000/100 Mbps RJ-45 connector for the Ethernet management connection. With EZSwitchSetup, this port supports switch IP address discovery and configuration, which eliminates the need to attach a serial cable to configure the switch IP address.
- One internal e-USB module provides 16 GB of persistent storage for increased serviceability, and error logging functionality.
- One USB Type-A (External) port for firmware upgrades, log downloads, and configuration backups.
- One Mini-USB (External) serial console management operating at 115200 b/s.
- Two hot-swappable redundant integrated 650W AC power supplies and fan assembly field-replaceable units.
- A total of 56 hot-pluggable SFP+ optical transceiver slots.
- A total of 56 bicolor LEDs (green/amber) to indicate the status of each SFP+ port.
- One green LED to indicate valid system power.
- One bicolor (green/amber) LED to indicate the system status.
- Two Ethernet LEDs: one green LED to indicate link speed of 1000 Mbps and one green LED to indicate activity.
- SEEPROM for switch identification.
- Console port. One RS-232 (UART), for terminal access by way of mini-USB.
- SEEPROM for switch identification, encryptable SEEPROM, and secure encryptable SEEPROM.

Note: The License ID is protected by the secure encryptable SEEPROM.

- Secure boot.
- Real-time:
 - Power monitoring.
 - Voltage monitoring.
 - Fan monitoring, including airflow direction.
 - Digital thermometers for temperature monitoring.
- Real-time clock (RTC) with battery.

Software license options

The IBM SAN56B-8 switch uses a capacity-based Ports on Demand (POD) license method. Starting with FOS 10.0.x, the following features no longer require a separate license:

- FICON CUP
- Brocade Trunking
- Extended Fabrics

- Integrated Routing

For more information, see this web page.

Port-side view of the device

The following figure shows the port-side view of the IBM SAN56B-8 Fibre Channel switch. The following components are identified:

1. USB, Ethernet, and Console LEDs and Connectors
2. SFP+ FC (Upper and Lower) Ports (0–55)
3. Switch ID Pull-out Tab




Figure: Port-side view of the IBM SAN56B-8

The following illustration shows the console connectors and LEDs on the port-side view of the IBM SAN56B-8 Fibre Channel switch:

1. Management Ethernet Activity LED
2. Management Ethernet Port
3. Management Ethernet Link LED
4. USB Type-A Port
5. UART Mini-USB Serial Console Port
6. 2-56 Threaded Holes for the retention of thumb screws of the console cable
7. System Power LED
8. System Power LED




Figure: IBM SAN56B-8 console connectors and LEDs

The following diagram shows how the ports are numbered on the IBM SAN56B-8 Fibre Channel switch:




Figure: IBM SAN56B-8 port numbering diagram

Non-port side view of the device

The following figure shows the non-port-side view of the IBM SAN56B-8 Fibre Channel switch. The following components are identified:

1. #6-32 for Screw Mounting of the Ground Cable
2. Ground Marking
3. Power Supply and Fan Assembly 2
4. Power Supply and Fan Assembly 1
5. Captive Screw
6. Handle
7. Power Supply and Fan Assembly Status LED
8. Power Supply Receptacle




Figure: Nonport-side view with AC power supply and fan assembly units

More information

For more information about the documentation for the IBM b-type Gen 8 SAN56B-8 (8634-F56), see this web page.

Gen 7 to Gen 8 migration

Note: Upgrading a Gen 7 director to Gen 8 is not possible. See this document.

IBM b-type Gen 8 transceiver module overview

The IBM SAN512B-8, IBM SAN256B-8, and IBM SAN56B-8 platforms are based on the 10th-generation Brocade Condor 6 ASIC. Brocade Condor 6 ASIC platforms (Gen 8) must use specific transceivers that are listed in the following table.

Table: IBM SAN512B-8, IBM SAN256B-8, and IBM SAN56B-8 supported transceivers

Description	IBM SKU	IBM Part Number	IBM Feature Code	SAN512B-8 & SAN256B-8 Port Blade	SAN512B-16 & SAN256B-16 Port Blade
128G FC SFP+ SWL (1 pack)	IB-000500	003UK893	2132	X	X
128G FC SFP+ SWL (8 pack)	IB-000501	003UK897	2133	X	X
OSFP SWL 100m (ICL Blades only)	IB-000502	003VG669	2134	X	
OSFP LWL 2Km (ICL Blades only)	IB-000503	003VG663	2135	X	
64G FC SFP+ SWL (1 pack)	IB-000462	003GM425	2630	X	X
64G FC SFP+ SWL (8 pack)	IB-000463	003GM427	2631	X	X

Gen 8 Fabric Operating System (FOS) overview

Gen 8 Fabric Operating System (FOS) v10.0 introduces a broad set of new features and enhancements. It is a modern operating system that provides a solid foundation to simplify management, improve performance, increase automation, and strengthen security. FOS 10.0 builds on previous releases by delivering enhanced scalability, operational efficiency, and support for modern SAN architectures.

Autonomous SAN

Enterprise data centers continue to push virtualization at higher density and support a much broader set of workloads. More virtual machines (VMs) are being deployed on the server end, and more storage LUNs are provisioned on storage arrays.

At the same time, adoption of all-flash storage has accelerated. Customers are increasingly deploying next-generation NVMe-based storage to reduce latency and improve application performance. Virtualization, flash storage, and automation tools enable faster deployment of applications and services while supporting higher-performance workloads.

The unprecedented number of application and service interactions also increased the complexity, risk, and instability of mission-critical operations. For example, increasing workload virtualization leads to a higher number of application flows, but reduces the visibility of the flows.

At the same time, storage virtualization increases the number of smaller Fibre Channel (FC) frames and the potential for congestion. As a result, IT organizations require storage networks that deliver ultra-low latency, higher-capacity bandwidth, and greater reliability.

To meet service-level agreements (SLAs) or achieve service-level objectives (SLOs), IT administrators also need new tools that can help ensure nonstop operations, quickly identify potential points of congestion, and maximize application performance, while simplifying administration.

In the past, the process of problem detection and mitigation relied on the diligence, skill, and experience of the fabric administrator. Traditionally, administrators relied on manually searching for and tracing error conditions. The speed at which data flows change in modern fabrics is frequently quicker than an experienced admin can identify the issue and its cause. The operational cost and potential disruptive impact on critical applications are no longer acceptable in today's data center.

To help administrative automation and improve network uptime, Broadcom developed Autonomous SAN technology that offers self-learning, self-optimizing, and self-healing capabilities. Autonomous SAN is realized through Fabric Vision technology that combines built-in capabilities in IBM b-type Gen 7 and Gen 8 platforms, Brocade Fabric OS (FOS), and Brocade SANnav Management Portal features. Fabric Vision provides detailed monitoring, alerting, response, and mitigation, which vastly improves the fabric administrator's insight and response time.

Self-learning

Brocade technology proactively monitors millions of I/O performance and behavior data points through integrated network sensors to gain deep insight into the environment. These data points are transformed into actionable intelligence that can monitor and alert when there are any abnormal changes.

Through the following capabilities, an administrator can identify individual applications and their performance characteristics across the fabric. They also can identify the performance of the various devices that comprise the fabric: the switches, hosts, and targets:

- I/O Insight

Proactively monitors I/O performance and behavior through integrated network sensors to baseline application performance and ensure operational stability. By combining the instrumentation of I/O Insight with the ability to self-learn the traffic flows, the Brocade autonomous SAN technology can generate performance and health metrics on each component and the applications to monitor for changes and alert the administrator with MAPS.

- Monitoring and Alerting Policy Suite (MAPS)

Provides an easy-to-use solution for policy-based threshold monitoring and alerting. MAPS proactively monitors the health and performance of any SCSI or NVMe storage infrastructure to ensure application uptime and availability. By using pre-built, rule-based, and policy-based templates, MAPS simplifies fabric-wide threshold configuration, monitoring, and alerting.

- Automatic Flow Learning

Provides automatic learning of all traffic flows from a specific host to storage across the SAN fabric. With this information, an administrator can automatically identify resource contention or congestion that is affecting application performance.

- Fabric Performance Impact

Uses predefined MAPS policies to automatically detect and alert administrators to varying levels of congestion severity, as well as identify credit-stalled devices such as misbehaving slow-drain devices and oversubscribed ports that can impact network performance. This feature pinpoints the devices that are causing or affected by port congestion and can automatically quarantine misbehaving devices to help maintain optimal network performance.

Self-optimizing

By using actionable intelligence gathered from self-learning capabilities, Brocade Fibre Channel SANs automatically apply priorities for specific data traffic to help ensure performance levels and monitor for traffic pattern shifts.

Learning traffic behavior enables the network to make smarter decisions about traffic prioritization, congestion avoidance, and adjustment to ensure optimal network performance for applications and storage. When something does change, the Brocade autonomous SAN technology isolates the port traffic for the misbehaving device to a virtual channel in the fabric and allow all other traffic to go around to maintain optimal performance.

In addition, Brocade SANnav Management Portal includes automated operational activities, such as infrastructure deployment and provisioning, to expedite IT services. Self-optimizing includes the following features:

- Adaptive Traffic Optimizer: Automatically classifies and segregates traffic with similar characteristics into dedicated flow groups to optimize performance in common SAN environments. By identifying and isolating specific traffic flows, ATO minimizes resource contention and prevents individual workloads from adversely affecting overall SAN performance, resulting in more predictable and efficient network operation.

- Advanced traffic shaping: Ensures application performance by proactively monitoring and actively shaping traffic.
- REST APIs: Eliminates human errors and performance effects through open DevOps.
- SAN Automation and Ansible: Gain cloud-like SAN orchestration for optimizing administration resources through Brocade REST API automation technology.

Self-healing

When potential disruptions are detected, the network automatically mitigates or resolves issues without manual intervention. This process is done by proactively monitoring the network to automatically identify abnormal or unexpected infrastructure behavior and then taking immediate action. These actions include alerting the end devices of the issue through a notification and signaling process within a SAN.

End devices can automatically adjust traffic or fail over to a healthy path to mitigate impact until the underlying infrastructure issue can be resolved. In addition, these capabilities detect and automatically reconfigure fabric mis-configurations that fall outside best practices. Self-healing includes the following features:

- Fabric congestion notification: Automatically detects congestion and notifies end devices to automatically mitigate congestion.
- Slow drain device quarantine (SDDQ): Automatically quarantines credit-stalled devices to prevent the misbehaving device from affecting the rest of the traffic.
- Automatic actions: Ensure data delivery with automatic failover from physical or congestion issues, such as port decommissioning, port toggling, and port fencing.
- COMPASS: Automates configuration compliance checking against predefined best-practice baselines, detects configuration drift, and helps administrators identify and correct out-of-compliance configurations.
- SANnav Management Portal: Reduces troubleshooting steps with built-in best practice recommendations to quickly resolve issues.

For more information about how self-healing SAN capabilities minimize the application performance impact, see [Broadcom Autonomous Self-healing SANS](#).

Automation

Most, if not all, IT administrators have first-hand experience in the growing complexity of managing enterprise infrastructure, including SANs.

IBM and Brocade are in the unique position to spot and understand the effect of automation, which helps organizations get more from their SAN infrastructure.

IBM and Brocade offer a combination of SAN automation with RESTful APIs and a SAN management platform to help organizations drive greater efficiency from their SANs. This is accomplished by using various means, including the following examples:

- IBM and Brocade SAN automation is provided with a multilayer architecture.
- RESTful API support on switches and management tools.
- Brocade's Ansible management framework, which eliminates repetitive tasks, simplifies management, and orchestrates across the full SAN infrastructure.

Why organizations should embrace automation

Organizations should embrace SAN automation for the following reasons:

- Reducing human error and streamlining operational processes has never been more crucial. As organizations move to digitize and adapt to new workloads, data availability, processing time, and agility to provision applications on-demand become the lifeblood of the business. This demands a more efficient and expedient infrastructure

management approach, leaving no room for human error. As a result, storage administrators must be freed from repetitive manual tasks, such as configuration management, reporting, documenting inventory, and troubleshooting. Instead, IT organizations need SAN automation to help them automate and orchestrate repetitive tasks, significantly improve efficiency, and decrease the risk of operational mistakes.

- Demand for more accurate and frequent infrastructure reports is on the rise. It is not only IT managers who need information about storage performance, use, and forecasting; business stakeholders also are asking for and expecting this data on-demand.

Business stakeholders increasingly expect infrastructure reports to be available on demand rather than waiting for administrators to generate them manually. Information must be accessible as frequently as business requirements dictate and available at the click of a button. Automation provides a level of responsiveness that traditional manual storage management processes cannot achieve. It also enables reports and dashboards to be tailored to specific audiences, ensuring that stakeholders receive timely, accurate, and relevant information aligned with their roles and responsibilities.

- SAN configuration management must be streamlined. As enterprise applications generate increasing volumes of data and virtualized environments continue to expand, deploying and configuring servers, storage, and network resources has become more complex and time-consuming than ever.

By streamlining SAN operations through automation, application provisioning workflows are simplified across hypervisor, network, and storage. As a result, they deliver agility and responsiveness to meet dynamic business demands.

- IT service delivery is not always able to respond at the pace that the business requires. As organizations become increasingly dependent on high-quality IT services to support proactive and agile decision-making, bottlenecks in service delivery must be identified and eliminated quickly. Achieving this goal without increasing headcount or expanding SAN-related capital expenditures requires a more efficient approach. SAN automation provides a practical solution by streamlining operational processes, reducing manual effort, and accelerating service delivery. As a result, organizations can improve IT agility while more closely aligning IT services with rapidly changing business requirements.
- Consistent configuration validation is a must. Manual configuration changes are occurring with greater frequency as enterprises diversify their IT architectures in general, and their storage architectures specifically. SAN automation ensures validation of consistent configuration parameters across the different SAN fabrics to facilitate troubleshooting of frequent alerts without reliance on manual intervention.

Brocade automation solutions use RESTful APIs to facilitate solutions architecture, share best practices, and get to production status faster. Brocade's SAN automation brings automation to the switch level that helps reduce the cost and complexity of managing storage systems.

For more information about how to set up and manage IBM b-type switches by using REST API, see the [Brocade Fabric OS REST API Reference Manual](#).

For more information about choosing your approach by using RESTful, Ansible, or Python, and to create an automation plan, see [SAN Automation for Dummies](#). This document also guides you through the process of how to clone and use the GitHub repositories.

FOS 10.x features

Brocade Fabric OS (FOS) 10.0 introduces support for 128 Gbps Fibre Channel hardware (Gen 8) alongside significant software modernization, automated fabric management, and simplified licensing.

In addition to these features, new features were added, including the following example:

- Supported generations: FOS 10.x exclusively targets Brocade Gen 7 (64G) and Gen 8 (128G) platforms.

- All-inclusive license bundling: All applicable feature licenses for Gen 7 and Gen 8 switches are permanently enabled by default in FOS 10.0.x without requiring manual license key installation. The only exceptions are Ports on Demand (POD) capacity expansion and Brocade TruFOS Certificates, as these features require separate management.
- Gen 8 acceleration: Built to drive the 128 Gbps line rate with auto-sensing support for 16G, 32G, 64G, and 128G Fibre Channel link speeds.
- NVMe & AI workload optimization: Optimized packet handling reduces switching latency by up to 50% relative to older generations, accelerating high-density flash arrays and AI training/inference data pipelines.

Quantum-safe security and network hardening

The following features are part of the quantum-safe security and network hardening:

- Quantum-resistant cryptography: Introduces support for post-quantum cryptographic capabilities, including quantum-resistant digital signature algorithms used to help protect software authenticity and system integrity.
- Unified IPFilter policies: Administrators can now define combined IPv4 and IPv6 rules under a single, unified IPFilter policy rather than managing distinct lists.
- Hardware & firmware validation: Enhanced anti-tampering checks dynamically verify software integrity during boots and upgrades to mitigate malware risks.

Autonomous SAN and embedded AI analytics

The following features are part of the autonomous SAN:

- Self-learning infrastructure: Uses embedded SAN AI telemetry to establish baseline operational behavior and automatically detect storage flow anomalies.
- Automated congestion avoidance: Continuously balances traffic across virtual channels dynamically to eliminate slow-drain bottlenecks without manual intervention.

MAPS and optics diagnostic enhancements

The following features are part of diagnostic enhancement:

- Dynamic SFP thresholding: Monitoring and Alerting Policy Suite (MAPS) dynamically calculates warning/alarm thresholds based on the optical transceiver's internal EEPROM, eliminating the need for vendor-specific static rules.
- Extension optical support: Extends MAPS monitoring capabilities to Ethernet SFPs on Extension platforms (voltage, current, temperature, RX/TX power).
- Detailed alerts: RASlog and SNMP events now embed transceiver serial numbers, part numbers, and actual vs. expected electrical metrics directly within the alert payload.

In this section, we describe some of these features. For a full list of enhancements, see the [Brocade Fabric OS Administration Guide](#).

Port identifiers

Port identifiers (PIDs, also called *fabric addresses*) are used by the routing and zoning services in FC fabrics to identify ports in the network. All devices in a fabric must use the same PID format. When you add equipment to the SAN, you might need to change the PID format on older equipment.

Many scenarios cause a device to receive a new PID; for example, unplugging the device from one port and plugging it into a different port as part of fabric maintenance; or changing the domain ID of a switch, which might be necessary when merging fabrics; or changing compatibility mode settings.

Some device drivers use the PID to map logical disk drives to physical FC counterparts. Most drivers can change PID mappings dynamically (also called *dynamic PID binding*) or use the WWN of the FC disk for mapping (also called *WWN binding*).

Some older device drivers behave as though a PID uniquely identifies a device; they use *static PID binding*. These device drivers must be updated, if possible, to use WWN binding or dynamic PID binding instead because static PID binding creates problems in many routine maintenance scenarios.

Fortunately, few device drivers still behave this way. Many current device drivers enable you to select static PID binding and WWN binding. Select static PID binding only if there is a compelling reason and only after you evaluated the effect of doing so.

Unified addressing mode

Unified addressing mode (UAM) provides a single addressing mode for an 8-bit, 10-bit, or shared area. It also provides other addressing modes for different combinations of platform and logical switch type. UAM simplifies these addressing mode combinations and features into a single unified addressing mode.

UAM is introduced in FOS 9.0.0 and is the default address allocation mode for modern platforms, including the IBM SAN512B-8 Director, IBM SAN256B-8 Director, and IBM SAN64B-8 switch.

The following table describes how UAM works for FOS features.

Feature	UAM mode	Pre-UAM mode
F_Port	A port is assigned only a base area.	A port is assigned only a base area.
F_Port with NPIV	For F_Ports with NPIV fewer than 63 devices, port is assigned only a base area. If the 64th NPIV device is received in F_Port, a new area is assigned to the existing base area. Whenever a port exhausts an area after 64 extra devices log in, another area is assigned. The maximum number of NPIVs supported is 255 devices.	The 10-bit addressing mode uses the 8-bit area ID and the borrowed upper 2 bits from the AL_PA portion of the PID. Areas 0x00 - 0x8F use only 8 bits for the port address and support up to 255 NPIV devices. Areas 0x90 - 0xFF use an extra 2 bits from the AL_PA for the port address. Therefore, these ports support only 64 NPIV devices per port.
F_Port Trunking	Same as F_Port and F_Port with NPIV. In FOS 9.0, UAM allocates more areas to F_Port trunking and Emulex HBA trunking.	F_Port trunking uses an 8-bit area in VF mode enabled logical switch. Default switch in VF or non-VF mode uses existing port address.
EX_Port	EX_Port in a base switch is assigned a 10-bit area.	EX_Port in a base switch uses a port address per the current switch addressing mode.
EX_Port on ICL Ports	When an ICL port is configured as an EX_Port, a new area is assigned to the ICL port from the pool. This is not applicable for ICL ports that are not configured as an EX_Port. From the FOS 9.0 release, when an ICL port is configured as an EX_Port, the new areas are assigned in both VFmode and non-VF mode.	When an ICL is configured as an EX_Port, the switch is assigned an 8-bit area in VF. EX_Port over ICL is not supported in non-VF mode.
EX_Port Trunking	In UAM mode, an EX_Port trunk behaves in a masterless manner both in the base switch and in non-VF mode.	In pre-UAM mode, an EX_Port trunk behaves in a masterless manner only in VF mode.
FCoE	In UAM mode, FCoE address assignment is simple. Whenever Ethernet Flexport requires an area, it is assigned an area from the pool of unused areas.	In pre-UAM mode, whenever Ethernet Flexport requires an area, it is assigned an area from the pool of unused areas. For default switches in VF mode or non-VF mode, the switches are converted to a 10-bit area mode when FCoE enodes are configured.
Port Address Binding	In platforms that are enabled in UAM mode, the port address binding feature is supported in VF and non-VF mode.	In pre-UAM mode, the port address binding feature is supported in only logical switches in VF enabled mode.
Persistent PID	In platforms that are enabled in UAM mode, the persistent PID feature is supported in VF and non-VF mode.	In pre-UAM mode, the persistent PID feature is supported in only logical switches in VF enabled mode.

Table: Difference between UAM and pre-UAM mode

Fabric zone locking

Before FOS 9.x, the level of zoning transaction protection against multiple concurrent clients existed at the local switch level only. For example, if a user on a switch starts a zone transaction by performing a zone create or edit

operation, and if another user attempts to perform a zone create or edit or commit operation on the same switch, the second user is locked out and prevented from performing their operation (see the following figure). This lock out occurs because the first user opened the transaction first.




Figure: Zone locking on a single switch

Although local switch protection exists, no fabric-wide locking mechanism is available. For versions earlier than FOS 9.x, the fabric-level detection support (Concurrent Transaction Detection) consists of only a warning mechanism that warns users if a transaction is open on a remote switch. However, it does not prevent users from committing changes if concurrent transactions exist across the fabric.

Furthermore, this fabric-wide detection is limited in scope to CLI zoning operations only. For example, if a user on one switch opens a zone transaction, and if a user on another switch in the fabric issues a zone edit or commit command, a warning is received that a transaction is occurring in the fabric. However, a zone commit operation can still be issued from either switch, which overwrites any pending zone transaction on other switches (see the following figure).




Figure: Fabric-wide zone commit overwrite in FOS 8.2.x

The Zone Fabric Locking feature provides zoning transaction protection at both the individual switch level and across the entire fabric (see the following figure). This feature is enabled by default. When a zone edit or zone commit operation is in progress anywhere in the fabric, additional zone edit or commit operations cannot be initiated on the same switch or on another switch until the current transaction completes or the default lock timeout period of five minutes expires.

At the start of a zoning transaction, the switch requests a fabric-wide lock to prevent conflicting zoning changes. The Fabric Lock Failsafe Timer, which is configurable and applies to the entire fabric, protects against orphaned locks that might result from interrupted zoning operations. When a fabric lock is established, a failsafe timer is started on all participating switches. If the zoning transaction is not completed before the timer expires, the lock is automatically released, allowing other zoning operations to proceed.




Figure: Fabric zone locking in FOS 10.0.0

When the failsafe timer expires, the open zone transaction is not stopped. If the same user attempts to resume the transaction by performing another edit or commit operation after the zone fabric lock expires, the transaction is allowed and the fabric lock is restarted.

If a different user attempts to start a new transaction after the first user's transaction timer expires, the transaction is allowed and the first user's transaction is stopped before the second user's transaction starts.

Consider the following points:

- If a fabric-lock is active in the fabric, the Fabric Lock Failsafe Timer timeout value cannot be modified.
- The fabric zone lock is a fabric-wide setting. Therefore, the fabric must be stable before performing zone edit, commit, and abort operations.
- If a zone transaction is stopped, the fabric zone lock is cleared.
- When the timer expires, the lock is cleared locally and on each individual switch in the fabric.
- When the virtual fabric function is used, the lock timeout value is independent of each logical switch and can have different values.
- Fabric zone locks can be preempted by zone merges.
- The Zone Fabric Locking feature is supported on switches with FOS 9.0.0 or later. Down-level switches do not participate. However, FOS 9.0.0 switches that own the zone fabric lock prevent updates from down-level switches.
- On remote switches, another 30-second buffer is added to the configured lock timeout value of the lock principal switch.

The fabric-wide locking mechanism ensures that any zone transaction that is started on any switch by any interface (CLI, REST, WebTools, SANnav, or Target Driven Zone) is not overwritten or stopped by another interface's transaction (except for zone merges) while the lock timer is active.

FICON logical switch

With the FICON logical switch feature, you can create a FICON logical switch. You can also modify a logical switch as a FICON logical switch. The FICON logical switch is supported on all platforms that support Virtual Fabrics in FOS 8.1.0 or later. A FICON logical switch can be created independent of enabling FMS mode. A FICON logical switch cannot be the default switch in the chassis, nor can it be the base switch of the chassis.

FICON logical switch provides auto-configuration of the following components, which simplifies the process of establishing a logical switch in the FICON environment:

- Insistent Domain ID (IDID) mode is set to on.
- Routing policy is set to device-based routing.
- Area mode is set to zero-based addressing.
- SCC security policy is auto-defined with only its own WWN.
- The created SCC policy is activated.
- The fabric-wide consistency policy is set to strict SCC policy.
- High Integrity Fabric (HIF) mode is enabled.
- FMS mode is enabled if FICON control unit port (CUP) license is installed on the chassis.

Management tools

The available management tools are described in this section.

WebTools

Brocade WebTools is a graphical user interface (GUI) that is embedded in the FOS firmware. It enables administrators to monitor and manage single or small fabrics, switches, and ports. New WebTools are Java-Free and are designed based on modern technology HTML5.

WebTools is started directly from a web browser or from SANnav Management Portal. The system requirements are described next.

WebTools system requirements

Before starting WebTools, verify that your workstation uses a supported operating system and web browser.

WebTools does not require a license. It is installed on the switch when you install FOS.

Supported operating systems

WebTools supports the following operating systems:

- Red Hat 8.0 and 8.1
- Windows 10 Pro
- Windows 2019

Supported web browsers

The following browsers can be used to access WebTools:

- Chrome version 96 and later
- Firefox version 95 and later
- Microsoft Edge version 97 and later

Starting WebTools

WebTools is started directly from a web browser or from SANnav Management Portal. You can start WebTools on any workstation with a compatible operating system and Web browser installed.

If the switch is configured with logical fabrics, you can log into any of the logical fabrics for which you have permission.

To start directly from a web browser, open your browser, enter the IP address of the switch, and press Enter. You can use HTTP or HTTPS; for example, <http://10.77.77.77> or <https://10.77.77.77> (see the following figure).


The image shows the WebTools Element Manager login screen. At the top is a logo consisting of a blue gear and several vertical bars in blue, yellow, and red. Below the logo, the text "WebTools" is in a large, dark blue font, and "Element Manager" is in a smaller, dark blue font. There are three input fields stacked vertically: the first contains the text "admin", the second contains a series of dots representing a password, and the third contains the number "128". Below these fields is a large blue button with the word "Login" in white text.

Figure: WebTools login screen

Enter the user name, password, and logical switch name or fabric ID (FID).

For the first switch login, the default username is `admin` and the default password is `password`. WebTools prompts you to change the default password.

If you are logging in to a Virtual Fabrics-enabled platform and you do not specify a logical switch, you are logged in to the default logical switch, which uses fabric ID `128`. For non-VF platforms, the FID option is not displayed.

Note: If you start from SANnav Management Portal, you might not need to log in, depending on the SANnav single sign-on configuration.

For information about how to use WebTools, see [Brocade Fabric OS Web Tools User Guide](#).

Security

Security entails authentication and encryption to restrict access and protect data from unauthorized access. Brocade products support a wide range of authentication, encryption, and management tools to protect fabrics and data from unauthorized access, including the following examples:

- Authentication: Authentication protocol support includes CHAP, DH-CHAP, FCAP, IKE, IPsec, RADIUS, and PEAP/MS-CHAP for RADIUS.

Note: TACACS+ was deprecated in Fabric OS 9.2.1 and is planned for removal in a future Fabric OS release. Existing deployments can continue to use TACACS+, but RADIUS is the recommended authentication framework for new deployments.

- Encryption: Brocade platforms support AES-based encryption for extension and management connections and can use IPsec to secure IP-based communications. These technologies provide high-performance encryption to help protect data in transit.
- In-flight encryption over ISLs: Brocade X6 directors with Gen 6 Fibre Channel port blades, Brocade X7 directors with Gen 6 or Gen 7 Fibre Channel port blades, Brocade X8 directors, and Brocade G820, G720, G630, and G620 switches support in-flight encryption for traffic traversing inter-switch links (ISLs). This capability helps protect sensitive data from unauthorized access while it is transmitted within and between data centers, including over long-distance links. Data-at-rest and data-in-flight encryption are complementary technologies that address different security requirements, and both might be necessary to meet regulatory or compliance objectives.
- Secure Boot: A switch validates the integrity and authenticity of the FOS boot image to establish a hardware-based root of trust through the manufacturing supply chain.
- Brocade Fabric OS (FOS) 10.0 introduces a comprehensive quantum-safe security architecture tailored for Gen 8 (128 Gbps) and Gen 7 Fibre Channel infrastructure. This update directly targets the threat of quantum computing and the emergence of "Harvest Now, Decrypt Later" (HNDL) attacks—where adversaries capture encrypted SAN traffic today to decrypt it in the future using Cryptographically Relevant Quantum Computers (CRQCs).

For more information, see section 10, "Storage area network security".

Adaptive Traffic Optimizer

Adaptive Traffic Optimizer (ATO) in Fabric OS (FOS) 10.0 is an evolution of Brocade's SAN traffic isolation architecture, designed specifically for Gen 8 (128 Gbps) Fibre Channel hardware (for example, Brocade G820 and X8 Directors).

While earlier Gen 7 Traffic Optimizer implementations relied primarily on static profiles (sorting flows into Performance Groups strictly by destination port speed or protocol), ATO on Gen 8 dynamically adapts to live fabric conditions in real time without manual administrative intervention.

The Adaptive Traffic Optimizer (ATO) includes the following key benefits:

1. Dynamic performance group customization:
 - Real-time environment profiling: Rather than using rigid system templates, ATO actively monitors the active SAN environment and dynamically tailors Performance Groups (PGs).
 - Optimal resource usage: FC fabrics have a finite number of Virtual Channels (VCs) and buffer credits per ASIC. ATO calculates optimal VC and buffer-credit allocations based on live device mix, IO profile variations, and link speed distributions.

2. Speed and protocol coexistence (16G to 128G):

- With Gen 8 hardware supporting speeds up to 128 Gbps alongside backward compatibility for 16G, 32G, and 64G, speed differentials create severe risk of head-of-line (HoL) blocking:
- Speed-matched virtual channels: ATO automatically segregates 128G high-speed flows from slower 16G, 32G, and 64G traffic flows, reducing head-of-line (HoL) blocking and improving overall fabric efficiency.
- SCSI versus FC-NVMe isolation: ATO dynamically distinguishes bursty, ultra-low-latency NVMe frames from standard Fibre Channel SCSI frames. This guarantees that NVMe queues maintain minimal latency without getting queued behind deeper SCSI buffer chains.

3. Autonomous congestion & slow-drain containment:

- Credit stall analytics: ATO tracks frame credit replenishment rates at the ASIC level.
- Dynamic quarantine: When a destination host or storage array experiences a credit stall or slow-drain behavior, ATO dynamically shifts that specific flow into a degraded/low-priority Virtual Channel.
- Non-impacted flow isolation: Neighboring healthy flows sharing the same physical Inter-Switch Link (ISL) or port continue to transmit at full line rate without experiencing buffer starvation.

4. Zero-touch operational model:

- Built-in enablement: ATO runs autonomously in FOS 10.0 on Gen 8 hardware. No manual zone mapping, QoS tagging, or profile binding is required.
- Continuous rebalancing: As new workloads scale up, host speeds are upgraded, or array paths fail over, ATO dynamically adjusts Virtual Channel boundaries to maintain balanced buffer allocation across the fabric.

Comparison: Gen 7 Traffic Optimizer versus Gen 8 Adaptive Traffic Optimizer

Attribute	Gen 7 Traffic Optimizer (FOS 9.x / 10.0)	Gen 8 Adaptive Traffic Optimizer (FOS 10.0)
Classification Model	Profile-based (Sys_TrafOpt_Version1/ Sys_TrafOpt_Version2)	Dynamic / Self-Learning
VC & Buffer Allocation	Static allocation per PG	Real-time adaptive reallocation
Top FC Link Speed	64Gbps	128Gbps
Congestion Handling	Rule-triggered Slow Drain Quarantine (SDDQ)	Automated flow-level credit stall isolation
Manual Configuration	Minimal (profile selection optional)	Fully Autonomous

Gen 6 platforms in FOS 9.0 continue the established behavior and maintain compatibility with Gen 8 platforms in the same fabric. If a Gen 8 switch or blade is connected to a port on a Gen 6 platform and the traffic flow source port is on the Gen 8 switch, the traffic flow is speed-based that is classified within the Gen 8 switch until reaching the port on the Gen 6 platform. Then, the traffic flow is handled with established behavior starting from the Gen 6 port.

Many years ago, Brocade Virtual Channels (VC) addressed this issue with great success and significant performance gains. Resources that affected performance include bandwidth, egress queues, ingress buffers, and QoS.

The implementation features many distinct egress queues that established virtual channels across the physical link, each channel having its own set of autonomous buffer credits. The bandwidth is shared, but flows are independent. When traffic characteristics are uniform, VCs offer an expedited pathway. Great benefit is gained by sorting traffic into VCs that are based on the most pertinent characteristic, which is speed.

Normally, VCs are a fundamental component of Fabric OS and are automatically managed by the switch. As Fibre Channel speeds have increased to 128 Gbps with Gen 8 platforms, SAN environments often contain a mix of 16, 32, 64, and 128 Gbps devices. These significant speed differences can introduce flow-control inefficiencies and increase the risk of head-of-line (HoL) blocking. Adaptive Traffic Optimizer (ATO) addresses this challenge by automatically segregating 128 Gbps traffic from lower-speed 64, 32, and 16 Gbps traffic flows, helping to maintain efficient flow control and deliver more predictable performance across the fabric.

For years, Brocade VCs operated within a relatively narrow bit rate of 1, 2, 4, 8, and 16 gigabit (8 - 15 Gbps) by use many logical independent paths. VCs are assigned by a hash called PID2VC, which results in a pseudorandom assignment of each flow to a VC. To a degree, this statistically mitigated the interference of slower flows impeding faster flows. Also, optionally faster flows can be manually assigned to a high QoS VC and slower flows to a low QoS VC to prevent interference.

As Fibre Channel technology has advanced to 128 Gbps with Gen 8 platforms, SAN fabrics increasingly contain a mix of 16, 32, 64, and 128 Gbps devices. Combined with the growth of NVMe, all-flash arrays (AFAs), and host virtualization, these mixed-speed environments can create flow-control inefficiencies and head-of-line (HoL) blocking. Adaptive Traffic Optimizer improves virtual channel utilization by automatically classifying and isolating traffic flows based on speed and workload characteristics, reducing contention and maintaining predictable performance across the fabric.

Speed-based classification

Flows are classified into one of several Performance Groups based on the destination speed and protocol. Automatic speed classification isolates traffic flows into different speed-based Performance Groups and eliminates speed mismatches that cause performance interference when lower speed devices collide with higher speed devices that are connected in the same fabric.

Slow drain device quarantine isolation

A slow-drain device, also sometimes referred to as a *credit-stalled device*, is a misbehaving device that promptly stops returning R_RDY signals (buffer credits) to the switch, which causes the switch to stop sending frames to the device. A credit-stalled device that stops returning credits to the switch for hundreds of milliseconds or more causes frame drops and link resets. This issue typically occurs when the end device is not working per its specification. An impaired flow that is associated with credit-stalled devices is automatically mapped to a low-priority Performance Group: PG_low_qos_1 or PG_low_qos_2.

Compatibility with QoS zones and CS_CTL

Gen 7 platforms remain compatible with Adaptive Traffic Optimizer when QoS is enabled. If QoS High or QoS Low priorities are defined, traffic flows are assigned to the corresponding QoS Performance Groups, ensuring that existing traffic prioritization policies are preserved.

If CS_CTL is enabled on a port, traffic flows retain their CS_CTL-based classification and priority handling. In this case, the flows are not automatically classified into ATO Performance Groups.

Compatibility with previous platforms

Adaptive Traffic Optimizer is designed to operate in mixed-generation SAN fabrics. Gen 8 switches can apply traffic classification and optimization to supported flows while maintaining interoperability with connected Gen 6 and Gen 7 platforms, providing a seamless migration path as fabrics evolve.

Traffic Optimizer and QoS mode

Traffic Optimizer and QoS mode are enabled by default on Gen 8 platforms. Advanced traffic classification capabilities rely on QoS services to identify, prioritize, and isolate traffic flows, helping maintain predictable performance during periods of congestion.

It is *not* recommended to disable QoS mode on any ports. If QoS mode is disabled on a port, Traffic Optimizer's advanced classification do not work on the port, and the port's traffic classification falls back to older PID to VC mapping.

Adaptive Traffic Optimizer considerations for FCR

Adaptive Traffic Optimizer is supported only in Fibre Channel Routing (FCR) edge-to-edge fabric topologies and is not supported in backbone-to-edge fabric topologies. To enable ATO across an FCR environment, all switches and FC routers in both the edge and backbone fabrics must be running Fabric OS 9.0 or later.

ATO across FC routers is supported only when Virtual Fabrics is disabled in the backbone fabric. If Virtual Fabrics is enabled in the backbone fabric, ATO traffic classification and optimization are not supported across FC routers.

When a traffic flow in a metaSAN traverses a platform that does not support ATO, the flow reverts to standard Fibre Channel traffic-handling behavior while maintaining interoperability with the remainder of the fabric.

For more information about the operating modes of backbone fabrics, see [Brocade Fabric OS Administration Guide, 10.0.x](#).

If a traffic flow in a metaSAN traverses a switch platform that does not support Adaptive Traffic Optimizer, the flow reverts to standard Fibre Channel traffic-handling behavior at that point while maintaining interoperability across the remainder of the fabric..

Fabric notifications

Fabric Notifications is a facility for proactively notifying devices in a fabric of performance impacting behaviors that affect the normal flow of I/O. Fabric impacting behaviors (FPI), such as signal integrity degradation, frame loss, or congestion, are detected by the fabric or participating devices. Signaling by way of FC Primitives, and notifications by way of Extended Link Services (ELS) are then used to communicate the type and severity of the FPI to the affected members of the fabric.

Fabric Notifications include the following components:

- Registered capabilities

A set of ELS commands exchanged between two ports is used to negotiate the range of Fabric Notification capabilities supported by each device.

- Congestion detection signaling

Hardware-generated primitive signals, supported on Gen 7 and Gen 8 platforms, are sent toward the point of congestion to indicate that the queue of frames awaiting transmission has reached a level at which fabric performance might be affected.

- FPI notifications (FPINs)

FPINs are software that is generated by using ELS commands. If a device registers to receive them, administrators can be notified of events that might affect fabric performance.

These components allow members of the fabric to indicate their ability to interpret and react to signals and notifications. Participating devices can respond by taking actions to mitigate and reduce the effects of congestion, frame loss, and signal integrity degradation. The FOS is the primary generator of signals and notifications because of its awareness of I/O flows, physical layer impacts, and device responses.

The relationship between the event detection, signaling, and notification operations is shown in the following figure.




Figure: Fabric Notifications signaling and notification operations

When an event is detected (1), the fabric starts the suitable response. If frame transmission is stalled, signal events are generated (2) toward the point of congestion (that is credit-stalled or oversubscribed device). If a transmission failure or sustained congestion event occurs, notification events are generated (3) and distributed to peer devices (4).

The fabric determines the distribution of the notification event from the zone member information of the triggering N_Port. If a transmission failure or software-based congestion signaling occurs, notification events are sent to the attached N_Port (5) to engage the device in mitigating activities.

IBM b-type switch Fabric Notifications solution

Fabric Notifications improves I/O behavior and helps prevent performance degradation by informing end devices about fabric conditions that can affect application performance. The IBM b-type Fabric Notifications solution combines hardware intelligence in the Fibre Channel switching ASICs, Fabric OS software services, and management capabilities provided by MAPS and SANnav to deliver a fabric-wide framework for detecting, reporting, and helping remediate performance-impacting events.

When a fabric issue is detected, detailed operational information is recorded, including the affected initiator-target-lun (ITL) relationships, event location, timestamp, and reason codes. This information simplifies troubleshooting and helps administrators more quickly identify and resolve the root cause of performance issues.

IBM b-type Fabric Performance Impact Notification

IBM b-type Fabric Performance Impact Notification (FPIN) enables switches to detect fabric conditions such as congestion, link integrity issues, frame-delivery problems, and other events that can affect application performance. When such conditions are detected, registered devices can be notified proactively, allowing them to take corrective actions, such as path failover, workload redistribution, or path optimization, before application performance is adversely affected.

FPIN provides a standardized mechanism for communicating fabric performance information to participating devices. To receive these notifications, a device must register with fabric services. Once registered, the device can receive notifications about conditions that affect a connection, a path, or overall fabric performance.

FPIN supports both hardware-assisted congestion signaling and software-based notifications. Hardware-assisted notifications are delivered directly from the switch to the host adapter, enabling a rapid response to congestion-related events. Software notifications are delivered through Fibre Channel services and processed by the adapter driver or operating system, providing additional diagnostic and operational information.

One advantage of hardware-assisted notifications is reduced reaction time because the host adapter can respond directly to changing fabric conditions. Software-based notifications provide additional context that can help systems and administrators make informed decisions about path selection, workload placement, and error recovery. Together, these notification mechanisms help improve application resilience, reduce the impact of congestion, and maintain optimal SAN performance.

Hardware-assisted signals are implemented as Fibre Channel primitives. Because primitives are not frames, they do not consume buffer credits and are not affected by buffer-credit depletion. In contrast, software-based notifications are transmitted as Extended Link Service (ELS) frames and therefore can be affected by congestion and buffer-credit limitations.

The way in which notifications are processed is determined by the adapter and operating-system vendors. Some implementations might simply log the notification, whereas others can initiate corrective actions. The specific actions taken in response to an FPIN event are vendor dependent. FPIN notifications can identify conditions such as congestion, link integrity issues, frame drops, and other fabric performance impacts.

FPIN notifications can be generated and consumed by any supported device in the SAN, including hosts, storage systems, and switches. By enabling end-to-end visibility of fabric conditions and facilitating automated responses, FPIN helps organizations move toward a more autonomous and self-managing SAN infrastructure.

Congestion signal primitives

Congestion signal primitives are ASIC-based hardware notifications supported on IBM b-type Gen 7 and Gen 8 platforms. These primitives are embedded within the Fibre Channel link encoding and are transmitted directly to connected Fibre Channel devices to provide immediate notification of congestion conditions.

Congestion signal primitives notify participating devices when degraded performance conditions, such as oversubscription or credit stalls, are detected. Fabric Notifications hardware continuously monitors fabric conditions and can respond immediately by signaling the attached device when congestion is identified.

Before a device can receive congestion signal primitives, the connected ports must exchange capabilities indicating support for hardware-assisted congestion signaling. Platforms that do not support congestion signal primitives can still participate in Fabric Notifications through software-based FPIN notifications that are delivered as Extended Link Service (ELS) frames.

Because congestion signal primitives are Fibre Channel primitives rather than frames, they do not consume buffer credits and are not affected by buffer-credit depletion. This capability enables rapid notification of congestion events and allows supported devices to react more quickly to changing fabric conditions.

Extended link service notifications

Extended Link Service (ELS) is a Fibre Channel protocol that enables switches, hosts, and storage devices to exchange capabilities, register services, and communicate operational and diagnostic information. The Fabric Notifications solution uses ELS-based Fabric Performance Impact Notifications (FPINs) to complement hardware-assisted congestion signal primitives and provide end-to-end visibility into fabric conditions.

FPINs provide a standardized mechanism for communicating congestion, peer congestion, link integrity issues, frame-delivery problems, and other fabric performance impacts to participating devices. Supported hosts, storage systems, and switches can generate and process FPIN notifications when conditions affecting application performance are detected.

A notification can be generated by any participating device or switch that detects a fabric performance issue, including congestion, link integrity degradation, or frame-delivery problems. The notification is processed consistently regardless of whether it originates from a switch port or an end device.

Fabric Notifications and MAPS work together to provide proactive monitoring and notification capabilities. MAPS policies can detect conditions such as congestion, frame-delivery issues, and link integrity problems and automatically trigger the corresponding FPIN notifications. These notifications enable registered devices to take corrective actions before application performance is affected.

FPIN notifications are delivered to registered in-zone devices that subscribe to the corresponding notification type. Notification delivery is independent of the zoning model being used and operates consistently across traditional and peer zoning environments. In addition, FPIN notifications can be propagated between switches, allowing registered devices throughout the fabric to receive notifications relevant to the paths and resources they use.

Hardware-assisted congestion signaling and ELS-based FPIN notifications are complementary technologies. Congestion Signal primitives provide immediate, low-latency hardware notifications between directly connected devices, while ELS-based FPIN notifications provide broader fabric-wide visibility and richer diagnostic information. Together, these capabilities help improve application resilience, accelerate problem identification, and support more automated and adaptive SAN operations.

Exchange diagnostic capabilities and register diagnostic functions

Exchange Diagnostic Capabilities (EDC) and Register Diagnostic Functions (RDF) are ELS services that are associated with Fabric Notifications (see the following figure). EDC is used between devices to exchange diagnostic capabilities. RDF is used by devices to register for specific notifications. The fabric does not transmit any notifications (Primitive or ELS) until a device is registered to receive them. Devices cannot deal with every type of problem; therefore, they cannot register for every available notification.




Figure: ELS Exchange Diagnostic Capabilities and Register Diagnostic Functions

FPIN notifications

A device can register to receive the following types of notifications:

- **Congestion:** Congestion notification descriptors are sent to devices connected to congested ports. Causes can include oversubscription, credit stalls, and buffer credit depletion. These notifications help identify devices that are contributing to or experiencing congestion.
- **Peer Congestion:** Peer Congestion notification descriptors are sent upstream from congested ports to registered in-zone devices. These notifications inform receiving devices about the presence and cause of congestion, enabling them to take corrective action if supported by the host or storage platform.
- **Link Integrity:** Link Integrity notifications are sent to affected devices and registered in-zone peers when a port experiences link integrity problems. MAPS alerts related to CRC errors and Invalid Transmission Words (ITWs) can trigger Link Integrity FPIN notifications. These notifications help hosts and storage systems identify degraded paths and, when supported, steer I/O traffic toward healthier paths.
- **Delivery:** Delivery notifications are sent to the registered originator when a SCSI command or status frame is dropped. The notification enables the host to react immediately rather than waiting for an upper-layer protocol timeout, which can significantly reduce recovery time.

Congestion signal

Congestion signals provide an immediate hardware-based feedback mechanism between directly connected devices on a Fibre Channel link. When congestion is detected, congestion signal primitives are transmitted within microseconds, providing rapid notification that transmission resources are becoming constrained. Congestion signals operate on a link-by-link basis.

Gen 7 and Gen 8 FC ASICs support direct hardware-based congestion signaling between connected devices. When congestion conditions are detected, Fabric Notifications can complement these hardware signals by generating FPIN Peer Congestion notifications for registered in-zone devices, enabling additional I/O optimization and remediation actions.




Figure: Congestion signal and notification for over-subscription

Congestion notification

FPIN Congestion notifications are valuable information for end-devices, which facilitates optimized I/O scheduling; that is, consider slower transfer rates or issue serial read I/Os.

The ELS Congestion Notification is the software equivalent of the Congestion Signal primitive and sent to end-devices that support notifications. An ELS congestion notification for over-subscription means that too much data was requested, and that the fabric was unable to accommodate the request without congestion. An ELS congestion notification for credit stall indicates to the end-device "you are stuck - fix the situation".

In general, congestion notifications indicate why exchange completion times might exist.

Congestion peer notification

FPIN Peer Congestion notifications are sent to registered in-zone peers of devices that are experiencing congestion. These notifications help identify congestion conditions within the SAN and provide participating devices with information that can be used to optimize traffic flow and reduce the impact of congestion.

In environments that support speed-aware traffic management, devices can use congestion information to adjust traffic distribution and make more efficient use of available bandwidth. In addition, hosts and storage systems can select alternative paths to avoid congested resources and improve overall I/O performance.

Peer Congestion notifications can also be used by multipathing software and other traffic-management functions to distribute workloads across multiple available paths. This approach helps improve path utilization, minimize bottlenecks, and maintain consistent application performance.

An example of a Peer Congestion notification workflow is shown in the following figure. When congestion is detected, the switch generates an FPIN notification and distributes it to registered in-zone devices. Supported hosts, storage systems, and management applications can then use this information to make informed decisions about path selection, workload placement, and congestion remediation.




Figure: Peer congestion notification flow across fabric zones

Under normal operations, as shown in the following figure, IBM SANnav management tool and the Emulex SAN Manager report minimal to no congestion.




Figure: Normal SAN operations without congestion

The workloads generally grow over time and become larger than the hardware can support. As shown in the following figure, this issue has a ripple effect on the entire infrastructure, which causes the congestion and results in major performance issues.




Figure: Over-subscription causing SAN congestion

IBM b-type Gen 8 SAN detects these types of F_Port congestion issues and sends FPIN Congestion Notification to the Emulex HBA on the AIX server. The Emulex SAN Manager (the management software to manage HBAs on the AIX server) is enabled to set congestion policy.

As shown in the following figure, depending on the choice of the policy between conservative, moderate, and aggressive, the action is taken to address the endpoint congestion issue. For the bully server, the workload is still too large; therefore, administrative attention is required.




Figure: Addressing congestion using IBM SANnav and Emulex SAN Manager

Link integrity

A complete solution requires end-devices (N_Port) to notify the fabric (F_Port) of receive-side errors of which the fabric otherwise is unaware. A link is a two-way street and FPIN Link Integrity notifications enable the identification of degraded paths in each direction, from N_Port → F_Port and F_Port → N_Port.

In notification F_Port → N_Port, HBAs register to receive FPIN notifications of interest; for example, Link Integrity. When F_Port experiences incoming errors because of faulty electronics, marginal optic, dirty connector, damaged fiber cable, and so on, a Link Integrity notification is sent to the HBA letting the host know of the transmission problem. In turn, the HBA forwards notifications to the MPIO driver, which can distinguish between a device failure and physical path failure. Link Integrity notifications are used by end-devices for troubleshooting links and by MPIO for selecting the best path.

In notification N_Port → F_Port, the HBA registers with the fabric to receive and send FPIN notifications. If the HBA sees incoming errors, it notifies the fabric that its connection to the end-device is erratic. An FPIN Link Integrity notification that is sent by N_Port to F_Port (specifically, the Fabric Controller) is distributed to the in-zone devices as well. The fabric logs the notification information so that the SAN administrator can take action.

Registered in-zone end-devices must be notified of link integrity problems. MAPS F_Port alerts for Cyclic Redundancy Check (CRC) and Invalid Transmission Word (ITW) are translated into Link Integrity FPINs and sent to the end-devices. Notifications integrate with the operating system device manager and notifying initiators and targets of link integrity issues along a data path enhances debugging and recovery efforts from multiple perspectives.

MPIO handles FPIN events and manages path selection. On a new pathway, MPIO drivers communicate with the appropriate service to manage outstanding and pending I/Os.

End-devices determine their best course of remediation, such as immediate failover to an alternative path within the MPIO environment. Link Integrity problems include the following examples:

- Slow host application performance can be traced to CRC errors on a storage port.
- Slow array performance can be traced to host port ITWs.

An example of how Link Integrity notifications operate in a real-world environment, using an Emulex HBA on an IBM AIX server connected to an IBM b-type SAN, is shown in the following figure.




Figure: Normal FC link integrity operations with balanced multi-path

Under normal operations, as shown in the previous figure with the FC links working, data is balanced between the different paths to the storage system. Over time, links might not work as well as they should.

As Fibre Channel links age or become impaired, they can remain operational while experiencing intermittent errors and degraded performance. When I/O traffic traverses such a degraded path, application performance can be adversely affected because data transfers are delayed or retransmitted (see the following figure). Rapid detection and remediation of these conditions are essential to maintaining application performance and, in some cases, application availability.




Figure: Effect of a degraded FC link on data path performance

With the Link Integrity feature, when a degraded but operational path is detected, a Fabric Performance Impact Notification - Link Integrity (FPIN-LI) is sent to the host HBA, as shown in the following figure. The Emulex HBA passes the notification up the protocol stack to the AIX multipathing software, enabling corrective actions such as selecting a healthier path while maintaining application availability.




Figure: FPIN Link Integrity notification sent to host HBA upon detection

As shown in the following figure, after the notification is received by the host HBA, the multipathing software uses information provided by the IBM b-type SAN fabric to make intelligent path-selection decisions and rebalance traffic away from the impaired link, improving availability and minimizing performance degradation.




Figure: MPIO rebalancing traffic away from a bad link after FPIN-LI

SCSI command delivery

When the fabric discards a command or status frame, Fabric Notifications notifies the originator of the failure by sending an FPIN Delivery notification, as shown in the following figure. Regardless of whether the frame is discarded by an E_Port or an F_Port, the originator is notified. The notification immediately triggers SCSI I/O error handling without requiring a timeout, helping reduce recovery time, avoid lengthy delays, and minimize application disruption.




Figure: FPIN SCSI command delivery failure notification

The Brocade's Fabric Notifications solution addresses each of these issues through a technology system of innovation. Various notifications with problem-specific descriptors are sent to or received from attached end-devices. In some cases, they are propagated to other directors or switches to notify peer end-devices.

Notifications are limited to registered in-zone devices and the device must register for each notification type it wants to receive.

Fabric Notifications complements MAPS technology and was introduced by Broadcom in FOS 9.0.

IBM SANnav Management Suite overview

Introduction and product family

IBM SANnav Management Suite is the next-generation management platform for IBM b-type Fibre Channel SAN environments. It provides centralized management, monitoring, troubleshooting, performance analysis, reporting, and automation capabilities through a modern browser-based interface.

Unlike previous SAN management applications that relied on Java-based clients, SANnav uses a browser-based architecture that provides a more intuitive user experience and streamlines deployment and maintenance. SANnav is designed to simplify SAN administration while supporting environments that range from a few hundred ports to large enterprise deployments with thousands of switches and ports distributed across multiple locations.

The SANnav software family consists of two products:

- IBM SANnav Management Portal: the primary SAN management application
- IBM SANnav Global View: a higher-level visibility layer across multiple Portal deployments

For the most current feature descriptions, supported platforms, known issues, and limitations, see the latest [SANnav Release Notes](#).

SANnav Management Portal

SANnav Management Portal enables administrators to perform the following tasks:

- Discover and monitor SAN fabrics
- Manage zoning configurations
- Monitor performance and capacity
- Analyze SAN health
- Monitor events and alerts
- Manage switch configurations
- Generate reports
- Track inventory information
- Troubleshoot SAN issues

A single SANnav Management Portal instance supports environments containing up to 15,000 physical SAN ports. For environments that exceed this size, multiple SANnav Management Portal instances can be deployed.




Figure: SANnav Management Suite architecture

SANnav Global View

SANnav Global View provides management visibility across multiple SANnav Management Portal instances.

Global View enables administrators to perform the following tasks:

- Consolidate information from multiple SANnav servers
- Monitor geographically dispersed SAN environments
- View summary dashboards
- Aggregate inventory information
- Track events across multiple SANnav instances
- Navigate directly into individual Management Portal servers

SANnav Global View does not replace SANnav Management Portal. Instead, it provides a higher-level management layer that consolidates information across multiple Portal deployments.

Note: SANnav Global View is compatible only with SANnav Management Portal instances running the same version. Global View requires a separate installation and license from Management Portal; both products cannot be installed on the same host or VM.

Architecture and key features in SANnav 3.0

SANnav 3.0 represents a major software release that introduces significant architectural changes in addition to support for the latest IBM b-type SAN platforms and Fabric OS releases.

SANnav 3.0 includes the following major enhancements:

- Support for Gen 8 Fibre Channel hardware platforms and blades
- Support for Fabric OS 10.0.0x new features and enhancements
- New software architecture: Docker replaced with Podman (rootless container operation)
- Updated server deployment and operating system support

- Enhanced installation, upgrade, and migration including disaster recovery
- SANnav Command Center (sannavcc) improvements
- SELinux support in enforced or permissive mode
- Quantum-safe cryptography support (FOS 10.0.0x)
- Configuration Policy Management (CPM) new features and enhancements
- Switch Configuration Backups enhancements
- Zoning new features and enhancements
- Health Summary Dashboard — Entity Profiles (new)
- Host and Storage Troubleshooting context (new)
- Fault management enhancements
- Webhooks and ServiceNow integration (new)
- Discovery enhancements
- User management enhancements
- External REST API interfaces (new)
- Flow Management enhancements
- Inventory enhancements
- Logical Fabrics enhancements
- Extension Tunnels and Circuits enhancements
- Topology enhancements
- Usability and UI/UX improvements

Transition from Docker to Podman

All versions of SANnav prior to v3.0 required a privileged Linux user (root or sudo) to install, upgrade, manage, and run the software. This was because the Docker daemon (`dockerd`) required root to operate.

SANnav 3.0 replaces Docker with Podman, which offers the following benefits:

- Daemonless operation (no root daemon required)
- Improved security through rootless container support
- Better RHEL integration (Podman is packaged with RHEL)
- Simplified lifecycle management

This architectural change reduces administrative complexity and aligns SANnav with current Red Hat Enterprise Linux containerization practices. Because of this change, upgrades from SANnav 2.x to SANnav 3.x require specific migration procedures. Administrators should review the [SANnav Management Portal Installation and Upgrade Guide](#) before planning upgrades.




Figure: SANnav 3.x replaces Docker with Podman and introduces rootless container operation

What's new in SANnav Global View 3.0

SANnav Global View 3.0 adopts the same Docker-to-Podman architecture change as Management Portal. SANnav Global View 3.0 includes the following additional enhancements:

- New software architecture: Docker replaced with Podman
- SANnav Command Center (sannavcc) enhancements
- Custom location for authorized keys SSHD file
- Security enhancements
- User management enhancements including OpenLDAP support
- Inventory enhancements
- IANA time zone support
- Usability and UI/UX improvements

Quantum-safe security

SANnav 3.x includes several security enhancements that extend beyond the container architecture change:

- Podman-based container architecture (rootless, daemonless)
- SELinux support in enforced or permissive mode (targeted mode only; MLS/minimal not supported)
- TLS-secured communications between clients, servers, and switches
- Enhanced authentication support (LDAP, RADIUS, TACACS+)
- Improved certificate management including support for third-party CA-signed certificates
- Quantum-safe code-signing verification using the Leighton-Micali Signature (LMS) algorithm to verify the integrity of SANnav software packages before installation
- Support for modern Fabric OS security capabilities including quantum-safe cryptography (FOS 10.0.0x)

Organizations should review their security requirements and align SANnav deployment practices with internal security standards.

Supported Fabric OS versions

SANnav Management Portal 3.0 supports the following Fabric OS releases:

- Fabric OS 10.0.0x
- Fabric OS 9.2.2x
- Fabric OS 9.2.1x or later
- Fabric OS 9.1.x or later
- Fabric OS 9.0.x or later
- Fabric OS 8.2.3e or later

Note: Gen 5 hardware platforms with Fabric OS versions other than 8.2.3e are managed by SANnav, however customer support is limited.

Supported hardware platforms

Generation	Speed	Model
Gen 8 Fixed-Port Switches	128 Gbps	IBM SAN56B-8 56-port Switch
Gen 8 Directors	128 Gbps	IBM SAN256B-8 4-slot Director
Gen 8 Directors	128 Gbps	IBM SAN512B-8 8-slot Director
Gen 7 Fixed-Port Switches	64 Gbps	IBM SAN24B-7 24-port Switch
Gen 7 Fixed-Port Switches	64 Gbps	IBM SAN64B-7 56-port Switch
Gen 7 Fixed-Port Switches	64 Gbps	IBM SAN128B-7 128-port Switch
Gen 7 Directors	64 Gbps	IBM SAN256B-7 4-slot Director
Gen 7 Directors	64 Gbps	IBM SAN512B-7 8-slot Director
Gen 6 Fixed-Port Switches	32 Gbps	IBM SAN24B-6 Switch
Gen 6 Fixed-Port Switches	32 Gbps	IBM SAN64B-6 Switch
Gen 6 Fixed-Port Switches	32 Gbps	IBM SAN128B-6 v2 Switch
Gen 6 Fixed-Port Switches	32 Gbps	IBM SAN18B-6 Extension Switch
Gen 6 Directors	32 Gbps	IBM SAN256B-6 Director
Gen 6 Directors	32 Gbps	IBM SAN512B-6 Director
Gen 5 Fixed-Port Switches	16 Gbps	IBM SAN24B-5 Switch
Gen 5 Fixed-Port Switches	16 Gbps	IBM SAN96B-5 Switch
Gen 5 Fixed-Port Switches	16 Gbps	IBM SAN42B-R Extension Switch
Gen 5 Directors	16 Gbps	IBM SAN384B-2 Director
Gen 5 Directors	16 Gbps	IBM SAN768B-2 Director

Kafka data streaming support

SANnav uses Kafka to provide real-time data streaming from switches running Fabric OS 9.0.x or higher. Data streaming enables collection of high-frequency performance statistics, FCIP performance and error statistics, and flow statistics.

Metric Type	Minimum Fabric OS Version
FC port metrics	Fabric OS 9.0.x or higher
FCIP metrics	Fabric OS 9.0.x or higher
Ethernet port, GigE port, and chassis metrics	Fabric OS 9.1.x or higher

Deployment and system requirements

Hardware and OS requirements

Before deploying SANnav, verify that the target server satisfies the required hardware and software prerequisites.

Deployment Size	Managed Ports	CPU Cores	CPU Sockets	CPU Speed	Memory	Storage
Small	Up to 3,000	16 cores	2 (minimum)	2,000 MHz	48 GB	600–700 GB
Large	Up to 15,000	24 cores	2 (minimum)	2,000 MHz	96 GB	1.2 TB

The following general recommendations apply:

- Dedicated deployment host (no other applications installed)
- MTU size of at least 1,500 bytes
- Reliable NTP synchronization
- Static network configuration
- Static MAC address (MAC address changes invalidate the SANnav license)

Note: SANnav Management Portal and SANnav Global View cannot be installed on the same host or VM. They can be installed on separate VMs within the same physical host if sufficient resources are available.

Because platform support changes over time, consult the latest [SANnav Release Notes](#) for current server and operating system support matrices.

Supported browsers

SANnav Management Portal and SANnav Global View are accessed through a standard web browser. No client software installation is required. Supported browsers include Google Chrome, Mozilla Firefox, and Microsoft Edge. For optimal client performance, the machine running the browser should have at least 16 GB of RAM.

Note: If accessing the SANnav client from a Remote Desktop session, user interface performance might be degraded.

SANnav 3.x introduces updated operating system requirements:

- **VM and bare-metal:** RHEL 9.6; RHEL 9.7 and 9.8 with conditional support. ESXi 8.x is the recommended hypervisor.
- **OVA:** Rocky Linux 9.6 is preconfigured in the OVA image. ESXi 8.x is required.

Always verify support for current operating system versions using the latest [SANnav Release Notes](#) and [SANnav Management Portal Installation and Upgrade Guide](#).

Deployment models

SANnav can be deployed using either of the following methods.

VM or bare-metal deployment

In this model, SANnav is installed on a dedicated physical server or a VMware virtual machine. This deployment option provides maximum flexibility and allows administrators to integrate SANnav into existing enterprise infrastructure standards.

OVA deployment

The SANnav appliance is delivered as an Open Virtual Appliance (OVA). Benefits include simplified deployment, a preconfigured operating system and storage layout, reduced installation effort, and faster time to deployment. The OVA model is generally recommended for organizations that want the quickest implementation and do not require extensive operating system customization.




Figure: SANnav deployment models

Installation planning checklist

Before installing SANnav, administrators should verify the following prerequisites:

- **Server capacity:** sufficient memory, CPU, disk capacity, and network bandwidth
- **Time synchronization:** all SAN switches and SANnav servers synchronized with the same NTP source; critical for event correlation, historical reporting, flow monitoring, and troubleshooting
- **Network connectivity:** DNS functionality, name resolution, required firewall connectivity, and SAN switch accessibility

- Security requirements: authentication requirements, LDAP integration, certificate requirements, and corporate security policies

Detailed installation procedures are intentionally not reproduced in this publication because they change more frequently than Redbooks publication cycles. See the latest [SANnav Management Portal Installation and Upgrade Guide](#) and [SANnav Global View Installation and Upgrade Guide](#) for current deployment procedures.

Network and firewall requirements

SANnav uses multiple ports for communication with SAN switches, browsers, authentication servers, syslog servers, SNMP services, and external integrations. The following table lists the ports that must be open in the firewall. These ports are automatically added to the IP tables when the SANnav server is running, so they do not need to be opened manually if the firewall runs on the SANnav server itself.

Port	Transport	Direction	Communication Path	Description
2222	TCP	Both	Client↔Server, Server↔Switch	Internal SSH server port. Open the replacement port if customized during installation.
25	TCP	Outbound	Server→SMTP Server	Email notifications without SSL or TLS.
80	TCP	Both	Client↔Server, Server↔Switch	HTTP port. Not used if HTTP-to-HTTPS redirection is disabled.
161	UDP	Outbound	Server→Switch	SNMP performance data collection.
162	UDP	Inbound	Switch→Server	SNMP trap port.
389	TCP	Outbound	Server→LDAP Server	LDAP authentication without SSL. Required only when using AD or CA LDAP without SSL.
443	TCP	Both	Client↔Server, Server↔Switch, Server↔vCenter	HTTPS secure access. Open the replacement port if customized during installation.
465	TCP	Outbound	Server→SMTP Server	Email notifications with SSL.
514	UDP	Inbound	Switch→Server	Syslog port. Not required if port 6514 is used.
587	TCP	Outbound	Server→SMTP Server	Email notifications with TLS.
636	TCP	Outbound	Server→LDAP Server	LDAP authentication with SSL. Required only when using AD, CA, or OpenLDAP with SSL.
3269	TCP	Outbound	Server→LDAP Global Catalog	LDAP Global Catalog port for authentication.
6514	TCP	Inbound	Switch→Server	Secure syslog port. Not required if port 514 is used.
9443	TCP	Outbound	Server→Dell SCG	Required only for Dell Technologies Call Home Center (SCG Gateway).

Port	Transport	Direction	Communication Path	Description
18082	TCP	Inbound	Switch→Server	Avro schema registry HTTPS port (Fabric OS 9.0.1 and higher). Required for Kafka streaming.
19094	TCP	Inbound	Switch→Server	Secured Kafka streaming port (IPv4 switches).
19095	TCP	Inbound	Switch→Server	Secured Kafka streaming port (IPv6 switches). Not required if port 19094 is used.

If using the Disaster Recovery feature, additional ports must be opened. See the [SANnav Management Portal Installation and Upgrade Guide](#) for details.

If an external authentication server (LDAP, RADIUS, or TACACS+) or email server (SMTP) is configured, verify that the SANnav server can reach the relevant ports listed above. Because firewall requirements change between releases, administrators should consult the latest [SANnav Management Portal Installation and Upgrade Guide](#) when planning firewall configurations.

Licensing and subscriptions

SANnav Management Portal and SANnav Global View are licensed separately. Each installed instance generates a unique Server UID that is tied to the host or VM. You cannot use the same license on a cloned VM.

SANnav licensing is subscription-based. The licensing process includes the following steps:

1. Obtaining the Server UID from the SANnav Licensing page
2. Obtaining a transaction key from IBM when purchasing a license
3. Generating the license file at the Broadcom Licensing Portal using the Server UID and transaction key
4. Applying the license XML file to SANnav within 30 days of generation

Note: The license file must be applied within 30 days of being generated, independent of the subscription expiration date. If not applied in time, the license must be regenerated at the [Broadcom Licensing Portal](#).

The following IBM part numbers apply for 1-year renewals:

Product	IBM Part Number
SANnav Management Portal Base, 1-Year Renewal	9240 B03 ANYY
SANnav Management Portal Enterprise, 1-Year Renewal	9240 E03 ANYZ
SANnav Global View, 1-Year Renewal	9240 G03 ANYO

All initial SANnav subscriptions begin when the order is processed by Broadcom, not when the license file is generated or applied.

SANnav provides progressive warnings as the license approaches expiration:

Time Before Expiry	Behavior
90–30 days	Warning message at login with option to dismiss for 7 days
29 days to expiry	Warning message displayed daily; cannot be dismissed
After expiry	Login blocked; SANnav UI inaccessible

Note: Although SANnav blocks UI access when the license expires, the backend server continues to collect data from the managed environment. No data is lost when a valid license is applied.

SANnav can be configured to automatically retrieve and activate renewed licenses from the [Broadcom Licensing Portal](#). This feature requires an active internet connection from the SANnav server to <https://api.broadcom.com/swpoemservice/>. Auto-retrieval can be enabled or disabled via the SANnav Command Center (`sannavcc`).

Licenses can also be managed in the following ways:

- Renewed: purchase a renewal through IBM or Broadcom sales channels
- Regenerated: if the 30-day application window expires, regenerate at the [Broadcom Licensing Portal](#)
- Rehosted: to move SANnav to a different server without purchasing a new license

Detailed licensing procedures are documented in the [SANnav Management Portal Installation and Upgrade Guide](#) and the [Broadcom Licensing Portal](#) documentation.

Administration and management interfaces

CLI management: SANnav Command Center (`sannavcc`)

The SANnav Command Center (`sannavcc`) is the primary CLI tool for managing the SANnav server without a browser. It is located in the `<install_home>/bin` directory and is run as the `sannavmgr` user.

Using `sannavcc`, administrators can perform the following tasks:

- Check SANnav server status
- Start, stop, and restart SANnav services
- Install SANnav prerequisites
- Install or upgrade SANnav
- Configure ports, swap location, and other server settings
- Enable or disable license auto-retrieval
- Manage SSL certificates
- Uninstall SANnav

To check server health at any time, run the `check-sannav-status.sh` script from `<install_home>/bin`. Example output from a healthy server:

SANnav server is healthy.
All the services are currently in running state.

SANnav monitors disk usage every 10 minutes with the following alert thresholds:

Threshold	Severity
70% disk used	Warning
80% disk used	Error
90% disk used	Critical

Element management — IBM SANnav WebTools

WebTools is the browser-based element manager built into each Fabric OS switch. It can be launched directly from a browser or from within SANnav Management Portal. You can start WebTools on any workstation with a compatible operating system and web browser installed.

Note: WebTools does not require a license. It is installed on the switch when you install Fabric OS. WebTools can be enabled or disabled in FOS.

Launching WebTools

If the switch is configured with logical fabrics, you can log in to any of the logical fabrics for which you have permission.

Note: To ensure that the latest version of WebTools is loaded in the browser, close the previous version of the WebTools browser and clear the browser cache before launching the new version.

From a web browser: Open a browser, enter the IP address of the switch, and press Enter. You can use HTTP or HTTPS, for example: <http://10.77.77.77> or <https://10.77.77.77>

Note: Starting with FOS v10.0.1, users can open a page in a new tab using Ctrl+Click or Cmd+Click. This action opens the specified page directly without requiring re-authentication.

From SANnav Management Portal: On the SANnav Inventory page, click the down arrow to the right of the switch and select **View in WebTools** from the action menu.

Note: You can also launch WebTools from the SANnav Management Portal Health Summary dashboard. For detailed information, see the [SANnav Management Portal User Guide](#).




Figure: Starting WebTools from SANnav Management Portal

Note: Record the firmware version of the switch. The corresponding version of WebTools launches, which might not be a version that is supported by this guide.

Enter the username, password, and logical switch name or fabric ID (FID). For information about launching WebTools (SSO or non-SSO) from SANnav, see the [SANnav Management Portal User Guide](#). If you are logging in to a virtual fabrics-enabled platform and you do not specify a logical switch, you are logged in to the default FID of the switch.

Note: Starting with FOS v10.0.1, the Logical Switch ID column is displayed regardless of whether virtual fabrics is supported or not supported on the platform.

For the first switch login, the default credentials are as follows:

- Username: `admin`
- Password: `password`




Figure: WebTools Element Manager login screen

WebTools prompts you to change the default password on first login. If you are logging in to a Virtual Fabrics-enabled platform and you do not specify a logical switch, you are logged in to the default logical switch, which uses fabric ID 128. For non-VF platforms, the FID option is not displayed. If you start from SANnav Management Portal, you might not need to log in, depending on the SANnav single sign-on configuration. Click **Login**.

WebTools user interface

After you are familiar with the basic components of WebTools, you can quickly start monitoring and managing your switch.




Figure: WebTools user interface overview

The WebTools UI consists of the following components:

1. **Navigation bar**: contains links to feature pages
2. **Logical switch drop-down**: displays links to other FIDs in the device
3. **Profile menu**: displays the logged-in user name, logout link, and dark theme toggle
4. **Title bar**: provides the page title and includes sub-navigation menus and buttons to take action within the page
5. **Content section**: displays the switch image and its respective details

The browser's title bar displays the switch name and switch IP address. Some pages also include a filter bar for searching and filtering the displayed content.




Figure: WebTools filter bar

The **Settings** page contains navigation options on the left side of the page. Clicking each item on the left navigation bar displays additional options. For example, to configure user accounts, click **Settings** → **Security Policies** → **User Management**.




Figure: WebTools Settings page

Some tables have an action menu that you can access by clicking the down arrow in the rightmost column. Click this arrow to display additional actions that you can perform on the associated object.




Figure: Using the down arrow to display additional actions in WebTools

Clicking **View** in the action menu opens a detail page for that object. You can also open the detail page by clicking the name of the object in the table. The detail page displays additional information about the object and might contain additional actions that you can perform.




Figure: WebTools detail page for a switch port

Note: Clear the browser cache before launching a new version of WebTools to ensure the latest version loads correctly.

For more information about the SANnav WebTools product, see the [SANnav Management Portal User Guide](#).

Initial setup and configuration workflow

After SANnav Management Portal is installed and licensed, administrators should perform the following configuration tasks to begin monitoring and managing the SAN environment effectively. These tasks establish fabric visibility, event notification, security, and operational monitoring capabilities.

Discover SAN fabrics

The first task is to discover the SAN infrastructure. SANnav requires connectivity to at least one seed switch in each fabric. During discovery, SANnav collects information about switches, fabrics, ports, hosts, storage systems, and inter-switch links, building a complete inventory view of the environment. After discovery is completed, administrators should verify that all expected devices and fabrics are visible and operating normally. For more information, see *Discovering a Fabric* in the [SANnav Management Portal User Guide](#).

Configure event notifications and SNMP

To ensure timely notification of operational issues, configure SANnav event notification services. SANnav supports email notifications, syslog forwarding, SNMP traps, and Call Home integration. Configure the SMTP mail server and Call Home, register the SANnav server as a syslog and SNMP trap recipient, and forward SNMP traps and syslog messages. For more information, see *Call Home and ESRS* and *Event Management* in the [SANnav Management Portal User Guide](#).

Configure user access and authentication

Administrators should configure user accounts, roles, and Areas of Responsibility (AORs) according to organizational requirements. SANnav supports role-based access control and can integrate with external authentication services such as LDAP, RADIUS, and TACACS+. Add users, roles, and AORs, configure a password policy, and set up external authentication. Password policies and account management settings should also be reviewed to align with security standards. See *Security* in the [SANnav Management Portal User Guide](#).

Configure MAPS monitoring and health policies

Monitoring and Alerting Policy Suite (MAPS) policies provide proactive health monitoring for SAN environments. Administrators should review the default policies and customize thresholds as necessary to match operational requirements. You can create a policy or import policies from a JSON or XML file — for example, you can import MAPS policies that were exported from Brocade Network Advisor. Properly configured monitoring policies help identify congestion, link issues, latency problems, and other abnormal conditions before they affect application performance.

Configure performance and telemetry collection

Verify that telemetry and performance-data collection are functioning correctly across all managed switches. SANnav uses streaming telemetry and performance statistics to provide health monitoring, capacity analysis, and troubleshooting capabilities. Confirm that performance data is being collected and displayed for switches, ports, hosts, and storage devices.

Create dashboards and reports

SANnav includes several predefined dashboards that provide visibility into fabric health, inventory, events, and performance metrics. Administrators can also create custom dashboards tailored to operational requirements. Custom views and reports help teams focus on key performance indicators and simplify ongoing monitoring activities.

Important: Following the completion of the initial configuration tasks, administrators should review dashboards, inventory views, alerts, and health summaries to confirm that fabrics, switches, hosts, and storage resources are being discovered and monitored correctly. Any configuration, discovery, or telemetry issues should be addressed before SANnav is used for ongoing production monitoring.

SANnav Management Portal user interface

By using the SANnav Management Portal, you can manage one or more SAN fabrics in multiple locations. After you are familiar with the basic components of SANnav, you can quickly start monitoring and managing your fabrics.

The basic layout of the SANnav user interface is shown in the following figure.



Figure: Basic layout of the SANnav user interface

The following components are shown in the previous figure:

1. **Navigation bar:** Contains links to feature pages. The SANnav link displays the Configurations and Settings page.
2. **Search:** Click the icon to display the search box, select a context on which to search from the list, and enter the search term in the field.
3. **Profile menu:** Click the icon to display more links for changing user preferences, displaying the SANnav version, and logging out.
4. **Subnavigation bar:** Provides the page title and optional item count within parentheses. Also includes buttons and menus to take actions within the page. The subnavigation bar is the main way to navigate within a feature.
5. **Filter bar:** Allows you to filter the display based on columns, fabrics, network scope, and customized filters.
6. **Expandable sidebar:** Provides an area where you can save selected inventory items for investigation later.

Dashboards and reports

SANnav provides several default dashboards that provide visibility into fabric health, inventory, events, and performance metrics. Dashboards provide summary status and performance information, from which you can drill down to get detailed views. By using filters and tags, you can sort and search your inventory to find the information that you want. A highly flexible reporting infrastructure allows you to generate custom graphical or tabular reports.

Detail pages

Clicking a fabric name, switch name, or port name in a table opens a detail page for that object. The detail page displays more information about the object and can include tasks that you can perform.




Figure: SANnav detail page for a switch

The detail page is different, depending on the context. For example, the detail page for a fabric on the Inventory page is different from the detail page for the same fabric on the Discovery page.

Tables

In tables, if an entry is truncated (indicated by an ellipsis at the end of the entry), hover over the entry to see the full text.




Figure: Hovering the mouse over truncated table text

Many tables include an action menu that you can access by clicking the down arrow in the rightmost column. Click this arrow to display more tasks that you can perform on the associated object.




Figure: Using the down arrow to display more actions

On the right side of the subnavigation bar, the **More** button often contains a **Bulk Select** option, which allows you to select and perform operations on one or more items at the same time. For example, you can turn on monitoring for multiple fabrics, or you can create a tag and apply it to multiple switches.




Figure: Bulk Select option in the subnavigation bar

SANnav Global View user interface and multi-site dashboard

After you are familiar with the basic components of SANnav Global View, you can quickly start monitoring SANnav Management Portal instances. When you first log in to SANnav Global View, you see the Summary dashboard — the single dashboard provided with Global View.

The basic layout of the SANnav Global View user interface is shown below.




Figure: SANnav Global View Summary dashboard

The following components are shown in the previous figure:

1. **Navigation bar:** Contains links to feature pages. The SANnav link displays the page where various settings and configurations can be performed.
2. **Profile menu:** Displays links for changing user preferences, displaying the SANnav version, and logging out.
3. **Subnavigation bar:** Provides the page title and optional item count within parentheses. Also includes buttons and menus to take actions within the page.
4. **Filter bar:** Allows you to filter the display based on columns, portal scope, and customized filters.

Health states and dashboard widgets

The SANnav Global View dashboard provides comprehensive visibility across multiple SANnav Management Portal instances. The Global View dashboard auto-refreshes every 15 minutes and provides the following widgets:

- **Health summary:** shows the count of fabrics, switches, hosts, and storage entities in each health state
- **Switch Health:** displays switch health grouped by Portal instance, firmware version, or model
- **Port Usage Summary:** summarizes port utilization across all managed environments
- **Alerts:** consolidates alerts from all connected Management Portal instances

These widgets indicate the status of fabric, switch, host, and storage entities managed by the Portals added to SANnav Global View. The health summary widgets are categorized by health state: Healthy, Degraded, or Poor. The health state for each object is determined by its health score, which is calculated from a set of predefined factors. For example, a fabric health score is influenced by whether a link is down or redundant paths are missing.

Fabric Health: 20 Poor						
 20 Items						
Name ^	Score	Switch Status	Events	Configurations	Portal	
Fabric_A	64	-30	-2	-4	New York DC	Show Details Show in SANnav Portal
FabricTest_27	70	-30	0	0	M	
FcoE	68	-30	0	-2	M	
FIDnine	70	-30	0	0	Munich DC	
Close						

Figure: SANnav Global View health states and scores

SANnav dashboards allow you to gather more information by drilling down into the dashboard widgets. By default, the scope of SANnav Global View is set to All Portals (all Management Portal instances). You can choose one or more Management Portal instances to change the scope.





Figure: SANnav Global View Summary dashboard with All Portals scope

Inventory, events, and settings

Click **Inventory** to display inventory information about fabrics, switches, switch ports, chassis, host and storage devices, and host and storage ports across all Management Portal instances.




Name	Type	FID	Fabric	Health	Model	Port Count	Portal
A__15643...	Switch	128	FabricTest...	Poor	Brocade X6-8	221	Munich DC
AG_17-mo...	Access Gate...	-	FCIP-104	Healthy	Brocade 6505	24	San Jose DC
AG_17-mo...	Access Gate...	-	Untitled N...	Healthy	Brocade 6505	24	Untitled New Portal

Figure: SANnav Global View Inventory page

Click **Events** to show all application events that stem from a user or system action. A system action is triggered under specific situations, such as when a portal is disconnected. You can filter the events list by date range.

Note: Events generated within individual SANnav Management Portal instances are not displayed on this page.



Description	Module	Category	Username	Occurred Time	Actions	Probable Cause
Retrieving SANnav license for Serial Number GLOBAL00001 failed.	Management Server License	SANnav Application Event	-	Aug 20, 2024 11:22:09 IST	-	No License Query details matched the given criteria
SANnav GLOBAL00001 license expires in 25 days. Please renew soon.	Management Server License	SANnav Application Event	-	Aug 20, 2024 02:10:09 IST	-	-
Retrieving SANnav license for Serial Number GLOBAL00001 failed.	Management Server License	SANnav Application Event	-	Aug 19, 2024 11:22:09 IST	-	No License Query details matched the given criteria
SANnav GLOBAL00001 license expires in 26 days. Please renew soon.	Management Server License	SANnav Application Event	-	Aug 19, 2024 02:10:09 IST	-	-
Retrieving SANnav license for Serial Number GLOBAL00001 failed.	Management Server License	SANnav Application Event	-	Aug 18, 2024 11:22:09 IST	-	No License Query details matched the given criteria
SANnav GLOBAL00001 license expires in 27 days. Please renew soon.	Management Server License	SANnav Application Event	-	Aug 18, 2024 02:10:09 IST	-	-
Retrieving SANnav license for Serial Number GLOBAL00001 failed.	Management Server License	SANnav Application Event	-	Aug 17, 2024 11:22:09 IST	-	No License Query details matched the given criteria
SANnav GLOBAL00001 license expires in 28 days. Please renew soon.	Management Server License	SANnav Application Event	-	Aug 17, 2024 02:10:09 IST	-	-

Figure: SANnav Global View Events page

Click **SANnav** to discover Management Portals, manage filters, and select various configuration settings across several feature categories. The profile menu is where you can add your phone number and change your password. It is also where you can enable **Persist Last Column Selection**, wherein the system remembers your column customization, and **FICON Display**. When **FICON Display** is enabled, specific columns in the switch and switch port inventory pages are rearranged for FICON mode.




Figure: SANnav Global View Administrator User Preferences

Drill-down fabric health monitoring

The Fabric Health widget shows the number of fabrics in Healthy, Degraded, and Poor state. The least overall score across all objects in the group is shown in the center of the circular widget.

To investigate the health state of fabrics across all SANnav Management Portal instances:

1. Click **Dashboard & Reports** in the navigation bar.
2. Hover over a colored portion of the Fabric Health widget to display details of the fabrics with that state. For example, hover over the red portion to display a list of portals and the number of fabrics in each portal having a poor state.




Figure: SANnav Global View Fabric Health widget showing poor health state

3. To view more details, click the colored segment of the widget. For example, clicking the red segment displays a table showing health scores for fabrics in poor health.
4. Click **Show Details** in the drop-down menu of the fabric list to see the factors contributing to that health score. The **Show Details** option is available only if the score is less than 100.

Fabric Health: 20 Poor						
 20 Items						
Name ^	Score ⇅	Switch Status ⇅	Events ⇅	Configurations ⇅	Portal ⇅	
Fabric_A	64	-30	-2	-4	New York DC	Show Details Show in SANnav Portal
FabricTest_27	70	-30	0	0	M	
FcoE	68	-30	0	-2	M	
FIDnine	70	-30	0	0	Munich DC	
Close						

Figure: SANnav Global View Show Details option for fabric health score

A pop-up window displays the factors that cause deductions in the health score for this fabric.

- Click **Back** to return to the list of portals.
- Select **Show in SANnav Portal** from the drop-down menu to open the Management Portal instance responsible for the data you are viewing. If you are not logged in to this portal, credentials are required.

Note: You might need to allow pop-ups in your browser for the portal window to open.

Cross-portal switch health navigation

The Switch Health widget displays details of switch health across Management Portal instances. You can get a summarized view and details of switches based on firmware version, Management Portal instance, switch model, and product category.

To display switch health across multiple Management Portal instances:

- Click **Dashboard & Reports** in the navigation bar. The Switch Health widget appears in the bottom row of the dashboard, showing switches grouped by SANnav Portal instances. You can select a different grouping from the drop-down list.




Figure: SANnav Global View Switch Health widget grouped by Portal

2. Hover over the graph to display details. For the Firmware Version and Model categories, more bars may be displayed than are listed in the row names on the left. In this case, hover over an intermediate bar to display the details.




Figure: SANnav Global View Switch Health widget grouped by model

- To see the details of all switches in a specific category, click the bar chart of the widget, then select **Show Details**.

Product Health: Poor						
Munich DC						
16 Items						
Name	IP Ad...	WWN	State	Status	Status Reason	Fabric
A_1564...		10:00:C4:...	Online	Critical	Switch Status is CRITICAL. Contributors:*B...	FabricTest_27
stinger16		10:00:50:...	Unknown	Not Reach	Switch Status is CRITICAL. Contributors:*BAD_PWR (CRITICAL). *BAD_FAN (CRITICAL).	est_27
sw2		10:00:C4:...	Online	Marginal	Switch Status is MARGINAL. Contributors:...	Switched New ...
sw3		10:00:C4:...	Online	Critical	Switch Status is CRITICAL. Contributors:*B...	Fcoe

Close

Figure: SANnav Global View Product Health poor state with Show Details

4. Click **Close** to return to the dashboard.
 5. Select **SANnav Portal** from the drop-down list in the widget, click the bar graph, then select **Show in Portal** to display the switch inventory in the selected SANnav Management Portal instance. If you are not logged in to this portal, credentials are required. A separate window opens for the Management Portal instance.
-

Documentation and references

This section provides only an overview of SANnav architecture and deployment considerations. Detailed procedures are documented in the following resources.

Release documentation

- [SANnav Management Portal Release Notes](#)
- [SANnav Global View Release Notes](#)

Installation documentation

- [SANnav Management Portal Installation and Upgrade Guide](#)
- [SANnav Global View Installation and Upgrade Guide](#)

User documentation

- [SANnav Management Portal User Guide](#)
- [SANnav Management Portal Flow Management User Guide](#)
- [SANnav Global View User Guide](#)
- [SANnav REST API Reference Guide](#)

These documents should be considered the primary source for installation procedures, supported configurations, known issues, and operational guidance.

Summary

IBM SANnav Management Suite delivers a modern, browser-based platform for managing IBM b-type SAN environments at scale. SANnav Management Portal supports operational SAN management activities, including monitoring, configuration, reporting, and troubleshooting, while SANnav Global View provides consolidated visibility across geographically distributed SANnav Management Portal instances.

SANnav 3.x introduces significant architectural improvements, including the transition from Docker to Podman, support for current Fabric OS releases, enhanced security capabilities, and improved deployment flexibility. For detailed installation, configuration, and operational procedures, consult the latest [SANnav product documentation](#) and [release notes](#).

Fibre Channel SAN design

The shift to NVMe-ready Fibre Channel technologies introduces new requirements for storage area network design. This is true regardless of the technology that is used. If the expectation is to use the capacity and performance of these platforms, careful attention must be given to SAN design.

Dedicated storage network infrastructure that provides lossless, low-latency, deterministic, scalable, and high-performance storage services to the applications becomes critical. Storage administrators talk about "fan-in" or "fan-out" ratios for storage platforms, which are the number of servers or applications in the network that use a specific array or array port for their storage access.

Depending on the types of applications and their performance needs, that ratio might range anywhere from low single digits to 40–50 servers (hundreds of virtual machines). As with any provisioning scenario, the storage administrator is dealing with projections of how much capacity and performance any server or application is going to use. But application performance is variable; time of day, week, month, and seasonal or event-driven events cause spikes or drops in application demand.

Another consideration is that applications are rarely refreshed simultaneously. In fact, the more common scenario is that multiple generations of performance exist in the environment simultaneously. This is driven by the ability to take a service window to re-platform existing environments to new server and storage acquisitions and the fact that some established applications might not feature an environment that runs in current operating system releases or application versions.

You might have 10-year-old or older operating systems that are running connected with a host bus adapter (HBA) that is two or more older generations of technology behind. How then do you balance that still critical application against the needs and performance of the newer machines?

The answer to this question is a combination of topology, balanced provisioning, granular monitoring, and automated mitigation.

SAN design basics

This section provides high-level guidelines for designing a typical SAN. The focus is on best practices for core-edge or edge-core-edge fabrics. The discussion starts at the highest level, the data center, and works down to the port level, providing recommendations at each point along the way.

Topologies

A typical SAN architecture comprises edge devices, network devices, and cabling. Topology is usually described in terms of how switches are interconnected, such as collapsed-core, core-edge, and full-mesh. The recommended SAN topology to optimize performance, availability, management, and scalability is a tiered, core-edge topology. This approach provides excellent performance without unnecessary interconnections. At a high level, the tiered topology uses many edge switches for device connectivity and a smaller number of core switches for routing traffic between the edge switches. The following figure shows examples of these topologies for a single SAN fabric (Fabric A or Fabric B). In dual-fabric SAN deployments, which are recommended for high availability, the second fabric is implemented as an independent mirror image of the topology shown.




Figure: Three scenarios of tiered network topologies

The primary difference between these three scenarios is device placement, where devices are connected to the network, and the associated traffic flows.

Scenario A: A collapsed-core architecture localizes traffic to a single platform. On a platform, each ASIC performs local switching. The receiving ASIC either switches a flow out to the end device or sends it to the next closer ASIC. A collapsed-core can have small performance advantages for performance-critical, latency-sensitive workloads. However, local switching does not scale beyond an ASIC's port group. A collapsed-core architecture can also reduce manageability by having one platform per fabric; the overall number of fabrics determines manageability, and most SANs have two fabrics.

Scenario B: A core-edge topology separates storage and server connectivity, providing ease of management and greater scalability. A core-edge topology has only one fabric hop from server to storage, providing performance comparable to full-mesh while allowing greater scalability. Core-edge SAN architecture is the most deployed design.

Scenario C: A full-mesh fabric has no more than one hop between server and storage, assuming the server and storage are not connected to the same platform. Designing fabrics with UltraScale ICLs is an efficient way to save valuable FC ports. Using best-practice SAN design considerations, users can quickly build a large fabric with 3456 ports or more.

Collapsed-core topology

Collapsed-core topology is often employed for high-performance storage connectivity, particularly for critical applications. In a collapsed-core topology, storage and servers connect to a single ASIC on a switch or director, which enables local switching. Redundant A and B fabrics are still used. A collapsed-core topology is also a consideration that drives some administrators to place storage ports in the same port group as the server ports, whether on a blade or a switch.

The collapsed-core topology places initiators (servers) and storage (targets) on the same chassis and potentially on the same blade and ASIC. This configuration realizes several benefits, depending on the size of the environment. It is frequently used when customers migrate from multiple switches to a single, dual-core architecture where all initiators and targets fit. If future design requirements include increased capacity, moving to a core-edge architecture is generally more beneficial in the long term for scalability and ease of management.

Core-edge topology

The core-edge topology places initiators (servers) on the edge tier and targets (storage) on the core tier. For redundancy, each fabric, A and B, has two cores. Because servers and storage are on different switches, this topology provides easy management, outstanding performance, and minimal fabric latency. Most traffic traverses only one hop from edge to core.

Storage-to-storage traffic requires two hops if the second storage platform destination is not connected to the same core. The two cores within the same fabric can be connected if storage-to-storage connectivity is required. The disadvantage of a core-edge design is that storage and core-to-edge connections contend for expansion as the environment scales; however, director platforms provide flexibility by allowing ICLs for inter-switch connectivity, freeing ports for additional devices.

Full-mesh topology

Another alternative topology is a full-mesh architecture. Every switch has direct ISLs to every other switch in the fabric. A full-mesh topology is problematic for environments with many switches because ISLs consume valuable ports that could otherwise be used to connect servers and storage. Director platforms have ICLs, allowing exceptional bandwidth and scale between chassis without consuming front-end ports. A full-mesh topology ensures that no end device is more than one hop away from another.

A full-mesh topology allows servers and storage to be placed anywhere because communication between source and destination is no more than one hop. This design uses director-class switches with UltraScale ICL ports for interconnectivity to ensure maximum device port availability and utilization. Design this architecture with a minimum of two switches and up to nine switches in a full-mesh topology.

Capabilities to consider for your storage architecture

A considerable number of critical capabilities must be considered when choosing the right storage strategy. Some of the more important considerations are availability, scalability, performance, extensibility, and manageability. For modern SAN environments, this also includes support for high-performance, latency-sensitive workloads, transmission reliability, and deterministic end-to-end data flows.

Availability, redundancy, and resiliency

An essential aspect of SAN architecture is fabric resiliency and redundancy. The objective is to remove single points of failure. Resiliency is the network's ability to continue functioning after a failure, while redundancy describes the duplication of components, typically the entire fabric, to eliminate a fabric failure as a single point of failure.

Availability

Brocade fabrics have resiliency built into Fabric OS (FOS), which runs on Brocade platforms. FOS can repair and overcome failures. For example, the Fabric Shortest Path First (FSPF) routing algorithm computes a new path when a fabric topology changes, such as when a link goes offline, assuming a second path exists where fabric resiliency is required.

Redundancy

The key to high availability and enterprise-class availability is redundancy. Business continuance is provided through most foreseeable and unforeseeable events by eliminating an entire fabric as a single point of failure. At the highest level of fabric design, the complete fabric should be redundant, with two mirrored, entirely separate fabrics that do not share common SAN platforms, routers, or switches.

Servers and storage devices should be connected to both fabrics, commonly referred to as Fabric A and Fabric B, using multipath I/O (MPIO) so that data can flow across both fabrics seamlessly in an active/active or active/passive mode. MPIO ensures that an alternate path is available if the current path fails. Ideally, redundant fabrics are identical or based on the same switch platforms to ensure consistency of performance and delivery.

Often, Fabric A and Fabric B reside in the same location. However, two locations, such as local and asynchronous sites, or three locations, such as local, synchronous, and asynchronous sites, are frequently used to provide disaster recovery. In those designs, a separate replication SAN is constructed so operations can continue after a disaster, either at the remote location or after data has been retrieved.

Resiliency

Regardless of the physical geography, production environments typically use two production fabrics for redundancy and may also use two replication fabrics. Replication fabrics connect directly to the arrays' replication ports and do not connect to the production fabrics.

Recommendations for SAN design are to ensure application availability and resiliency through the following guidelines:

- Fabric redundancy to avoid a fabric becoming the single point of failure
- Resiliency built into each fabric to prevent a single point of failure
- Redundant connections from each host to each fabric
- MPIO-based failover from initiator to target
- Identical architectures and platforms in each fabric where possible
- Redundant ISLs, inter-fabric links (IFLs), and ICLs for inter-switch connectivity
- Separate storage, or core-tier, and server, or edge-tier, layers for independent expansion
- Core switches of equal or higher performance compared to the edge switches
- The highest-performance switch in the fabric defined as the principal switch where applicable

Scalability

IBM B-type Gen 8 and later products are designed with scalability in mind, recognizing that most installations continue to expand and that growth must be supported with few restrictions. Following the same basic principles as the network grows helps ensure that performance, availability, redundancy, and resiliency continue to meet enterprise requirements.

Evaluate the effect on topology, data flow, workload, performance, redundancy, and resiliency of the entire fabric whenever one of the following actions is performed:

- Adding or removing initiators, including changes in workload and provisioning.
- Adding or removing storage, including changes in provisioning and storage media type, such as increased deployment of flash-based storage.
- Adding or removing switches.
- Adding or removing ISLs and ICLs.
- Changing virtualization strategies, workload placement, storage strategies, or traffic-flow patterns.

If these design best practices are followed when the network is deployed, small incremental changes should not adversely affect the availability and performance of the network. However, if changes are ongoing and the fabric is not correctly evaluated and updated, both performance and availability can be jeopardized.

Consider the following points when reviewing the status of a production Fibre Channel network for redundancy and resiliency:

- Are there at least two physically independent paths between each source and destination pair?
- Are there two redundant production fabrics?
- Does each host connect to two different edge switches or equivalent fabric access points?
- Are edge switches connected to at least two different core switches?
- Are inter-switch connections composed of redundant trunks with at least two ISLs where appropriate?
- Does each storage device connect to at least two different edge switches or separate port blades?
- Are storage ports provisioned so that every host has at least two paths through which it can access LUNs?

- Are replication fabrics separated from production fabrics when disaster recovery replication is required?

Performance

IBM B-type Gen 8 technology substantially accelerates a data center environment, offering 50% lower latency than the previous generation, while increasing bandwidth with 128 Gbps links. By using NVMe-ready FC technology, IBM B-type Gen 8 maximizes the performance of NVMe storage and high-transaction workloads and simplifies the process of integrating higher-performing NVMe storage into the environment without a rip-and-replace process.

Review the following performance requirements:

- Host-to-storage port fan-in/out ratios
- Over-subscription ratios:
 - Host to ISL
 - Edge switch to core switch
 - Storage to ISL
- Size of trunks
- Routing policy and currently assigned routes: Evaluate for potential imbalances
- Use of FEC for all ISLs and connections to Gen 5 (if supported) and Gen 6 devices
- Watching for effects such as:
 - Poor storage performance
 - Overloaded hosts or applications
 - Distance issues over constrained long-distance links resulting from changes in usage, such as adding mirroring, or too many workloads
 - Deteriorating optics that result in declining signal strength and increased error rate

Manageability

The IBM B-type Gen 8 FC switches integrate with the IBM SANnav Management Suite, which provides centralized management and visibility across SAN environments. SANnav simplifies fabric administration through features such as configuration management, monitoring, reporting, and diagnostic tools that help administrators maintain SAN health and performance.

Designing a SAN

The SAN planning process is like any type of project planning and includes the following phases:

- Phase I: Gathering requirements
- Phase II: Developing technical specifications
- Phase III: Estimating project costs
- Phase IV: Analyzing Return on Investment (ROI) or Total Cost of Ownership (TCO), if necessary
- Phase V: Creating a detailed SAN design and implementation plan

When selecting which criteria to meet, engage users, server, and storage subject matter experts (SMEs), and other relevant experts to understand the role of the fabric.

Because most SANs tend to operate for a long time before they are renewed, consider future growth because SANs are difficult to redesign. Deploying new SANs or expanding existing ones to meet more workloads in the fabrics requires a critical assessment of business and technology requirements.

A focus on planning ensures that the SAN, once deployed, meets all current and future business objectives, including availability, deployment simplicity, performance, future business growth, and cost.

Design worksheets that assist in designing a SAN environment are provided as a reference for documenting assets and metrics for SAN projects.

A critical aspect for successful implementation that is often overlooked is the ongoing management of the fabric. Identifying systems-level SMEs for all components that make up the SAN, and adequate and up-to-date training on those components, is critical for efficient design and operational management of the fabric. When designing a new SAN or expanding an existing SAN, consider the following parameters:

Application virtualization

Consider the following application virtualization parameters:

- Which applications will run under a virtual machine (VM) environment?
- How many VMs will run on a physical server?
- Under what conditions will the VMs be migrated (business and non-business hours; is more CPU or memory needed to maintain response times)?
- Is there a need for solid-state storage media to improve read response times?
- Homogeneous and heterogeneous server and storage platforms:
 - Is auto-tiering in place?
 - Which FOS versions are supported in a multivendor storage environment?
 - What is the planned refresh cycle of servers and storage platforms (2 years or 3 years)?
- Sizing:
 - How many user ports are needed now?
 - How many devices will connect through an access gateway?
 - How many inter-switch links (ISLs)/Brocade UltraScale inter-chassis links (ICLs) are required to minimize congestion in the fabric?
 - What distances for ISL/ICL connections need to be supported?
 - Does the fabric scale out at the edge or the core?

Backup and disaster tolerance

Consider the following backup and disaster tolerance parameters:

- Is there a centralized backup? (This parameter determines the number of ISLs that are needed to minimize congestion at peak loads.)
- What is the impact of backup on latency-sensitive applications?
- Is the disaster solution based on a metro FC or Fibre Channel over IP (FCIP) solution?

Diagnostics and manageability

Consider the following diagnostics and manageability parameters:

- What is the primary management interface to the SAN (command-line interface [CLI], Brocade SANnav, or third-party tool)?
- How often will Brocade FOS and SANnav be updated?
- How is cable and optics integrity validated?
- Is support needed for adding Gen 8 switches into a Gen 8 fabric?
- Is support needed for storage technologies, such as NVMe over Fabrics?
- What device interoperability support is required?
- Is interoperability required for other technologies, such as UCS?

High-performance latency-sensitive workloads

Enterprises increasingly leverage low-latency, high-throughput flash arrays for demanding, performance-sensitive workloads. Brocade Gen 8 Fibre Channel is well suited for these workloads because it provides sub-microsecond switch latency, high-speed bandwidth, and accurate I/O instrumentation. Performance testing has demonstrated that all-flash arrays can significantly benefit from Gen 8 SANs.

Fibre Channel standards include forward error correction (FEC), RS(544,514), to ensure transmission reliability and highly deterministic data flows. A single FEC can correct up to fifteen 10-bit symbol errors and detect up to thirty 10-bit symbol errors, helping avoid frame retransmits when bit errors are corrected. For Gen 8, dual FEC is enabled by default.

For highly demanding workloads, a no-hop fabric connection through a one-ASIC switch, such as the Brocade G820, or local switching within an ASIC on a director port blade minimizes latency to sub-microsecond speeds, measured in nanoseconds. Local switching performs cut-through switching of Fibre Channel frames from the ingress port to the egress port when both ports are in the same port group. Some platforms, such as Brocade X7 and X8 directors and G730 switches, traverse up to three switching ASICs between data ingress and egress. Keeping host and storage connections within an ASIC port group can reduce latency by a couple of hundred nanoseconds.

For details about port groups and local switching, see the *Brocade Fabric OS Administration Guide* and the hardware installation guide for the appropriate product.

Switch interconnections

Switch interconnections

As mentioned previously, at least two of every element in the SAN should provide redundancy and improve resiliency. The number of available ports and device locality (server/storage tiered design) determines the number of ISLs needed to meet performance requirements.

ISL requirements for directors include a minimum of two trunks with at least two ISLs per trunk.

Each source switch should be connected to at least two other switches. In the following figure, each connection line represents at least two physical cables. Two physical connections provide redundancy for ports, optics, fiber patch cables, patch panels, and fiber infrastructure.

Redundant trunks on a director platform should be placed in varying port groups on different blades, as shown in the following figure.




Figure: Connecting devices through redundant fabrics

See the appropriate hardware manual to determine port groups for the various models of port blades. For more details, see the *Brocade Fabric OS Administration Guide*. Whichever method is decided upon, it is crucial to be consistent across the SAN.

For example, do not place ISLs on lower port numbers in one chassis, as illustrated in the left side of the following diagram. Instead, stagger ISLs across a different chassis, as illustrated in the right side of the diagram.




Figure: ISL connection placement

Brocade UltraScale ICL connectivity

Brocade UltraScale ICL technology preserves Fibre Channel ports for end device connections. The Brocade X8-8 and X8-4 platforms use Brocade UltraScale ICL with Octal Small Form-factor Pluggable (OSFP) ports. The Brocade X8-8 has two Brocade ICL8 X8-8 blades, each with eight OSFP ports, totaling sixteen OSFP ports per 8-slot director for ICL connectivity. In a Brocade X8-8 Director, each OSFP port has eight independent links providing a total of 128 ICL links. Each link connects to one of the eight core ASICs within the Brocade X8-8 Director.

The Brocade X8-4 has two Brocade ICL4 X8-4 blades, each with four OSFP ports, totaling eight OSFP ports per 4-slot director for ICL connectivity. In a Brocade X8-4 Director, each OSFP port has eight independent links providing a total of 64 ICL links. Of the eight links, two links connect to one of the four core ASICs within the X8-4 Director.

Both 100m and 2 km OSFP optics are available. ICL oversubscription with Gen 8 directors is 3.6:1 when the director is fully populated with 48-port blades.

The Brocade X8 platform uses OSFP optical transceivers for UltraScale Inter-Chassis Link (ICL) connectivity. The X8-8 director supports sixteen OSFP ports per chassis, and the X8-4 director supports eight OSFP ports per chassis. These OSFP ports provide high-bandwidth inter-chassis connectivity while preserving front-end Fibre Channel ports for device attachment.



Figure: UltraScale X8-8 ICL core-edge architecture

Brocade UltraScale ICL best practices

Each core ICL blade must be connected to an ICL blade in an edge chassis to achieve redundancy. In the following figure, each X8-8 director has 16 ICL ports, eight per ICL blade. There are four core directors and eight edge directors. Each edge director ICL blade has a link to a core director ICL blade for a total of 16 ICLs. Eight edge directors × two ICL blades each, for a total of 16 ICLs per core director.

Note: A pair of ICL links is used to connect an ICL blade for redundancy.




Figure: Minimum ICL connections between Brocade X8-8 directors

Full-mesh topology with UltraScale ICL

A full-mesh architecture provides a single hop between source and destination. Broadcom supports a nine-director ICL mesh with up to 100-meter distances using QSFPs or OSFPs over OM4 fiber. In the example shown in the following figure, nine X8-8 directors with eight 64-port blades each, support 4608 128G Fibre Channel end-device ports. UltraScale ICLs with two ICL links between each director, equals 200 Gb/s between directors.

Alternatively, using UltraScale ICL connectivity, if the full-mesh had five X8-8 directors instead of nine, there would be 2560 128G Fibre Channel end-device ports with 400 Gb/s between each director.




Figure: Nine-chassis UltraScale ICL full-mesh architecture

Note: Refer to the Scale-Out Architecture with Brocade UltraScale Inter-Chassis Links Design Guide for details. UltraScale ICL connections are considered a hop of no concern in a FICON fabric.

In a core-edge architecture, the edge switches should connect to two core switches using Brocade Trunks (BTs) with at least two ISLs. Each BT should be attached to a different blade. Redundancy requires a second mirrored fabric, and end devices must be connected to both fabrics using MPIO to manage active-active or active-passive flows and failover and fallback flows.

The following recommendations are for ISL and UltraScale ICL connectivity:

- There should be at least two mirrored fabrics.
- There should be at least two core switches per fabric.
- Every edge switch should have at least two BTs to each core switch.
- Create small trunks. Keep BTs to two ISLs unless high traffic volumes are anticipated. Small trunks ensure that losing a BT does not result in significant bandwidth loss.
- Plan BTs based on one BT being offline.
- Place redundant BTs on different blades.
- BTs form only within a port group; port groups are in an ASIC boundary.
- Cable length difference within a BT should be identical for optimal performance.

Use the exact cable length for UltraScale ICL connections.

Use ISL or UltraScale ICL connectivity between the same domains; mixing the two types of connectivity is not supported.

Use the same optic type on both ends of an ISL, selecting from the following options:

- Short-wavelength multi-mode optical fiber (SWL MMF)
- Long-wavelength single-mode fiber (LWL SMF)
- Extended-long-wavelength single-mode fiber (ELWL SMF)

Implementation planning for a Fibre Channel SAN

This publication is for technical IT architects who are directly or indirectly responsible for SAN design that is based on IBM B-type Gen 7 and Gen 8 FC SAN platforms. It describes many of the challenges that face SAN designers in greenfield and established storage environments. Although not intended as a definitive design document, this publication introduces concepts and guidelines to help you avoid potential issues that can result from poor design practices. This section describes best practices guidelines.

Device placement and traffic locality

Device placement is a balance between traffic isolation, scalability, manageability, and serviceability. Virtualization has dramatically optimized compute platforms, driving the need for high performance and improved scalability in storage networks. Frame congestion can become a severe concern if there are end-device issues.

Traffic locality

Designing device connectivity depends significantly on the expected data flow between devices. For simplicity, communicating hosts and targets can be attached to the same switch as shown in the following figure.




Figure: Hosts and targets attached to the same switch to maximize locality of data flow

This approach needs to scale. With Fibre Channel's high-speed, low-latency nature, attaching these host-target pairs on different switches does not mean performance is adversely impacted for typical workloads. With the current generation of switches, local switching is not required to gain performance or achieve low latency. Architects might want to switch traffic locally for mission-critical applications that depend on speedy response times.

In exceptional cases, multi-hop concerns might involve traffic congestion, specifically, inadequate inter-switch connectivity or concerns about proper resiliency as shown in the following figure. Often, these concerns can be mitigated by provisioning trunks/UltraScale ICLs.




Figure: Hosts and targets attached to different switches for ease of management and expansion

A less common scheme for scaling a core-edge architecture is dividing the edge switches into a storage/target tier and a host/initiator tier. This approach is called edge-core-edge and lends itself to easier management and expansion. End devices do not connect to the core. Host and storage platforms have different performance requirements, cost structures, and other factors that are accommodated by placing initiators and targets in different tiers.

Data flow considerations

Fan-in ratios and oversubscription

A critical aspect of data flow is the fan-in ratio, which is the oversubscription of initiator ports to target ports or edge devices to BTs. Alternatively, oversubscription can be viewed from the perspective of the storage array, referred to as the fan-out ratio. The ratio is the number of devices that share a single port, whether an ISL, BT, UltraScale ICL, or

target port. Logically, a BT is a single port. The ratio is always expressed from the single entity's point of view, such as 7:1 for seven hosts utilizing a single BT or a single storage port.

What is the optimum number of hosts that should connect to a storage port? This question seems straightforward; however, the situation becomes complex once you consider clustered hosts, VMs, workload characteristics, and the number of LUNs per server. Determining how many hosts to connect to a particular storage port can be narrowed down to three considerations: port queue depth, IOPS, and throughput. Throughput is the only network component; thus, does a simple calculation adding up the expected peak bandwidth for each host suffice?

In practice, it is improbable that all hosts perform simultaneously at their maximum level. The bandwidth of the HBA was considerably overprovisioned with traditional application-per-server deployments. However, the game changed radically with the advent of VMs. Conceptually, oversubscription is built into virtual machines to optimize server resources. To the extent that servers optimize their resource utilization, they should proportionately increase their port utilization. Fewer virtual machine ports can oversubscribe a target port compared to non-virtualized machines. Nonetheless, it is prudent to oversubscribe ports to balance cost and performance.

Another method is to assign host ports to storage ports based on the host's I/O capacity requirements. The intended result is a small number of high-capacity servers assigned to each storage port, effectively resulting in many low-capacity VM workloads distributed across the storage ports.

Port monitoring should determine the actual port utilization from fan-in and fan-out ratios, driving necessary adjustments. Ongoing monitoring provides valuable heuristic data for the successful expansion of storage and the efficient assignment of storage ports.

Planning for storage

Storage arrays have evolved dramatically over the last few years. Performance increased exponentially, capacities exploded, and more LUNs are supported than ever before. The performance and capacity of low-end arrays also improved.

Today's storage arrays can include the following characteristics:

- Tiers comprising all-flash or flash-enhanced storage media became commonplace for mid-range and enterprise-class environments.
- Some arrays time out and reset their ports if they do not receive acknowledgments from the connected host after specific intervals.
- Use of in-band FC for control purposes has increased, putting extra stress on FC port buffer usage.

The use of flash-enhanced storage not only reduces response times and increases I/O rates by up to ten times, but also improves host efficiency of the host platforms, allowing them to run at a higher level of utilization and execute more I/O commands rather than waiting on responses from slow disk media.

Adding a layer of flash can open bottlenecks and relieve congestion that result from long wait times. Still, storage array performance can degrade over time, which can be attributed to factors that include the following examples:

- Proliferation of VMs and the "I/O blender" effect, whereby I/O processes become increasingly non-sequential, and increase reads and writes to disk.
- Insufficient controller cache.
- Misaligned LUN configurations where operating system block sizes do not match well to those on block storage.
- Misconfigured control values, such as edge hold time and queue depths.
- Provisioning strategies can favor capacity over usage.

Design guidelines

Consider the following design guidelines:

- Ensure that network throughput capacity matches or exceeds the capabilities of storage devices.

- Use care if you deploy mixed arrays with different performance characteristics. Experience shows that depending on how it is used, it is easy for a Tier 3 storage array to affect the performance of arrays in the same fabric. Troubleshooting in these situations is difficult.
- Limit network switch and optic speeds to two generations at most, with a best practice of uniform speed to avoid target ports from slowing to accommodate lower-speed ports and creating a back-pressure situation.
- Enable Forward Error Correction for ISL connections and connectivity to Gen 5 devices that support it. FEC is on by default for Gen 6 networks and always should be used to ensure consistency of network transit times.
- Control the number of LUNs behind each storage port based on the type of usage they receive.
- Check on any special short-frame traffic to avoid frame congestion at array ports. It might be necessary to increase the number of buffers at the array port to accommodate the extra control traffic.
- Use advanced Brocade FOS threshold monitoring to monitor hosts and storage arrays to ensure that array ports do not reset because of a high-latency host and therefore do not adversely affect other connected hosts.
- ISL use increases with in-line virtualized storage. Ensure that enough ISL bandwidth exists to manage the increased load.
- The in-line storage heads might communicate through FC or generate more SCSI control frames to manage their attached storage, which can contribute to frame congestion. You might need to increase the number of buffers at the ports that connect to the storage controller to accommodate this behavior.
- It is much more difficult to determine initiators and targets with in-line virtualized storage. Because they are on the same switch, be careful about deploying tools, such as Port Fencing.

NVMe

The next consideration in terms of why the storage network must be rethought and redesigned is Non-Volatile Memory Express (NVMe).

At a device level, consider the following characteristics:

- **Density:** Modern NVMe devices offer substantially greater storage density than DRAM, making them suitable for high-capacity, performance-sensitive environments.
- **Latency:** NVMe is optimized for low-latency operations, helping accelerate application response times and improve overall system efficiency.
- **Bandwidth:** By using the latest PCIe technologies, NVMe delivers high throughput and scalable performance for demanding workloads.
- **Protocol efficiency:** NVMe employs an optimized command architecture with support for multiple queues and parallel operations, allowing storage devices to process I/O requests more efficiently.

Storage virtualization

Storage virtualization enables LUNs that are accessed by servers to be extracted into a logical layer of LUNs from the physical storage (typically storage arrays) on which they are stored. Virtualized LUNs that are disassociated from their storage allow for more flexible storage provisioning processes. Performance can also improve because the virtual LUNs can be striped across multiple storage arrays.

Two general types of storage virtualization are available: one uses an external controller (called in-line virtualization), and in the other, the virtualization occurs inside a storage array. In-line solutions are slightly more flexible because they can use physical storage from a variety of sources and vendors.

The following figure shows a typical implementation of an in-line virtualized storage solution.




Figure: Typical implementation of an in-line virtualization storage solution

The host or VM accesses storage by way of a storage controller through the storage network. The red arrows indicate data access to and from the storage controller. The storage controller typically controls all access to the physical storage (indicated by the blue arrows). This configuration creates a flexible storage solution because logical LUNs can be striped across several physical arrays to improve performance. Also, logical LUNs can be manipulated transparently to the host or VM.

The major benefit of this type of storage virtualization is that storage can now be provisioned in units of capacity (500 GB or 1 TB) rather than physical LUNs. This first step is toward viewing storage as a service instead of as physical units. VM provisioning now becomes less complex and easier to automate.

Monitoring

Consider the following points regarding monitoring:

- Use MAPS to track thresholds of critical parameters.
- FPI monitoring is useful in determining latencies that are associated with virtualized storage.
- Use Flow Vision to review high-traffic-usage flows and identify the source of bottlenecks and congestion.
- Consider running Flow Vision constantly for the most mission-critical applications to keep a historical record of application performance profile and intermittent irregularities.
- For "frequent-caller" applications, run Flow Vision on a regular basis where time permits to verify good health.

Scalability and performance

IBM B-type Gen 8 products are designed with scalability in mind, knowing that most installations continue to expand and that growth is supported with few restrictions. However, following the same basic principles that were outlined in previous sections, as the network grows ensures that the levels of performance and availability continue.

Evaluate the effect on topology, data flow, workload, performance, and perhaps most importantly, redundancy and resiliency of the entire fabric any time one of the following actions is performed:

- Adding or removing:
 - Changes in provisioning

- Storage
- Switches
- ISLs and ICLs
- Changes in virtualization (workload and storage) strategies and traffic flow patterns

If these design best practices are followed when the network is deployed, small incremental changes should not adversely affect the availability and performance of the network. However, if changes are ongoing and the fabric is not suitably evaluated and updated, performance and availability can be jeopardized.

Consider the following points when reviewing the status of a production FC network — reviewing redundancy and resilience:

- Are there at least two physically independent paths between each source and destination pair?
- Are there two redundant fabrics?
- Does each host connect to two different edge switches?
- Are edge switches connected to at least two different core switches?
- Are inter-switch connections composed of two trunks of at least two ISLs?
- Does each storage device connect to at least two different edge switches or separate port blades?
- Are storage ports provisioned so that every host has at least two ports through which it can access LUNs?
- Are redundant power supplies attached to different power sources?
- Are zoning and security policies configured to allow for patch or device failover?

Review the following performance requirements:

- Host-to-storage port fan-in/out ratios:
 - Host to ISL
 - Edge switch to core switch
- Routing policy and currently assigned routes (evaluate actual utilization for potential imbalances)
- Use of FEC for all ISLs and connections to Gen 5 (if supported) and Gen 6 devices
- Watching for latencies because of the following issues:
 - Poor storage performance
 - Overloaded hosts or applications
 - Distance issues over constrained long-distance links, which result from changes in usage, such as adding mirroring or too many workloads
 - Deteriorating optics resulting in declining signal strength and increased error rate
- In Gen 7 and Gen 8 networks, storage response latency can be base-lined and monitored continuously by using I/O Insight with MAPS. Correct latencies immediately; they can have a profound effect on the fabric.

Although IBM B-type Gen 8 SANs are designed to allow for any-to-any connectivity and they support provision-anywhere implementations, these practices can have an adverse effect on the performance and availability of the SAN if left unchecked.

The network must be monitored for changes and routinely evaluated for how well it meets wanted redundancy and resiliency requirements.

SAN design for critical workloads

With all-flash arrays (AFAs) becoming the standard in today's enterprise data centers and the transition to FC-NVMe AFAs beginning, critical business applications are increasingly dependent on consistent low-latency, high-throughput storage performance for demanding, performance-sensitive workloads.

When designing the SAN, it is important to consider the placement of critical workloads (in relation to the placement of storage), the fan-in ratio to storage ports, and the fan-in ratio to ISLs/trunks.

Whereas the IBM B-type Gen 8 SAN technology provides a suite of built-in measures to avoid interference of workloads that expose any kind of congestion behavior, such as Adaptive Traffic Optimizer, FPIN, MAPS, and SDDQ, protecting critical workloads is crucial. Ideally, the most demanding and critical workloads include dedicated storage array target ports (or even dedicated arrays) and the shortest possible path through the SAN. The purpose for this configuration is to avoid any interference of other less critical workloads that might expose behavior that results in congestion or back pressure, which can interfere and adversely affect the performance of critical workloads.

Placement of servers with business-critical workloads

With core-edge SAN designs, it is often advantageous to connect servers with critical workloads directly on the core with the storage array ports. This configuration works well when the number of business-critical workloads is easily defined to a specific subset of servers in the SAN and enough ports are available on the core to connect both storage and business-critical servers.

If the number of business-critical servers exceeds the number of available ports on the core for server connections, the business-critical servers must be connected on edge switches. The most common model is to use dedicated edge switches for business-critical servers and decrease the fan-in ratio of servers to ISLs.

An alternative model is to attempt to evenly distribute business-critical servers across edge switches with the assumption that the workload pressure evens out with the less critical and demanding workloads on other servers. Although this approach might seem logical, experience shows that throughout the lifetime of applications and the SAN, it is operationally more complex to ensure optimal performance for business-critical workloads by using this model.

Business-critical VMs

In today's data centers, it is not uncommon for some or all business-critical workloads to be running on VMs. Combining this with the digital business environment, the value (business criticality) and performance requirements for an application often change throughout the life cycle of the application. Therefore, it can be difficult (or impossible) to predict the requirements for an application and plan placement accordingly from the beginning; thankfully, hypervisors provide the ability to move VMs between hosts without downtime and to migrate storage when necessary.

To apply the same principle of server placement of (bare metal installed) business-critical applications to VMs, deploy dedicated hypervisor clusters that are connected on the core or on high-performance edge switches. Use these hypervisor clusters for business-critical and performance-sensitive VM workloads only.

Visibility into each VM's workload on the same data store (backed by LUN or NSID) can be achieved with VM Insight. VM Insight enables storage administrators to monitor VM-level application performance and set baseline workload behavior. This information can be used to quickly determine whether a storage fabric is the source of performance anomalies for applications running on VMs. Based on VM Insight, storage administrators can provision and plan placement in the SAN on application requirements and can fine-tune the infrastructure to meet service-level objectives.

Capacity planning

One way to avoid performance issues and downtime is with suitable capacity planning, which helps you get an overall view of your IT infrastructure's overall capacity.

Gathering requirements

The SAN project team must interview all stakeholders (IT application owners, finance, corporate facilities, IT lab administrators, storage and network administrators, and users) who have a vested interest in the project. This interview process applies equally to planning for new and updated SANs.

Application owners

As critical stakeholders, application owners care because everyone is measured on application uptime. Application outages are something that users notice, and they can have severe financial effect on a business. With a redundant or a resilient infrastructure, hardware outages are transparent to the user, and only SAN administrators must pay attention.

To better understand their requirements, questions can be asked of the application owners including the following examples:

- What is the business goal for this application? (Is it a database that multiple applications rely on for business transactions?)
- What are the availability requirements?
- Is the application latency sensitive?
- Are there peak periods of utilization or other traffic patterns?
- What are the IOPS requirements in terms of read/writes?
- What is the worst-case response time before an outage?
- Is the application running on a cluster?
- Is there application downtime that can be used for applying patches, software upgrades, and maintenance?
- Can the application run on a VM? If so, how many other VMs can coexist on the same physical hardware?

The business criticality of the application determines the SAN design and the DR strategy, including backup and recovery. If the application is mission critical, the infrastructure must be fully redundant, with no single point of failure for mainframe or distributed open systems architectures.

Server and storage administrators

After the application requirements are defined, identify physical server and storage on which the application and data are stored to determine the overall high-level architecture of the SAN, especially if this includes existing and new equipment.

Consider the following points:

- Gather information about the servers on which the applications are running (blade or rack, CPU, memory, HBA/ embedded FC switch, operating system level, operating system patch level, or HBA driver version).
- How many HBAs are in the rack servers?
- Does each server have single or multiple port HBAs?
- What is the wanted storage response time (latency)? Do longer latency times that result from a disk hit severely affect performance?
- For hybrid storage, how much storage is allocated to SSD?
- For HDD-based storage, what is the current cache use? Is there enough cache to meet required response times? What is the average drive use for HDDs and SSDs? Contact your drive vendor to identify response times based on use for sizing workloads.
- How much storage capacity is allocated to flash (more cache allows for more consistent response time) if using a hybrid array?
- Are storage tiers used in the environment? What is the policy that is used for migrating data? Are different tiers used for online storage? What is the RAID level used? This factor determines available disk space and performance for the application.
- How many FC ports exist in the array? Are the arrays front-ended by a storage virtualization controller? If so, what is the additional latency?
- What are the recommended fan-in and fan-out ratios for the arrays that are used for this application? What are the limits?
- Does a disaster recovery (DR) site exist? If so, how is it connected (dark fiber, FCIP, and so on)?
- What is the available or required bandwidth between the intra-site for DR?
- Can the storage infrastructure support DR with the extra load?

- What tools are used for mirroring and replication (host-based or array-based)? If host-based, was the failover tested? If so, was application uptime affected? If storage-based, was the failover tested? Did the LUNs appear on the active ports?
- Was application uptime affected?

Fabric details

This section provides example tables that you can use to identify dominant factors, including facilities that can affect the storage area network (SAN) design.

This appendix includes the following topics:

- A.1 Current fabrics
- A.2 Individual fabrics
- A.3 Device
- A.4 Metrics and effect on SAN design and performance
- A.5 Consolidated SAN requirements
- A.6 Application-specific information

Current fabrics

Use the following table to record the fabrics that are used.

SAN/fabric	Number of switches	Type of switch	Total ports	Domains	Number of servers	Number of storage devices	Log data	Notes
Fabric 1								
Fabric 2								
Fabric 3								
Fabric 4								
Fabric 5								

Individual fabrics

Use the following table to record the individual fabrics that are used.

SAN/fabric	Domain number	Serial number	Model	Speed	WWN	IP addresses	Brocade Fabric OS	Notes or M-EOS Version
Switch 1								
Switch 2								
Switch 3								
Switch 4								
Switch 5								

Device

Use the following table to record information about the devices that are used.

Servers and storage	Model	WWN	Alias	Zone	Version	Application Switches	Notes
Server 1							
Server 2							
Server 3							
Storage 1							
Storage 2							
Storage 3							

Metrics and effect on SAN design and performance

Use the following table to record the metrics and effects on SAN design.

Metric	Source	Effect
Servers in the SAN	Estimate/Brocade SAN Health	Normal operations
Host Level Mirroring	Estimate	Distance; ISL; Traffic levels
Clusters (MSFT, HACMP, NetApp) — Average number of nodes; Workload level	Estimate; Estimate: High/Med/Low	In-band heartbeat; Frame congestion; Host fan-in; Traffic isolation
Virtualization: VIO Server — Number of servers; Consolidation ratio	Estimate; Estimate	Frame congestion; Edge traffic increase per port; Server fan-in on target ports; Device latencies
Virtualization: VMware — Number of VMware servers; Consolidated ratio; Shared VMFS; DRS; RDM; I/O intensive	Estimate; Estimate; Yes/No; Yes (%) / No; Yes (%) / No; High/Med/Low; Yes/No	Frame congestion; Device latencies; SCSI reset

Consolidated SAN requirements

Use the following table to record the consolidated SAN requirements.

SAN requirements data (complete for each SAN)

Item	Value
Fabric information	
Target number of user ports per fabric	
Target number of total ports per fabric	
Target number of switches per fabric (number of switches/ switch type, total switches)	
Number of fabrics	
Number of sites in environment	
Topology (core-edge, ring, mesh, other)	
Maximum hop count	
Fabric licenses	
SAN device information	
Number/types of hosts and operating system platforms	
Number/types of storage devices	
Number/types of tapes	
Number/types of HBAs	
Other devices (VTL/deduplication appliance)	
Total number of SAN devices per fabric	
Customer requirement for failover/redundancy, reliability of SAN (multipathing software utilized)	
Application details	
SAN Application (Storage Consolidation, Backup and Restore, Business Continuity); Fabric management application(s)	
Performance	
Maximum latency (ms)	
Targeted ISL oversubscription ratio (3:1, 7:1, 15:1, other)	

Application-specific information

Use the following table to record application-specific information.

Backup/restore infrastructure

Section	System	Operating system version, patch level	HBW driver version
Servers			
Server 1/HBA			
Server 2/HBA			
Server 3/HBA			

Section	Vendor	Version	Patch
Backup software			

Section	Vendor	Model	Firmware
FC switch			
Brocade			
Storage			
Array 1			
Array 2			
Tape library			
Library			

Virtual Fabrics overview

Virtual Fabrics is an architecture that is used to virtualize hardware boundaries. Traditionally, SAN design and management is done at the granularity of a physical switch. Virtual Fabrics allows SAN design and management to be done at the granularity of a port.

Virtual Fabrics is a suite of related features that can be customized based on your needs. The Virtual Fabrics suite consists of the following specific features:

- Logical switch
- Logical fabric

- Device sharing

This section describes a logical fabric and logical switch types, and device sharing with Virtual Fabrics.

Note: Virtual Fabrics is the name of the suite of features. A logical fabric is a type of fabric that you can create by using the Virtual Fabrics suite of features.

If SNMP monitoring is for the default switch, you do not need the SNMP user to be configured as a user on the switch.

Important: SNMPv3 is required to manage Virtual Fabrics. If you use SNMPGET to poll the switches (for example HiTrack), the SNMPv3 user must be configured based on the logical switch that is being monitored.

If SNMP monitoring is for the default switch, you do not need the SNMP user to be configured as a user on the switch.

Logical switch

Traditionally, each switch and all the ports in the switch act as a single Fibre Channel (FC) switch that participates in a single fabric. The logical switch feature allows you to divide a physical chassis into multiple fabric elements. Each of these fabric elements is referred to as a *logical switch*. Each logical switch functions as an independent, self-contained FC switch.

Default logical switch

To use the Virtual Fabrics features, you must first enable Virtual Fabrics on the switch, which creates a single logical switch in the physical chassis. This logical switch is called the *default logical switch*.

The default logical switch initially contains all of the ports in the physical chassis.

The following figure shows a switch before and after enabling Virtual Fabrics. In this example, the switch has 10 ports, which are labeled P0 - P9.




Figure: Switch before and after enabling virtual fabrics

The following figure shows a Virtual Fabrics-enabled switch before and after it is divided into logical switches. Before you create logical switches, the chassis appears as a single switch (default logical switch). After you create logical switches, the chassis appears as multiple independent logical switches. All of the ports continue to belong to the default logical switch until you move them to other logical switches.




Figure: Switch before and after creating logical switches

The default logical switch always exists. You can add and delete other logical switches, but you cannot delete the default logical switch unless you disable Virtual Fabrics.

Logical switches and fabric IDs

When you create a logical switch, you must assign it a fabric ID (FID). The fabric ID uniquely identifies each logical switch within a chassis and indicates to which fabric the logical switch belongs. You cannot define multiple logical switches with the same fabric ID within the chassis.

As shown in the following figure, logical switches 2, 3, 4, and 5 are assigned FIDs of 1, 15, 8, and 20, respectively. These logical switches belong to different fabrics, even though they are in the same physical chassis. For example, you cannot assign logical switch 5 a fabric ID of 15 because logical switch 3 is assigned FID 15 in the chassis. The default logical switch is initially assigned FID 128. You can change this value later.




Figure: Fabric IDs assigned to logical switches

Port assignment in logical switches

Initially, all ports belong to the default logical switch. When you create logical switches, they are empty, and you must assign ports to those logical switches. As you assign ports to a logical switch, the ports are moved from the default logical switch to the newly created logical switch. A specific port can be in only one logical switch.

As shown in the following figure, the default logical switch initially has 10 ports, which are labeled P0 - P9. After logical switches are created, the ports are assigned to specific logical switches. Ports 0, 1, 7, and 8 are not assigned to a logical switch; therefore, they remain assigned to the default logical switch.




Figure: Assigning ports to logical switches

A specific port is always in only one logical switch. The following scenarios refer to the chassis after port assignment that is shown in Figure 4.

- If you assign P2 to logical switch 2, you cannot assign P2 to any other logical switch.
- If you want to remove a port from a logical switch, you cannot delete it from the logical switch, but must move it to a different logical switch. For example, if you want to remove P4 from logical switch 3, you must assign it to a different logical switch: logical switch 2, logical switch 4, or logical switch 1 (the default logical switch).
- If you assign a port to a logical switch, it is removed automatically from the logical switch it is in. If you assign P3 to logical switch 3, P3 is automatically removed from logical switch 2.
- If you do not assign a port to any logical switch, it remains in the default logical switch, as is the case with ports 0, 1, 7, and 8.

A logical switch can have as many ports as are available in the chassis. As shown in the figure, the chassis has 10 ports. You can assign all 10 ports to a single logical switch, such as logical switch 2; however, if you made this assignment, no ports are available for logical switches 3 and 4.

You can move only F_Ports and E_Ports from one logical switch to another. If you want to configure a different type of port, such as a VE_Port or EX_Port, you must configure them after you move them. Some types of ports cannot be moved from the default logical switch.

Logical switches and connected devices

You can connect devices to logical switches, as shown in Figure 5. In logical switch 2, P2 is an F_Port that is connected to H1. In logical switch 3, P4 is an F_Port that is connected to D1. H1 and D1 cannot communicate with each other because they are in different fabrics, even though they are both connected to the same physical chassis.




Figure: Logical switches that are connected to devices and non-Virtual Fabrics switches

You can also connect other switches to logical switches. In the previous figure, P6 is an E_Port that forms an inter-switch link (ISL) between logical switch 4 and the non-Virtual Fabrics switch. Logical switch 4 is the only logical switch that can communicate with the non-Virtual Fabrics switch and D2 because the other logical switches are in different fabrics.

The following figure shows the logical switch view of a chassis where virtual fabrics are enabled and the devices are connected. The logical switch that includes Fabric 128 is the default switch when virtual fabrics are enabled. Logical switches Fabric 1 (connected to Host) and Fabric 15 (connected to Storage), and logical switch Fabric 8 is connected to a non-Virtual Fabric switch and storage.




Figure: Logical switches in a single chassis belong to separate fabrics

Management model for logical switches

The operations that you can perform on a logical switch depend on the context you are in. Some operations affect only a single logical switch, and some operations affect the entire physical chassis.

All user operations are classified into one of the following chassis management operations that span logical switch boundaries:

- Logical switch configuration (creating, deleting, or modifying logical switches)
- Account management (determining which accounts can access which logical switches)
- Field-replaceable unit (FRU) management (slot commands, such as `slotShow`)
- Firmware management (firmware upgrade, high availability failover)

These operations are limited to the logical switch, such as displaying or changing port states. Logical switch operations include all operations that are not covered in the chassis management operations.

When you log in, you are assigned an active context, or active logical switch. This context filters the view that you receive, and determines which ports you can see. You can change the active context. For example, if you are working with logical switch 1, you can change the context to logical switch 5. When you change the context to logical switch 5, you see only the ports that are assigned to that logical switch. You do not see any of the other ports in the chassis.

The scope of logical switch operations is defined by the active context. When you are in the context of a logical switch, you can perform port, switch, and fabric-level operations, subject to role-based access control (RBAC) rules.

If you have permission to run chassis-level commands, you can do so, regardless of which logical switch context you are in.

Logical fabric

A *logical fabric* is a fabric that contains at least one logical switch.

The four fabrics that are shown in Figure 5 and Figure 6 are logical fabrics because they each have at least one logical switch. You can connect logical switches to non-Virtual Fabrics switches and to other logical switches. You connect logical switches to non-Virtual Fabrics switches that use an ISL, as shown in Figure 5.

You connect logical switches to other logical switches in the following ways:

- Using ISLs

- Using base switches and extended ISLs (XISLs)

Logical fabric and ISLs

The following figure shows two physical chassis that are divided into logical switches. ISLs are used to connect the logical switches with FID 1 and the logical switches with FID 15. The logical switches with FID 8 are each connected to a non-Virtual Fabrics switch. The two logical switches and the non-Virtual Fabrics switch are all in the same fabric, with FID 8.




Figure: Logical switches that are connected to other logical switches through physical ISLs

The ISLs between the logical switches are *dedicated ISLs* because they carry traffic only for a single logical fabric. In the following figure, Fabric 128 has two switches (the default logical switches), but they cannot communicate with each other because they have no ISLs between them and they cannot use the ISLs between the other logical switches.




Figure: Logical switches connected to form logical fabrics

Base switch and extended ISLs

One method to connect logical switches is to use extended ISLs (XISL) and base switches.

When you divide a chassis into logical switches, you can designate one of the logical switches as the *base switch*. The base switch is a special logical switch that provides shared fabric services and supports communication between logical switches.

A base switch features the following properties:

- ISLs that are connected through the base switch can be used for communication among the other logical switches.
- Base switches do not support direct device connectivity. A base switch can have only E_Ports, VE_Ports, or EX_Ports, but no F_Ports.
- The base switch provides a common address space for communication between different logical fabrics.
- A base switch can be configured for the preferred domain ID, just like a non-Virtual Fabrics switch.
- You can have only one base switch in a physical chassis.

A base switch can be connected to other base switches through a special ISL, which is called a *shared ISL* or *extended ISL* (XISL). An extended ISL connects base switches. The XISL shares traffic among different logical fabrics.

Fabric formation across an XISL is based on the FIDs of the logical switches. A logical switch joins the logical fabric with that same FID.

The following figure shows two physical chassis that are divided into logical switches. Each chassis has one base switch. An ISL connects the two base switches. The ISL is an extended ISL (XISL) because it connects base switches.




Figure: Base switches connected by an XISL

Logical ISLs

The logical ISLs are not physical connections. The only physical connection is in the base switch. Traffic that needs to go from any of the logical switches in Chassis 1 to the logical switch in Chassis 2 with the same Fabric ID (FID) is routed up to the base switch, across the ISL and down from the base switch in Chassis 2.

Traffic between the logical switches can now flow across this XISL. The traffic can flow only between logical switches with the same fabric ID. For example, traffic can flow between logical switch 2 in chassis 1 and logical switch 2 in chassis 2 because they both have FID 8. Traffic cannot flow between logical switch 2 in Chassis 1 and logical switch 3 in Chassis 2 because they have different fabric IDs and are, therefore, in different fabrics.

It is useful to think of the logical switches as being connected with logical ISLs, as shown in the following figure. The logical ISLs are not connected to ports because they are not physical cables. They are a logical representation of the switch connections that are allowed by the XISL.

Tip: Configure at least two XISLs for redundancy.




Figure: Logical fabric that uses ISLs and XISLs

You can also connect logical switches by using a combination of ISLs and XISLs, as shown in Figure 10. In this example, traffic between the logical switches in FID 128 can travel over the ISL or the XISL. Traffic between the other logical switches travels over the XISL only.

By default, the physical ISL path is favored over the logical path (over the XISL) because the physical path has a lower cost. This behavior can be changed by configuring the cost of the dedicated physical ISL to match the cost of the logical ISL.

Attention: If you disable a base switch, all of the logical ISLs are disconnected and the logical switches cannot communicate with each other unless they are connected by a physical ISL.

Base fabric

Base switch ports on different chassis can be connected to form a fabric, called a *base ** fabric*. As with other logical switches, the base switches must have the same FID to be connected. If the base switches have different FIDs, the link between the switches is disabled.

The base fabric follows normal routing policies. If physical connectivity is available, the base fabric maintains connectivity for the logical fabrics.

Logical ports

Logical ISLs are formed to connect logical switches. A *logical port* represents the ports at each end of a logical ISL. A logical port is a software construct only and does not correspond to any physical port.

Most port commands are not supported on logical ports. For example, you cannot change the state or configuration of a logical port. However, state change on a logical link is permitted by using the `lfcfg --lislEnable` command.

The worldwide name (WWN) for logical ports is in NAA=5 format, and uses the following syntax:

`5n:nn:nn:nz:zz:zz:zx:xx`

The NAA=5 syntax uses the following variables:

- `nnnnnn` is the Organizationally Unique Identifier (OUI).

- zzzzzz is the logical fabric serial number.
- xxx is the logical port number, in the range 0 - FFF.

Logical fabric formation

Fabric formation is not based on connectivity; rather, it is based on the FIDs of the logical switches. The fabric formation uses the following basic order:

- Base fabric forms.
- Logical fabrics form when the base fabric is stable.
- Devices begin recognizing one another.
- Traffic is started between the logical switches.

Account management and Virtual Fabrics

When user accounts are created, they are assigned a list of logical fabrics to which they can log in and a home logical fabric (home FID). When you connect to a physical chassis, the home FID defines the logical switch to which you are logged in by default. You can change to a different logical switch context by using `setcontext` command.

When you are logged in to a logical switch, the system prompts changes to display the FID of that switch. The following example prompts for when you are logged in to the default logical switch (FID = 128) and a user-defined logical switch (FID = 15):

```
switch:FID128:admin> switch:FID15:admin>
```

Setting up IP addresses for a logical switch

Each physical chassis has one common IP address that is shared by all of the logical switches in the chassis. You can set up individual IPv4 addresses for each logical switch.

IPv4 addresses that are assigned to individual logical switches are assigned to IP over Fibre Channel (IPFC) network interfaces. In Virtual Fabrics environments, a single chassis can be assigned to multiple fabrics, each of which is logically distinct and separate from one another. Each IPFC point of connection to a specific chassis needs a separate IPv4 address and prefix to be accessible to a management host.

For a management host to access logical switches, the host bus adapter (HBA) must connect with the common, shared IP address and the individual IPv4 addresses that are configured for each logical switch.

IPFC addresses are not handled by `configupload` or `configdownload`. The IPFC address of the default logical switch or a non-VF switch is stored on the WWN card or compact flash. This address does not display in a `configshow`. Non-default logical switch IPFC addresses display in a `configshow`. The `ipaddrshow` command displays all switch addresses and IPFC addresses configured in the chassis.

Use the following procedure to set up IP addresses for a logical switch:

- Connect to the switch and log in by using an account with admin permissions.
- Enter the `ipAddrSet -ls` command. Consider the following points:
 - To add an IPv4 address, use the `--add` parameter. Specify the network information in dotted-decimal notation for the Ethernet IPv4 address with a Classless Inter-Domain Routing (CIDR) prefix.
 - To delete an IPv4 address, use the `--delete` parameter.
 - Enter the `ipaddrshow` command to verify the result.

The following example shows how to set IP addresses with the CIDR prefix for logical switches with FID 1, 2, and 128 (default logical switch).

```

switch:FID128:admin> ipaddrset -ls 1 --add 192.0.2.11/24
IP address is being changed...Done.
switch:FID128:admin> ipaddrset -ls 2 --add 192.0.2.22/24
IP address is being changed...Done.
switch:FID128:admin> ipaddrset -ls 128 --add 192.0.2.0/24
IP address is being changed...Done.
switch:FID128:admin> ipaddrshow
SWITCH
Ethernet IP Address: 198.51.100.0
Ethernet Subnetmask: 255.255.255.0
Gateway IP Address: 198.51.100.1
DHCP: Off
IPFC address for virtual fabric ID 128: 192.0.2.0/24
IPFC address for virtual fabric ID 1: 192.0.2.11/24
IPFC address for virtual fabric ID 2: 192.0.2.22/24

```

The following example shows how to delete an IP address for the logical switch with FID 1.

```

switch:FID128:admin> ipaddrset -ls 1 -delete

```

Logical switch log in context

When you log in to a logical switch using SSH, the login shell uses the destination IPFC address to look up and set the corresponding logical switch context. When you log in from the serial console, the login shell always sets your default home logical switch context.

The IPFC that is assigned to a logical switch is managed by using the `ipAddrSet` command. When you use SSH from a remote host, the IPFC address must share the network (subnet mask) as the management interface.

The values of FCIP addresses and subnet masks are stored in the WWN card. For example, if the management interface IP address is 10.38.18.183 with subnet mask 255.255.240.0, you can create the IPFC address. See the following example.

```

ipaddrset -ls 8 --add 10.38.18.183/20

```

To delete the IPFC address, you can use the `ipaddrset -ls 8 -delete` command.

If you log in to a switch by using the management interface IP address, the home logical switch context is established. You can log in directly to a specific logical switch by using SSH with the configured IPFC address assigned to that logical switch. You can manage the permissions by using the `userConfig` command (see the following example).

```

switch> userconfig --show admin Account name: admin Description: Administrator Enabled: Yes

Password Last Change Date: Unknown (UTC) Password Expiration Date: Not Applicable (UTC) Locked: No

Home LF Role: admin

Role-LF List: admin: 1-128 Chassis Role: admin

Home LF: 128

Day Time Access: N/A

```

Role-LF is the list of logical switch contexts for which you have permission to log in over the IPFC address. *Home LF Role* is the default logical switch context when you do not have permission to log in to a specific logical switch context or over management interface.

The following example shows the result of the `ipaddrshow` command. Only one IPFC is mapped to FID 8.

```
switch> ipaddrshow SWITCH

Ethernet IP Address: 10.38.18.183 Ethernet Subnetmask: 255.255.240.0 Gateway IP Address: 10.38.16.1 DHCP:

IPFC address for virtual fabric ID 8: 10.38.18.183/20
```

The following conditions apply for default users:

- For admin users, the login is successful to the specified logical switch, and the corresponding logical switch context (FID 1-128) is automatically established. When you connect using SSH to the IPFC address assigned to a logical switch, the session is placed in that logical switch context. For example, when you connect using SSH to the IPFC address assigned to FID 8, the context is set to FID 8.

The following conditions apply for non-default users:

- For non-default users that are configured in the switch with or without the logical switch permission, if the user logs in with the IPFC address without the permission for the corresponding logical switch, only the default home virtual fabric of the user is set for the context.
- For non-default users that are configured in the switch with or without the logical switch permission, if the user logs in with the IPFC address with the permission for the corresponding logical switch, login is successful to the specific logical switch with the corresponding logical switch context set.

Supported platforms for Virtual Fabrics

The following platforms support the Virtual Fabrics:

- IBM Storage Networking SAN512B-7
- IBM Storage Networking SAN128B-7
- IBM Storage Networking SAN64B-7
- IBM Storage Networking SAN56B-8
- IBM Storage Networking SAN42B-R7
- IBM Storage Networking SAN512B-8

Some restrictions apply to the ports, depending on the port type and blade type. The following sections explain these restrictions.

Supported port configurations in the fixed-port switch

No restrictions exist on the ports in IBM SAN64B-7 switches.

The following rules apply:

- Any port can belong to any logical switch (including the base switch and default logical switch), with the exception that F_Ports cannot belong to the base switch.
- The default logical switch can use XISLs.
- The default logical switch can also be a base switch, except on Gen 7 directors and Gen 6 directors.

Virtual Fabrics interaction with other Fabric OS features

The following table lists some Fabric OS (FOS) features and considerations that apply when Virtual Fabrics are used.

Table 2: Virtual Fabrics interaction with Fabric OS features supported logical switch creation limits

Fabric OS features	Virtual Fabrics Interaction
Access Gateway	Virtual Fabrics is not supported on a switch if AG mode is enabled.
Configuration upload and download	Virtual Fabrics uses a configuration file that is different from the configuration file that is used to download system configuration parameters.
FCR or Integrated Routing service	All EX_Ports must reside in a base switch. You cannot attach EX_Ports to a logical switch that has XISL use enabled. You must use ISLs to connect the logical switches in an edge fabric.
FICON	Up to two logical switches per chassis can run FICON Management Server (CUP), but the FICON logical switch can use ISLs and XISLs.
ISL R_RDY mode	ISL R_RDY mode is supported in the base switch and logical switch.
Licensing	Licenses are applicable for all logical switches in a chassis.
Network time protocol (NTP)	The clock information for a chassis is maintained by the default logical switch and not from the principal or primary FCS switch.
QoS	QoS VCs are maintained across the base fabric.

Limitations and restrictions of Virtual Fabrics

Before you use the Virtual Fabrics feature, be aware of the restrictions and limitations regarding QSFP ports and the maximum number of logical switches per chassis.

In the core blades, the four ports that belong to the same QSFP can be assigned to different logical switches. However, if one of the four ports belongs to the base switch, the remaining three ports cannot be assigned to any other logical switch.

The maximum number of logical switches per chassis varies depending on the switch model. The following table lists the supported platforms and the maximum number of logical switches (including the default logical switch) supported on each.

Table 3: Supported logical switch creation limits

Platform	Maximum logical switches per chassis a
IBM Storage Networking SAN512B-7	16
IBM Storage Networking SAN128B-7	16
IBM Storage Networking SAN64B-7	4
IBM Storage Networking SAN56B-8	4
IBM Storage Networking SAN42B-R7	4
IBM Storage Networking SAN512B-8	16

a. Numbers include the default switch and base switch.

Restrictions on the use of XISLs

The `Allow XISL Use` option of the `configure` command allows a logical switch to use XISLs in the base switch and any standard ISLs that are connected to that logical switch.

The following restrictions apply when XISLs are used. XISL use is not permitted when the logical switch meets any of the following conditions:

- Has ICL ports
- Is the default logical switch
- Is a base switch
- Is an edge switch for an FC router

When the logical switch is configured as an FCR edge switch, XISL use is not allowed while the logical switch is enabled. If the logical switch is disabled or has not yet joined an edge fabric, you *can* allow XISL use; however, fabric segmentation occurs when the logical switch is enabled or connected to an edge fabric.

Restrictions on moving ports

FC ports cannot be moved if any one of the following features is enabled:

- Long distance
- QoS
- F-Port buffers
- F_Port trunking - E_Port credits - SIM_Port
- Encryption

Before moving VE_Ports, you must remove the VE_Port tunnel configuration.

For more information about how to run a command in a different logical switch, change the fabric ID of a logical switch, or change a logical switch to base switch, see [FOS \(Fabric OS\) downloads and documentation](#).

Device sharing

The behavior of device sharing in the fabric depends on whether Virtual Fabrics is enabled or not enabled in the base switch. If the Virtual Fabrics is not enabled, FCR behavior is unchanged. If Virtual Fabrics is enabled, a base switch is like a backbone switch and a base fabric is like a backbone fabric in an FCR routing context.

FCR and Virtual Fabrics

If Virtual Fabrics is enabled, the following rules apply:

- EX_Ports can be configured on the base switch only. When you enable Virtual Fabrics, the chassis is automatically rebooted. When the switch comes up, only one default logical switch is present, with the default fabric ID (FID) of 128. All previously configured EX_Ports are persistently disabled with the reason ExPort in non-base switch.

You must specifically create a base switch, move the EX_Ports to the base switch, and then enable the ports. If you move EX_Ports to any logical switch other than the base switch, these ports are automatically disabled.

If you want to change an EX_Port on the logical switch to be a non-EX_Port, you must use the `portCfgDefault` command. You cannot use the `portCfgExPort` command because that command is allowed on the base switch only.

- From FOS v9.0.x and later, EX_Ports in a base switch receive a 10-bit dynamic area address instead of a zero-bit address. An EX_Port in the base switch is assigned any 10-bit area from 0x0000-00xFFC0, with the following exceptions because these areas are reserved for internal lookups:
 - Areas 0x8000-8300
 - 0x8040-8340
 - Areas 0x8080-8380
 - Areas 0x80C0-83C0
 - 0xC500
 - 0xC540
 - 0xC580
 - 0xC5C0

EX_Ports can connect to a logical switch that is in the same chassis or in a different chassis. However, the following configuration rules apply:

- If the logical switch is on the same chassis, the EX_Port FID must be set to a different value than the FID of the logical switch to which it is connecting.
- If the logical switch is on a different chassis, no FID for any logical switch in the FC router backbone fabric can be the same as the FID of the logical switch to which the EX_Port is connecting.
- EX_Ports in FC routers and those in a base switch cannot connect to any edge fabric with logical switches that are configured to use XISLs. If you connect an EX_Port to an edge fabric, you must ensure that no logical switches with XISL use are enabled in that edge fabric. If any logical switch in the edge fabric allows XISL use, the EX_Port is disabled.

Because XISL use is disallowed, dedicated links must be configured to route traffic across switches in the same logical fabric.

- Backbone-to-edge routing is not supported in the base switch (see Figure 7 for an example of established FC routers that allow backbone-to-edge routing with Virtual Fabrics).
- All FC router commands can be run in the base switch context only.

- The `fcrconfigure --bbfid` command *cannot* be used when Virtual Fabrics is enabled.

Instead, use the `lscfg` command to configure the FID.

- From FOS v9.0.x onwards, if a port is configured as an EX_Port and if the administrator moves the port from the base switch to another logical switch, the operation fails. You must remove the EX_Port configuration before moving it to another logical switch.
- When a Virtual Fabrics-enabled switch is part of the backbone, only the base switch with active EX_Ports can route frames between edge fabrics. If other base switches are participating in the fabric and do not have active EX_Ports, those switches participate as transient switches to forward frames to the next hop until they can be routed to the intended remote edge fabric.

Note: If any other default or logical switch is introduced as a hop between edge fabrics, the newly introduced switch that is not the base switch does not merge with the base switch fabric. This issue causes frames to be dropped.

For more information about device sharing and FC-FC routing, see [FOS \(Fabric OS\) downloads and documentation](#).

FICON design and implementation

Mainframes remain the definitive platform for mission-critical enterprise applications. To ensure that all components in a mainframe SAN meet the reliability standards and consistent behavior that most mission-critical applications demand, more testing and qualification programs are implemented for all devices that are connected to an IBM Z® processor.

The Gen 8 Directors and switches provide the high-performance, intelligent foundation required for the enterprise AI era and mission-critical FICON environments. As data growth continues to accelerate, Gen 8 allows organizations to consolidate sprawling workloads onto fewer, more powerful systems, reducing physical footprint while increasing operational headroom.

The Gen 8 platform is anchored by three primary pillars:

- **Quantum-Safe Security:** Integrated hardware and software safeguards, including Post-Quantum Cryptography (PQC), protect sensitive workloads from both current and future threats.
- **AI-Powered Autonomy:** Embedded SAN AI technology modernizes infrastructure management by learning, adapting, and self-correcting to reduce operational burden and minimize human error.
- **Powerful Performance:** Industry-leading 128G throughput supports AI-driven workloads and massive device density with ultra-low latency.

FICON is an IBM channel protocol that uses Fibre Channel (FC) transport for communication between IBM Z systems and storage devices. Although it shares the same transport technology as open-systems Fibre Channel, FICON has unique operational requirements, error handling, and performance characteristics. IBM b-type Gen 8 switches running Fabric OS (FOS) 10.x support FICON environments while also providing support for modern SAN technologies such as NVMe over Fabrics (NVMe-oF).

Important support notes for Gen 8 platforms

The following hardware lifecycle statuses must be considered for all new designs:

- **Gen 8 platforms:** The Brocade X8-8 Director, X8-4 Director, and G820 switch are the officially supported and recommended platforms for new FICON deployments.

- Gen 7 platforms: Supported as N-1 platforms running FOS 9.2.2x or higher.
- Gen 6 platforms: End of Availability and frozen on Fabric OS (FOS) 9.2.x; lack new security patches or support for Gen 8 interoperability features.

Remember to plan the timing of your upgrade projects to ensure adequate time before the EOS dates. Also plan for the coexistence period when mixed-generation fabrics may be required.

Software requirements

All Gen 8 switches require a minimum of FOS v10.0.0a1 for 128G performance tuning. However, it is strongly recommended to upgrade to FOS 10.0.1 or higher to take full advantage of the latest security and AI-powered automation capabilities.

Always see the IBM Broadcom FICON Qualification Letters to determine exact supported configurations: [FICON Qualification Letters](#).

A major change to note is the deprecation of the prohibit/allow matrix (PDCM) in FOS v9.1; customers must now use zoning for FICON port separation.

Note: Gen 8 FICON environments can be managed through FOS CLI, WebTools, or SANnav. When SANnav is used, SANnav Management Portal 3.0.1x or later is required.

FICON logical switch (FLS)

The creation of a FICON Logical Switch (FLS) is strictly required. In FOS 10.x, the FICON Control Unit Port (CUP) license is included by default; therefore, Fabric Management Server (FMS) mode is automatically enabled. A FICON Logical Switch provides the following benefits:

- Addresses are defined in the Hardware Configuration Definition (HCD) or IOCDS. If addresses change, the IODF becomes invalid, leading to outages.
- Improves security.
- Ensures properly configured FICON SAN fabrics.
- Simplifies and expedites configuration.

In earlier versions of FOS, determining whether the port address was bound was challenging for non-expert users. A "bound" address means that the port address is reserved and does not change — even if a cable is moved or a port is changed. This is critical for FICON environments where the port address is defined in HCD.

Another reason for the FICON logical switch is that it is possible to configure a switch to meet all of the FICON requirements and then back out specific changes after the channels come online. Because a mainframe channel checks only the suitable security attributes at login (when the PCHID is enabled, not when the CHPIDs are configured online), there is the potential for a channel not to come online the next time it goes off and back online. For most customers, this is the result of a CEC IML or channel path maintenance. Although circumstances for changing the FC address that is assigned to a port on the switch are unlikely, if it does occur, the IODF is no longer valid, which results in device or channel errors. This potential problem is fixed with the defined FICON logical switch.

Control unit port in a logical switch

An independent channel that is connected to the default switch with Control Unit Port (CUP) defined is not necessary. Because both disk and tape CUP functions are now incorporated into the FLS, a separate logical fabric for CUP is no longer required.

If a separate logical fabric is created for disk or tape mirroring and is put into its own logical fabric for mirroring, a separate CUP is still available.

Since TS7700, tape mirroring when done with FC and Metro Mirror use FCP protocol, which means a dedicated FICON channel is no longer required. TS7700 mirroring over FC now uses FCP protocol.

Link address refresher

Unified Addressing Mode (UAM) is now the default address allocation mode in FOS 10.x for both Virtual Fabric (VF) and non-VF environments. In UAM, link addresses are allocated sequentially starting from 0x00, ensuring predictable and consistent port addressing.




Figure: Link address and FC address

Domain ID and fabric security

Because the default zoning mode is "No Access" in FOS 10.x, administrators must ensure device connectivity between FICON channels and control units by creating FICON zones.

In two-byte or single-byte addressing, when the channel logs in, the domain area that is returned from the switch is used for the domain ID and area in the link address. With 2-byte link addressing, the channel establishes its link address when it first logs into the SAN.

After 2-byte link addressing is defined for a channel, all link addresses for that channel must use 2-byte link addressing. It is not possible to mix 1-byte and 2-byte link addressing for the same channel.

If the high integrity settings are not configured on the switch, zOS sets the channel offline and moves it to the "invalid" state. Because the fabric security features (high integrity fabric (HIF)) are not checked by the channel when single-byte addressing is used, it is acceptable to not set HIF mode when single-byte addressing is in use.

Port address

The hexadecimal port address is the single byte link address or the least significant byte of a 2-byte link address. The port address is bound to the physical port. The bits for N-Port ID Virtualization (NPIV) allow multiple WWNs to log in to the same switch port. These bits determine to which logical entity frames are routed when NPIV is used.

A FICON channel builds the Arbitrated Loop Physical Address (ALPA) portion of the FC address by using the same ALPA where the channel logged in. Therefore, the link address is identical to the port address in FICON environments.

FICON SAN design

Before you start a design project, see the FICON SAN migration section for more information about the switch differences and to understand cabling implications for directors. Consider the following areas as you design your FICON SAN:

- Management: SANnav Management Portal 3.0.1x must be used for Gen 8 management, as Network Advisor is End of Support (EOS) and does not support FOS 9.x or above. Key management considerations include:
- Although FICON SANs can be managed by using the command-line interface (CLI), doing so is unusual. Nearly all FICON SAN administrators use a GUI management tool.

- Network Advisor does not support FOS 9.x or above. Therefore, Gen 6/7/8 switches are not supported by Network Advisor.
- The SANnav Management Portal is the replacement product for Network Advisor. The enterprise version is required whenever FICON management is required.
- FICON environments nearly always require the Enterprise version of SANnav.
- SANnav 3.0.1 or higher is required when managing Gen 6/7/8 switches and is recommended for managing FICON fabrics.
- SANnav Global View is a new software package that can be useful when managing multiple independent production sites or a Disaster Recovery site.

The following routing method requirements apply:

- All switches in a fabric must use the same routing method; therefore, an outage must be planned if the routing method is changed.
- Confirm whether all control units in your environment support FiDR.
- Exchange-based routing (FiDR) is the most efficient and resilient routing method. If FiDR is supported on all channels and control units, use EBR.
- Use 48 port cards and high availability: Gen 8 director architectures utilize the FC128-48 port blades, providing 48 front-end 128G SFP+ ports. To ensure high availability during a SAN migration, both your cable plant layouts and Host Channel Definition (HCD) changes must be carefully planned around this specific 48-port density.

Consider the following cascaded performance factors:

- Most FICON CHPID mapping tools that are used to aid in channel path planning are built around 32 port cards. CHPID mapping tools must be updated to account for 48 port blades.
- Gen 8 replaces the legacy static traffic optimizer with the Adaptive Traffic Optimizer (ATO). ATO automatically learns and adapts to SAN traffic patterns to optimize performance.
- HCD changes and cable plant layout must be considered during a SAN migration.
- Virtual channels are a part of Brocade's FC switching architecture. In cascaded environments, the virtual channels enable multiple conversations to share the same ISL without contention.
- Multiple zones with different quality of service can be deployed. This is an especially important consideration for mixed workload environments.
- Ultra-low latency

Brocade FC switches are built on application-specific integrated circuits (ASICs) that support local switching. Local switching refers to when the input and output ports of an FC frame are on the same ASIC. Local switching on Gen 8 provides sub-microsecond latency (as low as 580 ns with dual FEC active) when the channel and control unit connect to ports on the same ASIC (or on the same port blade in the FC128-48, which has a single ASIC).

Naming convention

Before the mandatory FICON logical switch was required, the difference between a chassis, switch, and fabric were subtle. Now that they are distinct, creating good naming conventions is important. Naming conventions are preferences that vary for each organization. The only recommendation is that each name is unique and meaningful. Chassis, fabric, and switch names can be as long as 64 characters.

Domain ID and switch ID

The recommended best practice is to make the hex value of the domain ID (DID) match the Switch ID in HCD. This recommendation reduces confusion and simplifies troubleshooting. HCD changes and cable plant layout must be considered during a SAN migration.

Supported hardware and software

Ensure all components are compatible. Your FICON fabric design should center around Gen 8 hardware and FOS 10.x as the strategic platform. Complete an inventory check to ensure that all hardware and software is supported and that interoperability between new and existing components is validated. The recommended best practice is to complete all required changes to the environment before other upgrades are made to compartmentalize changes.

Typical redundant pathing

The 8-way pathing architecture remains the best practice. With Gen 8 FC128-48 port blades, carefully map your channel and control unit paths to ensure symmetry and redundancy.

8-way pathing to DASD spread across four directors provides the highest level of redundancy. Although two paths are connected to the same director, those paths are on different port blades, ensuring that a single blade failure does not affect both paths to the same device.

A common design approach is to assign the first path group to ports on one set of port blades and the second path group to ports on a different set of port blades. The objective is path separation and management simplicity rather than ASIC locality. Although connecting channels and control units to the same port blade can enable local switching and slightly lower latency, this is optional and generally beneficial only for highly latency-sensitive I/O environments.




Figure: 8-way pathing architecture for FICON DASD

Cascaded performance and virtual channels

Gen 8 replaces the legacy static Traffic Optimizer with the Adaptive Traffic Optimizer (ATO). ATO automatically learns and adapts to SAN traffic patterns to optimize performance.

Virtual channels allow for multiple conversations over the same link virtually at the same time. Even when all traffic is the same priority, virtual channels prevent a single busy flow from monopolizing an ISL.

For most enterprises, striping CHPIDs and devices at a diagonal across the channel cards is adequate, rather than straight across, to distribute load more evenly.

The use of Quality of Service (QoS) zones may be preferable in some cascaded configurations, particularly in the following scenarios:

- ISL bandwidth is limited and there is a mix of tier 1 applications and backup or mirroring applications.
- Cascaded CTC paths are defined. CTC traffic is typically small block. Each block, regardless of size, uses at least one FC frame. Each FC frame requires a buffer credit per hop. Each FC frame also incurs the ASIC processing overhead for that switch hop. Cascaded CTC is often put into its own low QoS zone.

The Brocade Gen 8 Virtual Channel architecture manages traffic by allocating flows to specific QoS priorities across 32 dedicated Virtual Channels (VCs) per port. These 32 VCs are architected as follows:

- VC 0–15: Data Traffic (16 VCs)
- VC 16–23: Fabric Services (8 VCs)
- VC 24–27: Control Frames (4 VCs)
- VC 28–31: Other/Reserved (4 VCs)




Figure: Virtual channels

Planning for ultra-low latency

Planning for ultra-low latency takes more planning and is done only when the application demands the highest possible performance. Brocade ASICs are designed to use cut-through routing. When the header of a frame is read, the frame is forwarded to the destination port before the entire frame is received. This design minimizes latency.

Local switching on Gen 8 provides sub-microsecond latency (as low as 580ns even with dual FEC active) when the channel and control unit connect to the same port blade. This latency is additive when traffic is traversing multiple ASICs. All ASIC-to-ASIC switching is performed by way of a backplane, which adds latency.

In Gen 8, the architecture is significantly simplified and optimized:

- Single ASIC per port blade: The Gen 8 FC128-48 port blade features a single, high-performance Condor 6 ASIC that manages all 48 front-end ports.

- True local switching (Red Loop): Because there is only one ASIC per port blade, any port-to-port communication on the same blade is handled entirely within that single ASIC — no backplane traversal required. This is the optimal path for ultra-low latency.
- Backplane routing (Green path): Backplane transit is strictly reserved for inter-blade traffic moving from a port on one blade to a port on a different blade.




Figure: Local switching versus switching through the backplane


Channel paths: 32 versus 48 port cards

Because Gen 8 relies heavily on the 48-port FC128-48 blades, channels and control units that previously used different 32-port field replaceable units (FRUs) in older generations might now land on the exact same FRU. As shown in Figure 6-5, CHPID 10 features an entry link address of xx00 and control unit link addresses xx0F and xx1F. Similarly, CHPID 11 has an entry link address of xx20 and control unit link addresses xx2F and xx3F. Each is on their own FRU when connected to a 32-port card. Assuming sequential addressing, with the 48 port cards, CHPIDs 10 and 11 are on the same FRU and the control unit at link address 3F is on a different FRU than the CHPID.

Note: Starting with Fabric OS 10.0.1, Unified Storage Fabric (USF) is fully supported on Brocade X8 Directors. The right half of the FC128-48 blade (ports 24 through 47) consists of "AnyIO" ports, which are Ethernet-capable and support Brocade USF for IP storage connectivity at 10G, 25G, 50G, and 100G speeds.

FICON CHPID Mapping Legend:

■ xx0x Group A: Mapping Ports 00–0F ■ xx1x Group B: Mapping Ports 10–1F ■ xx2x Group C: Mapping Ports 20–2F



FC128-48 Port Blade (Gen 8)

Figure: Gen 5 32-port card versus X8 48-port card channel path examples

				17	2F		47	5F
				16	2E		46	5E
				15	2D		45	5D
				14	2C		44	5C
				13	2B		43	5B
				12	2A		42	5A
				11	29		41	59
				10	28		40	58
0F	1F	2F	3F	0F	27	3F	57	
0E	1E	2E	3E	0E	26	3E	56	
0D	1D	2D	3D	0D	25	3D	55	
0C	1C	2C	3C	0C	24	3C	54	
0B	1B	2B	3B	0B	23	3B	53	
0A	1A	2A	3A	0A	22	3A	52	
09	19	29	39	09	21	39	51	
08	18	28	38	08	20	38	50	
07	17	27	37	07	1F	37	4F	
06	16	26	36	06	1E	36	4E	
05	15	25	35	05	1D	35	4D	
04	14	24	34	04	1C	34	4C	
03	13	23	33	03	1B	33	4B	
02	12	22	32	02	1A	32	4A	
01	11	21	31	01	19	31	49	
00	10	20	30	00	18	30	48	
Slot 1		Slot 2		Slot 1		Slot 2		

Figure: Gen 5 32 port card versus X8 48 port card channel path examples

The IBM CHPID Mapping Tool was not updated to account for 48 port blades. To aid in link address planning, an Excel Workbook is available from your IBM technical representative.

Chassis configuration options using 48 port cards

Addresses can be assigned to port cards in any order, but sequential addressing is highly recommended because it keeps documentation simple and the switch addresses align with the physical port layout. With 48 port cards, an 8-slot director can contain as many as 384 ports. However, because link addressing is limited to 254 usable addresses per logical switch (0xFE and 0xFF are reserved for CUP), careful planning is required. See the following table.

Table: Chassis configuration options with 48 port cards

Option	Considerations
8-slot Director, one FICON logical switch	Although this configuration is not recommended, it is possible. A logical switch can be created by using only the lower 16 ports on each column of ports per blade, effectively making it a 32-port card. An 8-slot chassis can be ordered with only five cards for 240 ports or six cards for 288 ports.
8-slot Director, 2 FICON logical switches per chassis	Creating two FICON logical switches per chassis allows for the greatest density — as many as 384 usable ports per chassis, more than 50% increase over single logical switch. The recommended layout is to populate the first logical switch with port cards on the left half of the chassis and the second logical switch on the right half.
4-slot Director	The maximum port count is limited to 192 ports. If FRUs were considered in the channel path planning, a new CHPID map is likely required.

Two logical FICON switches per chassis

Creating two FICON logical switches per X8 chassis allows for maximum density, yielding over a 50% increase in port density compared to a single logical switch per chassis. Instead of 254 usable ports per chassis, as many as 384 usable ports can be available per chassis.

When CUP is enabled, link addresses xxFE and xxFF are logical addresses. If no ports are associated with those addresses, those link addresses do not appear in the fabric and are not visible to the channel.




Figure: Eight logical switches for 8-way pathing option to fully use all ports

ISL bandwidth sharing with logical switches

A base switch can be used to share UltraScale Inter-Chassis Links (ICLs) among multiple logical fabrics. Gen 8 directors support a base switch that allows multiple logical fabrics to share the same physical ICL infrastructure.

Logical switches can be used for several other purposes. The idea of the base switch allows multiple logical fabrics to share the same physical ICL. XISLs feature the following limitations:

- They are not visible through CUP.
- System automation tools cannot be used to manage or monitor the links.
- RMF or Disk Magic cannot access the 74-7 records on these links.
- The quality of service (QoS) feature cannot be used.
- More bandwidth planning and congestion management considerations are required.
- Over subscription of XISL bandwidth in one logical switch affects available bandwidth for other logical switches.

It is possible for some logical switches to use the base switch (XISLs) while other logical switches can be configured to use direct ISLs. The base switch is the recommended approach for cascaded FICON environments.




Figure: Logical switch options

FICON SAN implementation

The SANnav Management Portal 3.0.1x replaces Network Advisor. SANnav includes more features for troubleshooting and monitoring and is required for Gen 8 management.

Chassis-level management

Chassis management in SANnav is performed on the default switch, which operates on a default Fabric ID (FID) of 128. All chassis-level operations — firmware upgrades, hardware monitoring, and HA status — are visible through the default switch context.

If you are new to the use of logical fabrics for FICON, you must leave the default switch that is discovered for chassis management. The default switch is not a FICON logical switch and must not be used for FICON traffic.

FICON display mode setup

In SANnav User Preferences, set the **FICON Display** to Yes to view RNID data and specific FICON columns. Mainframe systems use hex numbers for addresses and identifiers. Similar to the FICON display mode in Network Advisor, SANnav has an equivalent feature. This display mode ensures that numbers are represented in hex and adds RNID columns to the switch and switch port inventory views.

Setting up the FICON display mode is done on a per-user basis and should be the first setting that a FICON SAN user makes when logging into SANnav for the first time.

Complete the following steps in SANnav 3.0.1x:

1. In the upper right, click the person icon. In the drop-down menu, select **User Preferences**.




Figure: Configuring user preferences

2. Scroll to the bottom of the next window and find **FICON Display** under the Tables section. If the FICON Display is Disabled, click **Edit**.





Figure: Editing FICON Display

3. Check the **Enable FICON Display** option to show FICON-related columns in the switch and switch port inventory. Most users also uncheck the **Persist Last Filter Selection**, **Persist Table Column Customization**, and **Show Compare Dialog on Zone Configuration Activate**.

4. Click **Save**.



A dialog box for selecting user preferences. It contains a checked checkbox labeled "FICON Display" with an information icon to its right. Below this is a label "Persist Last Column Selection" followed by a dropdown menu currently showing "Yes". At the bottom of the dialog are two buttons: a blue "Save" button and a grey "Cancel" button.

Figure: Selecting user preferences

Configuring FICON fabric

Create a FICON logical switch and assign unique domain IDs that map to the HCD Switch ID. Standard fabric parameters like zoning and security settings are applied consistently across the fabric. Use the following steps to create the logical switch:

1. Navigate to the SAN configuration window. Click **SANnav** → **SAN Configuration** → **Logical Fabric Management**.



Figure: Switch configuration

2. Create a FICON fabric:
 - a. Click **FICON Fabrics**.
 - b. Click **+** to create a FICON fabric.




Figure: Selecting FICON fabrics

3. Configure fabric parameters:

- Enter a name for the fabric.
- Enter a description for the fabric (optional).
- Enter an FID number for the fabric.
- Click **Add**.

Figure: Configuring fabric parameters

The fabric parameters are listed in the following table.

Table: Fabric parameters

Parameter	Description
Name	The name of the fabric, not the logical switch.
Description	The description of the fabric. When cascading, the recommendation is to add a meaningful description such as "tape backup for ..." or "disk mirroring for ...".
Tags	A SANnav database search tag used for generating filtered reports. Unrelated to the RNID tag. Creating a tag is optional and can always be added later.
Fabric ID	A required integer 1–128 that identifies logical fabrics. The default FID for the default switch is 128. FID 127 is commonly used for the base switch for ISL sharing. The FID must be unique for each logical switch in a chassis and must be the same for all logical switches in the same fabric.

4. Select the logical switch chassis.




☐ Name ^	IP Address ↕	Logical Switch Count ↕	VF Status ↕	Firmware ↕
☐ Brocade6510	10.124.5.186	1	Enabled	v8.2.1
☐ Brocade_DCX4S	10.124.5.183	1	Enabled	v8.2.1c
☒ BrocadeG620	10.124.5.187	1	Enabled	v8.2.1a
☒ fcs01enwdco_X6	10.124.5.180	1	Enabled	v8.2.1b

1 – Select chassis to create FICON switch(es) for this fabric.

Next Cancel

Figure: Selecting the logical switch

- Set the Domain IDs for each switch in the fabric. The domain ID (DID) is the highest byte of the port address.



Name ^	Domain ID ↕
BrocadeG620	0x01
fcs01enwdco_X6	0x02

Back Next Cancel

Figure: Setting the switch domain ID

- Add the ports. Most administrators click the Slot/Port # column to sort by port number.




Figure: Selecting the port to add

- Click **OK** to finish adding the ports.




Figure: Completing the addition of the ports

- Configure the remaining fabric parameters. These parameters are applied to all logical switches in the fabric. Select **Fabric Parameters** to display parameter options. Enter the wanted parameter options as listed in Table 7-3 and click **Save**.

Click on the pull down arrow next to “Fabric Parameters” to display these options.

▼ Fabric Parameters

RA TOV	10000	☐ Allow XISL use
ED TOV	2000	☐ Enable device probing
Maximum Hops	7	☐ Per-frame routing priority
BB Credit	16	☐ Suppress Class F traffic
Data Field Size	2112	☐ Long distance fabric
Routing Policy	Exchange Based ▼ Device Based Exchange Based	

☒ Activate

It doesn't matter what order you select the options in. Once all options are selected, click “Save”

Figure: Configuring the remaining fabric parameters

Table: Fabric parameter definitions

Parameter	Description
Long-distance fabric	Adds buffer credits to ISL (E-Port) links. Use in cascaded environments when ISL distance exceeds approximately 1 km for standard FICON traffic, or for shorter CTC links. Typically not set for campus cascading. For distances exceeding 10 km, other methods provide buffer credits.
Activate	When selected, logical switches and ports are enabled after the creation process completes. It is not safe to perform cabling duties with enabled SFPs; leave unchecked when cabling is incomplete.
Discover	Controls whether logical switches are automatically discovered by SANnav. Logical switches are discovered by default.
Rename	Controls whether logical switches are renamed according to the fabric naming convention. Switch names default to Switch_ followed by the FID number.
Disable unused ports	Disables ports that have no SFP installed or no device connected. It is not safe to have enabled SFPs with no connection.
Zone ports	Automatically creates a zone containing all non-ISL ports. Because the default zoning mode is "No Access" in FOS 10.x, a zone must be created for connectivity.
Zone configuration	Creates a zone configuration containing the zone created in the previous step.
Disable All-Access	Disables the All-Access zone policy for added security.

9. Discover the logical switches. Logical switches are discovered by default. The system administrator might disable automatic discovery, or a network error might cause discovery to fail. Verify that all logical switches appear in the SANnav inventory.
10. Rename the logical switches. Switch names default to Switch_ followed by the FID number. Rename logical switches to meaningful names aligned with your naming convention.
11. Disable unused ports. It is not safe to have lasers enabled on SFPs to which nothing is connected. In some instances, the open connector can reflect enough light back to the SFP to keep it active.
12. Zone the FICON ports. Because the default zoning mode is "No Access" in FOS 10.x, administrators must ensure device connectivity by creating FICON zones:
 - Put all ports into a single zone that is based on the port domain and index.
 - Do not put ICL ports in the zone.
 - Although it is ideal to not have any E-Ports in the zone when cascading, having the E-Ports in the zone does not cause problems.
 - Consider more zoning granularity. For example, separate mirrored ports with zoning.
 - Domain and port index zoning is always used for FICON.
 - To simplify zoning algorithms and to accommodate various switch form factors, a port index rather than a port number is used for FICON zoning.
13. Create a zone configuration with the zone, then activate the zone configuration, and disable the All-Access option. Complete the following steps to create a zone:

a. Select **Zoning**, then select **Zones**, then click **+**.




Figure: Creating a zone

b. Select the fabric. Zoning is applied to a fabric, not an individual switch. Because most FICON fabrics are single-switch fabrics, there is typically only one fabric to select.




Figure: Selecting the fabric to zone

c. Name the zone and configure the type of zone.

1 – Name the zone

2 – Optional – Add a description of the zone

Name: FICON_Zone

Description: All ports based on d.i

Tags:

Zone Type: Standard Zone

Fabric: FICON-XRC-z15

3 – Select “Standard Zone” for the Zone Type

4 – Click “Add”

Members: 0 Items

Zone Ali... Member Count Type Status Add

Figure: Name zone and configure zone type

d. Add the ports to the zone. Typically, all but the ISL ports are added to the zone.

Select discovered Devices/Ports

Enter manually

Alias: Domain, Port Index

1– By default zone by “Alias” will be selected. For a FICON zone, select “Domain, Port Index”

2– Select the ports to add to the zone. For most FICON fabrics, all ports will be selected which can be accomplished with a single check in the upper box by “Domain, Port”

3– Click the right arrow key to add the ports to the zone.

Alias	Domain, Port Index	P	F	S	Zone Ali
0X65, 0X74	sw14...	6500...	sw14...		
0X65, 0X75	sw14...	6515...	sw14...		
0X65, 0X76	sw14...	6517...	sw14...		
0X65, 0X77	sw14...	6503...	sw14...		
0X65, 0X78	sw14...	6511...	sw14...		

Figure: Adding the ports to the zone

e. Complete the zone creation and save it.

1– Click “OK” at the bottom of the page after adding ports to the zone.

2– Click “Save” at the bottom of the page after completing the zone definition.

OK

Save

Figure: Completing the zone creation

f. Add the zone to a zone configuration. A zone configuration is a collection of zones that is activated together.




Figure: Creating a zone configuration

g. For FICON users who defined only one zone, adding that zone to a zone configuration can seem redundant.




Figure: Defining the zone configuration

h. Add the zones to the zone configuration. In open systems environments, it is common to have multiple zone configurations; therefore, activating the zone configuration is optional. For most FICON environments, only one zone configuration is used, which is often activated when the zone configuration is created.

Add Zones

1 – Typically only one zone with all ports is added to the zone configuration but if a more granular zoning approach was used, select all zones that should be in the zone configuration.

☐	Name ^	Type
☒	FICON_Zone	Standard
☐	gen7-xrc	Standard

2 – Click "OK"

3 – Check "Activate" to activate the zone configuration.

4 – Click "Save"

Buttons: OK, Create New, Cancel, Activate, Save, Cancel

Figure: Saving and activating the zone configuration

i. For added security, disable the All-Access zone.

Zone Configurations ▾ Zones ▾ Zone Aliases Policies

1 – Select "Policies"

☐	Name ^	WWN
☐	Base_Fabric	10:00:C4:F5:7C:D9:DD:0E
☐	Fabric_B	10:00:C4:F5:7C:D9:DD:DB
☒	FICON-XRC-z15	10:00:88:94:71:E0:32:8C

2 – Select the fabric

3 – Use the pull down arrow by the fabric to select "Disable (No Access)"

Buttons: All Access, All Access, Disable (No Access)

Figure: Disabling the All-Access policy

Note: In Gen 8 platforms, 128G F-Ports are automatically enabled with Forward Error Correction via Turbo-Transmit and Sync (FEC-via-TTS). This provides enhanced error correction for legacy FICON channels at 128G. FEC-via-TTS is transparent to the channel and requires no configuration.

Managing FICON fabrics

Gen 8 hardware and FOS 10.x introduce SAN Fabric Intelligence (SAN FI), providing holistic, end-to-end visibility into SAN health and performance. SANnav Management Portal 3.0.1x introduces several new FICON-specific columns to improve troubleshooting.

Customizing the display

SANnav allows customization of host and storage port columns to show FICON RNID tags. Setting the FICON display mode defaults all column displays to the typical displays with RNID data that most mainframe administrators require.




Figure: Customizing the display

Viewing connections

Use the SANnav Inventory tab to view Host Ports (to see channels) and Storage Ports (to see control units).

1. Select **Inventory**.
2. Select **Host Ports** to see the channels, or **Storage Ports** to view the control units.




Figure: Viewing channel connections

The resulting display is shown in the following figure.

Link address is the first 2 bytes of the FC address

RNID data

FC A... ^	Type ⚙	Model ⚙	Manu... ⚙	Sequence # ⚙	RNID Tag ⚙
651300	N	T01	IBM	00000003F6E8	309c
651400	N	T01	IBM	00000003F6E8	309f
651500	N	994	IBM	0000000KNK21	0231

Figure: CHPID connectivity display

Similarly, control units can be monitored by clicking **Storage Ports** instead of Host Ports.

SANnav 3.0.1x introduces several new FICON-specific columns to improve troubleshooting. When viewing the Host Ports or Storage Ports inventory, the following columns are available: CHPID, CSS ID, MIF flag, link address, UID, and serial number.

Converting a production switch to a FICON logical switch

Converting a switch that is not already configured for FICON requires downtime. Ensure that new environments are deployed as FICON logical switches from the start to avoid this requirement.

The easiest way to check is by using CLI. Save the session output so that you can email it to your Brocade Business Partner or IBM representative for analysis. Use the following CLI commands to capture the current switch configuration:

```
switchshow
fabricshow
portaddress --show
```

Checking high integrity fabric and address mode

Validate that the address mode is mode 1 (zero based) and that the high integrity fabric (HIF) mode is set.

```
switch_30:FID30:admin> switchshow
```

switchName	switch_30
switchType	166.0
switchState	Online
switchMode	Native
switchRole	1
switchDomain	Principal
switchId	fffc01
switchWwn	10:00:c4:f5:7c:64:5b:62
zoning	OFF
switchBeacon	OFF
FC Router:	OFF
HIF Mode	ON
Allow XISL Use	OFF
LS Attributes	FID: 30, Base Switch: No, Default Switch: No, Ficon Switch: No, Address Mode 1]

Figure: Validating the address mode

If HIF Mode is not On or the Address Mode is not 1, an outage is likely needed to reconfigure the switch.

Validating address binding

Use the `portaddress --show` command to verify that all ports have an address successfully bound ("Y" must appear in the User bound column).

Example: Validating address binding

```
switch_30:FID30:admin> portaddress --show
Partition Address Mode :1
```

Index	Slot	Port	Area	Mode	User bound
72	9	8	0x0000	8-bit	Y
73	9	9	0x0200	8-bit	Y
74	9	10	0x0300	8-bit	Y
75	9	11	0x0400	8-bit	Y
76	9	12	0x0500	8-bit	Y

A Y must appear in the column for all ports.

Validating FICON logical switch

Run the `switchshow` command. `HIF Mode` must be `ON`, and the LS Attributes must show `FICON Switch: Yes`. The example shown in the following figure is a standard logical switch and therefore displays "FICON Switch: No".


```
switch_30:FID30:admin> switchshow
```

switchName	switch_30
switchType	162.0
switchState	Online
switchMode	Native
switchRole	Principal
switchDomain	1
switchId	fffc01
switchWwn	10:00:c4:f5:7c:16:8b:3e
zoning	OFF
switchBeacon	OFF
FC Router:	OFF
HIF Mode	ON
Allow XISL Use	OFF
LS Attributes	FID: 30, Base Switch: No, Default Switch: No, Ficon Switch: No, Address Mode 1]

Figure: Validating FICON logical switch

Note: Fabric OS v10.0.0a introduces cryptographic checksum verification for switch configuration backups. Checksums must match before a configuration restore is permitted, preventing tampered or corrupted backups from being applied.

Routing policy

This policy primarily affects load balancing in cascaded environments. Most FICON environments today support FICON Dynamic Routing (FiDR), which corresponds to Exchange-Based Routing (EBR) in FOS. The routing policy must be consistent across all switches in a fabric; however, each fabric operates independently.

A route is the set of switches and ports a path takes to deliver FC frames from the channel to the control unit. The routing policy determines how the route is selected.

Table: Routing methods and descriptions

Routing method	Description
Port Based Routing (PBR)	Deprecated. Routes are determined by a predefined map for each possible route in the SAN. The channel path always takes the same route. The only reason to configure PBR today is if a new switch must be connected to an older switch still configured for PBR. Setting PBR is not an option in SANnav; use the CLI.
Device-Based Routing (DBR)	Default. Routes are determined by usage. A round-robin algorithm selects the ISL or ISL trunk group. Routes are established dynamically as a channel to control unit path is needed, but once selected, the path remains the same. Provides better statistical load balancing across ISLs than PBR.
Exchange-Based Routing (EBR)	Analogous to FICON Dynamic Routing (FiDR). Strictly recommended for new Gen 8 installations when FiDR is supported by all channels and control units. Selects routes dynamically on a per-exchange basis (per CCW chain). The least-used ISL or ISL trunk group is selected at the time a new CCW chain is received. Results in the most efficient and dynamic use of ISL bandwidth.

Exchange-based routing and FiDR background

Mainframe channel protocol was developed for buffered control units that manage physical storage devices. Channel programs are composed of Channel Command Words (CCWs). Write commands include data. CCWs can be, and often are, strung together in a channel program.

Status is a bit flag in a byte returned by the control unit. The only status bits pertinent to this discussion are Channel End (CE) and Device End (DE). CE is sent by the control unit when it completes receiving and has processed the data. DE is sent by the control unit when the physical device has completed its operation.

In a serialized bit stream across FC, CCWs are divided into FC packets, and the packets are sent to their destination. All CCWs, and all FC packets in a channel program, are sent on the same exchange so all packets arrive in order.

When split status is returned to the channel, DE is returned on a different exchange. Because the protocol was designed for buffered control units that had already completed the operation, the channel program never had to account for receiving DE before CE.

FICON SAN migration

Getting a SAN Health report and a Brocade Support Link (BSL) report is the best way to determine migration tasks. It is also the best way to ensure that a support case is opened with the relevant data.

Important: Migrate to Gen 8 hardware and FOS 10.0.1 to ensure access to Frontier AI security updates and PQC compliance. Delaying migration increases security risk and limits access to new performance features.

Firmware

FICON switch configurations require FICON-supported FOS levels (FOS 10.0.0a/10.0.0a1 or FOS 10.0.1 for Gen 8, and 9.2.2c or higher for Gen 7). Firmware updates are an important part of maintaining the security of the SAN infrastructure. Before installing a Fabric OS update, verify that the firmware image was obtained from an authorized

IBM or Broadcom source and that its integrity and authenticity can be validated. Follow the applicable IBM and Broadcom procedures for verifying, downloading, and installing Fabric OS firmware. Keep firmware versions current according to the supported upgrade paths and compatibility requirements for the environment.

Attention: FCIP FICON Acceleration features (XRC, TAPE, Teradata, TinTlr, and DcvAck) are officially no longer supported in FOS 10.0.x. Customers relying on FCIP FICON Acceleration must evaluate their extension strategy before upgrading to FOS 10.x.

Attention: Fabric OS 10.x strictly enforces advanced security standards. All legacy cryptographic algorithms — including MD5, SHA1, and DSA — are permanently removed. If any switch in the fabric uses these algorithms, the firmware upgrade from FOS 9.2.x to FOS 10.x will be blocked.

Note: Fabric OS 10.x strictly enforces the presence of a valid Trusted FOS (TruFOS) certificate to authenticate firmware packages. Switches without a valid TruFOS certificate cannot be upgraded to FOS 10.x without first obtaining the certificate from Broadcom.

Note: It is strongly recommended to upgrade Gen 8 platforms to FOS v10.0.0a/v10.0.0a1 or higher immediately upon receipt. Gen 8 platforms should not be downgraded below FOS 10.0.0a.

Attention: Fabric OS 10.x completely obsoletes non-secure management protocols. If SNMPv1, FTP, HTTP, Telnet, TLSv1.1, or non-secure Syslog are configured or in use, the firmware upgrade from FOS 9.2.x to FOS 10.x will be strictly blocked.

Note: Environments with Gen 7 switches operating in AG mode with Slow Drain Device Detection (SDDD) configured should review the SDDD configuration before upgrading to FOS 10.x.

Attention: Fabric OS 10.x introduces several mandatory security enhancements based on the Principle of Least Privilege. Customer-created scripts stored on the switch may be automatically deleted during the upgrade process. Back up all custom scripts before upgrading.

FICON logical switch

Creating a specific FICON logical switch is strictly required. For FOS 10.0.x and later, the FICON Control Unit Port (CUP) license is permanently enabled by default. Management Server (FMS) mode is automatically enabled natively. A logical switch defined for FICON cannot be reconfigured as a non-FICON switch without a disruptive reconfiguration.

To avoid the need for an outage, it is highly recommended that all new switch configurations include virtual fabrics that are enabled with a FICON logical switch from the beginning.

Inter-chassis links

High-density cabling for Gen 8 infrastructure utilizes advanced optical form factors to achieve extreme port density. Gen 8 directors use OSFP (Octal Small Form-factor Pluggable) transceivers for UltraScale ICLs.

Connector technology — OSFP: Eight-lane transceivers used for X8 UltraScale ICLs. When connecting a Gen 8 (X8) director to a Gen 7 (X7) director, a breakout cable is required because the Gen 8 side uses OSFP and the Gen 7 side uses QSFP.

The following ICL connectivity and breakout requirements apply:

- X8 side: OSFP connector using MPO-16/APC (Angled Physical Contact), typically identified by a green connector body.
- X7 side: Two QSFP connectors using MPO-12/UPC (Ultra Physical Contact), typically identified by blue or aqua connector bodies.

Attention: Never attempt to plug an APC connector into a UPC port. The mismatched end-faces will cause permanent optical damage to both connectors and the transceivers.

Inter-Chassis Links (ICLs) are available in 100 m and 2 km options. The 100 m option requires OM4 or OM5 multi-mode fiber and an MPO/MTP connector. The 2 km option requires single-mode fiber.

It is recommended that both directors that are cascaded with ICLs are upgraded at the same time to avoid the need for purchasing new breakout cables or intermediate adapters.

Cabling and link diagnostics

Although wholesale single-mode cable replacement is generally not necessary in mainframe environments, it is not uncommon for cables to be reused from previous SAN generations. Cable hygiene is critical for high-speed optics.

Although multi-mode cable is rare in mainframe environments, some exceptions exist. The most common exception is for ICL connections within a datacenter, where the short distances make multi-mode practical.

As with any reuse of cabling, good cable hygiene is recommended, which must include cleaning all connection points.

Note: Because high-speed optics are highly sensitive to physical impairments, as a best practice, use Brocade ClearLink Diagnostics (D_Port) on all cables before and after deployment. D_Port testing validates cable plant quality including optical loss, distance measurements, and end-to-end link integrity.

Patch room and fiber plant

Migrating from products with different form factors or changing port count can affect documentation and patch room design. When migrating from Gen 7 to Gen 8 directors, the physical blade orientation changes from vertical to horizontal (in 8-slot models). All documentation templates and diagrams that include chassis slot numbering must be updated.

SFPs

As with all previous FC switching products from Brocade, the latest generation product supports SFPs of the same generation and one generation below. A Gen 8 switch port is capable of 128G; therefore, it can accept a 128G SFP+ or a 64G SFP+. A Gen 7 switch port is capable of 64G; therefore, it can accept a 64G SFP+ or a 32G SFP+. Each SFP supports three speeds: the current generation, one generation back, and two generations back. Consider the following points:

- A 128G optic can auto-negotiate to 128G, 64G, and 32G.

- A 64G optic can auto-negotiate to 64G, 32G, and 16G.
- A 32G optic can auto-negotiate to 32G, 16G, and 8G.

If an 8G connection is required (typically for older tape or DWDM), you need a 32G SFP+. Because 32G SFPs and 8G port speeds are not supported on Gen 8 products (the minimum supported speed on a Gen 8 FC128-48 blade is 16G), you need a Gen 7 switch or Gen 6 switch to accommodate older connections that require 8G. You can use the Brocade-branded 32G SFPs from your existing products in a Gen 7 or Gen 6 switch, but do so only if 8G support is required.

Note: Gen 8 platforms and Fabric OS 10.0.1 officially support 128G Long Wavelength (LWL) optics capable of 10km and 128G Extended Long Wavelength (ELWL) optics capable of 25km. Connectivity between ELWL and LWL optics is strictly not supported; administrators must ensure identical optic types are used on both ends of an ISL.

Most FICON SANs are single-switch (single-domain) fabrics. Channel paths can be striped across multiple SANs using a mix of Gen 6, Gen 7, and Gen 8 technologies. The fabrics do not need to run the same FOS version; each fabric can use any FICON-supported FOS release.

Note: Gen 5 reached End of Support on 30 April 2025 and should not be merged into new fabrics. Stand up new Gen 8 fabrics and swing the cables.

Note: Because separate FICON fabrics operate independently, each fabric can run a different FOS version, provided that the version is supported for FICON.




Figure: Mixed Gen 8 and Gen 7 Single Switch Fabrics

In cascaded environments, a non-VF switch can connect to a VF-enabled logical switch. However, all switches in the fabric must operate within the same logical fabric context; otherwise, the fabric will not form correctly.

Configuration requirements

The following configuration requirements must be met when operating mixed-generation fabrics:

- The FOS code must be at v10.0.0a/v10.0.0a1 or higher (with the FOS 10.1.0x target path recommended for security) in the Gen 8 switch. Gen 7 switches must be at FOS 9.2.2c or higher.
- All fabric parameters must match.
- For FICON environments, the routing policy is the most likely parameter that is different. While many legacy implementations used PBR or DBR, FiDR (EBR in FOS) is the recommended standard for new deployments.
- Default Zoning Mode mismatches will cause fabric segmentation. In FOS 10.x, the default zoning mode (DefZone) is strictly set to "No Access". When merging a Gen 8 switch with an older switch that uses "All Access", the E-Port will instantly segment due to a Zone Conflict. Ensure all existing switches are set to "No Access" or have an active zone configuration before merging.
- It is unlikely that other fabric parameters are different; however, if switch generations must be mixed in the same fabric, a careful review of all fabric parameters is required before the merge.




Figure: Sample mixed Gen 8 and Gen 7 cascaded fabrics

FMS and CUP

In Fabric OS 10.0.x and later, the FICON Control Unit Port (CUP) license is permanently enabled by default; therefore, FMS mode is automatically enabled. A FICON logical switch cannot be reconfigured as a non-FICON switch without a disruptive reconfiguration.

FICON port naming enhancement

In Fabric OS 10.0.0a, a new PortName option has been added to the `ficoncupshow` and `ficoncupset` commands. This allows administrators to assign descriptive names to FICON CUP ports for easier identification in management tools.

Gen 7 to Gen 8 fixed-port switches

The Gen 7 SAN64B-7 (Brocade G720) and the Gen 8 SAN56B-8 (Brocade G820) are the only FICON-supported fixed-port switches. When upgrading, the new switch must be placed into a FICON logical switch context.

Gen 7 to Gen 8 directors

When upgrading from Gen 7 directors (SAN512B-7/X7-8 and SAN256B-7/X7-4) to Gen 8 directors (SAN512B-8/X8-8 and SAN256B-8/X8-4), the following physical differences must be considered:

- In the Gen 7 X7-8 chassis, port blades are mounted vertically. In the Gen 8 X8-8 chassis, port blades are mounted horizontally.
- Cable routing and reach change because the blade orientation changes from vertical to horizontal, and the blade count per row increases.
- CHPID path mapping: The physical port locations have shifted from Gen 7 to Gen 8, meaning channel paths will land on entirely different physical ports after migration.
- Cable diagrams and documentation: All physical documentation templates and diagrams that include chassis slot numbering must be updated.
- Patch room design: Patch panel layouts designed to match the vertical symmetry of Gen 7 blades will need to be redesigned.
- Automation: Custom scripts and SANnav software alerts for automatic service ticket generation must be reviewed and updated.

Because of these profound physical differences, organizations should consider migrating fabrics non-disruptively by utilizing the new Gen 8 chassis in parallel with the existing Gen 7 chassis before decommissioning.




Figure: Gen 8 versus Gen 7 card slots

Consider the following points regarding physical upgrades:

- All documentation templates must be updated.
- Custom scripts and software to automatically generate service tickets or other monitoring activity might need to be updated.
- The patch room layout might need to be redesigned.
- Cabling methodologies likely will change.

Documentation and references

For more information, see the following resources:

- [Brocade FICON Resource page](#)
- [SANnav Management Portal User Guide](#)
- [SANnav Management Portal Release Notes](#)
- [SANnav Management Portal Installation and Upgrade Guide](#)
- [Brocade Transceiver Support Matrix](#)
- [High-Density Cabling Design Guide](#)
- [Brocade X8-8 Director Hardware Installation Guide](#)
- [Brocade X8-4 Director Hardware Installation Guide](#)
- [Brocade G820 Switch Hardware Installation Guide](#)
- [Brocade X7-8 Director Hardware Installation Guide](#)
- [Brocade X7-4 Director Hardware Installation Guide](#)
- [Brocade G720 Switch Hardware Installation Guide](#)
- [IBM FICON Qualification Letters](#)

Summary

IBM b-type Gen 8 hardware and Fabric OS 10.x represent the strategic platform for FICON SAN environments. The transition to Gen 8 delivers significant improvements in performance (128G throughput, sub-microsecond local switching latency), security (PQC, TruFOS, Podman-based management), and density (384 ports per chassis with two logical switches).

Key design principles for Gen 8 FICON environments include:

- Always deploy FICON Logical Switches from the start; converting existing switches requires downtime.
- Use Exchange-Based Routing (EBR/FiDR) for all new deployments where supported by channels and control units.
- Plan for 48-port blade addressing in HCD and CHPID mapping; update all documentation.
- Use SANnav Management Portal 3.0.1x or higher for all Gen 8 FICON management.
- Enable FICON Display mode in SANnav user preferences before beginning any management tasks.
- Validate HIF mode, address binding, and routing policy before placing a FICON logical switch into production.

Introduction to Brocade Unified Storage Fabric

Brocade Unified Storage Fabric (USF) is available on Brocade X7 and X8 directors with the 48-port blade (X7: FC64-48 and X8: FC128-48). Half of the ports on the blade (24) are AnyIO ports, which can be configured for either FC or Ethernet. There are no licenses required to use USF.

The following sections describe the use cases, advantages, and deployment models for USF.

Architectural challenges in dual-network storage environments

Modern enterprise data centers face continuous pressure to increase application density via hyper-converged compute infrastructures, server virtualization, and container frameworks. These infrastructures support diverse workloads that use both Fibre Channel (FC) and IP-based storage protocols, depending on storage tiering requirements and application criticality.

Traditionally, this split results in operating two disparate network topologies:

- Fibre Channel SANs: A dedicated network environment purpose-built for storage traffic, providing deterministic, lossless, and secure transport managed entirely by a storage operations team.
- IP/Ethernet Networks: A general-purpose network supporting a mix of data traffic types (for example, front-end access, server-to-server traffic, internet traffic) managed by a data center network team.

When enterprises scale block- or file-based IP storage over these shared data center networks, severe operational challenges regularly emerge. Contention between general application networks and latency-sensitive storage often yields unpredictable performance and availability. Furthermore, many IP network administrators are less familiar with the unique requirements of mission-critical storage traffic, which can cause routine maintenance activities or congestion events to adversely affect storage services. As a result, minor network maintenance routines or congestion points can escalate into catastrophic storage system disruptions.

The Brocade Unified Storage Fabric solution

Brocade Unified Storage Fabric (USF) mitigates these dual-network shortcomings by allowing enterprises to consolidate all storage traffic over existing, dedicated high-performance Fibre Channel infrastructure.

Note: Brocade USF is a cyber-resilient, dedicated storage network featuring integrated services that simultaneously support both traditional Fibre Channel protocols and IP-based storage traffic on a single physical fabric.

USF isolates Fibre Channel and IP storage streams through a logically separated architecture inside the director chassis. It leverages logical fabrics, logical switches, and software-defined AnyIO ports to guarantee that traffic flows in parallel without protocol translation or data blending between environments.

Supported storage protocols

With USF, the Brocade fabric supports the following storage protocols across both Fibre Channel and IP:

- Fibre Channel (FC):
 - FCP (SCSI)
 - FICON
 - FC-NVMe
- IP:
 - iSCSI
 - NVMe/TCP
 - NAS

Key benefits of USF

Organizations running business-critical applications on IP storage that use USF can achieve the following benefits:

Management and provisioning simplification

With USF, a single team owns the fabric infrastructure and associated tasks for management operations and provisioning. Customers emphasize the importance of having both the storage and SAN fabrics owned directly by the storage team to meet service level agreements (SLAs). Single-team ownership mitigates cross-organizational friction with network groups and minimizes application delivery bottlenecks.

Consolidated management of FC and IP storage is integrated within the fabric, mirroring traditional FC SAN management workflows to simplify configuration and administration.

Diagnostic and monitoring capabilities familiar to the storage team apply to both storage and networking environments, minimizing operational risk, complexity, and troubleshooting time.

Cyber resiliency by design

The Brocade infrastructure includes silicon root of trust hardware and software integrity validation on startup as well as periodically to provide a tamper-proof infrastructure.

While IP servers on general data center networks potentially expose their local storage subnets to lateral threats when a single device is compromised, USF securely isolates the storage infrastructure away from general internal and external IP data networks.

Enterprise reliability and performance

By extending Brocade fabric services to IP storage block-level protocols, such as iSCSI and NVMe/TCP, USF enables organizations to leverage redundant dual-fabric architectures and lossless transmission capabilities that are fundamental to high-end SANs, helping ensure 99.9999% availability.

IP storage logical switches provide loop-free networks (requiring no Spanning Tree Protocol) and do not permit BUM traffic (broadcast, unknown unicast, and multicast), which can negatively affect storage performance.

Hardware support and specifications

Brocade USF is supported in Fabric OS (FOS) v9.2.1 and later.

Both the X7 and X8 48-port blades designate 24 ports as AnyIO slots, while the remaining 24 ports function strictly for Fibre Channel transport. AnyIO ports support both Fibre Channel and Ethernet connectivity and can be reconfigured for either protocol as required. In accordance with this, the SFP changes personality to either Ethernet or FC depending on the port configuration.

Technical specification matrix

Feature	Brocade Gen 7 (Condor 5 ASIC)	Brocade Gen 8 (Condor 6 ASIC)
Switch models	X7-8, X7-4	X8-8, X8-4
Port blade	FC64-48	FC128-48
AnyIO ports	24 ports per blade (ports 16–23, 32–47)	24 ports per blade
Max port speed	64G FC / 25G Ethernet	128G FC / 100G Ethernet
FC speeds	8/16/32/64G FC	16/32/64/128G FC
IP speeds	10/25G Ethernet	10/25/100G Ethernet
MTU size	2k	9k (Jumbo Frame Support)
Minimum software version	FOS v9.2.1	FOS v10.0.1

Reference design and deployment

Direct attachment architecture

The direct attachment architecture is designed for smaller, tightly controlled environments, with server NICs and storage array controller ports connected directly to designated AnyIO ports on the IP storage (IPS) logical switch, which provides the IP storage services within the USF architecture. This simple deployment model facilitates the integration of additional storage capacity and can typically scale to several thousand endpoints.

Unified Storage Fabric




Figure: Direct Attach Architecture

ToR aggregation architecture

The most common and flexible deployment design separates storage traffic at the data center Top-of-Rack (ToR) switches. Separating storage traffic from the broader data center network at the first-hop ToR boundary helps protect storage services from network congestion and other data center network disruptions.

Uplinks to USF can be configured at Layer 2 or Layer 3 using port channels with a minimum of two physical connections per uplink.

When configured at Layer 3, USF uses static routes because the dual-fabric architecture and multipath storage connectivity do not require dynamic routing protocols. This approach simplifies deployment, troubleshooting, and ongoing operations while insulating the USF network from route changes and other events on the broader data center network.

USF's dual-fabric architecture aligns well with block-level multipathing and provides redundancy. Implementations commonly adopt the dual-fabric model used in Fibre Channel SAN environments, enabling concurrent path connectivity through Fabric A and Fabric B.




Figure: ToR Aggregation Architecture

Unified replication architecture

In replication deployments between data centers, USF provides an ideal solution when the storage replication ports are a mix of FC and IP. In these solutions, long distance ISLs are already in place for FC replication and can be provisioned for use by USF for IP storage replication.

On Gen 8 platforms, the long distance ISLs can be configured as XISLs and transport both FC and IP replication traffic on the same XISL.

On Gen 7 platforms, the ISLs used for IP replication traffic must be dedicated ISLs, separating the IP replication from FC replication traffic on the ISLs.




Figure: Unified Replication Architecture

Deployment

For details on deploying USF, see the [Brocade Fabric Administration Guide](#).

Migration to Gen 8

Gen 7 to Gen 8 migration

This section describes the process of migrating from Brocade Gen 7 platforms to Gen 8 platforms. The migration requires verification of Fabric OS compatibility, fabric interoperability, zoning consistency, and hardware readiness to ensure a smooth transition with minimal disruption to SAN operations.

Pre-migration assessment

Before starting the migration, assess the existing SAN environment to verify hardware, software, compatibility, zoning, and fabric readiness. This assessment helps identify potential issues and ensures a smooth migration to the target platform.

Migrating to FOS v10.0.0a

The supported upgrade paths to Fabric OS v10.0.0a are as follows:

Note: The upgrade path table is not available in the source document. See the Brocade Fabric OS Software Upgrade User Guide, 10.0.x for the full list of supported upgrade paths.

See the *Brocade Fabric OS Software Upgrade User Guide, 10.0.x* for detailed instructions on non-disruptive and disruptive upgrade procedures. In addition, review the *Fabric OS v10.0.0a Release Notes* to verify supported upgrade paths, hardware compatibility, interoperability requirements, and known issues before performing the upgrade.

Note: Gen 8 platforms do not support downgrade from FOS v10.0.0a to v10.0.0. This does not apply to Gen 7 platforms.

Inventory and baseline

Before making any changes to the SAN environment, collect and document the current switch, fabric, zoning, and configuration information. This baseline provides a point of reference for validating the migration, troubleshooting issues, and restoring the environment if necessary. Capturing this information also helps verify compatibility and identify potential migration risks.

Run the following commands on the Gen 7 switch to collect baseline information:

```
switchshow
chassisshow
fabricshow
version
licenseshow
cfgshow
portshow
```

Capture and document the following information:

- Current FOS version (must be compatible with the target Gen 8 FOS release)
- Domain ID assignments
- ISL topology

- Zoning database checksum
- Port types (F, E, EX, and VE ports)

Back up the current configuration and support data:

```
configupload
supportsave
```

FOS 10.x compatibility planning

Before integrating Gen 8 switches into an existing SAN fabric, verify Fabric OS interoperability and supported upgrade paths by consulting the applicable Fabric OS Release Notes, Software Upgrade Guide, and Hardware/Firmware Support Matrix. Fabric interoperability is dependent on the specific FOS levels running on each platform, and supported combinations can vary by release.

Ensure that all Gen 7 and Gen 8 switches are running Fabric OS versions that are explicitly supported for coexistence. When required, upgrade existing switches to a supported FOS level before establishing ISLs or attempting a fabric merge.

Important: Gen 8 platforms require Fabric OS 10.x or later. See the applicable release documentation to determine the minimum supported FOS level for interoperability with existing fabric components.

- All switches must run compatible FOS before fabrics can merge.
- Mixed-gen fabrics are supported only within the validated FOS matrix.

The following rules apply:

- Upgrade Gen 7 to the minimum FOS version required for Gen 8 coexistence.
- Gen 8 platforms require Fabric OS 10.x or later. Verify the installed Fabric OS version before deployment.

Gen 8 initial deployment

Before integrating the Gen 8 switch into the existing SAN fabric, perform the initial chassis configuration. This setup establishes the switch identity, management connectivity, and basic operating parameters required for fabric integration and ongoing administration.

```
ipaddrset
chassisname
configure
```

ISL integration

After the Gen 8 switch has been deployed and configured, establish Inter-Switch Links (ISLs) between the Gen 7 and Gen 8 switches to enable fabric connectivity and communication. Before creating the ISLs, verify that Fabric OS versions are compatible, zoning information is synchronized, and all fabric merge requirements are met. Proper ISL integration ensures that devices can communicate across the fabric and supports a smooth migration of workloads to the Gen 8 platform.

Create ISL

Create Inter-Switch Links (ISLs) between the Gen 7 and Gen 8 switches to establish connectivity and enable communication across the SAN fabric. ISLs allow fabric services, zoning information, and device information to be shared between switches, providing the foundation for a seamless migration and fabric integration.

```
portcfgtrunkport <port> 1
portenable <port>
```

Verify

After completing the base chassis configuration, verify that the management IP address, chassis name, switch status, and Fabric OS version are configured correctly before proceeding with fabric integration. This validation helps identify configuration issues early and ensures that the switch is ready for deployment.

```
islshow
trunkshow
```

Fabric merge conditions

Before connecting the Gen 8 switch to the existing fabric, verify that all fabric merge requirements are satisfied. Fabric merge issues can prevent switches from joining the fabric and may result in segmentation or operational disruptions. Ensuring that fabric parameters, zoning information, and domain IDs are configured correctly helps enable a successful fabric merge and seamless SAN operation.

The following parameters must match:

- Fabric name
- Zoning database (checksum or identical config)
- Principal switch election
- Domain ID uniqueness

Run the following command to check fabric status:

```
fabricshow
```

Zoning synchronization

Before connecting the Gen 8 switch to the production fabric, verify that the zoning configurations are consistent across both environments. Zoning synchronization ensures that hosts and storage devices maintain proper access after the fabric merge and helps prevent segmentation or connectivity issues. If zoning differences are detected, use one of the following methods to resolve them.

If a zoning mismatch is detected, use one of the following options:

- Option A: Push zoning from Gen 7

```
cfgsave
cfgenable <cfg>
```

- Option B: Clear Gen 8 zoning

```
cfgclear
cfgdisable
```

Note: The zoning databases on all switches must match exactly before a fabric merge can occur successfully. Partial zoning merges are not supported.

Workload migration strategy

After the Gen 8 switch has been integrated into the fabric and all validation checks have been completed, migrate host and storage workloads to the new platform. A phased migration approach minimizes risk, allows issues to be identified and resolved incrementally, and helps maintain application availability throughout the migration process.

Non-disruptive method

This method uses the existing fabric connectivity to move workloads gradually from Gen 7 to Gen 8 without requiring a complete outage. Hosts and storage connections are migrated one at a time, enabling validation of each move before proceeding with the next workload.

- Keep the ISL active
- Migrate host cables gradually

Complete the following steps:

1. Disable old Gen 7 port:

```
portdisable <port>
```

2. Move cable to Gen 8.
3. Enable Gen 8 port:

```
portenable <port>
```

4. Validate:

```
nsshow  
nodefind <WWN>
```

Validation and monitoring

After workloads have been migrated to the Gen 8 platform, validate fabric operations and monitor switch performance to confirm that the migration was successful. Review port status, traffic statistics, error counters, and fabric events to identify any issues that could affect connectivity, performance, or stability. Continued monitoring helps ensure that hosts and storage devices are operating as expected in the new environment.

Key commands:

```
porterrshow  
portstatsshow  
portperfshow  
fabriclog --show
```

Decommission

After all workloads have been successfully migrated and validated on the Gen 8 platform, decommission the legacy Gen 7 switch. Before removing the switch from service, confirm that all hosts, storage devices, and ISLs have been migrated, and that no active traffic is using the legacy infrastructure. Proper decommissioning helps prevent unintended disruptions and ensures a clean transition to the new environment.

Complete the following tasks:

1. Remove ISLs.

2. Verify no devices are logged in through the Gen 7 switch.
3. Disable switch ports, if required.
4. Back up final configurations and support data.
5. Disable the switch.
6. Remove the switch from the environment following site procedures.

Practical runbook

This section provides a consolidated, step-by-step workflow for planning, executing, validating, and completing the migration from a Gen 7 to a Gen 8 environment. Use this runbook as a high-level checklist to help ensure that all prerequisite tasks, migration activities, and post-migration validation steps are completed in the proper sequence.

Complete the following runbook steps:

1. Back up configurations and collect baseline information.
2. Verify Fabric OS compatibility and upgrade requirements.
3. Deploy and configure the Gen 8 switch.
4. Validate base chassis settings and connectivity.
5. Align zoning databases and fabric parameters.
6. Establish and verify ISLs.
7. Validate successful fabric merge.
8. Migrate hosts and storage incrementally.
9. Monitor fabric health and performance.
10. Validate application and storage connectivity.
11. Decommission the legacy Gen 7 switch.

Gen 6 to Gen 8 migration: Key differences

This section highlights the key considerations when migrating from Gen 6 to Gen 8 platforms. Unlike Gen 7 platforms, Gen 6 switches do not support Fabric OS 10.x and require additional planning to ensure a smooth migration. Understanding these differences helps minimize migration risks and avoid Fabric OS compatibility and interoperability issues.

Major differences compared to Gen 7

Before planning a Gen 6 to Gen 8 migration, it is important to understand the architectural, hardware, and software differences between the two generations. These differences can affect migration planning, Fabric OS requirements, interoperability considerations, performance expectations, and feature availability. The following comparison highlights the key distinctions that should be evaluated before beginning the migration process.

Feature	Gen 6	Gen 8
Max Speed	32G	128G
ASICs	Condor 4	Condor 6
FOS baseline	8.x	10.x required

Fabric OS 10.x is not supported on Gen 6 platforms. However, Gen 6 platforms running a supported Fabric OS 9.x release can interoperate with Gen 7 and Gen 8 platforms, allowing mixed-generation switches to coexist within the same SAN fabric, subject to the interoperability requirements of the applicable Fabric OS release.

Fabric interoperability between Gen 8 platforms and earlier generations depends on the specific hardware platforms and Fabric OS versions in use. Before planning a migration, consult the applicable Fabric OS Release Notes, Fabric OS Software Upgrade Guide, and Hardware/Firmware Support Matrix to verify supported interoperability, coexistence requirements, and upgrade paths.

Cisco to Brocade migration

This section describes the considerations and approaches for migrating from a Cisco MDS SAN environment to a Brocade SAN fabric. Because Cisco MDS and Brocade switches cannot participate in the same native Fibre Channel fabric, migration requires either a physical migration approach or another supported interoperability method as documented by Broadcom and Cisco. Careful planning is required to migrate zoning, connectivity, and workloads while minimizing disruption.

Note: Cisco MDS and Brocade switches cannot participate in the same native Fibre Channel fabric. Migration requires either a physical migration (recommended) or a supported interoperability method as documented by Broadcom and Cisco. See the applicable interoperability documentation for supported connectivity options and restrictions.

Migration models

The appropriate migration approach depends on factors such as SAN size, acceptable outage windows, operational risk, business continuity requirements, and whether the target SAN architecture differs from the source environment.

Forklift migration (clean cutover)

A forklift migration replaces the existing Cisco SAN with a Brocade SAN during a planned maintenance window. This approach is often suitable when significant architectural changes are planned or when the existing environment is being retired completely. A forklift migration includes the following advantages and considerations:

- Advantages:
 - Simpler migration process
 - Easier implementation of new SAN designs
 - No requirement to operate parallel environments
- Considerations:
 - Typically requires a larger outage window
 - Higher risk due to the scope of change

Phased migration

A phased migration moves hosts and storage incrementally into the new Brocade SAN environment. This approach allows validation after each migration step and can reduce operational risk.

Phased migration includes the following advantages:

- Reduced migration risk
- Smaller maintenance windows
- Easier rollback of individual migration activities

Phased migration includes the following factors:

- More complex planning and coordination
- Requires temporary coexistence of old and new environments

Note: If the migration includes a significant SAN redesign — such as transitioning between full-mesh and core-edge topologies — the most appropriate migration model should be evaluated based on the target architecture and operational requirements.

Pre-migration data collection (Cisco)

Before beginning the migration, collect and document the current Cisco MDS fabric configuration. This information provides a baseline for migration planning, zoning conversion, validation, and troubleshooting. Capturing the existing environment helps ensure that host-to-storage connectivity, zoning policies, and device configurations can be accurately recreated in the Brocade fabric.

Run the following commands on the Cisco MDS switch:

```
show flogi database
show zoneset active
show zone
show device-alias database
```

Map the following items:

- WWNs → aliases → zones

Zoning conversion

Before migrating workloads to the Brocade fabric, convert the existing Cisco zoning configuration to the Brocade zoning model. Although both vendors use Fibre Channel zoning, their configuration objects and management methods differ. Understanding how Cisco zoning elements map to Brocade zoning elements helps ensure that host-to-storage access is preserved during the migration and reduces the risk of configuration errors.

Cisco-to-Brocade zoning object mapping

Cisco and Brocade use different terminology and configuration objects for Fibre Channel zoning. Before migrating zoning configurations, understand how Cisco zoning components map to their Brocade equivalents. This mapping helps ensure that host-to-storage access is maintained and that the zoning configuration is recreated accurately in the target Brocade fabric.

Cisco	Brocade
zoneset	cfg
zone	zone
device-alias	alias

Manual zoning conversion

After collecting the Cisco zoning configuration, recreate the zoning objects in the Brocade fabric. Because Cisco and Brocade use different zoning constructs and management methods, the conversion process typically requires recreating aliases, zones, and configurations manually. Review all WWNs, aliases, zone memberships, and active

zonesets before enabling the migrated configuration. Validate the converted configuration before migrating production workloads.

Complete the following migration process steps:

1. Create Brocade aliases that correspond to Cisco device aliases.
2. Recreate zones using the appropriate host and storage members.
3. Create a Brocade configuration (`cfg`) containing the migrated zones.
4. Save and enable the configuration.
5. Verify zoning and device visibility before migrating workloads.

Consider the following example commands and verification steps:

- Cisco:

```
zone name ZONE1 vsan 10
  member pwwn 10:00:00:00:00:01
  member pwwn 50:00:00:00:00:01
```

- Brocade:

```
allicreate "host1","10:00:00:00:00:01"
allicreate "storage1","50:00:00:00:00:01"
zonecreate "ZONE1","host1;storage1"
cfgcreate "PROD_CFG","ZONE1"
cfgenable "PROD_CFG"
cfgsave
```

- Verification:

```
switch:admin> alishow
alias: host1
  10:00:00:00:00:01
alias: storage1
  50:00:00:00:00:01

switch:admin> zoneshow
zone: ZONE1
  host1
  storage1

switch:admin> cfgshow
Defined configuration:
  cfg: PROD_CFG
  ZONE1
Effective configuration:
  cfg: PROD_CFG
  ZONE1
```

Host migration process

After the Brocade fabric has been configured and validated, migrate hosts from the Cisco fabric in a controlled and incremental manner. Migrating one host at a time allows connectivity and application access to be verified before proceeding with additional hosts. This approach minimizes risk, simplifies troubleshooting, and helps maintain service availability throughout the migration.

Complete the following steps for each host:

- Verify multipathing is functioning correctly.
- Disable the Cisco switch port.
- Move the cable to the Brocade switch.
- Enable the Brocade switch port.
- Verify the host successfully logs in to the fabric.
- Confirm storage visibility and application connectivity.
- Use the `nsshow` command as in the following example to verify that the host has successfully logged in to the fabric after being connected to the Brocade switch. Confirm that the host WWN appears in the Name Server database and that the attached storage devices are visible as expected.

```
switch:admin> nsshow
```

Type	Pid	COS	PortName
N	010200	3;	10:00:00:00:00:00:00:01

The presence of the host WWN in the Name Server database indicates that the host has successfully logged in to the fabric.

Using Brocade SAN Health for migration planning

Successful SAN migrations require careful planning, validation, and documentation. It is important to understand the current SAN environment and establish a baseline before making any changes. Brocade SAN Health provides a point-in-time snapshot of the SAN environment that can be used for migration planning, validation, troubleshooting, and post-migration comparison.

Brocade SAN Health is a free assessment tool that gathers switch, fabric, zoning, and configuration information from supported Brocade SAN environments. The tool analyzes the gathered information and generates reports that help administrators assess SAN health, inventory assets, validate best practices, and identify potential issues before initiating a migration. SAN Health performs the following tasks:

- Taking inventory of devices, switches, firmware versions, and SAN fabrics
- Capturing and displaying historical performance data
- Comparing zoning and switch configurations against best practices
- Assessing performance statistics and error conditions
- Alerting administrators to relevant technical support bulletins (TSBs)
- Producing detailed reports (Microsoft Excel) and diagrams (Microsoft Visio)

Because it provides a point-in-time snapshot of the SAN environment, SAN Health is invaluable for change tracking, migration planning, and post-migration validation.

For more information about Brocade SAN Health, including download and usage instructions, [see the Brocade SAN Health Diagnostics Capture page](#).

SAN migration factors

Today's SAN administrators are faced with the need for more storage capacity and speed. They require high performance and redundant SAN networks that can meet their demands and scale for growth.

To accommodate these new requirements, SAN administrators often must migrate from and upgrade their storage network. Migrating from older SANs to newer SAN technology requires a plan that includes design, configuration, and implementation processes along with a post-migration analysis.

This section describes those migration factors that must be considered for the migration plan. The use of SAN Health allows you to gather information to help make the decisions for a smooth and seamless transition.

In addition to knowing about the infrastructure, the SAN design must be reviewed to help incorporate future growth, backup and recovery or business continuity projects, and server or storage upgrades.

Consider the following example factors:

- Fabric topology and network inventory
- Application failover considerations
- Storage failover considerations
- Review of topology changes
- Identification of critical servers, storage arrays, and applications
- Developing a zone configuration export and import strategy
- Ensuring that correct licensing sets are included
- Planning for software interoperability
- Preparing and updating SAN or storage diagram to meet new requirements

Optics

Optics play a critical role in SAN connectivity and performance. Before deploying or migrating to new switches, verify that all transceivers are supported on the target platform and are compatible with the required link speeds, distances, and cable types. Validating optic compatibility in advance helps prevent link initialization failures, performance issues, and interoperability problems during the migration.

Complete the following optics validation tasks:

1. Verify transceiver compatibility with the target switch platform.
2. Confirm supported link speeds and distances.
3. Validate fiber cable type and requirements.
4. Check for supported and qualified optics in the Brocade Transceiver Support Matrix.
5. Verify optic health and link status after installation.

See [Brocade® Transceiver Support Matrix](#)

Storage area network security

Security has always been a key advantage of Fibre Channel (FC) storage area networks (SANs). FC SANs are often deployed as isolated or air-gapped networks, reducing exposure to external threats. In addition, Fibre Channel SANs are typically deployed as dedicated and segmented networks, which can reduce exposure to unauthorized access. Zoning and LUN masking further restrict a host's access to designated storage resources, helping to limit the potential impact of a compromised host.

However, FC SANs are not immune to security threats. Mission-critical applications depend on SAN infrastructures, making it essential to protect both the network and the data it carries.

Securing an FC SAN requires a layered approach that addresses multiple security challenges across the infrastructure. Implementing security controls at each layer helps maintain the confidentiality, integrity, and availability of SAN resources over time.

In this section, we examine each security layer, beginning with management security and progressing inward to the layers that protect access to and security of the data itself.

Fibre Channel SAN security layers

The following figure shows the layers of FC SAN security.




Figure: FC SAN security layers overview

Frontier AI and quantum decryption: Post-quantum cryptography

Legacy networks cannot contain adversarial AI tools that exploit and move laterally at machine speed, nor can they protect against future quantum decryption. To insulate your mission-critical data, your storage must reside on an unassailable foundation. Fibre Channel is widely regarded as one of the most secure storage networking technologies in the data center. Next-generation IBM Gen 8 Fibre Channel delivers a highly resilient data fabric. The network architecture combines AI-threat secure defense with quantum-safe cryptography to defend your most critical data assets today and tomorrow.

The compounding threat: Frontier AI and quantum decryption

Autonomous AI tools can launch nonstop, automated attacks that learn and exploit undetected vulnerabilities. Left unchecked, these intelligent agents can silently exfiltrate vast troves of enterprise data for future decryption — "Harvest now, decrypt later." By pairing an AI-threat secure network architecture with quantum-safe cryptography, IBM Fibre Channel ensures a complete defense lifecycle:

- **Neutralize AI reasoning and scanning:** Restrict the physical network fabric to deny autonomous tools the visibility they need to map the environment, render custom exploit code useless, and halt machine-speed breaches.
- **Defeat quantum horizons:** Ensure that even if a breach is attempted, your data remains completely hidden, insulated, and designed to resist current and emerging cryptographic threats into the post-quantum future.

This defense is built upon the most secure network in the data center for a reason. Traditional data center IP/Ethernet networks are inherently routable, highly interconnected, and plagued by a sprawling attack surface exposed to lateral reach contagion. Fibre Channel is a completely non-routable, dedicated storage layer. This architectural isolation means that IBM fabrics remain structurally invisible to network-layer threat actors and autonomous scanning tools. Taking security one step further, IBM proactively utilizes frontier AI tools to continuously scan, detect, and remediate vulnerabilities before they can impact production infrastructure. Layering AI-threat secure and quantum-safe defenses onto this isolated, AI-fortified architecture creates a data vault that significantly reduces attack exposure by today's adversarial AI and tomorrow's quantum technology. The result is an enterprise storage infrastructure that achieves the highest compliance and data defense standards possible with zero compromises to performance.

IBM Gen 8 Fibre Channel guardrails against frontier AI and quantum decryption

Physical isolation

Uses a non-routable, point-to-point architecture to naturally segment storage traffic, eliminating the lateral attack paths common to traditional IP data center networks.

Hardened attack surface

Implements strict principle-of-least-privilege access controls within the Fabric OS (FOS) software by eliminating root-level accounts, neutralizing hijacking and malware exploits.

Hardware and software validation

Embeds silicon root of trust and utilizes unique Trusted FOS Certificates to continuously verify authenticity, blocking counterfeit components or malicious firmware.

Proactive vulnerability mitigation

IBM's PSIRT maintains a robust vulnerability intake process utilizing multiple scanning tools and penetration testing methods. Leverages advanced frontier AI models and automated penetration testing to actively isolate and patch code flaws before they ever hit production.

IBM Gen 8 and FOS 10.0.x security enhancements

While IBM Gen 7 focused heavily on introducing foundational cyber resilience and self-healing automation, IBM Gen 8 shifts its primary security focus toward future-proofing the SAN against emerging cryptographic threats — specifically targeting the era of AI-driven attacks and quantum computing.

The changes to security include several architectural and operational differences.

Quantum-safe cryptography versus standard encryption

IBM Gen 7 relies on traditional, modern cryptographic standards such as AES-256 for data-in-flight encryption and standard RSA/ECC for secure boot and code signing. While highly secure against current threats, these algorithms are theoretically vulnerable to future quantum computers.

IBM Gen 8 includes Post-Quantum Cryptography (PQC). It integrates quantum-resistant 256-bit encryption and advanced cryptographic protocols natively into the hardware and Fabric OS (FOS). This prevents *harvest now, decrypt later* attacks, where bad actors steal encrypted data today with the intent of decrypting it once quantum computing matures.

Zero trust and access architecture

IBM Gen 7 uses standard secure management access controls, role-based access control (RBAC), and validated Secure Boot via a hardware Root of Trust to prevent unauthorized firmware modifications.

IBM Gen 8 formalizes a strict Principle-of-Least-Privilege (PoLP) architecture directly into the fabric design. Gen 8 drastically minimizes the default attack surface by strictly isolating control planes, segmenting management domains, and ensuring that connected devices or administrators only have the minimum necessary permissions required to operate.

AI-powered threat and anomaly detection

IBM Gen 7 introduced self-learning and automation primarily geared toward SAN health and reliability. It tracks telemetry data, utilizes Fabric Performance Impact Notifications (FPINs), and flags physical layer anomalies or congestion that could be exploited for Denial of Service (DoS) situations.

IBM Gen 8 enhances this with embedded SAN AI technology and Fabric Intelligence. The Gen 8 AI doesn't just watch for structural failures; it models baseline traffic behaviors contextually down to the Virtual Machine (VM) level. This

allows it to intelligently detect subtle pattern deviations that could indicate data exfiltration, ransomware behavior, or hijacking attempts, initiating self-healing protocols autonomously to quarantine compromised zones.

Summary comparison

Security Feature	IBM Gen 7	IBM Gen 8
Cryptographic Status	Standard Enterprise Encryption (AES-256 / RSA)	Quantum-Safe / Post-Quantum Cryptography
Access Control Philosophy	Standard RBAC & Secure Boot	Principle of Least Privilege (PoLP) Minimizing Attack Surfaces
Threat Intelligence	Telemetry-based automated health & congestion alerts	Embedded SAN AI Engine for behavioral anomaly & malware tracking
Hardware Foundation	Hardware Root of Trust & Secure FOS	Enhanced Hardware Root of Trust with PQC-validated firmware chains

Gen 7 secures your SAN against the threats of today's digital landscape by ensuring hardware integrity and automated fabric resilience. Gen 8 builds directly on top of that foundation, upgrading the cryptographic core to make your storage infrastructure resilient against the weaponization of AI and quantum computing.

FOS 10.0.x default secure switch configuration and hardening

Updated security baseline rules including mandatory SHA-512 password hashing, dual lockout enforcement (`adminlockout` & `adminlockoutconsoleaccess`), SSH key regeneration (RSA 2048 / ECDSA P521), 900s SSH rekeying, and blocking unencrypted syslog and file transfer operations.

The following section describes the secure switch configuration and hardening requirements that must be resolved before you can upgrade from FOS 9.2.2x to FOS 10.0.0x.

When upgrading from Brocade Fabric OS (FOS) 9.2.2 to FOS 10.0.0a, Broadcom enforces strict pre-upgrade checks. If any unsupported, legacy, or non-secure configurations are active, the firmware upgrade process will be explicitly blocked.

Brocade Fabric OS (FOS) 10.0.x pre-upgrade requirements checklist

This section details mandatory authentication, cryptographic, protocol hardening, and configuration prerequisites required prior to upgrading to Brocade Fabric OS (FOS) 10.0.x. Non-compliant configurations blocks the upgrade process.

Authentication and certificates

Certificates include the following prerequisites:

- TruFOS Certificate Requirement

A valid, unexpired TruFOS Certificate must be installed on the switch/director prior to initiating the upgrade.

- SMTPS CA Certificate

If Secure SMTP (SMTPS) is enabled, you must import the required SMTPS CA certificate before upgrading. The upgrade will block if SMTPS is enabled without this certificate.

Cryptographic and hashing algorithms

Review the supported algorithms and implement changes as necessary:

- Reconfigure auth hashing (SHA1/MD5 to SHA256)

SHA1 and MD5 hashing protocols are unsupported in FOS 10.0.x. You must reconfigure authentication hashing to SHA256:

```
authutil --set -h sha256
```

- Reset custom SSH host key algorithms & Unsupported Ciphers

If custom SSH host key algorithm configurations or legacy/unsupported cryptographic ciphers are active, the upgrade fails. Revert to supported cipher groups using:

```
seccryptocfg --apply -group <supported_group>
```

Note: For more information on `seccryptocfg` and supported groups, see the *Brocade Fabric OS 10.0.0 Command Reference Guide* and search for `seccryptocfg`.

Protocol and service hardening

Review and implement required protocol changes:

- Disable non-secure switch mode

Switches operating in non-secure mode must have non-secure mode disabled prior to starting the upgrade.

- Remove non-secure syslog servers

FOS 10.0.x blocks upgrades if non-secure Syslog servers are configured. Remove any non-secure Syslog destinations:

```
syslogadmin --remove --ip <syslog_ip>
```

- Update SMTPS TLS protocol version

Configurations using legacy TLSv1 for secure SMTP are not permitted. Reconfigure SMTPS to use a supported version (such as TLS 1.2).

Note: For more information on SMTPS TLS Protocol Versions, see the *Brocade Fabric OS 10.0.0 Command Reference Guide* and search for `seccryptocfg`, `seccertmgmt`, and `smtps`.

Other blocking pre-upgrade constraints

Broadcom checks the following security and group configurations prior to upgrading:

- Custom SFP and Circuit groups

Delete custom SFP or Circuit logical groups using the following command:

```
logicalGroup --delete
```

- Paused SFP Objects

Restart any paused SFP monitoring objects using the following command:

```
mapsConfig --config restart -type sfp -members
```

- FCIP Tunnel Features

Disable FICON features on FCIP tunnels, including:

- `ficon-xrc`
- `ficon-tin-tir`
- `ficon-tape-pipelining`
- `ficon-tera-data`
- `ficon-device-acking`

Management interfaces

This section describes the management interfaces provided by Fabric OS (FOS) and the security mechanisms that can be used to protect them. Because management interfaces provide administrative access to the switch, they should be secured as part of an overall defense-in-depth strategy.

Access to the management network should be restricted to authorized administrative systems and networks. Common practices include isolating management traffic by using dedicated VLANs or management networks and limiting access to only the management protocols and services required in the environment. This approach reduces the attack surface and helps prevent unauthorized access to the switch management interfaces.

IP Filter policies in FOS can be used to permit or deny traffic to the management interfaces based on defined access rules. An IP Filter policy acts as a packet-filtering firewall for the management interfaces, permitting or denying traffic according to configured rules. Secure switch management is performed through SSH and HTTPS. Administrators should configure IP Filter policies to allow access only from authorized management hosts and networks and to restrict access to only the services required for administration. In dual-stack environments, IPv4 and IPv6 access should be restricted consistently according to the organization's security requirements.

In current Gen 8 systems, unsecured management protocols such as Telnet and HTTP are disabled by default and are not intended for switch management. Although Telnet and HTTP entries might still appear in the default IP Filter policy, they are configured with deny rules. The default IP Filter policy permits secure management access through SSH and HTTPS. Administrators should further restrict access by allowing connections only from authorized management networks and by blocking any additional services that are not required in their environment.

The IP Filter policy should be used in environments where strict control of fabric access is required. As with other ACL policies, it is important to regularly review and maintain the IP Filter policy to ensure that it remains aligned with security and operational requirements.

Displaying the IP filter

To display the current IP Filter in use, run the `ipfilter --show` command, as shown in the following example:

```
IBM_8635_F88:FID128:admin> ipfilter --show

Name: default_ipv4, Type: ipv4, State: defined

Rule  Source IP  Protocol  Dest Port  Action
----  -
1     any          tcp       22         permit
2     any          tcp       23         deny
3     any          tcp       80         deny
4     any          tcp       443        permit
5     any          udp       161        permit
6     any          udp       123        permit
7     any          tcp       600 - 1023 permit
8     any          udp       600 - 1023 permit

Name: default_ipv6, Type: ipv6, State: defined

Rule  Source IP  Protocol  Dest Port  Action
----  -
1     any          tcp       22         permit
2     any          tcp       23         deny
3     any          tcp       80         deny
4     any          tcp       443        permit
5     any          udp       161        permit
6     any          udp       123        permit
7     any          tcp       600 - 1023 permit
8     any          udp       600 - 1023 permit

Name: unified, Type: any, State: active

Rule  Source IP  Protocol  Dest Port  Action
----  -
1     any          tcp       22         permit
2     any          tcp       23         deny
3     any          tcp       80         deny
4     any          tcp       443        permit
5     any          udp       161        permit
6     any          udp       123        permit
7     any          tcp       600 - 1023 permit
8     any          udp       600 - 1023 permit
```

Note: By default, Fabric OS does not permit cleartext Telnet (TCP port 23) or HTTP (TCP port 80).

In Fabric OS (FOS) v10.0.0, Broadcom introduced the Unified IP Filter policy as the new default architecture, replacing the legacy dual-policy model used in FOS v9.x and earlier.

The primary differences between the Unified IP Filter policy and the legacy IPv4 and IPv6 IP Filter policies are related to enforcement, activation behavior, and firmware compatibility:

Feature	Legacy IPv4 & IPv6 Policies (FOS 9.x and earlier)	Unified IP Filter Policy (FOS 10.0.0+)
Active Policy Structure	Requires two separate active policies running simultaneously (one for IPv4 and one for IPv6).	Consolidates rules into a single unified policy active by default.
Mutual Exclusivity	Dual IPv4 and IPv6 active policies coexist without disabling one another.	Mutually exclusive with legacy policies. Activating a unified policy automatically disables legacy v4/v6 policies, and vice versa.
Traffic Handling During Mode Transition	Switching between IPv4 policies does not affect active IPv6 traffic rules (and vice versa).	If you switch from Unified to a non-unified (legacy) IPv4 policy, IPv6 traffic is completely blocked (INPUT/OUTPUT) until an IPv6 policy is explicitly activated.
FOS Downgrade Restrictions	Standard upgrade/downgrade rules apply.	Downgrades are blocked if a unified policy is currently active or defined in configuration.

Role-based access controls

One way to provide limited accessibility to the fabric is through user roles. FOS includes predefined user roles, each of which can access a subset of the CLI commands. These roles are known as role-based access controls (RBACs), which are associated with the user's login credentials.

When you log in to a switch, your user account is associated with a predefined role or a user-defined role. The role that your account is associated with determines the level of access you have on that switch and in the fabric.

Default accounts

FOS offers predefined accounts:

- Admin
- Factory
- User

Although the factory account is reserved for development and manufacturing purposes, the passwords for all default accounts must be changed and secured during the initial installation and configuration of each switch.

Note: The root account was removed in FOS 9.1 and later releases and cannot be activated.

From Brocade Fabric OS (FOS) 9.0 to FOS 10.0, Broadcom significantly tightened security and account access controls, primary among which was the complete removal of raw root access.

The following table describes the major changes to default accounts and account policies across these release trains:

Updates to FOS 9.0 include several key account changes.

Removal of the root account (FOS 9.1+)

Before FOS 9.0.x, the traditional Linux root account was accessible for advanced troubleshooting, low-level system access, and vendor debugging.

For version FOS 9.1.x through FOS 10.0.x, the root user is disabled and removed from functional access to mitigate malware, hijacking, and unauthorized kernel access.

Although the root account can appear in account listings, it is disabled and cannot be used for login.

Introduction of the maintenance account

To replace root for vendor service procedures, FOS introduced the maintenance account.

The default password for maintenance is set to `password`.

Unlike the old root account, logging into maintenance or invoking debug/kernel mode requires token-based authentication supervised by the service provider (Broadcom/Dell/HPE). You generate a challenge token via `serviceshell --generate` and receive a temporary response token from vendor support.

Mandatory default password changes forced at first login

In earlier FOS generations, default passwords (such as admin/password, user/password) could persist indefinitely if not manually updated.

For FOS 9.0+ & FOS 10.0 and later: First-time login to default accounts (admin, user, maintenance, etc.) forces an immediate password update.

Note: The new password cannot match the factory-default password.

Disabling and deprecating insecure default accounts

Default accounts: Unused or legacy service accounts (for example, default SNMP users and legacy factory diagnostic accounts) are disabled by default.

Factory default lockdown: Switches shipped with FOS 9.2 or later adhere to Default Secure profiles out-of-the-box, enforcing local security policies right from initial setup.

Default account comparison

Default Account / Feature	FOS 9.0	FOS 9.1 / 9.2	FOS 10.0
root Account	Present / Accessible	Disabled (No privileges/ access)	Disabled / Removed
maintenance Account	Introduced as default	Replaces root (Requires token auth)	Active (Requires vendor token auth)
Mandatory Password Change	Enforced at initial setup	Enforced at initial setup	Enforced at initial setup
Default Crypto / Security	Standard settings	Introduced Default Secure templates	Default Secure enforced out-of-box
Default SNMP Users	Can be disabled	Disabled by default	Disabled by default

Note: Recovering lost passwords might require significant effort and fabric downtime and assistance from Support resources.

User accounts

In addition to the default permissions that are assigned to the roles of admin, factory, and user, FOS supports up to 252 other user accounts on the chassis. These accounts expand the ability to track account access and audit administrative activities.

Each user account is associated with the following attributes:

- Virtual fabric list: Specifies the Virtual Fabric that a user account is allowed to log in to.
- Home virtual fabric: Specifies the Virtual Fabric that the user is logged in to, if available. The home Virtual Fabric must be a member of the user's Virtual Fabric list. If the fabric ID is not available, the next-lower valid fabric ID is used.
- LF permission list: Determines functional access levels within the bounds of the user's Virtual Fabrics.
- Chassis role: Similar to switch-level roles, but applies to a different subset of commands.

Fabric OS supports remote users through one of the following external Authentication, Authorization, and Accounting (AAA) services:

- Lightweight Directory Access Protocol (LDAP)
- RSA SecurID (RSA)
- Federated Authentication (FA)

Note: When the AAA server is configured as RSA, the communication between the server and the switch is through the REST API over HTTPS.

When configured for remote authentication and authorization, a switch becomes a LDAP, RSA client, or FA resource server. In any of these configurations, authentication records are stored in the authentication host server database. Account name, assigned permissions, and time-accounting records for each user are stored as per the following table.

Multi-factor authentication

Multi-factor authentication (MFA) is a multi-step account login process that requires users to enter more information than just a password. For example, in addition to a password, users might be asked to enter a code that is sent to their email, answer a secret question, or scan a fingerprint.

FOS supports MFA with the following services:

1. RSA with SecurID via REST API with HTTPS
2. FA

RSA SecurID provides the capability to verify the identity of users using MFA through the REST API over HTTPS. The users are provided with a time-based rotating numeric passcode. This passcode with the user PIN (PIN + OTP) is used for the user authentication. Fabric OS uses this mechanism to support the MFA user authentication for access to FOS.

Note: A one-time password (OTP), also known as a one-time PIN, one-time authorization code (OTAC) or dynamic password, is a password that is valid for only one login session or transaction, on a computer system or other digital device. In case of FA, the MFA can be enabled on the identity provider.

Federated authentication

FOS supports federated authentication (FA) performed by an Identity Provider (IDP). IDP manages the applications, user identities, credentials, and MFA configurations. The user or application attempts to log in to the switch using a valid access token that is issued by the IDP.

Note: The FOS switch can either fetch IDP public keys or the keys can be imported by the user to perform token validation.

Remote user configuration on the switch

At least one LDAP, RSA, or FA server must be configured before you can enable the AAA service. You can configure the AAA service even if it is disabled on the switch. You must be logged in as admin or securityAdmin to configure the remote authentication service.

LDAP, RSA, or FA configuration is chassis-based configuration data. On platforms that contain multiple switch instances, the configuration applies to all instances. The configuration is persistent across reboots and firmware downloads.

Security protocols

Security protocols provide endpoint authentication and communications privacy by using cryptography. Typically, the user is authenticated to the switch while the switch remains unauthenticated. This authentication means that you can be sure that you know with whom you are communicating.

The security protocols supported by FOS are listed in the following table.

Protocol	Description
HTTPS	HTTPS is a Uniform Resource Identifier scheme that is used to indicate a secure HTTP connection. Web Tools supports the use of Hypertext Transfer Protocol over SSL (HTTPS). Quantum safe cryptography can be configured for HTTPS.
LDAP	Lightweight Directory Access Protocol (LDAP) with Transport Layer Security (TLS) uses a certificate authority (CA). With the import of signed certificates, you can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL)/ TLS technology with LDAP.
RSA	RSA SecurID REST API-based authentication over HTTPS. The CA root certificate for the RSA server can be imported using the seccertmgmt command.
FA	FOS communicates to the IDP through the REST API over HTTPS protocol. If the signing authority does not exist in the FOS CA bundle, then the CA certificate that signed the IDP server certificate is required for import using the seccertmgmt command.
SCP	Secure Copy (SCP) is a means of securely transferring computer files between a local and a remote host or between two remote hosts, using the Secure Shell (SSH) protocol. Configuration upload and download support the use of SCP.
Secure Syslog	Secure syslog requires importing syslog CA certificates using the seccertmgmt command.
SFTP	Secure File Transfer Protocol (SFTP) is a network protocol for securely transferring files on a network. Configuration upload and download support the use of SFTP.
SNMPV3	Simple Network Management Protocol version 3 (SNMPv3) is used by network management systems to monitor and manage network-attached devices. SNMPv3 provides authentication, message integrity, and optional encryption to help secure management communications.
SMTPS	Secure SMTP (SMTPS) provides secure email delivery for MAPS notifications by using TLS encryption and authentication between the switch and the configured mail relay server.
SSH	Secure Shell (SSH) is a network protocol that allows data to be exchanged over a secure channel between two network devices. Encryption provides the confidentiality and integrity of data. SSH uses the public-key cryptography to authenticate the remote network device and to allow the remote network device to authenticate the user, if necessary. Quantum safe cryptography can be configured for SSH.
SSL	Fabric OS uses Secure Socket Layer (SSL) to support LDAP, RSA, FA, SYSLOG, SMTPS, KAFKA, and HTTPS. A certificate must be generated and installed on each switch to enable the HTTPS server communication. This protocol supports 128-bit encryption by default. TLS v1.2 is recommended for a stronger encryption. HTTPS, SMTPS, and FA supports TLS v1.3.

Access control lists

Between Brocade Fabric OS (FOS) 9.0 and FOS 10.0, changes to Access Control Lists (ACLs), which encompass policies such as FCS (Fabric Configuration Server), DCC (Device Connection Control), SCC (Switch Connection Control), and IP Security Access Lists, fall into several key architectural, security, and hardware compatibility categories.

In FICON environments, ACLs help protect fabric integrity by controlling switch membership, device connectivity, and management access. Although FICON logical switches (FLSs) provide the logical separation required for mainframe workloads, ACL policies are configured and managed independently and should be implemented as part of the overall security architecture. The specific ACL requirements depend on the fabric topology, operational requirements, and security policies in effect.

The following list describes the primary changes:

- Improved policy management and consistency across physical and logical fabrics.
- Continued support for Fabric Configuration Server (FCS), Device Connection Control (DCC), and Switch Connection Control (SCC) policies to control fabric membership and device access.
- Improved support for modern hardware platforms, including Gen 8 directors and switches.
- Expanded support for IPv6 and unified IP filtering capabilities for management network access control.
- Better integration with virtual fabrics and FICON logical switch environments.

Hardware architecture and platform constraints

Gen 8 hardware integration: FOS 10.0 introduces support for Brocade Gen 8 platforms alongside Gen 7. ACL security policy processing natively extends across Gen 7 and Gen 8 fixed-switch, director, and extension platforms without requiring different ACL syntax or structural changes.

Gen 6 platform retirement: Gen 6 platforms reached their final supported software levels within the FOS 9.x family. Consequently, ACL policy deployment and management in FOS 10.0 are limited to supported Gen 7 and Gen 8 switches and directors.

In-band management security controls (FCS and DCC)

FOS utilizes specific built-in ACL policy types (such as Fabric Configuration Server / FCS and Device Connection Control / DCC) to govern in-band management and fabric access:

Extension Platform (FCIP) ACL support: While fixed-port switches and directors supported full ACL policies in FOS 9.0.0, complete ACL feature parity for Gen 7 Extension Switches such as the Brocade 7810/7850 matured in FOS 9.2.0 and carries forward into FOS 10.0.

Automation-friendly ACL management: In FOS 9.x, Broadcom introduced REST API support for security configurations, which expanded further into FOS 10.0. ACL management (activating, modifying, distributing FCS and DCC policies) shifted away from legacy Java-based interfaces toward structured REST payloads and SANnav orchestration.

FICON considerations: Within FICON logical switch (FLS) environments, ACL policies are configured independently of the logical switch definition. Organizations commonly use FCS, DCC, and SCC policies to control switch membership, restrict device connectivity, and maintain fabric integrity across single- and dual-director topologies.

WebTools and management interface ACLs

Removal of legacy WebTools: FOS 9.0 replaced the old Java-based WebTools interface with a Java-free HTML5 version. By FOS 10.0, legacy Java utilities and deprecated embedded tools like the legacy embedded SSH Client are removed entirely. ACL policy creation (DCC, FCS, SCC) is managed through the modern WebTools interface or CLI/REST API.

Tightened default security base

Although standard IP management ACL functionality (`secpolicymanager` / `ipfilter`) operates similarly between 9.x and 10.x, the base system surrounding network ACLs has been significantly tightened:

- Default secure policy profiles: FOS 10.0 defaults to stricter security profiles out of the box, requiring explicit ACL entries or TLS/crypto configurations for any out-of-band management traffic trying to traverse management interfaces.
- Deprecation of weak protocol rules: IP filter ACL rules permitting unencrypted legacy management protocols (HTTP, Telnet) are disabled by default in FOS 10.0 security baselines in favor of HTTPS and SSH.

ACL policies

Access control lists (ACLs) are used to provide network security through policy sets. FOS provides several ACL policies, including a Switch Connection Control (SCC) policy, a Device Connection Control (DCC) policy, a Fabric Configuration Server (FCS) policy, and an IP Filter.

SCC policy

Use the SCC policy to restrict FC switches that can join the fabric. Only switches that are specified in the policy can join the fabric. All other switches fail authentication if they attempt to connect to the fabric, which results in the respective E_Ports being segmented because of the security violation.

Use the SCC policy in environments where you need strict control of fabric switch members.

The SCC policy can prevent switches from participating in a fabric; therefore, it is important to regularly review and properly maintain the SCC policy.

DCC policy

Use the DCC policy to restrict devices that can attach to a single FC port. The policy specifies the FC port and one or more worldwide names (WWNs) that can connect to the port. The DCC policy comprises all of the DCC policies that are defined for individual FC ports.

Not every FC port must have a DCC policy, and only ports with a DCC policy in the active policy set enforce access controls. A port that is present in the active DCC policy allows only WWNs in its respective DCC policy to connect and join the fabric. All other devices fail authentication when they attempt to connect to the fabric, which results in the respective F_Ports being disabled because of the security violation.

Use the DCC policy in environments where you need strict control of fabric device members.

Because the DCC policy can prevent devices from participating in a fabric, it is important to regularly review and properly maintain the DCC policy.

FCS policy

Use the FCS policy to restrict the source of fabric-wide settings to a single FC switch. The policy contains the WWNs of one or more switches, and the first WWN (online) in the list is the primary FCS. If the FCS policy is active, only the primary FCS can make or propagate fabric-wide parameters. These parameters include zoning, security (ACL) policy databases, and other fabric settings.

Securing end points in SAN infrastructure

The SAN infrastructure consists of HBAs, switches, and storage devices. Each of these components offer security features that must be enabled to maximize security.

Several options are available regarding HBAs. You can use Single Initiator Zones to limit the number of LUNs that might be compromised if the server is compromised.

On the storage device, LUN masking must be implemented to ensure that only those LUNs zoned to the host bus adapter (HBA) are accessible through the associated storage port. DH-CHAP can also be enabled to authenticate devices when they join the fabric.

Additional security controls should be implemented to reduce the potential impact of a compromised device. Single-initiator zoning should be used in conjunction with LUN masking to limit device visibility and restrict access to only the required storage resources.

On FC switches, unused ports should be permanently disabled to reduce the attack surface. Active F_Ports should also be configured so that they cannot transition to E_Ports, preventing unauthorized interswitch links (ISLs) from being established and helping maintain fabric integrity.

Additional security controls are available for Fibre Channel switches. Unused switch ports can be permanently disabled. Administrators can also prevent active F_Ports from becoming E_Ports (ISLs) by configuring the ports to disallow E_Port operation.

By using a DCC policy, administrators can control which devices are allowed to log in to specific switch ports based on their WWNs. Similarly, an SCC policy can be used to control which switches are permitted to join the fabric based on switch WWNs. These controls are widely used in High Integrity Fabric (HIF) environments and are commonly implemented in FICON deployments to help maintain fabric integrity and restrict unauthorized access.

Between Brocade Fabric OS (FOS) 9.0 and FOS 10.0, Broadcom significantly modernized the security model for SAN endpoints (HBAs, servers, and storage target ports) and switch hardware. The focus shifted from optional security controls to zero-trust hardware enforcement, hardened endpoint authentication, and strict cryptographic standards.

Endpoints and HBAs (N_Port and F_Port security)

FC-SP-2 and modernized DH-CHAP authentication

FOS 9.0: Supported FC-SP (Fibre Channel Security Protocol) and DH-CHAP authentication for N_Ports (HBAs and storage arrays connecting to switch F_Ports). However, legacy cryptographic hashes like MD5 were still widely tolerated.

FOS 10.0: Enforces FC-SP-2 standards with SHA-256 / SHA-512 cryptographic hashing for DH-CHAP endpoint authentication. Legacy MD5 hashing is disabled or strongly deprecated, forcing HBAs and storage targets to authenticate using modern key exchange algorithms before logging into the fabric (FLOGI).

FC-NVMe endpoint security and isolation

Enhancements: As flash/NVMe storage adoption grew across the 9.x and 10.x release trains, FOS strengthened endpoint isolation for FC-NVMe workloads.

Impact: Tightened Fabric Zoning Lock mechanisms prevent dynamic FC-NVMe host/target discovery exploitation, ensuring HBAs cannot probe or map unauthorized NVMe namespaces or storage arrays on shared fabric switches.

Storage array and target device controls

Target rate limiting and flow guard security

FOS 9.0: Flow-based monitoring (Fabric Vision / MAPS) monitored oversubscription and congestion from storage arrays.

FOS 10.0: Deep integration between storage endpoint security and MAPS (Monitoring and Alerting Policy Suite) allows automatic quarantine or port-fencing of compromised, noisy, or misbehaving storage targets.

Restricted LUN and WWN masking and zoning integrity

4 MB zoning database: FOS 9.0 increased the maximum zoning database size to 4 MB, enabling support for larger and more complex SAN environments while maintaining centralized management and consistent enforcement of zoning policies across the fabric.

Transactional zoning locks: FOS 10.0 enforces transactional zoning locks that help prevent concurrent or unauthorized zoning modifications through management interfaces, including REST APIs. These controls help maintain zoning integrity and prevent unintended changes that could affect host-to-storage access.

Switch hardware and platform integrity

TruFOS certificates and counterfeit protection

FOS 9.0: Introduced digital license certificates and signature checks.

FOS 10.0 / Gen 8: Brocade TruFOS Certificates are strictly enforced. The switch verifies the cryptographic authenticity of the operating system, optics, and firmware updates. Unauthorized firmware or counterfeit hardware modules are blocked at boot time via hardware-rooted Secure Boot.

ISL encryption and switch-to-switch (E_Port) security

In-flight encryption: AES-256 encryption on Inter-Switch Links (ISLs) is updated to support modern cipher suites. In FOS 10.0, session key renegotiation for encrypted links between switches is hardened to protect against replay attacks.

Overall management and protocol baseline

FOS 9.2 introduced — and FOS 10.0 standardizes — out-of-the-box Default Secure configurations.

Protocol hardening: Legacy, unencrypted management protocols (Telnet, HTTP) and weak SNMP users are disabled by default. All REST API calls, WebTools sessions, and SSH connections require TLS 1.3 and elliptic curve Diffie-Hellman (ECDH) key exchanges.

ISL in-flight encryption and compression

IBM b-type Gen 8 FC platforms support in-flight compression and encryption at a port level for local and long-distance ISLs. In-flight data compression is a useful tool for saving money when bandwidth caps or bandwidth usage charges are in place for transferring data between fabrics. Similarly, in-flight encryption enables a further layer of security with no key management impact when transferring data between local and long-distance data centers besides the initial setup.

Between Brocade Fabric OS (FOS) 9.0 and FOS 10.0, changes to In-Flight Encryption (IFE) and In-Flight Compression (ICOMP) focus on scaling these features for 128Gb/s Gen 8 Fibre Channel architecture, enforcing modern cryptographic standards, and streamlining key exchange protocols.

Gen 8 hardware scaling and port density

Expanded ASIC throughput: In FOS 9.x (Gen 6/Gen 7 platforms), IFE and ICOMP were bound by 32G/64G ASIC buffer and port group limits. FOS 10.0 introduces native support for Gen 8 platforms (such as 128G switches/blades like the G820 and FC128-48).

Increased bandwidth capabilities: Under FOS 10.0, individual port groups support up to 256 Gb/s of combined encryption and compression bandwidth per ASIC port group, allowing up to 32 ports per group to run in-flight security/compression at native rates without needing auxiliary blades or external appliances.

Mixed-generation interoperability: FOS 10.0 natively supports encrypted/compressed Inter-Switch Links (ISLs and IFLs) operating between Gen 7 and Gen 8 switches running 10.x.

Cryptographic hardening and key exchange

Modern authentication protocols only: To establish an encrypted ISL (E_Port / EX_Port), the switches must authenticate prior to enabling encryption. While FOS 9.0 still tolerated legacy DH-CHAP hashing variants, FOS 10.0 strictly enforces modern, high-bit hashing algorithms (DH-CHAP with SHA-256 / SHA-512) and FCAP (FC-SP Certificate Authentication Protocol).

Removal of weak cipher suites: FOS 10.0 deprecates older 128-bit or legacy IPSec/AES key exchange modes on ISL encryption pipelines in favor of AES-256 GCM (Galois/Counter Mode), ensuring protection against modern replay attacks over long-distance links.

Auto-key rotation improvements: Session key renegotiation for in-flight encryption on long-haul ISLs operates with tighter background key rotation cycles in FOS 10.0, avoiding frame drops during session re-keying.

Enabling in-flight ISL encryption increases latency because frames must be processed by the ASIC before transmission. In-flight compression also requires additional processing, although the reduced volume of transmitted data may partially or fully offset the associated latency, depending on the workload.

Implementing DH-CHAP for ISL encryption

The following example shows how to implement DH-CHAP for ISL encryption:

1. Verify whether DH-CHAP secrets exist in the switch database:

```
switch:admin> secauthsecret --show
Chap secret key database does not exist.
```

2. Run `secauthsecret --set` to set DH-CHAP authentication and create secrets on each switch:

```
switch:admin> secauthsecret --set
This command is used to set up secret keys for the DH-CHAP authentication.
The minimum length of a secret key is 8 characters and maximum 40 characters.
Setting up secret keys does not initiate DH-CHAP authentication.
If switch is configured to do DH-CHAP, it is performed whenever a port or a switch is enabled.
Warning: Please use a secure channel for setting secrets.
Using an insecure channel is not safe and may compromise secrets.
Following inputs are specified for each entry.
WWN for which secret is being set up.
Peer secret: The secret of the peer that authenticates to peer.
Local secret: The local secret that authenticates peer.

Press Enter to start setting up shared secrets >
Enter WWN, Domain, or switch name (Leave blank when done): 10:00:8j:94:7f:40:90:20
Enter peer secret:
Reenter peer secret:
Enter local secret:
Reenter local secret:
Warning: Keeping same peer and local secrets can compromise the security of your network.
To change, reenter WWN and secrets.
Enter peer WWN, Domain, or switch name (Leave blank when done):
Are you done? (yes, y, no, n): [no] y
Saving data to keystore... Done.
```

3. Validate that switch secrets are now stored in the switch database:

```
switch:admin> secauthsecret --show
WWN              DIId  Name
10:00:8j:94:7f:40:90:20  60   switchAname_DH-CHAP_example
```

4. Enable DH-CHAP as the authentication method:

```
switch:admin> authutil --set -a dhchap
Authentication is set to dhchap.
```

5. Set the DH Group to 4:

```
switch:admin> authutil --set -g "4"
switch:admin> authutil --show
AUTH TYPE      HASH TYPE      GROUP TYPE
fcap,dhchap    sha256,sha1,md5  4
```

6. Activate the configuration:

```
switch:admin> authutil --policy -sw active
Warning: Activating the authentication policy requires either DH-CHAP secrets or PKI
certificates depending on the protocol selected. Otherwise, ISLs will be segmented
during next E-port bring-up.
ARE YOU SURE (yes, y, no, n): [no] yes
Auth Policy is set to active.
```

7. Validate the authentication configuration:

```
switch:admin> authutil --show
AUTH TYPE      HASH TYPE      GROUP TYPE
fcap,dhchap    sha256,sha1,md5  4
Switch Authentication Policy: ACTIVE
Device Authentication Policy: OFF
```

8. Configure ports to use encryption for ISLs on local and remote switches:

```
switch:admin> portdisable 1/23
switch:admin> portcfgencrypt --enable 1/23
switch:admin> portenable 1/23
```

9. Validate the port's configuration:


```

switch:admin> portcfgshow 1/23
Area Number:          23
Octet Speed Combo:    1(32G|16G|8G|4G)
Speed Level:          8G
AL_PA Offset 13:      OFF
Trunk Port             ON
Long Distance          OFF
VC Link Init           OFF
Locked L_Port          OFF
Locked G_Port          OFF
Disabled E_Port        OFF
Locked E_Port          OFF
ISL R_RDY Mode         OFF
RSCN Suppressed        OFF
Persistent Disable     OFF
LOS TOV mode           0(OFF)
NPIV capability        ON
QOS Port               AE
Port Auto Disable:     OFF
EX Port                OFF
Mirror Port            OFF
SIM Port               OFF
Credit Recovery        ON
F_Port Buffers         OFF
E_Port Credits         OFF
Fault Delay:           0(R_A_TOV)
NPIV PP Limit:         126
NPIV FLOGI Logout:     OFF
CSCTL mode:            OFF
TDZ mode:              OFF
D-Port mode:           OFF
D-Port over DWDM:      OFF
Compression:           OFF
Encryption:            ON
10G/16G FEC:           ON
16G FEC via TTS:       OFF

```

Implementing FCAP for ISL encryption

The following process shows an alternative way to enable FCAP for ISL encryption:

1. Generate switch CSR files:

```
switch:admin> seccertmgmt generate -csr fcap
```

2. Export switch CSR files to CA for signing:

```
switch:admin> seccertmgmt export -csr fcap
```

Example with complete customer-supplied input:

```
seccertmgmt export -csr fcap -protocol scp -ipaddr 10.1.1.1 -login hostsusername \
  -password hostspassword -remotedir /home/fcap/fcap_files
```

3. Import CA certificate to each switch:

```
switch:admin> seccertmgmt import -ca -client fcap
```

Example with complete customer-supplied input:

```
seccertmgmt import -ca -client fcap -protocol scp -ipaddr 10.1.1.1 \  
-remotedir /home/fcap/fcap_files -certname ca.pem \  
-login hostsusername -password hostspassword
```

4. On the host, verify switch CSRs:

```
openssl req -text -noout -verify -in FcapSw.csr
```

5. On the host, CA signs switch CSRs:

```
openssl ca -in FcapSw.csr -out FcapSw.pem
```

6. Verify CA signed switch CSRs:

```
openssl x509 -in FcapSw.pem -text -noout
```

7. Validate filename permission on host:

```
chmod 777 FcapSw.pem
```

8. Import CA signed switch CSRs:

```
switch:admin> seccertmgmt import -cert fcap
```

Example with complete customer-supplied input:

```
seccertmgmt import -cert fcap -protocol scp -ipaddr 10.1.1.1 \  
-remotedir /home/fcap/fcap_files -certname FcapSw.pem \  
-login hostsusername -password hostspassword
```

9. Validate switch database:

```
switch:admin> seccertmgmt show -all
```

10. Enable FCAP authentication:

```
switch:admin> authutil --set -a fcap  
Authentication is set to fcap.  
switch:admin> authutil --set -g "4"  
DH Group was set to 4.  
switch:admin> authutil --set -h sha256  
Hash is set to sha256.
```

11. Configure ports to use ISL encryption on local and remote switches:

```
switch:admin> portdisable 1/23
switch:admin> portcfgencrypt --enable 1/23
switch:admin> portenable 1/23
```

12. Validate the port's configuration:

```
switch:admin> portcfgshow 1/23
Area Number:          23
Octet Speed Combo:    1(32G|16G|8G|4G)
Speed Level:          8G
AL_PA Offset 13:      OFF
Trunk Port            ON
Long Distance         OFF
VC Link Init          OFF
Locked L_Port         OFF
Locked G_Port         OFF
Disabled E_Port       OFF
Locked E_Port         OFF
ISL R_RDY Mode        OFF
RSCN Suppressed       OFF
Persistent Disable    OFF
LOS TOV mode          0(OFF)
NPIV capability       ON
QOS Port              AE
Port Auto Disable:    OFF
EX Port               OFF
Mirror Port           OFF
SIM Port              OFF
Credit Recovery       ON
F_Port Buffers        OFF
E_Port Credits        OFF
Fault Delay:          0(R_A_TOV)
NPIV PP Limit:        126
NPIV FLOGI Logout:    OFF
CSCTL mode:           OFF
TDZ mode:             OFF
D-Port mode:          OFF
D-Port over DWDM:     OFF
Compression:          OFF
Encryption:           ON
10G/16G FEC:          ON
16G FEC via TTS:      OFF
```

Consider the following points regarding in-flight encryption and compression:

- Encryption and compression are supported on E_Ports and EX_Ports.
- ISL ports must be set to long-distance (LD) mode when compression is used.
- Twice the number of buffers must be allocated if compression is enabled for long-distance links because frame sizes might be reduced by approximately half.
- If both compression and encryption are used, enable compression first.

Implementing ISL encryption across multiple ISLs between the same switch pair requires that either all ISLs be configured for encryption or none of them.

The supported number of encryption- and compression-enabled ports varies according to the platform generation, hardware architecture, and link speed. Earlier platforms imposed limits on a per-ASIC basis. See the platform-specific documentation for the applicable encryption and compression capacity limits.

Consider the following points regarding Virtual Fabric (encryption and compression):

- E_Ports in user-created logical switches, the base switch, and the default switch can support encryption and compression.
- Encryption and compression are supported on XISL (Extended Inter-Switch Link) ports but not on LISL (Logical Inter-Switch Link) ports. LISLs are internal links that connect logical switches through an XISL and are used for Virtual Fabrics communication.
- If encryption or compression is enabled on a port, it must be disabled before the port can be moved from one logical switch (LS) to another.

Policy database distribution

Security Policy Database Distribution provides a mechanism for controlling the distribution of each policy on a per-switch basis. Switches can individually configure policies to accept or reject a policy distribution from another switch in the fabric.

In addition, a fabric-wide distribution policy can be defined for the SCC and DCC policies with support for strict, tolerant, and absent modes. This setting can be used to enforce whether the SCC or DCC policy must be consistent throughout the fabric.

The Policy Database Distribution features the following modes:

Strict

All updated and new policies of the specified type (SCC, DCC, or both) must be distributed to all switches in the fabric, and all switches must accept the policy distribution.

Tolerant

All updated and new policies of the type specified (SCC, DCC, or both) are distributed to all switches (FOS V6.2.0 or later) in the fabric. However, the policy does not need to be accepted.

Absent

Updated and new policies of the type specified (SCC, DCC, or both) are not automatically distributed to other switches in the fabric. Policies can still be manually distributed.

Together, the policy distribution and fabric-wide consistency settings provide a range of control on the security policies from little or no control to strict control. Strict mode is required for FICON environments that require High Integrity Fabrics be enabled.

Cryptographic and security standards

Fabric OS firmware uses the Brocade Fabric OS FIPS Cryptographic Module 9.1.1 library to perform cryptographic functions. The module must be used in a Federal Information Processing Standards (FIPS) compliant operational environment along with the proper device configuration.

Fabric OS firmware uses the Brocade Fabric OS Common Criteria (CC) standards with Fabric OS 9.1.1 Network Device Collaborative Protection Profile (NDcPP) to perform cryptographic functions.

Fabric OS version 9.1.1b8 is in accordance with the BSI (Bundesamt für Sicherheit in der Informationstechnik) EAL2 (Evaluation Assurance Level 2) certified configuration.

For more information about other topics in this subject area that are out of the scope of this document, such as encrypting data that is at rest and in flight, see [Brocade Fabric OS Administration Guide, 10.0x](#).

IBM b-type Gen 8 Fabric OS administrative tools and diagnostics

This section provides an overview of the embedded management and diagnostic tools available in Fabric OS (FOS) 10.

Command-line interface overview

The command-line interface (CLI) enables administrators to monitor and manage individual switches, ports, and entire fabrics from a standard workstation. In Fabric OS 10, the CLI can be accessed remotely over the management Ethernet network by using Secure Shell (SSH) or locally through the switch serial console. SSH is the preferred management interface because it provides encrypted communications and secure authentication.

Fabric OS 10 enhances management security by removing support for several legacy and non-secure protocols and cryptographic mechanisms, including Telnet, HTTP, FTP, SNMPv1, non-secure syslog, SHA-1, and MD5. Administrators should use secure alternatives such as SSH and HTTPS for switch management and monitoring.

Fabric OS uses role-based access control (RBAC) to control access to administrative operations. Each feature is associated with one or more RBAC roles that determine whether a user can execute commands, modify switch settings, or view command output. When creating user accounts, ensure that users are assigned the roles required for their job responsibilities. To determine the RBAC role required for a specific command, see the Role-Based Access Control section of the Brocade Fabric OS Administration Guide, 10.0.x. CLI commands and options are case-sensitive. For detailed command syntax and usage information, see the [Brocade Fabric OS Command Reference Manual, 10.0.x](#).

Connecting to Brocade Fabric OS through the serial port

A serial console connection provides local access to Fabric OS and is commonly used during initial switch setup, network configuration, password recovery, and other maintenance procedures when network access is unavailable. On Brocade X8 Directors, administrators can connect to either control processor (CP0 or CP1) through the available serial ports.

For detailed cabling requirements, terminal emulator settings, and platform-specific procedures, see the [Brocade Fabric OS Administration Guide, 10.0.x](#).

Connecting to Brocade Fabric OS through SSH

SSH is the preferred method for remote command-line administration of Fabric OS. The switch must have a configured management IP address and network connectivity before an SSH session can be established. If network connectivity is not available, use the serial console to perform the initial configuration.

To establish an SSH session:

```
ssh admin@<switch-ip-address>
```

After authentication, the CLI prompt displays the switch name and the currently logged-in user.

For detailed connection procedures and authentication options, see the [Brocade Fabric OS Administration Guide, 10.0.x](#).

Web Tools overview

Brocade Web Tools is a graphical user interface (GUI) embedded in the Fabric OS firmware. It enables administrators to monitor and manage small fabrics, switches, and ports. Web Tools can be launched directly from a standard web

browser or from the SANnav Management Portal. Additionally, Web Tools supports firmware migration to FOS v10.x from FOS v9.2.2x.

Note: Web Tools will always show English language irrespective of the browser or operating system language setting.

Enabling and disabling Web Tools

In FOS, Web Tools is enabled by default. Starting with FOS v10.x, administrators can enable or disable Web Tools using the `mgmtapp` command or via REST URI.

If Web Tools is disabled and a user attempts to launch it from a browser or SANnav, the following error message is displayed:

"Web Tools interface is disabled in Fabric OS."

Launching Web Tools

Web Tools can be launched on any workstation with a compatible operating system and web browser. If the switch is configured with Logical Fabrics, users can log in to any authorized logical fabric based on their account permissions.

To launch Web Tools directly from a web browser:

1. Open a supported web browser.
2. Enter the HTTPS protocol followed by the switch's IP address. For example: `https://10.77.77.77`

Note: Starting with FOS v10.0.1, users can open a page in a new tab using Ctrl+Click or Cmd+Click. This action opens the specified page directly without requiring re-authentication.

For more information about configuring and using Web Tools, see the [Brocade® Fabric OS® Web Tools User Guide, 10.0.x](#), available from the Broadcom TechDocs website.

Fabric Vision overview

Fabric Vision technology is an advanced hardware and software architecture that combines capabilities from FOS, b-type Gen 6, 7 and 8 devices, and SANnav. This suite of tools can be used to help administrators address problems before operations are affected and accelerate new application deployments. Fabric Vision technology features are described in the following sections.

Monitoring and Alerting Policy Suite

The Monitoring and Alerting Policy Suite (MAPS) is the primary health monitoring and alerting framework in IBM b-type Fabric OS. MAPS continuously monitors the SAN by comparing operational metrics against predefined or user-defined policies. When monitored values exceed configured thresholds, MAPS can generate alerts or automatically perform corrective actions, helping administrators identify and address problems before they affect application availability.

MAPS simplifies SAN operations by replacing manual monitoring with policy-based automation. Instead of continuously reviewing logs and performance statistics, administrators can define policies that monitor the health of switches, ports, traffic flows, environmental sensors, and other critical resources. This proactive approach improves operational efficiency while reducing the risk of service disruptions.

MAPS is commonly used for the following purposes:

- Monitoring switch and port health
- Detecting congestion and slow-drain devices
- Monitoring environmental conditions and hardware components
- Tracking traffic flow and performance metrics
- Automatically fencing, quarantining, or toggling ports when policy thresholds are exceeded

MAPS is a licensed feature that requires the Fabric Vision license. This publication assumes that the required license has been installed and activated.

Fabric OS includes several predefined policies that provide different monitoring sensitivities:

Policy	Description
dflt_aggressive_policy	Uses the most stringent thresholds and is intended for environments that require rapid detection and response.
dflt_moderate_policy	Provides a balanced set of thresholds suitable for most production environments.
dflt_conservative_policy	Uses more tolerant thresholds to reduce unnecessary alerts in environments where occasional errors are acceptable.
dflt_base_policy	Supports the subset of monitoring functions available without a Fabric Vision license.

Administrators can customize these predefined policies or create their own rules, groups, and policies to meet specific operational requirements. MAPS configurations are persistent and remain available after switch reboots or high-availability failovers.

MAPS can be managed through either the Fabric OS CLI or SANnav Management Portal. SANnav simplifies administration by providing centralized policy management across multiple switches, allowing administrators to create, modify, distribute, and activate MAPS policies from a single interface. This centralized approach is particularly beneficial in large SAN environments where maintaining consistent monitoring policies across multiple fabrics is essential.

For detailed information about Web Tools installation, configuration, and administration, see the [Brocade® Fabric OS® Web Tools User Guide, 10.0.x](#).

Flow Vision

Flow Vision is a Fabric Vision technology that provides advanced traffic monitoring, analysis, testing, and diagnostic capabilities for Fibre Channel SAN environments. It helps administrators gain visibility into application traffic flows, validate fabric infrastructure, and troubleshoot performance issues with minimal impact on production workloads.

Flow Monitor

Flow Monitor is a Fabric Vision feature that provides visibility into application traffic flows across the SAN. It collects Fibre Channel, SCSI, and NVMe statistics, enabling administrators to monitor workload behavior, identify congestion, troubleshoot performance issues, and perform capacity planning. Flow Monitor supports both predefined and user-defined flows based on flow characteristics such as source and destination devices, ingress and egress ports, and protocol types. On Gen 7 and Gen 8 platforms running Fabric OS 10.0.x, automated flow monitoring can dynamically discover application flows and collect performance statistics without disrupting production workloads.

Flow Monitor includes the following key benefits:

- Monitoring end-to-end SAN traffic flows
- Identifying bottlenecks, congestion, and slow-drain devices

- Collecting SCSI and NVMe performance metrics
- Supporting capacity planning and workload analysis
- Providing nonintrusive traffic visibility across the fabric

For detailed configuration procedures and command syntax, see the Brocade Fabric OS Administration Guide, 10.0.x.

Flow Generator

Flow Generator is a Fabric Vision testing feature that generates synthetic Fibre Channel traffic to validate SAN connectivity, optics, ISLs, and switch hardware under controlled load conditions. It enables administrators to verify infrastructure readiness and evaluate performance before introducing production workloads.

Unlike traditional testing methods that require hosts and storage devices, Flow Generator can use simulated devices to generate traffic at line rate, allowing SAN components to be exercised under realistic operating conditions.

Flow Generator supports the following use cases:

- Validating new SAN deployments before production use
- Testing optics, cables, and inter-switch links (ISLs)
- Verifying fabric performance and throughput
- Troubleshooting connectivity or infrastructure issues
- Evaluating the impact of configuration changes

Notes:

- SANnav does not manage Flow Generator.
- Fabric OS 10.x supports Flow Generator on Brocade Gen 7 and Gen 8 platforms.

Important: Flow Generator should not be used in active production environments. Generated traffic can consume significant bandwidth and may affect application workloads that share the same links.

Flow Generator uses simulation-mode (SIM) ports. These ports behave similarly to standard F_Ports but are used exclusively for testing. Traffic generated by SIM ports remains within the fabric and does not reach production hosts or storage devices.

For detailed configuration and operational procedures, see the Brocade Fabric OS Administration Guide, 10.0.x.

Flow Mirror

As SAN environments grow in size and complexity, administrators increasingly require nonintrusive methods for analyzing traffic without disrupting production workloads. Flow Mirror addresses this requirement by allowing selected traffic flows to be copied and analyzed while normal I/O operations continue uninterrupted.

Flow Mirror supports the following use cases:

- Capture application traffic for detailed analysis
- Investigate intermittent performance or latency issues
- Analyze traffic associated with congestion events
- Troubleshoot slow-drain devices
- Observe traffic patterns without affecting application operation

Flow Mirror supports two mirroring methods:

- CPU Flow Mirroring (CFM)

CPU Flow Mirroring copies selected frames to the switch CPU for analysis. This method is useful for capturing small amounts of traffic for troubleshooting and diagnostic purposes. CPU Flow Mirroring supports a maximum rate of 256 frames per second.

- Local Flow Mirroring (LFM)

Local Flow Mirroring copies traffic to a designated analyzer port connected to an external packet capture or protocol analysis tool. This approach allows detailed examination of Fibre Channel traffic without affecting the original flow. The available bandwidth is limited by the speed of the mirror destination port.

Flow Mirror operates on user-defined flows and can mirror either ingress or egress traffic. Mirrored flows can be configured as active or inactive and can operate in either unidirectional or bidirectional mode.

A common use case is troubleshooting a slow-drain device. By mirroring traffic associated with the affected device, administrators can analyze frame behavior and identify conditions contributing to congestion or performance degradation.




Figure: Flow Mirror CFM and LFM operation diagram

For detailed information about Fabric OS administration, configuration, security, and monitoring, see the [Brocade® Fabric OS® Administration Guide, 10.0.x](#).

D_Port testing

Brocade ClearLink Diagnostic Port (D_Port) is a Fabric OS diagnostic feature that allows administrators to temporarily convert a Fibre Channel port into a diagnostic port for validating links and testing hardware before it is placed into production. D_Port testing can be performed through either the Fabric OS CLI or SANnav Management Portal.

Using Brocade ClearLink technology, D_Port testing helps verify the health of switches, adapters, optics, cables, and inter-switch links (ISLs). This capability is particularly useful when deploying new hardware, replacing SFPs, or troubleshooting suspected link problems.

The following tests are supported:

- Electrical loopback testing
- Optical loopback testing
- Link traffic testing
- Link latency and distance measurement

A common use case is validating newly installed optics and cabling between two switches before the ports are enabled as production E_Ports. If the diagnostic tests complete successfully, the ports can be placed into service with confidence. If failures are detected, corrective action can be taken before the issue affects application traffic.

Note: D_Port testing uses Brocade ClearLink® proprietary technology. For detailed procedures, supported configurations, and command syntax, see the [Brocade Fabric OS Troubleshooting and Diagnostics Reference Manual](#).

Starting D_Port testing

D_Port testing can be initiated from either the Fabric OS CLI or SANnav Management Portal.

The following example shows how to start a D_Port test from the CLI:

```
switch:admin> portcfgdport --enable 1/0

switch:admin> porttestdport --start 1/0

switch:admin> porttestdport --show 1/0

Port 1/0:
Electrical Loopback Test ..... PASS
Optical Loopback Test ..... PASS
Link Traffic Test ..... PASS
Link Latency Test ..... PASS
```

SANnav also provides a graphical interface that allows administrators to select one or more ports, configure diagnostic parameters, execute the tests, and review the results. This approach simplifies D_Port testing across multiple switches and ports from a centralized management interface. For more information, see the [Brocade Fabric OS Troubleshooting and Diagnostics Reference Manual](#).

Fabric performance impact

Fabric Performance Impact (FPI) is a Fabric Vision capability that uses predefined MAPS policies to automatically detect, classify, and report SAN congestion conditions. FPI helps administrators identify issues such as slow-drain devices, credit-stalled devices, and oversubscribed ports that can adversely affect application performance.

By continuously monitoring E_Ports and F_Ports, FPI identifies both the devices causing congestion and the devices affected by it. FPI can also automatically quarantine credit-stalled devices to prevent buffer credit starvation from propagating through the fabric.

FPI is particularly useful in large SAN environments where congestion-related issues can be difficult to detect and isolate. By providing early warning and automated corrective actions, FPI helps maintain application performance and improve overall fabric resiliency.

For detailed configuration procedures and policy settings, see the Brocade® Fabric OS® Administration Guide, 10.0.x.

I/O Insight

I/O Insight is a Fabric Vision analytics capability that provides application-level visibility into storage traffic by monitoring SCSI and NVMe I/O activity within the SAN. Using integrated network sensors, I/O Insight continuously collects performance and latency metrics without disrupting production workloads.

The collected data is integrated with MAPS, enabling administrators to define thresholds and alerts that help identify performance anomalies before they affect application service levels. By monitoring IOPS, I/O latency, response times, and outstanding I/O activity, I/O Insight provides valuable information for troubleshooting performance problems and establishing application performance baselines.

I/O Insight includes the following key benefits:

- Monitoring application and storage performance at the I/O level
- Detecting abnormal latency and performance degradation
- Supporting SLA monitoring and performance baseline creation
- Providing visibility into both SCSI and NVMe workloads
- Assisting with storage tuning and capacity planning
- Simplifying the identification of application and virtual machine performance issues

I/O Insight enables administrators to proactively manage SAN performance and quickly identify the root causes of application slowdowns before they affect business operations.

For detailed configuration procedures and supported platforms, see the Brocade® Fabric OS® Administration Guide, 10.0.x.

Adaptive Traffic Optimizer

Adaptive Traffic Optimizer (ATO) is an automated traffic management capability available on Brocade Gen 8 platforms running Fabric OS 10.x. ATO continuously analyzes fabric conditions and dynamically optimizes traffic distribution to improve congestion management and maximize the utilization of available buffer and virtual channel resources.

Unlike earlier optimization approaches that relied on predefined profiles, ATO automatically learns the characteristics of the storage environment, including device speeds, protocols, and congestion conditions. Based on these observations, ATO dynamically adjusts Performance Groups (PGs) to provide optimal workload isolation and traffic handling without requiring administrator intervention.

ATO includes the following key benefits:

- Automatic optimization of traffic flows across the fabric
- Improved congestion management
- Better utilization of virtual channel and buffer resources
- Dynamic adaptation to changing workload patterns
- Reduced administrative effort through policy-free optimization
- Enhanced support for mixed SCSI and NVMe environments

ATO is particularly beneficial in dynamic environments where workload characteristics frequently change. By continuously adapting to current fabric conditions, ATO helps maintain consistent performance without requiring manual tuning.

For detailed information about Performance Groups, traffic classification, and configuration options, see the Brocade® Fabric OS® Administration Guide, 10.0.x.

SAN Fabric Intelligence

SAN Fabric Intelligence (SAN FI) provides a unified view of server, storage, and SAN infrastructure information, helping administrators quickly identify areas that require further investigation. Available through the Fabric OS CLI, SAN FI consolidates information from multiple sources into easy-to-consume dashboards and provides a common framework for monitoring and troubleshooting across the fabric.

By correlating data from switches, hosts, storage systems, virtual machines, MAPS, and Flow Vision, SAN FI can significantly reduce the time required to identify and resolve operational issues.

Common SAN FI use cases include the following examples:

- Viewing SAN infrastructure, including storage systems, physical servers, virtual machines (VMs), and their relationships
- Viewing summarized SAN telemetry and monitoring information, including MAPS events, Flow Vision data, port statistics, and congestion metrics
- Identifying performance issues through congestion analytics
- Identifying physical-layer issues through MAPS violations
- Analyzing fabric paths between source and destination devices using path enumeration
- Visualizing all available fabric paths between endpoints
- Comparing fabric latency and device latency
- Identifying congestion microbursts that may affect application performance

Path enumeration

Path enumeration provides visibility into all logical paths between a source and destination within the SAN. Examples include paths between a virtual machine and a storage LUN, a host port and a target port, or two switches. By displaying the complete path through the fabric, administrators can identify routing issues, congestion points, and latency bottlenecks that may affect application performance.

Fabric latency and device latency

SAN FI distinguishes between fabric latency and device latency to help isolate the source of performance problems.

Fabric latency is the delay introduced by the SAN infrastructure itself, including switches, inter-switch links (ISLs), and fabric services.

Device latency is the delay introduced by endpoint devices, such as servers, HBAs, storage controllers, or storage arrays.

This distinction helps administrators determine whether an application slowdown is being caused by the SAN fabric or by the attached storage or host systems.

Congestion microbursts

A congestion microburst is a short-duration spike in traffic that temporarily exhausts available buffer credits and creates localized congestion within the fabric. Although these events may last only milliseconds, they can significantly affect application response times and are often difficult to detect using traditional monitoring methods.

SAN FI helps identify congestion microbursts and correlate them with affected devices and fabric paths. This visibility enables administrators to investigate intermittent performance problems that might otherwise go undetected.

SAN FI vCenter integration

SAN Fabric Intelligence (SAN FI) integrates SAN infrastructure information with VMware virtualization metadata. Fabric OS 10.0.0 supports a single VMware vCenter instance, whereas Fabric OS 10.0.1 and later supports up to four vCenter instances. Support for multiple vCenters allows SAN FI to correlate information across multiple VMware environments, providing end-to-end visibility between virtual machines, hosts, storage resources, and SAN infrastructure.

A SAN FI-enabled fabric can be configured to communicate with up to four VMware vCenter instances. Before adding multiple vCenters, all switches in the fabric must be running Fabric OS 10.0.1 or later, and the vCenter root CA certificate must be imported into the switch trust store.

The following example illustrates the process of adding a vCenter to SAN FI:

```
switch:admin> sanfi --config -addvcenter vc01.example.com \  
-username administrator@vsphere.local  
  
vCenter Password: *****  
  
Successfully added vCenter:  
vc01.example.com
```

Use the fully qualified domain name (FQDN) of the vCenter server whenever possible to simplify certificate validation and improve configuration stability.

SAN FI dashboards

SAN FI provides several dashboards that help administrators analyze inventory, performance, health, and congestion information across the fabric:

Dashboard	Description
Object Summary Dashboard	Displays inventory, performance, and health information for SAN objects, including fabrics, switches, ports, servers, virtual machines, and storage resources.
Congestion Dashboard	Summarizes fabric congestion events and assists with identifying recurring congestion patterns.
MAPS Violations Dashboard	Displays the most significant MAPS events and allows administrators to quickly identify physical-layer or operational issues.
Fabric Path Enumeration Dashboard	Displays available logical paths between source and destination objects along with path latency, errors, and congestion information.

Troubleshooting with SAN FI

SAN FI provides workflow-driven troubleshooting that helps administrators quickly isolate physical-layer, congestion-related, and performance issues.

A typical troubleshooting workflow includes the following steps:

- Reviewing the Fabric Summary dashboard to identify areas requiring further investigation
- Examining switch-level summaries to identify affected switches
- Reviewing port-level information to isolate affected ports
- Examining MAPS violations and congestion information to determine root cause

Using this workflow, administrators can rapidly identify issues such as CRC errors, invalid transmission words (ITWs), congestion events, and excessive latency.

Fabric performance impact notifications

Fabric Performance Impact Notifications (FPINs) are Fibre Channel Extended Link Service (ELS) notifications that communicate fabric performance issues to participating devices. When supported by both the fabric and endpoint devices, FPINs enable hosts and storage systems to react intelligently to congestion, signal integrity issues, and frame delivery problems.

The following table includes the primary FPIN descriptor types:

Descriptor	Description
Congestion Descriptor	Indicates congestion conditions such as credit stalls or oversubscription.
Peer Congestion Descriptor	Notifies devices that share a zone with a congested device.
Link Integrity Descriptor	Reports physical-layer issues, including CRC and Invalid Transmission Word (ITW) errors.
Delivery Descriptor	Indicates that a SCSI-FCP command frame was dropped within the fabric.

FPINs improve end-to-end performance awareness by extending fabric intelligence beyond the SAN switches and into connected hosts and storage systems.

FICON

FICON (Fiber Connection) is the IBM mainframe channel architecture used to connect IBM Z systems to storage devices through a Fibre Channel SAN. Fabric OS provides specialized capabilities for FICON environments, including FICON logical switches, management controls, security policies, and diagnostic functions designed to meet the performance, availability, and operational requirements of mainframe workloads. Proper configuration of FICON fabrics is essential to ensure predictable behavior, high availability, and compliance with IBM Z operational requirements.

FICON requirements

The following are requirements for configuring FICON on Brocade switches and directors:

- A FICON logical switch (FICON LS) must be used in a FICON environment. Do not connect a FICON channel or device to a non-FICON LS port. The default switch cannot be changed to a FICON LS; therefore, Virtual Fabrics must be used.
- Configuring a domain ID: Give all switches a unique domain ID.
- Configuring an SCC policy: An SCC policy is created when you create a FICON LS. The initial SCC policy contains only the current native domain. Before the new domain can join an existing fabric, the new SCC policy must be updated to contain the switches in the existing fabric's SCC policy.
- Port addresses: Ports in a FICON LS must have an address bound to them before they can be enabled. When moving ports into a FICON LS, use the `portAddress --bind` command to bind port addresses before enabling the port.

Note: You cannot bind to port addresses 0xFE and 0xFF in a FICON LS.

FICON best practices

Follow these best practices for configuring FICON on Brocade switches and directors:

- Switch ID

Make the switch ID in the IOCP or HCD the same as the hexadecimal equivalent of the domain ID; although not required, this makes it easier to correlate switches to the IOCP. For example, if the domain ID is 32 (0x20), the switch ID should be 20.

- Firmware

- Use only qualified firmware for FICON; see Brocade Fabric OS Target Path.
- All versions of platform firmware that Brocade SANnav Management Portal is managing should be in the SANnav repository.
- All fabric platforms should be at the same code level.
- When upgrading Fabric OS more than two levels, plan the upgrades so that no fabric switch is more than one release apart.
- Always check the switch's firmware version before putting the switch into production.
- If a firmware upgrade is required on a new switch installation, upgrade the firmware before making any configuration changes.

- CHPID and control unit placement

Place CHPID and control unit connections on the same blade when possible, to minimize the impact of a blade failure. Using the same blade port group, which is optional, maximizes Brocade local switching to minimize latency and optimize performance. CHPID and control unit placement should be done in conjunction with your IO Configuration Data Set (IOCDS) planning.

- Default Switch

Leave the default switch FID at 128.

- Port Fencing

Use the default MAPS aggressive policy (`dflt_aggressive_policy`) to set the port fencing parameters. For more information about customizing MAPS policies for FICON, see SAN and Fabric Resiliency Best Practices for IBM b-type Products.

- WWN-based persistent port area IDs (PIDs)

Do not enable WWN-based persistent PIDs in FICON environments.

- Zoning for FICON

The default zoning mode changed from All Access to No Access in Fabric OS v10.0.x. To ensure device connectivity between FICON channels and FICON devices, Brocade recommends that you create one or more zones for FICON that match the HCD/IOCDS configuration, add the zones to a zoning configuration, and then activate the zoning configuration. Discrepancies in the default zoning modes can cause E-Ports to segment.

Note: Default zoning is not supported with FICON when using a Base Fabric.

Preparing a switch for FICON

Follow these steps to verify and prepare a switch for use in a switched point-to-point FICON environment:

1. Connect to the switch, and log on using an account assigned the admin role.
2. Verify that the proper management IP addresses have been configured.
3. Verify that the switches can be seen by your management tool, for example SANnav.

4. Verify that the switches have the correct version of Fabric OS.
5. Add feature keys, if applicable.
6. Enter the `configUpload` command to save a baseline of the switch configuration.
7. To create a FICON logical switch, see [Creating a FICON logical switch](#).

Currently supported platforms

The following platforms are currently supported for Fabric OS 10.0.0 for the IBM z14 and later systems:

Platform	Type
Brocade X8-8 Director	Gen 8 Director
Brocade X8-4 Director	Gen 8 Director
Brocade G820 Switch	Gen 8 Switch
Brocade X7-8 Director	Gen 7 Director
Brocade X7-4 Director	Gen 7 Director
Brocade G720 Switch	Gen 7 Switch
Brocade 7850 Extension Switch	Gen 7 Extension

To determine the FICON support for Fabric OS 10.0.x, see the FICON qualification support letters on the [Broadcom Mainframe SAN page](#).

IBM b-type Gen 8 SAN automation

Most, if not all, IT administrators have first-hand experience in the growing complexity of managing enterprise infrastructure, including SANs.

IBM b-type Gen 8 provides comprehensive automation capabilities that help organizations improve the efficiency, consistency, and scalability of SAN operations.

IBM b-type Gen 8 offers a combination of SAN automation with RESTful APIs and a SAN management platform to help organizations drive greater efficiency from their SANs.

This greater efficiency is realized through several capabilities:

- IBM b-type Gen 8 SAN automation is provided with a multilayer architecture.
- RESTful API on directors and switches, and on management tools.
- IBM b-type Gen 8's Ansible management framework, which is designed to eliminate repetitive tasks, simplify management, and orchestrate across the full SAN infrastructure.

Supported hardware and software

Software support

The Fabric OS (FOS) REST API is supported on IBM b-type Gen 7 devices running FOS 9.1 or later and IBM b-type Gen 8 devices running FOS 10.0 or later. These versions include a set of CLI commands that control the REST session management and API functionality.

Hardware support

The FOS REST API supports all Brocade Gen 7 Fibre Channel devices that are supported by Fabric OS v9.1.0 and later and Brocade Gen 8 devices that are supported by Fabric OS v10.0.0 and later.

Motivation to automate

Organizations should embrace SAN automation for the following reasons:

- Reducing human error and streamlining operational processes has never been more crucial.

As organizations move to digitize and adapt to new workloads, data availability, processing time, and agility to provision applications on-demand become the lifeblood of the business. Achieving these business objectives demands a more efficient and expedient infrastructure management approach, which leaves minimizing the potential for human error.

As a result, storage administrators must be freed from repetitive manual tasks, such as configuration management, reporting, documenting inventory, and troubleshooting.

Instead, IT organizations need SAN automation to help them automate and orchestrate repetitive tasks, which significantly improves efficiency, and decreases the risk of operational mistakes.

- Demand for more accurate and more frequent infrastructure reports is on the rise.

Increasingly, not only IT managers but also business stakeholders require access to storage performance, utilization, and forecasting information. They expect this data to be readily available and accessible on demand. No one wants to wait for a slot when storage administrators can allocate time to produce a report. This information must be available as frequently as business demands dictate, on demand. Automation not only provides this kind of responsiveness that traditional manual storage management processes cannot deliver, but it also can be customized so that each stakeholder is getting more accurate data that is aligned to their responsibility.

- SAN configuration management must be streamlined.

With more enterprise applications demanding access to more data and with more virtual machines, deploying and configuring servers, storage, and the network became more time-consuming and more complex than ever. By streamlining SAN operations through automation, application provisioning workflows are simplified across hypervisor, network, and storage, which delivers agility and responsiveness to meet dynamic business demands.

- IT service delivery is not always as responsive as the business demands.

As organizations become increasingly dependent on high-quality IT services to support proactive and agile business decision-making, any bottlenecks in IT service delivery must be identified and resolved quickly. At the same time, organizations are under pressure to achieve these goals without increasing the number of storage administrators or significantly expanding SAN-related capital expenditures. As a result, SAN automation has become a critical enabler of IT agility, allowing organizations to streamline operations, improve efficiency, and align IT service delivery more closely with rapidly evolving business requirements.

- Consistent configuration validation is a must.

Manual configuration changes are becoming more frequent as enterprises diversify their IT infrastructures and adopt increasingly complex storage architectures. SAN automation helps maintain consistent configuration settings across multiple SAN fabrics, reducing the risk of configuration drift and human error. By enforcing standardized policies and automating routine tasks, SAN automation also simplifies troubleshooting and

accelerates the resolution of frequent alerts, minimizing the need for manual intervention and improving overall operational efficiency.

IBM b-type Gen 8 automation solutions use RESTful APIs to facilitate solution architecture, share best practices, and get to production status faster.

REST API overview

The Brocade Fabric OS (FOS) REST API provides a programmable web services interface for IBM b-type Gen 8 SAN environments, enabling centralized management of SAN switches across a fabric. The API uses standard HTTP methods to perform create, read, update, and delete operations on fabric configuration data, allowing administrators and applications to automate common management tasks.

In addition to configuration management, the API supports provisioning, monitoring, status collection, and validation operations through the YANG data model defined in the YANG 1.1 RFC. Unlike NETCONF implementations, however, it does not include a YANG-based datastore management framework. The API is hosted by an Apache web server embedded within Fabric OS, providing a secure and standards-based interface for integrating SAN management with automation and orchestration tools.

The RESTful API approach helps you think of a network device as a web server. By using standard web-based tools, an automation can send and receive transactions to or from a network device just as it sends transactions to and from a website. Therefore, transactions occur over a secure connection by using HTTP rules to handle the exchange. The data appears in the form of XML or JSON, depending on the RESTful API services that are implemented on the networking device.

To interact with a SAN (or other) device, review its RESTful API reference to learn, among other things, what Uniform Resource Identifiers (URIs) you must use. URIs are identifiers that can be used as part of a web address. According to the documentation, the following URL is used for accessing a listing of the zones in the active configuration:

```
GET <base_URI>/running/zoning/defined-configuration/
```

The model that is used to represent state and configuration information is expressed in a modeling language called YANG. YANG describes the structure of the different elements that are inside the model. It is used to describe whether each element is read-only or read/write.

It also describes the type of data that the element can hold, such as string or integer, and it shows the relationship among various elements, the other nested elements they contain, their peer elements, and the parent elements that contain them. A segment of the description of a zone in YANG is shown in the following example:

```

list zone {
  key "zone-name";
  description
    "List of the members in the zone. The members can only be identified as a WWN,
    domain, index, or a zone alias.";

  leaf zone-name {
    type zoning-name-type;
    description
      "The zone name.";
  }

  leaf zone-type {
    type zone-type-type;
    description
      "The zone type. Note that target zone types cannot be created or modified
      (only deleted).";
  }

  container member-entry {
    description
      "The zone member.";
    leaf-list entry-name {
      type zone-member-type;
      min-elements 1;
      description
        "List of the members in the zone. The members can only be identified as a
        WWN, domain, index, or zone alias.";
    }

    leaf-list principal-entry-name {
      when "../zone-type=1 or ../zone-type=2";
      type zone-member-type;
      min-elements 1;
      description
        "List of the principal members in the peer zone. The members can only be
        identified as a WWN, domain, index, or zone alias.";
    }
  }
}

```

Ordinarily, more information goes into a YANG module, such as revisioning and governance information; this listing omits them for brevity. Therefore, the YANG description is complete, but it is also verbose. Although this precision is necessary when interacting with the model programmatically, it is sometimes useful to get a global view of the abstraction that is provided by the model to see how the data is structured.

An open source tool called `pyang` can parse the YANG model and produce a tree that represents the elements in the model. The listing includes information about each element, such as whether it is read-only or read/write, a list, optional, or nested.

A representation of the zoning model in tree form is shown in the following example:

```

module: IBM-b-type-zone
+--rw IBM-b-type-zone
+--rw defined-configuration
|   +--rw cfg* [cfg-name]
|   |   +--rw cfg-name      zoning-name-type
|   |   +--rw member-zone
|   |       +--rw zone-name* zoning-name-type
|   +--rw zone* [zone-name]
|   |   +--rw zone-name      zoning-name-type
|   |   +--rw zone-type?     zone-type-type
|   |   +--rw member-entry
|   |       +--rw entry-name* zone-member-type
|   |       +--rw principal-entry-name* zone-member-type
|   +--rw alias* [alias-name]
|   |   +--rw alias-name      zoning-name-type
|   |   +--rw member-entry
|   |       +--rw alias-entry-name* union
+--rw effective-configuration
+--rw cfg-name?      zoning-name-type
+--rw checksum?      string
+--rw cfg-action?    uint8
+--rw default-zone-access? uint8
+--ro db-max?        uint32
+--ro db-avail?      uint32
+--ro db-committed?  uint32
+--ro db-transaction? uint32
+--ro transaction-token? uint32
+--ro db-chassis-wide-committed? uint32
+--ro enabled-zone* [zone-name]
+--ro zone-name      zoning-name-type
+--ro zone-type?     zone-type-type
+--ro member-entry
+--ro entry-name*     union
+--ro principal-entry-name* union

```

Simple automation example

In this example, `<base_URI>` is defined as `http://<switch-ip-address>/rest`. To create a login session with a switch in the fabric, run the following command as a single line:

```
curl -X POST -v -u admin:password http://10.18.254.37/rest/login
```

In this command, `curl` is a command-line utility that is used to transfer data to and from a server by using various protocols, including HTTP and HTTPS. The command authenticates to the switch by using the specified user credentials and establishes a REST API session. The POST method is used to send the login request to the switch's REST API endpoint.

Consider the following points:

- `-X POST` specifies the POST HTTP method (instead of GET).
- `-v` specifies verbose output to access the authorization string in the header of the response that is used in the next step.
- `-u admin:password` specifies the credentials to use.

The last parameter is the URI for curl to use to log in. (The URI value is described in the RESTful API reference.)

This command establishes the session that is used for the following commands. The following example is a trace of its execution:

```

* Trying 10.18.254.37...
* Connected to 10.18.254.37 (10.18.254.37) port 80 (#0)
* Server auth using Basic with user 'admin'
> POST /rest/login HTTP/1.1
> Host: 10.18.254.37
> Authorization: Basic YWRtaW46cGFzc3dvcmQ=
> User-Agent: curl/7.47.0
> Accept: */*
>
< HTTP/1.1 200 OK
< Date: Wed, 31 Jan 2018 16:01:24 GMT
< Server: Apache
< Authorization: Custom_Basic YWRtaW46eHh40jNkYTl1ZmM3NzMxYjk4OGU2ODg1YzZkMGRjNWJlMzMyNjBhZDYxZThkOWQ2MWMx
< Cache-Control: no-cache
< X-Frame-Options: DENY
< Content-Secure-Policy: default-src 'self'
< X-Content-Type-Options: nosniff
< X-XSS-Protection: 1; mode=block
< Connection: close
< Transfer-Encoding: chunked
< Content-Type: application/yang-data+xml

```

Next, you perform a GET of the URI to return the current configuration by using the `Custom_Basic` value that is returned from the login for authentication:

```

curl -v -H "Authorization: Custom_Basic YWRtaW46eHh40jNkYTl1ZmM3NzMxYjk4OGU2ODg1YzZkMGRjNWJlMzMyNjBhZDYxZT
http://10.18.254.37/rest/running/zoning/defined-configuration

```

By default, curl uses the GET method, so you do not need to specify it. `-H "Authorization: Custom_Basic YWR...jA="` is the authentication and session identifying string that is returned in the previous command. `-H` places the string into the GET request header, as shown in the following trace:

```

* Trying 10.18.254.37...
* Connected to 10.18.254.37 (10.18.254.37) port 80 (#0)
> GET /rest/running/zoning/defined-configuration HTTP/1.1
> Host: 10.18.254.37
> User-Agent: curl/7.47.0
> Accept: */*
> Authorization: Custom_Basic YWRtaW46eHh40jNkYTl1ZmM3NzMxYjk4OGU2ODg1YzZkMGRjNWJlMzMyNjBhZDYxZThkOWQ2MWMx
>
< HTTP/1.1 200 OK
< Date: Wed, 31 Jan 2018 16:09:39 GMT
< Server: Apache
< Cache-Control: no-cache
< X-Frame-Options: DENY
< Content-Secure-Policy: default-src 'self'
< X-Content-Type-Options: nosniff
< X-XSS-Protection: 1; mode=block
< Connection: close
< Transfer-Encoding: chunked
< Content-Type: application/yang-data+xml

```

```

<?xml version="1.0"?>
<Response>
  <defined-configuration>
    <cfg>
      <cfg-name>CFG_FABRIC_A</cfg-name>
      <member-zone>
        <zone-name>CLUSTER1</zone-name>
        <zone-name>Z_AIXHOST_FCS2_VMAX01_SN1234_9F0</zone-name>
        ...
      </member-zone>
    </cfg>
    <alias>
      <alias-name>esx66_5d3d00</alias-name>
      <member-entry>
        <alias-entry-name>10:00:8c:7c:ff:5d:3d:00</alias-entry-name>
      </member-entry>
    </alias>
  </defined-configuration>
</Response>
* Closing connection 0

```

The results appear as an XML data segment that is structured according to the YANG model. Therefore, see the YANG model together with the Fabric OS REST API documentation. The Fabric OS REST API documentation includes the YANG models that define the REST resources, their attributes, and supported operations. The latest documentation is available from the Broadcom support portal for the applicable Fabric OS release.

The RESTful API manual is available from the Broadcom website. Having retrieved the zoning information from the fabric, you close the session by using the CLI command (the results are omitted to save space):

```

curl -v -H "Authorization: Custom_Basic YWRtaW46eHh40jNkYTl1ZmM3NzMxYjk40GU2ODg1YzZkMGRjNWJlMzMyNjBhZDYxZT
http://10.18.254.37/rest/logout

```

In version 8.2.2 or later of FOS, the REST API sessionless operation allows you to provide authentication credentials directly for each GET request. Essentially, FOS REST API sessionless operation completes the login, GET operation, and logout as one complete request. You can use only basic authentication formats for REST API sessionless operation, which includes plain text or Base64.

The following example shows a GET request that uses plain text authentication:

```

curl -u admin:password http://10.155.2.190/rest/running/IBM-b-type-media/media-rdp

```

Ansible as an alternative

The previous section showed an example of an approach that uses a procedural methodology. The workflow starts at the beginning, runs a series of steps, and then ends. Most traditional programs work this way.

Ansible takes a declarative approach. Rather than provide steps, Ansible describes each of the hosts in an inventory. The description appears in a document that is called a playbook.

For example, rather than provide steps to install a specific application, Ansible describes a host state where the application is installed. When you run the playbook, Ansible takes no action if the application is installed. If the application is not installed, Ansible calls installation routines so the host is brought into the desired state without requiring the administrator to write any specific steps.

In the realm of storage networks, the use of a declarative language means you can describe switches and fabrics where, for example, a zone is configured with the suitable hosts and storage arrays. When you run the Ansible playbook, those zones are defined as needed, and the hosts and storage arrays are added to them, if necessary.

With some other declarative automation utilities, it is necessary to install an agent on each host that the utility manages. This agent retrieves the commands from a command center and runs them on the local host.

Ansible is unique in that it does not require agents. To make switch state changes, Ansible establishes a Secure Shell session to a proxy and sends it a small Python script. The script carries out the necessary operations by using the switch API and removes itself from the host.

You need two different skill sets to successfully implement an Ansible solution. First, you must understand the most common playbook operations. These operations are coded and installed for use by the playbooks. As vendors announce support for Ansible, they also provide script libraries for the most common tasks. If a task exists that is required but is not available in the official Ansible distribution, the open source community might provide code for that task in publicly available repositories.

Second, you must understand your business needs to provide ongoing playbook development. The person that is maintaining the playbooks does not need to be a programmer and does not need to know how remote system operations occur. That person only needs to know the desired outcomes and must construct playbooks in YAML, which is a human-readable data serialization language that is used by Ansible. An Ansible playbook is shown in the following example:

```
---
- hosts: edgeSwitches
  vars_files:
    - ../fos_passwords.yml
  gather_facts: False
  tasks:
    - name: run fos commands
      IBM-b-type_fos_command:
        switch_login: "{{switch_admin_account}}"
        switch_password: "{{switch_password}}"
        switch_address: "{{switch_ip_address}}"
        command_set:
          - command: alicreate "SampleAlias1", "10:23:45:67:76:54:32:10"
```

The three dashes at the beginning are part of the YAML specification. The `hosts` section identifies automation target switches. You can keep sensitive information in a separate file, as demonstrated by the `fos_passwords.yml` line. The name of a variable in double braces, such as `{{switch_password}}`, indicates variable substitution. The variable file that is specified in `vars_files` indicates where to find external variables.

SANnav REST API

SANnav Management Portal and SANnav Global View provide end-to-end visibility into enterprise SANs. These tools detect, analyze, and take action based on SAN behavior and performance, which helps administrators get to the root of problems faster and remediate them fully. Storage administrators can troubleshoot across the storage fabric in as little as 30 seconds. This ability is significantly faster than with any other shared storage infrastructure architectures.

The SANnav REST API provides functions that complement what is found in the FOS REST APIs. The SANnav REST API includes support for the following SANnav features:

- List the fabrics that are managed by the SANnav server.
- List the seed and principal switch data for each fabric.
- List the switches (members) in the fabric.
- Display the FCR topology information for routing topologies, such as edge-to-edge, backbone-to-edge, and edge-to-backbone.

- Configure event-forwarding.
- Acknowledge or unacknowledge events.
- Display a list of filtered events.
- Perform an inventory search.

Conclusion

SAN automation is a critical component of IT modernization and digital transformation because it enables organizations to manage storage-related operations more efficiently without increasing administrative staffing levels or expanding storage infrastructure capital expenditures. By reducing manual tasks, improving consistency, and accelerating operational processes, SAN automation helps transform network storage into a strategic business asset.

IBM b-type Gen 8 demonstrates a strong commitment to SAN automation through its advanced management capabilities, programmable interfaces, and integration with modern orchestration frameworks. This chapter introduces these automation capabilities at a high level and highlights their business and operational value. Readers who want to deploy and implement SAN automation solutions should consult the IBM b-type Gen 8 product documentation for detailed configuration instructions and examples.

IBM Storage SAN design guidelines and best practices

This section provides Fibre Channel SAN design guidance for attaching IBM Storage FlashSystem and IBM DS8000 systems to an IBM b-type Gen8 fabric. The section is prescriptive, stating the recommended design and reason for it and assuming the reader is an experienced storage or SAN administrator. Generic SAN design topics, such as fabric topologies, inter-switch links (ISLs) and trunks, and zoning fundamentals, are covered in Section 5. Virtual fabrics are covered in Section 6. This section adds product-specific design rules and recommendations that the previous sections do not cover.

FlashSystem host-to-storage path design

SAN best practices for Fibre Channel connectivity include the following guidelines:

- Use two independent fabrics to eliminate single points of failure.
- Cross-connect hosts and storage systems across both fabrics to ensure redundancy and continuous access during component failures.

A cross-connected design ensures that each host and storage controller is connected to both fabrics, eliminating single points of failure and enabling continuous access during component outages. This approach ensures that no single component failure disrupts host access to storage and is further enhanced by symmetric connectivity. For a simple illustration of a cross-connected SAN design, see the diagram below. In the figure, the IBM SAN768B-8 represents Fabric A, and the IBM SAN256B-8 represents Fabric B.




Figure: Cross-connected SAN design

Symmetric connectivity ensures that equivalent ports across controllers are connected to the same fabric, enabling predictable failover behavior, easier troubleshooting, and consistent performance during both normal operation and failure conditions.

Note: A key objective of this design is predictability, ensuring that failures result in consistent and well-understood behavior across the environment.

In the following picture, the ports on controller A and controller B are symmetrically connected.❖

Fabric Topology




Figure: Symmetric connectivity

Although SAN design becomes more complex as additional devices are introduced, the fundamental goal remains the same: eliminating single points of failure.

An IBM FlashSystem deployment consists of one FlashSystem control enclosure and two node canisters (controllers). Each host that consumes block storage from the array should maintain redundant, balanced FC connectivity to both controllers through each fabric.

The recommended host attachment pattern is summarized as follows:

- Two HBAs per host, with one connecting to SAN Fabric A and the other to SAN Fabric B, delivering fabric redundancy to the host edge.
- Controller ports are cabled in an alternating-fabric pattern. Each FlashSystem node canister has at least four FC ports and, depending on the model, can be expanded to eight, 16, or 24 ports per canister. Regardless of the port count, cabling alternates between Fabric A and Fabric B across the canister's ports, ensuring every canister is symmetrically present in both fabrics. The recommended pattern for a four-port-per-canister configuration is shown in Section 13.1.1 below. The same alternating pattern repeats for higher port counts.
- Four logical paths per LUN: The host sees its volumes through four active paths, two via Fabric A and two via Fabric B, one per canister. Even on systems with eight, 16, or 24 ports per canister, the host's path count is

determined by the number of target ports to which the volume is mapped via portsets, not by the number of physically cabled ports.

This four-path design is the baseline best practice for a single-system configuration. It tolerates the loss of any single HBA, switch, ISL, controller port, or controller without losing access to the volume. This keeps the multipath topology simple and easy for the operating system's multipath driver to manage.

The four-path design ensures continuous access under the following failure conditions:

- HBA failure: Remaining paths through the second HBA continue serving I/O.
- Fabric failure: Paths through the surviving fabric remain active.
- Controller canister failure: Paths automatically fail over to the surviving canister.
- Port failure: Alternative ports on the same fabric continue to serve traffic.

These predictable behaviors simplify troubleshooting and ensure consistent availability.

These connectivity guidelines establish a consistent and predictable foundation for FlashSystem deployments, which is implemented through the alternating-fabric cabling pattern described in the following section.

Alternating-fabric cabling pattern

The cabling between an IBM FlashSystem and the two SAN fabrics should follow a consistent, alternating pattern. Odd-numbered ports on each canister connect to Fabric A, and even-numbered ports connect to Fabric B. This convention can be reversed as long as it is applied consistently.

This approach creates a symmetric topology, ensuring that each canister has an equal presence in both fabrics. As a result, the loss of a single fabric does not isolate any controller.

When additional ports are added—for example, upgrading from four to eight ports per canister—the existing pattern can be extended without requiring any changes to the existing cabling.

The following table shows the recommended cabling pattern for a baseline four-port-per-canister configuration.

Table: Recommended alternating-fabric cabling for a 4-port FlashSystem canister

Canister	Ports to Fabric A	Ports to Fabric B
Node canister 1 (N1)	P1, P3	P2, P4
Node canister 2 (N2)	P1, P3	P2, P4

The same alternating pattern continues for 8-, 16-, and 24-port configurations: odd ports (P1, P3, P5, P7) connect to Fabric A, and even ports (P2, P4, P6, P8) connect to Fabric B. The alternating pattern provides the following practical and cumulative benefits:

1. Symmetric fabric presence: Every canister has the same number of ports in Fabric A as in Fabric B. Therefore, the loss of either fabric leaves both canisters equally reachable.
2. Predictable failure semantics: Because the cabling pattern is uniform across canisters, the same failure (for example, an even-numbered port failure or a Fabric B outage) has the same effect everywhere, making troubleshooting and runbooks simpler.
3. Compatibility with auto zoning: FlashSystem auto zoning requires that the same FC I/O port ID on every node connect to the same fabric. The alternating pattern shown in the previous table satisfies this requirement.
4. Clean expansion: When a customer adds a port expansion, such as moving from four to eight ports per canister, the existing cabling does not need to be re-engineered; the new ports simply continue the alternation.

- Support for portset design: Storage Virtualize portsets, used to group target ports for host mapping and auto zoning, can be cleanly defined along fabric lines. For example, a portsets can contain only Fabric A ports or P1 and P2 from every canister because the cabling makes those groupings physically meaningful.

Note: Although the host only sees four paths per LUN in the standard configuration, the additional cabled ports in 8, 16, or 24-port deployments are not idle. They are used to host different port sets, serve different host groups, separate workload tiers (see Section 13.4), and provide spare capacity for future host onboarding and storage partition migration. Cabling more ports than a single host's path count does not violate the four-path rule; rather, it allows storage administrators to apply the rule consistently across many hosts.

NPIV and equivalent ports

FlashSystem uses N_Port ID Virtualization (NPIV), an industry-standard technology that allows a single physical FC port to present multiple WWPNs to the fabric. NPIV enables FlashSystem's fast, fabric-level failover and is the reason alternating-fabric cabling is mandatory.

Each physical FC I/O port on a FlashSystem node presents three WWPNs, each dedicated to a specific traffic type:

- Physical WWPN (pWWPN): storage-to-storage: The underlying physical port identity. It carries storage-to-storage traffic, including Policy-Based HA, Policy-Based Replication, FlashSystem grid inter-system traffic, and external storage virtualization (backend). It does not accept host I/O.
- SCSI host vWWPN: A virtual WWPN dedicated to SCSI-FC host attachment. SCSI hosts log in here.
- NVMe host vWWPN: A virtual WWPN dedicated to NVMe-over-FC host attachment. NVMe hosts log in here.

The split is visible in `lstargetportfc`. For each FC I/O port ID the command lists three rows: the pWWPN (virtualized no, host_io_permitted no), the SCSI vWWPN, and the NVMe vWWPN (both virtualized yes, host_io_permitted yes).

Example: `lstargetportfc` showing the three WWPNs of FC I/O ports 1 and 2 on node 1

```
IBM_FlashSystem:IBM340102:> lstargetportfc
id WWPN                port_id host_io_permitted virtualized protocol fc_io_port_id
1  500507681011D278  1      no                  no          scsi        1
2  500507681015D278  1      yes                 yes          scsi        1
3  500507681019D278  1      yes                 yes          nvme        1
4  500507681012D278  2      no                  no          scsi        2
5  500507681016D278  2      yes                 yes          scsi        2
6  50050768101AD278  2      yes                 yes          nvme        2
...
```

The `host_io_permitted no` on the pWWPN row is the system enforcing the NPIV access model: hosts cannot do I/O to the physical WWPN, only to the protocol-matched virtual WWPN.

How NPIV protects host I/O during canister events

Without NPIV, a canister failure or planned maintenance shutdown causes the WWPNs of the affected canister to log out of the fabric. The hosts lose those paths and must rely on the operating system's multipath driver to detect the failure, mark the paths down, and redirect I/O to the surviving canister's WWPNs. With NPIV enabled, the host vWWPNs do not disappear when their canister goes offline. Instead, they migrate to the partner canister and log back in to the same fabric. Thus, the target WWPN that the host is zoned to remains continuously present. The host never sees its target leave the fabric because failover happens at the fabric layer, not in the host multipath stack. This eliminates a class of multipath-related outages during code upgrades and maintenance.

The equivalent-ports rule

A host vWWPN does not fail over to an arbitrary port on the surviving canister. Rather, it moves to the same FC I/O port number on the partner canister. The vWWPNs of N1 port 1 fail over to N2 port 1, the vWWPNs of N1 port 3 fail over to N2 port 3, and so on. For the migrated vWWPN to remain reachable in the fabric where the host already has paths, both equivalent ports must be cabled to the same fabric. For example, N1 port 1 and N2 port 1 are cabled to Fabric A, and N1 port 2 and N2 port 2 are cabled to Fabric B. This continues for every port number across every canister. This is the equivalent-ports rule, which makes the alternating-fabric pattern a hard requirement. The cabling in the previous table satisfies this rule by design. If equivalent ports are cabled to different fabrics, NPIV failover will break when it is needed most: when the vWWPN migrates to the partner port. However, that port is in the wrong fabric, so the host's existing zoning cannot reach it.

Implications for zoning

NPIV must be reflected in the construction of zones. These implications carry through to every zoning topic in this section:

- Zone hosts to the protocol-matched host vWWPN, never the pWWPN: SCSI hosts are zoned to the SCSI vWWPN, and NVMe-over-FC hosts are zoned to the NVMe vWWPN.
- Zone storage-to-storage traffic to the pWWPN: PB-HA, PB-R, grid inter-system, and external virtualization zones use the physical WWPN, which does not accept host I/O.
- Use WWPN zoning, not switch-port (hard) zoning: Switch-port zoning includes every WWPN logged in through the port, including the pWWPN and both vWWPNs, which breaks the NPIV access model.
- Peer zoning and auto zoning support NPIV naturally: Peer zones are WWPN-based by construction, and FlashSystem auto zoning, zones hosts to the correct host vWWPNs automatically once the FC I/O ports are cabled in compliance with the equivalent-ports rule.

Path count design

Design the number of paths for the host-to-volume connections to these targets:

- Single system: Four paths: Two paths through Fabric A (one per canister) and two through Fabric B (one per canister).
- High-availability configuration (two systems, PB-HA): Eight paths: A PB-HA pair is two independent systems, so four canisters in total. Each system contributes the same 4 paths as a standalone system (two per fabric, one to each of its canisters); two systems give $4 \times 2 = 8$ paths. 4 through Fabric A and 4 through Fabric B, one to each of the four canisters.

Do not exceed these path counts. Beyond four paths in a standard configuration, or eight paths in a high availability configuration, the host multipath layer provides no additional resiliency because all failure domains, including the HBA, fabric, controller port, canister, and enclosures are already covered.

Although additional paths are technically possible, they introduce unnecessary overhead for path probing, failover processing, and bus rescanning. Excessive path counts can increase operational complexity and often lead to slower failover times, as well as prolonged system boot and device discovery.

Note: It is normal and often beneficial to cable more than four ports per canister. These additional ports are used to support other hosts, separate workload tiers, carry replication traffic, or provide spare capacity for future growth. The path-count guideline applies to the number of paths presented to a single host for a given volume, not to the total number of physically cabled ports in the system.

This is not only a design-time rule; it is also monitored by default. IBM Storage Insights includes an Advisor that issues a warning when a host exceeds the recommended path count.

The advisory restates the rule used here: limit the number of paths managed by the host multipathing software to four, even though eight is the maximum supported.

The advisory lists the affected hosts and their measured path count.

Hosts have more than the maximum number of supported paths to a volume

Storage systemFS5200-3-89

Device typeFlashSystem 5200 - 4662

Reporting hardware4662-6H2

Learn more

Overview

Recommendations

Based on our general experience, it is best to limit the total number of paths from any host to the system. Limit the total number of paths that the multipathing software on each host is managing to four paths, even though the maximum supported is eight paths. Following these rules solves many issues with high port fan-outs, fabric state changes, and host memory management, and it improves performance.

The most important reason to limit the number of paths that are available to a host from the system is for error recovery, failover, and failback purposes. The overall time for handling errors by a host is reduced. In addition, resources within the host are greatly reduced when you remove a path from the multipathing management.

The table below shows the hosts that have more than the maximum number of logins for a single host to a single IO group.

For more details refer to the chapter "Hosts" in the Redbook "IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines".

Affected resource

Export report

id	io_group	name	path_count
0	io_grp0	VMHOST01	16
1	io_grp0	VMHOST02	16
3	io_grp0	VMHOST05	16

Close

Acknowledge

Figure: IBM Storage Insights Advisor warning for hosts exceeding the recommended path count

Fan-in ratios and noisy neighbors

Fan-in refers to the number of host initiator ports that share access to a single FlashSystem target port. Because a target port is a finite shared resource, excessive I/O activity from one host can consume a disproportionate share of the available bandwidth and processing capacity, increasing latency for all other hosts using the same port. This behavior is commonly known as the *noisy neighbor* effect. Activities such as backup operations, database checkpoints, or heavily utilized virtual machines can increase response times for unrelated business-critical workloads that share the same target port.

There is no single correct ratio. Plan fan-in by workload tier:

Table: Example fan-in classification by workload tier

Workload tier	Example workloads	Fan-in target
Tier 1: Latency Critical	OLTP databases, ERP, mission-critical applications, VDI	Conservative (for example, up to ~2:4:1)
Tier 2: General Production	Application servers, mid-tier databases, file services, general VMs	Moderate (for example, ~6:10:1)
Tier 3: Non-Critical	Dev/test, archive, reporting	Higher (for example, ~12:16+:1)

Apply these practices:

- Count and document fan-in per target port before cabling, and review it as the environment grows.
- Reserve a subset of target ports for tier-1 workloads and keep their fan-in conservative.
- Pin backup and bulk-copy hosts to dedicated ports so sustained sequential traffic does not share a port with latency-sensitive OLTP.
- Distribute initiators evenly across target ports; do not zone every host to the first two ports because they were easiest to cable.
- Monitor with IBM Storage Insights and Fabric Vision (MAPS, Flow Vision) and rebalance based on measured per-port utilization, queue depth, and latency.

High availability design

High availability is achieved through active-active architectures that ensure continuous access to data even in the event of system or site failures, with automatic failover and no manual intervention.


Policy-based high availability and policy-based replication are partnerships between two independent FlashSystem systems. Policy-based high availability provides synchronous active-active access across two sites, while policy-based replication provides asynchronous disaster recovery. Traffic between the two systems crosses the SAN and must be isolated from host I/O. The minimum requirement is that the two systems can see each other over Fibre Channel. The rules below are redundancy and performance recommendations on top of this requirement:

- For policy-based high availability configurations use at least two ports per controller for partnership traffic, with one port per fabric. A single port is a single point of failure for the partnership, and a single fabric path caps replication throughput.
- Do not share partnership ports with host attachment or external storage virtualization.
- Isolate partnership traffic with portsets and partnerfcportmask. If the same physical SAN carries both host and partnership traffic, use virtual fabrics with dedicated inter-site ISLs for the partnership fabric.

- Size each fabric for 100% of peak partnership write traffic. If one fabric fails, all partnership writes will fall onto the survivor, so sizing each fabric for only half is not redundant. Trunk the inter-site links rather than using independent ISLs.

IBM FlashSystem grid

The IBM FlashSystem grid can federate up to 32 FlashSystem or SAN Volume Controller systems into a single managed platform. The unit of workload placement is the storage partition, which groups volumes, volume groups, host objects, mappings, and replication and snapshot policies into a single entity. This entity can be migrated between member systems without host disruption. Grid management runs over IP, while the SAN carries partition migration traffic and host I/O between systems. For grid construction and partition management procedures, see *Unleashing the Power of IBM Storage FlashSystem grid*, SG24-8586.



Grid architecture overview

A FlashSystem grid has one coordinator system and up to 31 member systems. The coordinator creates the grid and handles configuration; once the grid exists, it can be managed and monitored from any member. The grid control plane is IP-based, systems are loosely clustered over the management network, so the grid itself imposes no SAN requirement. The SAN is needed only for the data plane: partition migration between members and PB-R/PB-HA replication.

The following member eligibility criteria apply:

- Single I/O-group systems (one control enclosure) with standard topology and replication layer. Systems with multiple I/O groups cannot join.
- IBM Storage Virtualize 8.7.3.x or later is recommended for all members (9.1 raises the grid limit from 8 to 32 systems).
- Supported models: FlashSystem 5200/5300/5600, 7200/7300/7600, 9100/9200/9500/ 9600, C200, and SVC SA2/SV2/SV3. FlashSystem 50x5 is not supported.
- A grid can contain mixed models and code levels.

Storage partitions

In a grid, the unit of placement and migration is the storage partition. This workload-centric object bundles an application's data with all the necessary logical configurations to serve it, including volumes, volume groups, host objects, host-to-volume mappings, replication policies, and snapshot policies. A grid can support up to 32 storage partitions. Because the partition contains both back-end data and front-end, host-facing configurations, migrating a partition between systems transfers the entire workload without disruption, eliminating the need for host reconfiguration, rezoning, or application downtime. Storage partitions serve hosts that connect over SCSI, FC, or iSCSI. NVMe-over-FC hosts continue to use traditional volumes and volume groups and are outside the partition construct.

Inter-system connectivity

The partition migration process automatically establishes a Fibre Channel partnership to copy partition data between the source and target systems. The only hard requirement is FC visibility between the systems. For redundancy and performance, follow these guidelines:

- Use at least two ports per controller for inter-system connectivity, one per fabric.
- Do not share inter-system ports with host attachment or external storage virtualization. While no migration or PB-HA/PB-R replication is running, the ports are idle, so sharing appears to work. However, during a migration, latency occurs when the shared host or backend I/O is used, and the migration runs slower. Dedicated ports keep performance predictable during migration windows.
- Use the same FC I/O port IDs for the inter-system role on every member so that zoning stays uniform, and so that every member is a valid migration source or target.
- Size inter-system ISLs for migration plus any concurrent replication, and verify utilization with Fabric Vision before adding members or partitions.

Storage partition mobility

Partition migration requires two things of the SAN design: inter-system zoning and host zoning to candidate targets:

- Inter-system partnership zones: Create one peer zone per fabric that contains the WWPNs of the dedicated inter-system ports of the grid members. When a new member joins, add its inter-system ports to the existing peer zone. Make one edit per fabric, regardless of grid size. Use storage-to-storage WWPNs in these zones. Auto zoning does not create array-to-array zones.
- Host zones to candidate targets: After migration, hosts perform input/output (I/O) on the target system. Therefore, a host must already be zoned to every member to which its partition can migrate.

Path count during migration

The host sees the source and target paths simultaneously during migration because the volumes carry the same identifiers. Thus, the multipath layer treats them as one device. With four steady-state paths per system, this results in eight paths, which is within multipath limits and the reason not to exceed four steady-state paths per system. After the commit step removes the source copy, the host returns to four paths against the target. Confirm the host path state before committing.

IBM FlashSystem auto zoning

Auto zoning enables the FlashSystem to create and maintain its own host zones on the B-type Gen8 fabric, eliminating the need for a SAN administrator to create them manually. This simplifies the initial installation of the system and the day-to-day host provisioning workflow, while also removing the need for the storage and SAN teams to hand over zoning information for every host onboarding and change.

Auto zoning integrates IBM Storage Virtualize with Brocade Target-Driven Zoning (TDZ). When a host is associated with an auto zoning-enabled FC portset, the FlashSystem instructs the fabric to create the necessary peer zones and storage ports as principals, as well as host ports as members, eliminating the need for manual configuration on the switch.

Auto zoning produces standard peer zones that behave identically to manually created peer zones.

Note: Auto zoning requires Brocade Fabric OS (FOS) 9.1.1 or later. The TDZ feature must also be enabled by the fabric administrator on switch ports to which the FlashSystem target ports are connected. TDZ does not need to be enabled on host ports. Auto zoning only takes effect after TDZ is enabled on the relevant target-side ports.

Auto zoning architecture

The principal device is the storage array target port, and the member devices are the host initiator ports. The principal device can communicate with all member devices, and each member device can communicate with the principal. However, member devices cannot communicate with one another. This arrangement provides the same level of isolation as a single-initiator zoning model, where hosts are prevented from communicating directly with other hosts. At the same time, it simplifies administration by requiring only a single peer zone for each target port rather than a separate zone for every initiator-target pair. The result is a standards-based peer zone that behaves the same as one that was manually created while being easier to manage and maintain.

Portset fabric requirements

Auto zoning only operates on user-defined FC portsets, which must present symmetric, balanced connectivity:

- Symmetric fabric connectivity is required. Every FC I/O port ID in an auto-zoning-enabled portset must connect to the same fabric on every node. For example, FC I/O port 1 from all nodes connects to Fabric A. The alternating-fabric cabling of 13.1.1 satisfies this requirement.
- Balanced across fabrics: A portset should contain an equal number of target ports from each redundant fabric.
- Host ports should be on the matching fabric. Host initiator ports must be connected to the same fabric as the target ports in the portset with which the host is associated. Auto zoning creates peer zones only between ports that share a fabric.

Target-driven zoning

Auto zoning requires IBM Storage Virtualize 9.1.x or later and Brocade Fabric OS 9.1.1 or later. TDZ must be enabled on switch ports to which the FlashSystem target ports connect, not host ports. Auto zoning takes effect only after TDZ is enabled on the target ports.

Example: Enabling TDZ on the target-side switch ports and confirming the mode

```
SWG620:FID128:admin>
SWG620:FID128:admin> portcfgtdz --show "8,9,16,17"
  Port      Mode
  =====
      8      ON
      9      ON
     16      ON
     17      ON
SWG620:FID128:admin>
```

Host identification with FDMI

Auto zoning uses the Fabric Device Management Interface (FDMI) to identify hosts based on the data that initiators register with the fabric. Rather than manually entering WWPNs, a host can be created on the FlashSystem using its FDMI host name. You can inspect the fabric's FDMI database per port to confirm what the initiator registered.

Example: Inspecting registered FDMI host data on a target-side port

```
Local HBA database contains (port 8):
50:05:07:68:13:92:45:3a
  Port attributes:
    FC4 Types: FCP
    Supported Speed: 16 32 64 Gb/s
    Port Speed: 32 Gb/s
    Port Name: 50:05:07:68:13:92:45:3a
    Port Symb Name: IBM-SV_9.1.2.1_78E4V18_1_Port66_SCSI_HOST
    FC4 Active Type: FCP
  HBA attributes:
    Manufacturer: IBM
    Model: 03JK911
    Model Description: 03JK911:IBM 4x64Gb FC PCIe HIC
    OS Name and Version: IBM-SV-V1_REPLICATION_0x204fac08a74_AutoZoning_CL
```

Auto zone naming and prefixes

Auto-created zones follow the format `IBM_SV_<target_port_WWPN>_00`. Before creating the first auto-zoning-enabled portset, you can set a site-specific prefix with the command `chsystem -autozoneprefix`. Once an auto-zoning-enabled portset exists, the prefix is locked. Set a meaningful prefix, such as the system name, to make zones easy to identify in the fabric.

Example: Setting and confirming the auto zone name prefix

```
[19:07:48] node1:~ # lssystem | grep auto_zone_prefix
auto_zone_prefix FS7600
```

With the prefix set, the zones auto zoning creates carry that prefix and contain the target WWPN together with the host WWPNs as a peer zone.

Example: Auto-created peer zones, shown with the FS7600 prefix

```

zone: FS7600_IBM_SV_50050768138244d9_00
      50:05:07:68:13:82:44:d9; 10:00:00:10:9b:0f:db:f2;
      10:00:00:62:0b:3d:06:06
zone: FS7600_IBM_SV_500507681382453a_00
      50:05:07:68:13:82:45:3a; 10:00:00:10:9b:0f:db:f2;
      10:00:00:62:0b:3d:06:06
zone: FS7600_IBM_SV_50050768139244d9_00
      50:05:07:68:13:92:44:d9; 10:00:00:10:9b:0f:db:f1;
      10:00:00:62:0b:3d:06:05
zone: FS7600_IBM_SV_500507681392453a_00
      50:05:07:68:13:92:45:3a; 10:00:00:10:9b:0f:db:f1;
      10:00:00:62:0b:3d:06:05

```

Operational considerations

Once TDZ is enabled on the switch ports, follow these onboarding steps:

1. Create an auto-zoning enabled FC portset with the appropriate policy.
2. Add auto-zoning-capable FC I/O ports to the port set. The `fabric_auto_zone_capability` field in `lspportfc` must show "yes."
3. Create the host object associated with the portset, either by FDMI host name or by host port WWPNs.
4. Monitor `lshostzone` to see if the `applied_config` moves from `zone_needed` to `zone_status zoned`. Finally, confirm that the host is online in `lshost`.

Example: Creating an auto-zoning portset and host with the one_to_one policy

```

svctask mkportset -porttype fc -autozoneenabled yes -name azps1

# Add an auto zoning capable FC I/O port (check lspportfc first)
svctask addfcportsetmember -portset azps1 -fcioportid 5

# Create the host, associated with the portset
svctask mkhost -fcwwpn 2100000E1EC36DF4 -portset azps1 \
    -protocol fcscsi -type generic -name azhost

# Verify: applied_config zone_needed -> zone_status zoned
lshostzone
# Verify: host status online
lshost

```

Zoning policies

Policies are set per portset and control how initiator WWPNs are matched to target ports:

- `one_to_one`: Each initiator is zoned with one target port per canister within its own fabric, with initiators distributed evenly across the available target ports. Use it as the standard policy where host WWPNs are visible in the fabric at host creation.
- `one_to_one_all_fabrics`: Each initiator is zoned with one target port from every fabric, even though it has a single active host path from the fabric it is connected to. Use it where initiator WWPNs are not visible on all fabrics during initial discovery: SAN-boot scenarios and NPIV WWPN configurations.

Both policies produce balanced peer zones consistent with the path-count targets. The policy is selected at portset creation with the `autozonepolicy` parameter.

Verification and troubleshooting

Confirm host zoning with `lshostzone`: the host moves from `zone_needed` to `zoned`, after which `lshost` shows it online. If a host stays at `zone_needed`, check that TDZ is enabled on the correct target-side ports that the portset is balanced across fabrics, and that the host initiator is cabled to the same fabric as the portset target ports.

Note: For new deployments using IBM Storage Virtualize 9.1.x or later and FOS 9.1.1 or later, enable TDZ on the target ports and set auto zoning as the default host zoning method. Before creating the first auto-zoning-enabled portset, set the auto zone prefix. Manually create classic zones only for inter-system replication, grid inter-system traffic, and third-party storage - auto zoning is limited to FlashSystem host-to-storage zoning.

Partnership port and zoning design

Partnership ports require a specific zoning pattern. Zone two ports from each node or canister in the local system's I/O group, and two ports from each canister in the remote system's I/O group. If dual-redundant fabrics are available, zone one port from each node across each fabric. This ensures that each fabric carries an independent path between every node pair, delivering the greatest fault tolerance. No other FC port on any node should carry a remote (partnership) zone. Confining partnership zoning to dedicated ports prevents host or back-end WWPNs from unintentionally being exposed to inter-system communication. In practice, this results in two to four partnership ports per node, usually one per redundant fabric. Visibility is limited to the I/O groups actually required on the partnered system. This zoning pattern applies identically whether the partnership carries PB-HA, asynchronous PB-R, or synchronous PB-R traffic. The zone construction does not change with the replication type; only the Fibre Channel versus IP requirement changes. For the complete partnership zoning requirements, see [Zoning for partnerships in IBM Docs](#).

IBM DS8000 host attachment overview

The IBM DS8000 is an enterprise storage system designed for IBM Z and large Open Systems environments. The current generation is the DS8000 G10 (10th generation) family. It is built on POWER9+ processors and includes the following models:

- DS8A10 (A00 rack-mount and A01 single-frame) for medium-sized environments
- DS8A50 (A05) and DS8A80 (A08) multiframe models for medium and large environments For more information about the performance and architecture of the DS8000, see the [IBM Storage DS8A00 Architecture and Implementation Guide \(SG24-8559\)](#).

Hardware concepts that shape SAN design

Three DS8000 hardware concepts drive host attachment design:

- I/O enclosures, paired left and right: I/O enclosures are numbered left (even: 0, 2, 4, 6) and right (odd: 1, 3, 5, 7). Spreading host connections across both sides is the basis of DS8000 availability, a whole-enclosure failure must not remove all of a host's paths.
- Host adapters (HAs): All ports on the G10 family are 32 Gigabit Fibre Channel (GFC), encryption-capable host adapters. Each four-port adapter port is independently configurable as FCP (Open Systems) or FICON (IBM Z), and supports up to 509 logins including host and Copy Services logins. Host port totals are model-dependent: the DS8A10 (A00/A01) provides up to 64 host ports, and the DS8A50 (A05) and DS8A80 (A08) provide up to 128 host ports.
- I/O port IDs: Each FC port has a stable, persistent WWPN and an I/O port ID of the form `I0xyz` - enclosure, HA slot, and port. The `lsioport` command lists every port with its WWPN, type (FCP or FICON), speed, I/O enclosure, and HA card; use it to document the port plan.

```
dscli> lsioport -l
```

Date/Time: June 20, 2026 11:05:30 PM TRT IBM DSCLI Version: 7.8.58.11 DS: IBM.2107-00000000										
ID	WWPN	State	Type	topo	portgrp	Speed	Frame	I/O	Enclosure	HA Card
=====										
I0000	5005076309001324	Communication established	Fibre Channel-LW SCSI-FCP	0	0	32 Gb/s	1	1		1
I0001	5005076309005324	No light detected	Fibre Channel-LW FICON	0	0	32 Gb/s	1	1		1
I0002	5005076309009324	No light detected	Fibre Channel-LW FICON	0	0	32 Gb/s	1	1		1
I0003	500507630900D324	No light detected	Fibre Channel-LW FICON	0	0	32 Gb/s	1	1		1
I0030	5005076309031324	Communication established	Fibre Channel-LW SCSI-FCP	0	0	32 Gb/s	1	1		4
I0031	5005076309035324	No light detected	Fibre Channel-LW SCSI-FCP	0	0	32 Gb/s	1	1		4
I0032	5005076309039324	No light detected	Fibre Channel-LW FICON	0	0	32 Gb/s	1	1		4
I0033	500507630903D324	No light detected	Fibre Channel-LW FICON	0	0	32 Gb/s	1	1		4

Figure: Lsioport output documenting I/O ports, WWPNs, and protocol (DSCLI)

Spreading host attachments

Spreading is the DS8000 equivalent of the IBM FlashSystem alternating-fabric discipline. Apply these rules:

- Minimum two host connections per host, to different HAs in different I/O enclosures. A single HA or single I/O enclosure must never be a single point of failure for a host.
- Balance connections across the left and right I/O enclosures: Where a host's connection count is at or below the number of I/O enclosures, assign an equal number to left (0, 2, 4, 6) and right (1, 3, 5, 7) enclosures.
- Keep paths per volume small, four is a good target: Some open systems hosts limit paths and volumes. In general, do not exceed four paths per volume to achieve the best balance of availability and performance.
- Isolate I/O-intensive workloads onto dedicated HAs: If one workload is large-block sequential or otherwise dominant, assign it its own HA cards so that it does not saturate ports shared with latency-sensitive workloads.

Access control model

The DS8000 controls host access through host attachments and volume groups rather than NPIV:

- Host connection: Identifies one host HBA by its WWPN.
- Host attachment: A set of host connections from the same host system, managed collectively. Up to 8,192 host attachments per system.
- Volume group: The LUN masking construct is the set of FB volumes that a host can access. One volume group can be assigned to multiple host connections, but one host connection can only map to one volume group. This is how volumes are shared between hosts.

A host attachment may carry a port mask that restricts the DS8000 I/O ports to which the HBA may log in. However, the preferred practice is to leave the mask open to allow access to all ports and control access in the fabric with zoning.

DS8000 Copy Services port isolation

DS8000 remote replication, including Metro Mirror, Global Mirror, Global Copy, and their combinations, carries traffic between DS8000 systems via FC I/O ports. This traffic must be isolated from host I/O. On the DS8000, isolation occurs at the I/O port and HA level. Apply the following best practices:

- Use separate I/O ports for Copy Services and host attachment: Do not use the same DS8000 I/O port for host I/O and remote replication. Shared replication ports carry sustained inter-site write traffic that interferes with host latency.
- Prefer dedicated HA pairs for Copy Services: If enough adapters are available, dedicate entire HA cards, or ideally, HA pairs, to Copy Services. This ensures that replication and host workloads do not compete for resources on the same adapter.
- Spread Copy Services connections across at least two HAs in different I/O enclosures. Replication paths require the same enclosure-level redundancy as host paths.

- Account for replication port logins: Each FC port supports up to 509 logins, including PPRC initiator and target ports. Include replication in the login budget when planning shared infrastructure.

DS8000 zoning specifics

DS8000 zoning follows the single-initiator peer-zoning practice common to all enterprise storage. There are two DS8000-specific points:

- There is no NPIV, so zones target the physical I/O port WWPN directly.
- Host-port access control belongs in the fabric and in volume groups, not in the host-attachment port mask.

Zoning

The recommended approach for host zoning is peer zoning. In a peer zone, the DS8000 I/O port WWPN is designated as the principal, and the host HBA WWPNs are designated as non-principals. Each non-principal can communicate with the principal, but non-principals cannot communicate with each other. This provides the same host isolation as single-initiator zoning—hosts never see each other—but in a single, maintainable zone object per target port rather than one zone per initiator-target pair. Fewer zone objects mean a smaller zone database and shorter change windows:

- Build one peer zone per DS8000 I/O port, designating the I/O port WWPN as the principal and the host HBAs that require it as non-principals.
- Use WWPN zoning rather than port-number zoning so that access follows the device, not the cable.
- Zone to the DS8000 physical I/O port WWPN; the I/O port WWPN from lsioport is the zoning target.
- Embed the I/O port ID in the zone alias so the physical port is identifiable from the zone name.

The older practice of single-initiator zoning, which involves one host HBA per zone with one or more DS8000 target ports, remains valid and supported. Many established environments use it. The rationale behind this method is RSCN containment. When an HBA joins or leaves the fabric, it issues a Registered State Change Notification (RSCN) that reaches every device in the same zone but does not cross zone boundaries. Therefore, a single-initiator zone confines the disruption caused by a flapping or rebooting HBA to that one host. Peer zoning achieves the same containment with far fewer zones, which is why it is now preferred.

Host access control

It is possible to restrict which DS8000 I/O ports an HBA may use through the host-attachment port mask, but the preferred practice is to leave host access open to all available DS8000 HA ports and control access in two places: the fabric (zoning, which determines which DS8000 ports the host can see) and the volume group (LUN masking, which determines which volumes the host can access). This keeps access control within the designated layers and avoids the creation of an additional, easily overlooked masking point within the host attachment.

Separate FCP and FICON

DS8000 I/O ports are individually configurable as FCP or FICON, but isolate the two protocols onto separate HA cards where possible. This keeps the Open Systems (FCP) and IBM Z (FICON) environments from sharing adapter resources and simplifies both zoning and problem determination.

FICON attachment is an IBM Z topic with its own design rules, logical paths per port, control-unit images, and IBM Z host channel planning. Zone FICON host channels to FICON-configured DS8000 ports in single-initiator zones, the same containment principle as FCP. An IBM Z host typically uses four paths to a DS8000. For FICON fabric design, see Chapter 7; for FICON performance, see SG24-8501 and SG24-8559.

zHyperLink is a separate, ultra-low-latency IBM Z connection (response time as low as 16 microseconds) that complements FICON rather than replacing it. It is point-to-point between the IBM Z server and the DS8000 and is not part of the Fibre Channel zoning model, so it requires no zoning. Plan it as a direct-attach link alongside the FICON paths, not as a fabric device.

Endpoint security: In-flight encryption

IBM Fibre Channel Endpoint Security (IFCES) encrypts data in flight on the FICON and FCP links between IBM Z and the DS8000, and authenticates the endpoints so that only authorized server and storage devices exchange data. On the G10 family the 32 GFC host adapters are encryption-capable, and IBM states the encryption runs without measurable performance impact. IFCES is a SAN-security design choice rather than a zoning rule. It does not change how zones are built but it belongs in the design plan: decide up front which IBM Z-to-DS8000 links require in-flight encryption and endpoint authentication, and confirm the host adapters and IBM Z channels are configured for it. It applies independently of the operating system or access method in use.

Appendix A - Optics and cables

IBM offers a wide range of optics and cables for IBM b-type SAN switches and directors.

When purchasing optics for IBM b-type SAN switches (Brocade OEM products), you order them using IBM part numbers and feature codes; however, the optics themselves are physically and electronically branded as Brocade products.

The following explains the relationship among IBM ordering, hardware branding, and CLI identification:

1. Ordering & Support (IBM side)

- IBM Feature Codes: Optics are ordered through IBM's catalog via IBM Feature Codes (for example, Feature Code 2132 for 128Gbps SW SFPs or specific 32G/64G/128G optic bundles).
- IBM Warranty & Support: Transceivers ordered directly from IBM are 100% qualified, validated, and covered under your IBM switch hardware maintenance contract.

2. Physical & Firmware Branding (Brocade side)

- Physical Labeling: The optical transceivers carry Brocade logos, Brocade part numbers (for example, 57-1000xxx-xx), and vendor serial numbers. You will rarely see an IBM logo on the actual SFP or QSFP body.
- Switch EEPROM (`sfps show`): When you run the `sfps show <slot/port>` command in Fabric OS (FOS), the switch reads the internal EEPROM of the optic. The output displays:
 - Vendor Name: Brocade (or BRCD)
 - Vendor PN: Brocade's internal optic part number

3. Compatibility & Third-Party Optics Rules

- FOS Qualification: Brocade Fabric OS (FOS) strictly enforces optic verification. The switch checks for genuine Brocade EEPROM signatures.
- Third-Party / Unapproved Optics: Non-Brocade or generic optics might trigger warnings in FOS and, depending on the switch generation and FOS security settings, the optic may be disabled.
- IBM Technical Support: If a SAN port experiences link errors while using third-party optics not coded for Brocade/IBM b-type switches, IBM Support might require you to replace the optic with a genuine Brocade-coded SFP before troubleshooting physical layer issues.

Transceivers

IBM offers different optics for different SAN switches and directors. Check the IBM announcement letter for a specific switch to see what is available for that specific switch. For example, see the announcement letter for the IBM Storage Networking SAN512B-8 Gen 8 director:

<https://www.ibm.com/docs/en/announcements/family-863502-storage-networking-san512b-8>

The following transceiver feature codes are available for the IBM Storage Networking SAN512B-8:

- Transceiver
 - (#2132) - SFP+, SWL, 128G, 1PK, SECURE
 - (#2133) - SFP+, SWL, 128G, 8PK, SECURE
 - (#2134) - OSFP, 100M, GEN8 ICL, 1PK
 - (#2135) - OSFP, 2KM, GEN8 ICL, 1PK
 - (#2136) - OSFP, 100M, GEN8 ICL, 4PK
 - (#2137) - OSFP, 2KM, GEN8 ICL, 4PK
 - (#2630) - SFP+, SWL, 64G, 1-PK, SECURE
 - (#2631) - SFP+, SWL, 64G, 8-PK, SECURE
 - (#2671) - SFP+, LWL, 64G, 10KM, 1-PK, SECURE
 - (#2672) - SFP+, LWL, 64G, 10KM, 8-PK, SECURE
 - (#2673) - SFP+ 64Gbps ELW 25km, SECURE

For specifications of Brocade Transceiver Modules, see the product briefs:

<https://www.broadcom.com/products/fibre-channel-networking/transceiver-modules>

To determine whether a specific transceiver is supported in a particular IBM b-type SAN switch or director, see the IBM Sales Manual or the Brocade Transceiver Support Matrix:

<https://docs.broadcom.com/doc/GA-MX-460>

IBM and Brocade product name mapping

The following table maps IBM product names to their corresponding Brocade product names:

IBM name	IBM MTM	Brocade name	Type of equipment
Gen 8 (128 Gbps) family			
IBM Storage Networking SAN512B-8	8635-F88	Brocade X8-8	SAN Director
IBM Storage Networking SAN256B-8	8635-F84	Brocade X8-4	SAN Director
IBM Storage Networking SAN56B-8	8634-F56	Brocade G820	SAN Switch
Gen 7 (64 Gbps) family			
IBM Storage Networking SAN512B-7	8961-F78	Brocade X7-8	SAN Director
IBM Storage Networking SAN256B-7	8961-F74	Brocade X7-4	SAN Director
IBM Storage Networking SAN128B-7	8969-P96/R96	Brocade G730	SAN Switch
IBM Storage Networking SAN64B-7	8960-P64/R64, 8969-P64/R64	Brocade G720	SAN Switch
IBM Storage Networking SAN24B-7	8969-P24	Brocade G710	SAN Switch
IBM Storage Networking SAN42B-R7	8969-R42	Brocade 7850	SAN Router
Gen 6 (32 Gbps) family			
IBM System Networking SAN512B-6	8961-F08	Brocade X6-8	SAN Director
IBM System Networking SAN256B-6	8961-F04	Brocade X6-4	SAN Director
IBM System Networking SAN128B-6	8960-F97/N97, 8960-F96/N96	Brocade G630 v2, Brocade G630 v1	SAN Switch
IBM System Networking SAN64B-6	8960-F65/N65, 8960-F64/N64	Brocade G620 v2, Brocade G620 v1	SAN Switch
IBM System Networking SAN24B-6	8969-F24, 8960-F24	Brocade G610	SAN Switch
IBM System Networking SAN18B-6	8960-R18	Brocade 7810	SAN Router

Fiber-optic cables

IBM offers multimode and single-mode fiber-optic cables. Check the IBM announcement letter for a specific switch to determine which cable options are available. For example, see the announcement letter for the IBM Storage Networking SAN512B-8 Gen 8 director:

<https://www.ibm.com/docs/en/announcements/family-863502-storage-networking-san512b-8>

The following cable feature codes are available for the IBM Storage Networking SAN512B-8:

- Cable
 - (#5811) - OM4 Cable LC/LC 8m
 - (#5812) - OM4 Cable LC/LC 8m - 4 Pack
 - (#5820) - OM5 Cable LC/LC 5m
 - (#5821) - OM5 Cable LC/LC 20m
 - (#5822) - OM5 Cable LC/LC 5m - 8 Pack
 - (#5823) - OM5 Cable LC/LC 20m - 8 Pack

IBM does not offer fiber-optic cabling solutions for QSFP, OSFP, and Double Density (DD) transceivers. Instead, refer to the Brocade High-Density Cabling Design Guide for recommended third-party fiber-optic cable options for OSFP, QSFP, and DD optics. For details about cable specifications and cable management, see the High-Density Cabling Design Guide:

<https://docs.broadcom.com/doc/HD-Cabling-DG>

In the Brocade High Density Cabling Design Guide Appendix D "Equipment List", you will find specific cable vendors and part numbers for QSFP to QSFP, QSFP to 4xSFP, SFP-DD to SFP+, SFP-DD to SFP-DD, SFP-DD to QSFP, OSFP to OSFP, OSFP to 8xSFP+, OSFP to 2xQSFP, and LCx2 to LCx2 connections.

Power cords

IBM also offers power cords for all different geographical locations. Check the IBM announcement letter for a specific switch to see what is available for that specific switch. For example, see the announcement letter for the IBM Storage Networking SAN512B-8 Gen 8 director:

<https://www.ibm.com/docs/en/announcements/family-863502-storage-networking-san512b-8>

The following linecord feature codes are available for the IBM Storage Networking SAN512B-8:

- Linecord
 - (#9800) - Linecord, US, CA, LA, Japan (NEMA L6-20P, IEC 60320 C19)
 - (#9802) - Linecord, EMEA (16 amp / 250 Vac w/earth pin CEE 7 / Sheet VII)
 - (#9831) - Linecord, Australia, NZ (15 amp / 250 Vac, AS/NZS 3112/2000)
 - (#9835) - Linecord, Korea (Feller Type V11G, CEE(7) V11, 16 A - 250V)
 - (#9897) - PDU Linecord, India
 - (#9899) - Linecord, India

Note: All power cords are provided by IBM and are not sourced from Broadcom-Brocade.

Abbreviations

Abbreviation	Expansion
AAA	Authentication, Authorization, and Accounting
ACL	Access Control List
AD	Active Directory
AES	Advanced Encryption Standard
AES-256 GCM	AES-256 Galois/Counter Mode
AFA	All-Flash Array
AG	Access Gateway
AI	Artificial Intelligence
ALPA	Arbitrated Loop Physical Address
APC	Angled Physical Contact
ASIC	Application-Specific Integrated Circuit
ATO	Adaptive Traffic Optimizer
BT	Brocade Trunking
BUM	Broadcast, Unknown Unicast, and Multicast
BSI	Bundesamt für Sicherheit in der Informationstechnik
CA	Certificate Authority
CC	Common Criteria
CFM	CPU Flow Mirroring
CHPID	Channel Path ID
CIDR	Classless Inter-Domain Routing
CID	Chassis Identifier
CLI	Command-Line Interface
CP	Control Processor

Abbreviation	Expansion
CPM	Configuration Policy Management
CR	Core Routing
CRC	Cyclic Redundancy Check
CRUD	Create, Read, Update, and Delete
CSR	Certificate Signing Request
CSS	Channel Subsystem
CTC	Channel-to-Control
CUP	Control Unit Port
DBR	Device-Based Routing
DCC	Device Connection Control
DHCP	Dynamic Host Configuration Protocol
DH	Diffie-Hellman
DH-CHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
DID	Domain ID
DoS	Denial of Service
DPS	Dynamic Path Selection
DR	Disaster Recovery
DWDM	Dense Wavelength Division Multiplexing
EAL2	Evaluation Assurance Level 2
EBR	Exchange-Based Routing
ECDH	Elliptic Curve Diffie-Hellman
EEPROM	Electrically Erasable Programmable Read-Only Memory
ELS	Extended Link Service

Abbreviation	Expansion
ELWL	Extended Long-Wavelength
EMI	Electromagnetic Interference
EOS	End of Support
ERP	Enterprise Resource Planning
E_Port	Expansion Port
EX_Port	Extended Expansion Port
F_Port	Fabric Port
FC	Fibre Channel
FCAP	Fibre Channel Certificate Authentication Protocol
FCP	Fibre Channel Protocol
FCoE	Fibre Channel over Ethernet
FCIP	Fibre Channel over IP
FC-NVMe	Fibre Channel NVMe
FCR	Fibre Channel Routing
FCS	Fabric Configuration Server
FDMI	Fabric Device Management Interface
FEC	Forward Error Correction
FID	Fabric ID
FiDR	FICON Dynamic Routing
FICON	Fiber Connection
FIPS	Federal Information Processing Standards
FLS	FICON Logical Switch
FMS	Fabric Management Server

Abbreviation	Expansion
FOS	Fabric OS
FPIN	Fabric Performance Impact Notification
FQDN	Fully Qualified Domain Name
FRU	Field-Replaceable Unit
FSPF	Fabric Shortest Path First
GFC	Gigabit Fibre Channel
GigE	Gigabit Ethernet
HA	High Availability
HBA	Host Bus Adapter
HCD	Hardware Configuration Definition
HIF	High Integrity Fabric
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
I/O	Input/Output
ICL	Inter-Chassis Link
ICOMP	In-flight Compression
IEC	International Electrotechnical Commission
IFE	In-flight Encryption
IOCDS	IO Configuration Data Set
IOPS	Input/Output Operations Per Second
ISL	Inter-Switch Link
IPS	IP Storage
IT	Information Technology

Abbreviation	Expansion
ITW	Invalid Transmission Word
JSON	JavaScript Object Notation
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LFM	Local Flow Mirroring
LISL	Logical Inter-Switch Link
LMS	Leighton-Micali Signature
LUN	Logical Unit Number
LWL	Long-Wavelength
M_Port	Mirror Port
MAC	Media Access Control
MAPS	Monitoring and Alerting Policy Suite
MFA	Multi-Factor Authentication
MPIO	Multipath I/O
MTM	Machine Type Model
MTU	Maximum Transmission Unit
MTTR	Mean Time to Resolution
N_Port	Node Port
NAA	Network Address Authority
NAS	Network-Attached Storage
NDcPP	Network Device Collaborative Protection Profile
NEMA	National Electrical Manufacturers Association
NETCONF	Network Configuration Protocol

Abbreviation	Expansion
NPIV	N_Port ID Virtualization
NTP	Network Time Protocol
NVMe	Non-Volatile Memory Express
OEM	Original Equipment Manufacturer
OLTP	Online Transaction Processing
OM4	Optical Multimode 4
OM5	Optical Multimode 5
OSFP	Octal Small Form-factor Pluggable
OTP	One-Time Password
OUI	Organizationally Unique Identifier
OVA	Open Virtual Appliance
PCHID	Physical Channel ID
PDU	Power Distribution Unit
PPRC	Peer-to-Peer Remote Copy
PQC	Post-Quantum Cryptography
QSFP	Quad Small Form-factor Pluggable
RADIUS	Remote Authentication Dial-In User Service
RBAC	Role-Based Access Control
REST	Representational State Transfer
RHEL	Red Hat Enterprise Linux
RMF	Resource Measurement Facility
RNID	Registered Node Identification
RSCN	Registered State Change Notification

Abbreviation	Expansion
SAN	Storage Area Network
SCC	Switch Connection Control
SCP	Secure Copy
SCSI	Small Computer System Interface
SDDD	Slow Drain Device Detection
SFP	Small Form-factor Pluggable
SFP+	Enhanced Small Form-factor Pluggable
SHA	Secure Hash Algorithm
SLA	Service Level Agreement
SMF	Single-Mode Fiber
SNMP	Simple Network Management Protocol
SSH	Secure Shell
SSL	Secure Sockets Layer
SSO	Single Sign-On
SVC	SAN Volume Controller
SWL	Short-Wavelength
TACACS+	Terminal Access Controller Access-Control System Plus
TCO	Total Cost of Ownership
TDZ	Target-Driven Zoning
TLS	Transport Layer Security
UAM	Unified Addressing Mode
URI	Uniform Resource Identifier
USB	Universal Serial Bus

Abbreviation	Expansion
USF	Unified Storage Fabric
VDI	Virtual Desktop Infrastructure
VE_Port	Virtual E_Port
VF	Virtual Fabrics
VM	Virtual Machine
VMFS	Virtual Machine File System
VMID	Virtual Machine ID
VTL	Virtual Tape Library
WWN	Worldwide Name
WWPN	World Wide Port Name
XISL	Extended Inter-Switch Link
XML	Extensible Markup Language
YAML	YAML Ain't Markup Language
YANG	YANG Data Modeling Language

Additional Resources

This document complements but does not replace IBM product documentation. For detailed API references, complete configuration options, and the latest updates, consult the following additional resources:

IBM Redbooks publications:

- [IBM b-type Gen 7 Installation, Migration, and Best Practices Guide \(SG24-8497\)](#)
- [IBM SANnav Management Portal v2.2.X Implementation Guide \(SG24-8534\)](#)

IBM documentation:

- [IBM b-type SAN Switch and Broadcom Switch Product Matrix 2](#)
- [IBM Storage Networking Fibre Channel SAN b-type Documentation](#)

Other related documentation:

- [Broadcom Fibre Channel Networking Documentation Portal](#)

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
United States of America

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.

AI Attribution

This work was primarily human-created. AI was used to make stylistic edits, such as changes to structure, wording, and clarity. AI was used to make content edits, such as changes to scope, information, and ideas. AI was prompted for its contributions, or AI assistance was enabled. AI-generated content was reviewed and approved. The following model(s) or application(s) were used: IBM Bob, Microsoft CoPilot and Google Gemini.