BROCADE
A Broadcom Company

Mainframe: Migrate Brocade® Gen 5 to Gen 7

White Paper

Table of Contents

Chapter 1: Before You Begin

1.1 Overview and Support

This paper discusses the following topics:

- Gen 5 to Gen 7 technical upgrade considerations for FICON fabrics
- New features in Brocade® Fabric OS® 9.x
- Upgrading Brocade Network Advisor to SANnav™

The most common reason to plan an upgrade is due to end-of-support (EOS) dates as products near end of life. Remember to plan the timing of your upgrade projects to ensure adequate time before the EOS dates.

Table 1: EOS Dates

Product	EOS[a]	Comments
z15 on Gen 5	October 31, 2021	Although support for these hardware platforms continues through April 30, 2025, for z15 connections, the last day of support is October 31, 2021. The replacement for the Brocade 6510 Switch is the Brocade G620 Switch or the Brocade G720 Switch. The replacements for the Brocade 8510-8 Director and Brocade 8510-4 Director are the Brocade X6-8 Director and Brocade X6-4 Director, or the Brocade X7-8 Director and Brocade X7-4 Director. Although the Brocade 7840 Extension Switch is still supported at this time, the replacement for the Brocade 7840 Extension Switch is the Brocade 7850 Extension Switch.
Gen 5	April 30, 2025	For FICON, the Brocade 8510-8 Director, Brocade 8510-4 Director, and Brocade 6510 Switch products will reach EOS on April 30, 2025. The exception is the Brocade 7840 Extension Switch, which will reach EOS on February 28, 2029. The replacements for the Brocade 8510-8 Director and Brocade 8510-4 Director are the Brocade X6-8 Director and Brocade X6-4 Director, or the Brocade X7-8 Director and Brocade X7-4 Director. The replacements for the Brocade 6510 Switch are the Brocade G620 Switch or the Brocade G720 Switch. The replacement for the Brocade 7840 Extension Switch is the Brocade 7850 Extension Switch.
Brocade Network Advisor	February 8, 2022	SANnav is the replacement product. Most customers plan the management system upgrade to coincide a few months before the end of their Brocade Network Advisor support contract.

a.) These are Brocade EOS dates. Although, these dates vary slightly by OEM, the dates do not vary by more than a few weeks.

1.2 Planning

As you prepare to migrate your environment from Gen 5 technology to Gen 7 technology, consider preparing for the transition to the upcoming Brocade Fabric OS (FOS) 9.x release. In particular, this paper recommends using the FICON Logical Switch feature available in FOS 8.1.x. This feature is optional in FOS 8.x, but it is required for FICON connectivity in FOS 9.0 and later. See the FICON Logical Switch section in this document.

The advent of the FICON logical switch is intended to simplify the process of setting up Brocade directors for FICON as well as FICON configuration and management. In earlier releases of FOS, many of the settings required manual configuration and were not validated during the process. This approach was prone to mistakes when configuring the FICON settings, which then required spending valuable time trying to determine the cause of the resulting connectivity errors. The Brocade High Integrity Fabric (HIF) feature limits manually induced errors, and the FICON logical switch was added to further minimize configuration errors by providing inherent configuration of the necessary parameters.

In addition to simplification, the FICON logical switch provides a mechanism to abstract the FICON director definition from the physical hardware, which is now capable of supporting more ports than the FICON director architecture permits. This mechanism mimics the abstraction that logical partitions (LPARs) provide on the mainframe and is a container for all FICON connectivity.

1.3 Terminology

- **Storage area network (SAN)** – SAN is one or more fabrics, connected HBAs, and storage.
- **Fabric** – Typically, a SAN consists of multiple fabrics and other components. Fabric refers to a set of interconnected Fibre Channel (FC) and FICON switches.
- **Inter-switch link (ISL)** – An ISL is the fiber connection between switches.

1.4 Gen 5 versus Gen 7

Two significant changes affect FICON environments:

- A 32-port card is no longer available
- A logical switch configured specifically as a FICON switch is required with FOS 9.0 and above

A 32-port card is no longer available in Gen 6 and Gen 7, and the control processor (CP) blades are now half-height blades that occupy the first slot in the chassis. This change affects the following:

- CHPID path mapping
- Cable diagrams and other documentation that include chassis slot numbering
- Supported SAN switch addressing modes

NOTE: As features, rather than speed, have become the more prominent reason for upgrading a SAN, switches are marketed by generation (Gen) instead of speed. Gen 5 is 16G FC capable, Gen 6 is 32G FC capable, and Gen 7 is 64G FC capable.

Although the FICON logical switch is not hardware dependent, Gen 7 switches require FOS 9.0 and above.

1.5 SAN Health®

A SAN Health® report is the best way to get a basic understanding of the existing SAN. Since SAN Health reports are confidential, remember to authorize your business partners and Brocade representative to receive a copy.

1.6 Extension

When this document was published, the current extension products supporting FICON were the Brocade SX6 Extension Blade for the Brocade X7-4 Director, Brocade X7-8 Director, and the Brocade 7850 Extension Switch. The Brocade SX6 Extension Blade is Gen 6 technology, and the Brocade 7850 Extension Switch is Gen 7 technology. The Brocade 7850 Extension Switch operates on FOS version 9.2.0 or later. Extension interoperability is supported when both platforms run the same version of FOS.

The Brocade 7810 Extension Switch does not support FICON; however, array-to-array replication uses FC, which is FCP, not FICON. To avoid FICON fabric requirements for array-to-array replication, consider moving replication to a separate and dedicated virtual fabric.

1.7 Gen 5 FOS Versions

The last code stream supported on Gen 5 is FOS 8.x. Gen 7 supports only FOS 9.x.

All SAN switches in a fabric must be no more than one major code revision apart. All switches operating with FOS 8.x and 9.x can participate in the same fabric. Switches operating with FOS 7.x cannot be included in the same fabric with FOS 9.x switches.

1.7.1 Allow/Prohibit Matrix (PDCM)

Deprecating the Allow/Prohibit Matrix, also known as the Prohibit Dynamic Connectivity Mask (PDCM), is planned for FOS 9.1. Customers who are using the PDCM should consider other methods, typically zoning, to block paths. The PDCM is being deprecated because all IBM software products that used the PDCM are no longer supported by IBM.

1.7.2 FICON Logical Switch Summary

- A FICON logical switch was optional in FOS 8.x; however, it is required in FOS 9.x.
- FICON-specific information and control are now exposed in the RESTCONF API.
- The Allow/Prohibit Matrix (also known as the PDCM) was deprecated in FOS 9.1.

1.7.3 FICON Logical Switch

Configuring Virtual Fabrics (VFs) is analogous to configuring LPARs on an IBM Z central processor complex (CPC). Mainframe customers often disable logical fabrics when configuring switches or leave VF enabled using only the default switch. For FOS 9.0 and above, VFs is required, and you will need to create a FICON logical switch.

The FICON logical switch was introduced in FOS 8.x, so customers planning to migrate to Gen 6 should plan on implementing a FICON logical switch during initial deployment rather than risking a future outage when upgrading. Upgrading is a concern because FOS 8.x will reach the end of support before FOS 9.x.

1.7.3.1 The Need for a FICON Logical Switch

Designating a logical switch as a FICON logical switch allows you to do the following:

- Improve security
- Ensure properly configured FICON SAN fabrics
- Simplify configuration

For the user, a FICON logical switch has the following advantages:

- **Address binding** – In prior versions of FOS, determining if a port address was bound was challenging for non-expert users. A bound address means that the FC address for the port cannot change. A bound address is important because the IBM Z channel subsystem builds FC addresses based on the link addresses defined in the Hardware Configuration Definition (HCD) tool.

  Although circumstances for changing the FC address assigned to a port on the switch are unlikely, should it happen, the I/O definition file (IODF) would no longer be valid, resulting in device or channel errors. This potential problem is fixed by defining a FICON logical switch.
- **Ensuring FICON switch requirements do not change** – It is possible to configure a logical switch to meet all FICON requirements and then back out certain changes after the channels have come online. A logical switch defined specifically for FICON use does not allow any changes to the logical switch that do not conform to the FICON requirements.

  Since a mainframe channel checks the proper security attributes only at login time, with a logical switch not specifically designated a FICON logical switch, the potential exists for a channel not to come online the next time it goes offline and back online. Note that login is handled by the physical channel (PCHID), not the logical channel path (CHPID). For most customers, a channel login is the result of a CPC initial machine load (IML) or channel path maintenance.
- **Address mode** – A number of addressing mode restrictions depend on the switch type; these restrictions were causing considerable field confusion and in some cases limiting options for customers. The FICON logical switch ensures that the proper address mode is used.

  The address mode is important because certain address modes use the lower byte, ALPA. Because of the way IBM Z builds FC addresses from link addresses, the lower byte must always be the same throughout the fabric.

1.7.3.2 Logical Switch Defined for FICON versus Logical Switch Configured for FICON

The FICON logical switch was introduced in FOS 8.1.0. Customers who deployed their FICON SAN fabrics with FOS 8.1.0 and above likely used this logical switch, but those who upgraded from an earlier version of FOS configured the fabric either manually or with the FICON configuration wizard. The difference between a logical switch defined for FICON and a logical switch configured for FICON is that a logical switch configured for FICON can be reconfigured such that it no longer meets the FICON requirements, whereas a logical switch defined as a FICON logical switch cannot be reconfigured to remove the FICON requirements.

1.7.3.3 Do I Need to Change My Current FICON SAN?

All existing FICON switch configurations running any FICON-supported version of FOS are supported with switches defined as FICON logical switches. All upgrade scenarios are supported, including a mix of logical switches defined as FICON and FCP. Remember that all switches in a fabric cannot be more than one major FOS version apart, so any switches on FOS 7.x must be upgraded to 8.x before any other switch in the fabric can be upgraded to 9.x.

Most FICON SANs are single-switch (single-domain) fabrics. Channel paths can be striped across SANs using a mix of Gen 5, Gen 6, and Gen 7 technology as long as they are correctly configured for FICON. The paths can be on different hardware using any version of FICON-supported FOS; they can have different addressing modes and routing methods.

When cascading, a non-VF-enabled switch can be connected to a VF-enabled logical switch as long as every switch in the fabric that has VF enabled has the same fabric ID (FID).

Using different hardware and configurations gives customers the flexibility of upgrading everything in the channel path simultaneously or during separate upgrade events. Doing so, however, should be a short-term migration strategy only because different configurations have different performance characteristics.

Making changes to an existing FICON SAN is only required when upgrading FOS to 9.0 or higher.

The following figure illustrates how different channel paths can be on different hardware.

Figure 1: Mixture of Gen 5, Gen 6, or Gen 7 Technology within a FICON Fabric

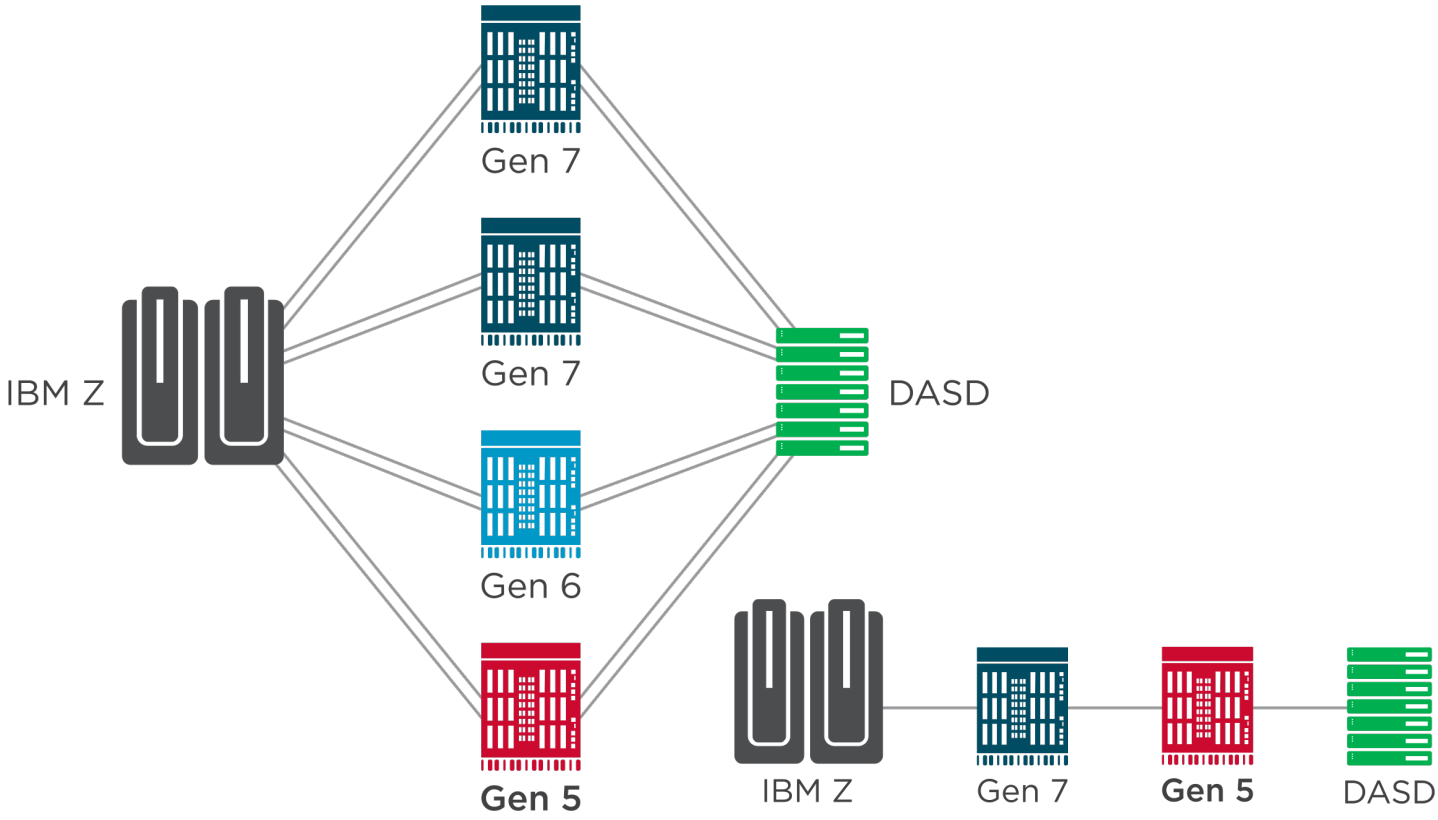

1.7.3.4 Is an Outage Required to Create a FICON Logical Switch?

Should the need arise to configure an existing logical switch to a FICON logical switch, use the following table. Fill in the *Your Environment* column to determine if an outage is required. The most common reason to change an existing switch is to upgrade the FOS code to 9.x.

Table 2: Determine If an Outage is Required

Parameter	Requirement for a Nondisruptive Conversion	Your Environment	Comments
Switch type	■ Brocade X6-8 Director or Brocade X6-4 Director ■ Brocade X7-8 Director or Brocade X7-4 Director		The Brocade 7850 Extension Switch, Brocade G620 Switch, and Brocade G720 Switch require an outage.
Logical switch (LS)	FICON LS (non-default LS)		VF must be enabled, and a LS other than the default logical switch must be properly configured for FICON.
Address mode	Zero based (mode 1)		Zero-based address mode is the addressing mode that ensures that the FC address assigned to a port is in compliance with the channel path link addressing. This mode was not always set previously because certain configurations were inherently in compliance, so you will need to check.
HIF	Must be enabled		All properly configured FICON fabrics should have HIF enabled. When HIF is enabled, insistent domain ID is enabled and the proper security policy is enabled and distributed throughout the fabric.
Port address binding	Must be in effect		All properly configured FICON fabrics should have port address binding in effect; however, checking is highly recommended.

1.7.3.5 FOS and Hardware Support for a FICON Logical Switch

The FICON logical switch was introduced in FOS 8.1.0 and is supported on all platforms. The recommended version of FOS for changing logical switches to FICON logical switches is FOS 8.2.1c or higher.

1.7.3.6 FICON Management Server (Also Known as CUP)

Enabling FICON Management Server (FMS) lets administrators define and configure the Control Unit Port (CUP) device online. All chassis-level alerts are reported to the CUP device on all logical switches, so a channel is unnecessary to dedicate to the default switch. All other actions, such as the display matrix (DM) commands and 24-7 record capturing, are limited to the logical switch where the channel with CUP is connected. If multiple logical switches are in use, you need to define CUP on at least one channel connected to each switch if CUP management is desired on more than one switch.

Other than common chassis components, each logical switch behaves as independent hardware. There is no sharing of information between one logical switch and another. When tools, such as System Automation or Disk Magic, are used to manage and monitor switches through CUP, a separate channel with CUP defined must be attached to each logical switch.

If you create a separate logical fabric for disk mirroring, a separate channel is required to monitor the replication ports. Since the CHPID is a FICON channel, the logical switch must be defined as a FICON logical switch. Defining a logical switch as FICON does not affect other protocols, such as FCP, used for mirroring. A FICON logical switch only adds security features and supports CUP.

Since TS7700 Grid uses IP, and Metro Mirror and Global Mirror use FCP, there is no need to configure a FICON logical switch. CUP is available only on FICON channels, and a dedicated FICON channel is required if CUP is desired. Since it is a FICON channel, the logical switch is configured to meet FICON requirements. If replication and FICON channel ports are in the same logical switch, isolate them with zoning.

1.8 SAN Management

Brocade Network Advisor does not support switches that are running FOS 9.0 and above. The EOS date for Brocade Network Advisor is February 8, 2022. Customers should consider upgrading to SANnav as soon as possible and must upgrade to SANnav if any of the products to be discovered are running FOS 9.0 or higher.

1.8.1 SANnav Management Portal

The SANnav Management Portal supports all switches and all versions of FOS that are supported.

Typically, customers plan their upgrade to SANnav to coincide with the end of their Brocade Network Advisor support contract. It is highly recommended to start the transition 90 days in advance. There is a free, fully functional, 90-day trial version of SANnav that customers can use for this purpose.

SANnav 2.1 or later is recommended for FICON environments.

See Chapter 4, Useful Resources, for information about SANnav training. As with any new software interface, there is a learning curve, but since SANnav is a complete redesign, there is consistency throughout and the learning curve should be minimal.

1.8.2 SANnav Global View

For customers with multiple instances of Brocade Network Advisor or new applications that require multiple management platforms, SANnav Global View provides single-pane-of-glass management for up to 20 instances of SANnav Management Portal. Previously with Brocade Network Advisor, each instance of Brocade Network Advisor was a completely separate management system.

1.8.3 Chassis-Level Management

All chassis management is performed on the default switch, which is identified by FID128. If you are new to logical fabrics for FICON, the default switch must be discoverable by SANnav in order to support chassis management and chassis-level alarms, such as a high-temperature warning.

1.8.4 Chassis Card Layout

The CP cards for Gen 7 directors were reduced to half-height size, and now both CPs fit in a single chassis slot. This change affects cable diagrams and patch layouts, which should be considered during Gen 5 to Gen 7 director migration.

Although both CPs physically occupy one chassis slot, they are still referred to as slot 1 and slot 2. High-speed optics generate more heat, and reducing the number of physical slots allows each slot more space, which improves air flow for better cooling.

Figure 2: Gen 5 versus Gen 6 or Gen 7 Card Slots

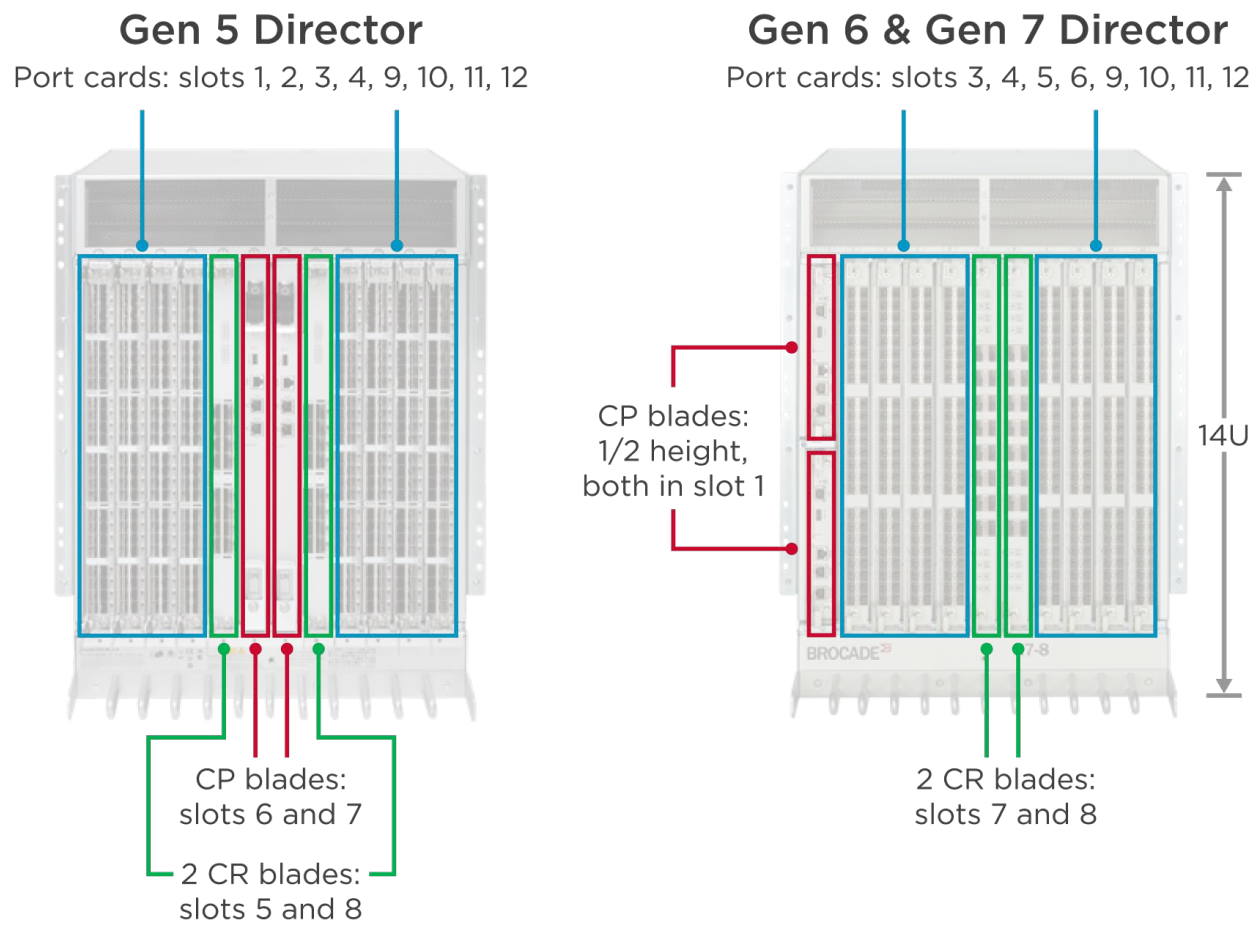

Chapter 2: HCD and Channel Path Planning

2.1 Link Address Overview

Link addressing has not changed; however, due to the change from 32-port cards to 48-port cards, link addresses land on different port cards when configurations are migrated from Gen 5 to Gen 7 directors. As a result, much of the design work is dedicated to CHPID path mapping, which requires an in-depth understanding of how link addresses relate to FC addressing. The following figure shows how the elements of the FC address map to the fields of a link address.

Figure 3: Deriving a Fibre Channel Address from a Link Address

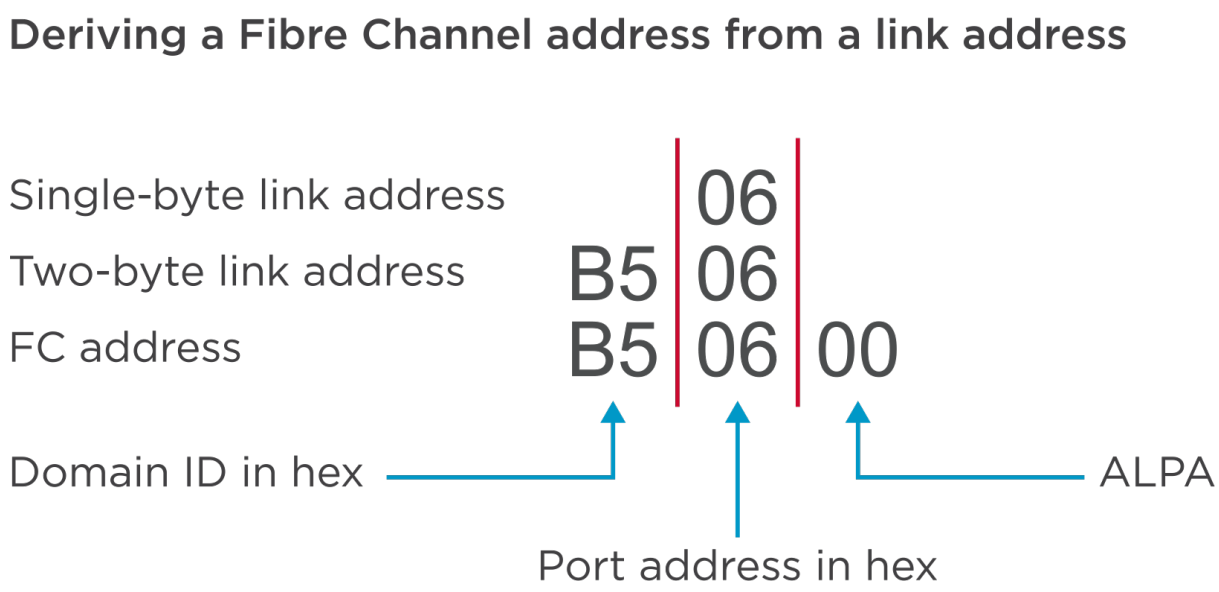

2.1.1 Domain ID

The hexadecimal domain ID maps to the most significant byte of a two-byte link address.

If the host configuration specifies single-byte link addressing, the channel assumes that the Domain ID field of the FC address is common for all port addresses configured for the FICON director; therefore, channels and control units must be connected to the same logical switch.

If the host configuration specifies two-byte link addressing, all port addresses accessed by that channel must specify two-byte link addresses. If two-byte link addresses are defined, all fabric security settings must be enabled on the switch to ensure that the high-integrity fabric requirements of the IBM Z host are met (see Fabric Security) regardless of whether cascading topologies are in use. So, for FOS 9.0 and above, and therefore for all Gen 7 switches, the channel must be connected to a FICON logical switch.

2.1.2 Port Address

The hexadecimal port address maps to the single-byte link address or the least significant byte of a two-byte link address.

2.1.3 ALPA

The hexadecimal arbitrated loop physical address (ALPA) field does not map to any part of the link address.

An arbitrated loop has never been supported by FICON and it is not supported on FC Gen 6 and later. Notably, the lowest byte of the FC address is still often referred to as the ALPA field. Today, the ALPA field is more commonly used to extend the FCP and N_Port ID Virtualization (NPIV) port address range.

The NPIV protocol provides a method for a FC device to associate multiple WWNs with the same switch port. This protocol determines which logical entity within the FC device is associated with I/O requests. This technique is used by virtual servers running on zLinux or zVM (using the FCP protocol) to share the same IBM Z host channel. Although the FICON protocol does not exploit the ALPA field, the value of the ALPA field must be the same for all FICON control units and channels and is handled by the IBM Z channel implementation.

For example, a FICON channel builds the ALPA portion of the FC address by using the same ALPA field value from the FC address allocated to the channel when it logs in to the fabric. Consequently, the ALPA portion of the FC address allocated to control units that log in to the fabric must have the same ALPA field value as the channel. In a logical switch defined for FICON, the ALPA portion of the FC address is always 00.

Other addressing modes use the upper two bits of the ALPA field to extend addressing beyond 256 ports. Although not used for FICON, zLinux can use these additional address bits.

2.2 Fabric Security

If the HIF settings are not enabled on the switch, the IBM Z host will place the channel in the offline state and indicate that the reason for the channel state is *Invalid Attach*. Once a channel is in the Invalid Attach state, the misconfiguration on the switch must be corrected, and the channel must be re-enabled to bring up the link. The PCHID (not the CHPID) is placed in the Invalid Attach state, so you cannot simply configure the CHPID offline and back online to clear the condition. To reset the condition, the PCHID state must be toggled from the System Element (SE).

Configuring two-byte link addresses is the most secure method to ensure that the HIF requirements are met. Since the channel does not check the HIF security features, two-byte link addressing is considered the best practice if single-byte addressing is specified.

Chapter 3: FICON SAN Design

3.1 Overview

Often, storage and processor upgrades are aligned with SAN upgrades. To help you make sure that all SAN requirements are well known before commencing any design planning (such as remembering to account for spare migration ports when upgrading storage and processors), the following checklist is provided for your reference.

Table 3: FICON Design Checklist

Area	Considerations
Management	Although a FICON SAN can be managed through the CLI, doing so is unusual. Nearly all FICON SANs use management software. Brocade Network Advisor does not support FOS 9.x. Therefore, Gen 7 switches are not supported by Brocade Network Advisor. SANnav Management Portal is the replacement for Brocade Network Advisor. The enterprise version is required whenever managing a director-class product or environments with more than 600 ports. Otherwise, the base version is adequate. FICON environments nearly always require the enterprise version of SANnav. Gen 7 SANnav 2.2 or higher is recommended for managing FICON fabrics. SANnav Global View is software used to manage multiple independent production sites or active-active backup sites. SANnav Global View provides a single-pane-of-glass for up to 20 instances of SANnav Management Portal. Typically, customers plan their upgrade from Brocade Network Advisor to SANnav to coincide with the end of their Brocade Network Advisor support contract. It is highly recommended to start the transition 90 days in advance.
FEC	It is recommended that FEC be enabled on 16G ports; FEC is enabled by default. The 32G FICON and faster connections inherently support FEC. TTS should only be enabled if a similarly TTS-capable and FEC enabled device is connected. Enabling or disabling FEC and TTS is disruptive to traffic.
Routing method	It is recommended that all switches in a fabric use the same routing method, so plan for an outage if the routing method needs to be changed. Do all control units in your environment support FiDR? If so, then exchange-based routing might be an effective performance enhancement for your environment. Exchange-based routing is the most efficient and resilient routing method, so if FiDR is supported on all channels and control units, exchange-based routing is recommended.
Using 48-port cards and high availability	HCD and cable plant consideration. Most FICON SANs use director-class products with 32-port blades. Furthermore, the CHPID mapping tool and other tools used to aid in channel path planning are built around 32-port cards. CHPID mapping based on 32-port cards will likely not work with 48-port cards.
Cascaded performance	HCD and cable plant consideration. Virtual channels have always been a part of the Brocade FC switching architecture. In cascaded environments, the virtual channel used for a path is determined by the port address, often referred to as the port ID or PID. Where control units are connected can affect channel I/O rate performance. Multiple zones with different quality of service can be deployed. This consideration is especially important for cascaded channel-to-channel (CTC) traffic.

Table 3: FICON Design Checklist (Continued)

Area	Considerations
Ultra-low latency	HCD and cable plant consideration. Brocade FC switches are built on ASICs that support local switching. Local switching is when the input and output are on the same ASIC. For most applications, the added design considerations to take advantage of local switching are not necessary. Low latency results in higher transaction rates. High-frequency trading applications are the most common reason to plan for channel paths to be on the same ASIC.
Naming convention	Before the availability of the FICON logical switch definition, the difference between a chassis, a switch, and a fabric was subtle and of little or no consequence to mainframe administrators. Although the technical differences mean little to a mainframe administrator, there are times when you will need to differentiate between a chassis and a logical switch, especially when creating multiple logical switches in the same chassis. Similarly, when cascading, you will need to differentiate the logical switches from the fabric. Naming conventions are preferences that vary for each organization. The only recommendation here is that you make each name unique and easily identifiable. This can be done by simply appending *_fab* at the end of fabric names and *_chassis* at the end of chassis names. Chassis, fabric, and switch names can be as long as 64 characters. Packing too much information into a name can make it difficult to easily identify the item with which the name is associated and fills up too much screen space. The SANnav database has a description field where information such as rack and aisle location can be added, so there is no need for overly complex names. Although there is no specific recommendation for naming conventions, the general advice is to keep names simple.
Domain ID and switch ID	HCD and cable plant consideration. The recommended best practice is to use the hex value of the switch's domain ID for the switch ID in HCD. This recommendation also implies that all FICON logical switches have a unique domain ID and that the switch ID always be two characters. Valid domain IDs are from 0x01 to 0xEF.
FICON Logical Fabric FID	FOS 9.2.x or later, Brocade platforms have VFs enabled and require one or more FICON LS in a FICON environment. Cascading two FICON logical switches creates a FICON logical fabric. Creating a new FICON LS on a platform requires assigning a FID; there is no default FID. On a platform, each logical switch requires a unique FID. The same FID must be used on each FICON logical switch being cascaded; if the FID is not the same, the logical fabric will not form.
Supported hardware and software	An inventory check is supported in the existing environment to ensure that all hardware and software are currently supported and to ensure the interoperability between new hardware and software being introduced to the fabric. The recommended best practice is to complete all required changes to the existing environment before commencing other upgrades so as to compartmentalize the changes.

3.2 Forward Error Correction

Forward error correction (FEC) provides a mechanism for reducing error rates during data transmissions over 16G FICON links. When FEC is enabled on a port, the sender adds systematically generated error-correcting code (ECC) to its data transmission. This mechanism allows the receiver to detect and correct errors without needing to get additional information from the sender. FEC encoding can correct a burst of up to 11-bit errors in every 2112-bit transmission. The error correction covers both frames and primitives. There is no loss of bandwidth or added transmission data rate overhead to the 16G FC link.

On 16G Brocade platforms, there is FEC itself and Transmitted Training Signal (TTS). FEC is enabled by default on Brocade 16G platforms and considered best practice. TTS is disabled by default. At 16G, FEC is optional; therefore, TTS was added to determine if both ends have FEC enabled; if not, FEC is disabled. FEC is enabled and not optional in 32G and higher FC/FICON therefore, there is no need for TTS.

TTS should only be enabled if a similarly TTS-capable and FEC enabled device is connected. TTS and encoding used on 16G FC and higher are not compatible.

- Fixed speed 16G FC, TTS is not compatible with the encoding used on 16G FC or higher at the peer port. The final port state will be *No_Sync*.
- Auto negotiation, the local port will try the TTS encoding; if the other end is not compatible, the local port will try 8-bit or 10-bit encoding (8G FC and lower). The port will settle at 8G FC port speed if the peer port supports it.

The 32G FICON and higher connections inherently support FEC. For DWDM, ISL connectivity to 32G and higher DWDM ports inherently have FEC. If the DWDM is transparent, it sends the inherent FEC. If the DWDM is protocol recognizing, it sends only the FC frame and uses its own error detection and correction.

NOTE: Enabling or disabling FEC and TTS is disruptive to traffic.

3.3 FICON Logical Fabric FID

FOS 9.2.x or later, Brocade platforms have VFs enabled and require one or more FICON LS in a FICON environment. Cascading two FICON logical switches creates a FICON logical fabric. Creating a new FICON LS on a platform requires assigning a FID; there is no default FID. On a platform, each logical switch requires a unique FID. The same FID must be used on each FICON logical switch being cascaded; if the FID is not the same, the cascaded logical fabric will not form.

3.4 Typical Redundant Pathing

The typical configuration uses 8-way pathing to the DASD spread across four directors. Although two paths are connected to the same director, the paths should be planned such that the channel and control units associated with a path do not share the same port cards. A common design approach is to plan the first path group to use the cards to the left of the core blades, which are in the center of the director, and the second path group to use the cards to the right of the core blades.

Figure 4: Typical 8-Way Pathing

Despite this separation, most FICON implementations that use directors with 32-port cards put both channel paths on the same switch. Both channel paths, therefore, use the same domain ID, so the high byte of the link addresses is the same for both paths. Only the lower byte of the link address is different.

To minimize the number of FRUs in a channel path, some organizations plan the CHPID and control unit connections to all reside on the same port card. For high-speed transaction processing that demands the lowest latency, some organizations plan the CHPID and control unit connections to be connected to ports that share the same ASIC.

3.5 Cascaded Performance and Virtual Channels

Virtual channels allow for multiple conversations over the same link virtually at the same time. Even when all traffic is the same priority, it is a good practice to balance traffic across virtual channels so that impacted devices have a minimal effect on other traffic. Traffic balancing does not have to be a scientific planning project. For most enterprises, simply striping CHPIDs and devices at a diagonal across the channel cards, rather than straight across the top or bottom of the card, is adequate. Due to the card design, the same virtual channel is used when striping across the top or bottom.

It might be desirable to use QoS zones in some cases. The two most common reasons follow:

- 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 comprised of short messages. Each message, regardless of size, utilizes at least one FC frame, and each frame requires one buffer credit; so CTC applications can starve other applications for buffer credits even though the link bandwidth is not fully utilized.
  - -Cascaded CTC is usually put into its own low QoS zone.

To aid in virtual channel planning, there is an Excel workbook available from your Brocade representative.

Figure 5: Virtual Channels

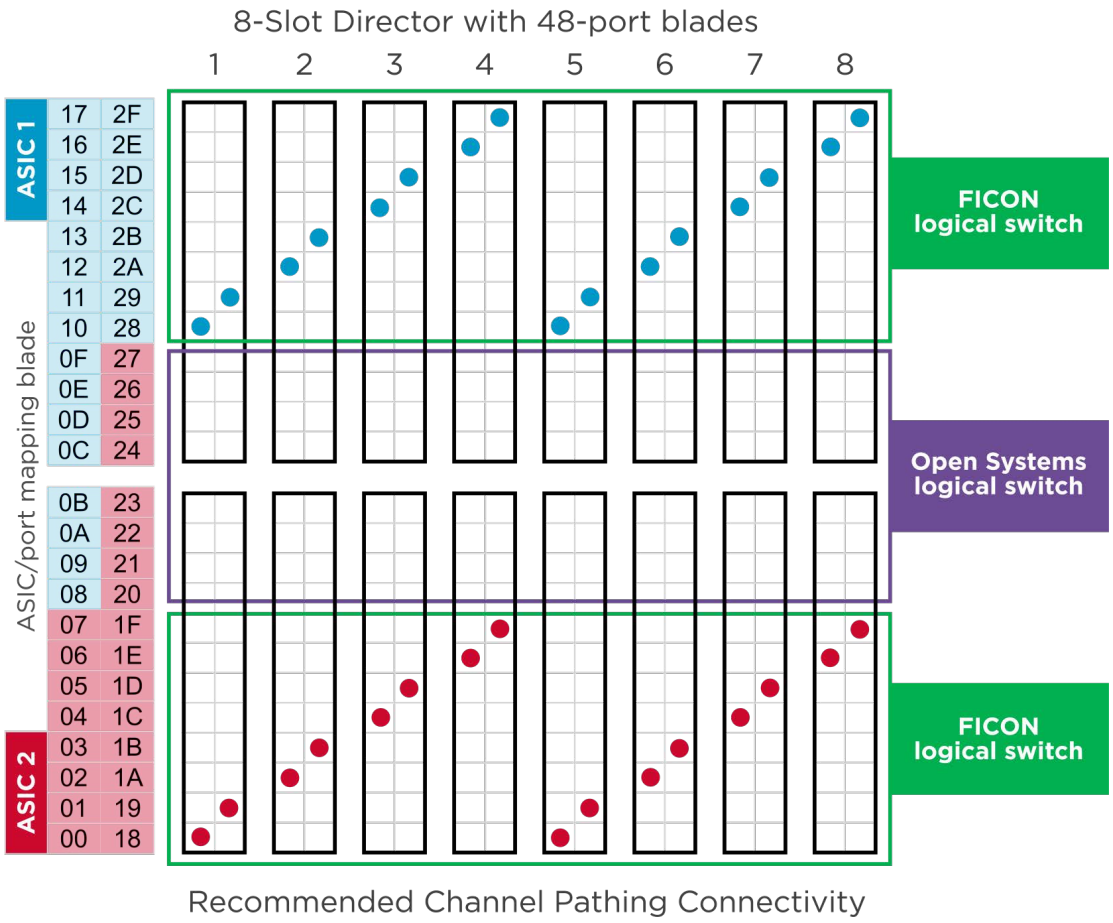

3.6 Ultra-Low Latency

Designing the configuration for ultra-low latency takes additional planning and should be done only when the application demands the highest-speed switching with the lowest latency.

Brocade ASICs are designed to use cut-through routing. As soon as the header of a frame is read, the frame is forwarded to the next port. Local switching occurs when the channel and control unit are connected to ports with the same ASIC. The switching latency within the Gen 7 ASIC is 460 ns. This latency is additive when traffic traverses multiple ASICs. All ASIC-to-ASIC switching is performed through a connection through the core blade, even if the ASICs are on the same port card. Since ASIC-to-ASIC traffic must traverse a core blade, a third ASIC is introduced in the path.

Figure 6: Local Switching versus Switching through the Backplane

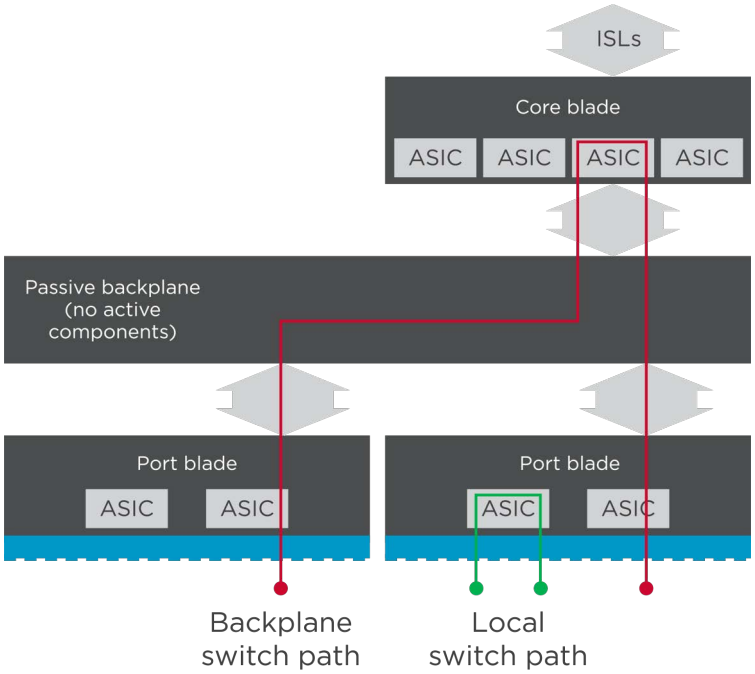

3.7 Channel Paths: 32-Port Cards versus 48-Port Cards

Because there are 48 ports per blade instead of 32, channel paths that once used the same FRU (blade) might no longer use the same FRU. Conversely, channel paths that use different FRUs might now be on the same FRU. For example, in Figure 7, on a 32-port card, CHPID 10 has an entry link address of xx00, and the control unit link addresses xx0F and xx1F. Similarly, CHPID 11 has an entry link address of xx20, and the control unit link addresses xx2F and xx3F. Each channel path is on its own FRU when connected to a 32-port card, as shown on the left in Figure 7. However, when moving to a 48-port card and assuming sequential addressing, 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.

Figure 7: 2-Port versus 48-Port Card Channel Path: Example

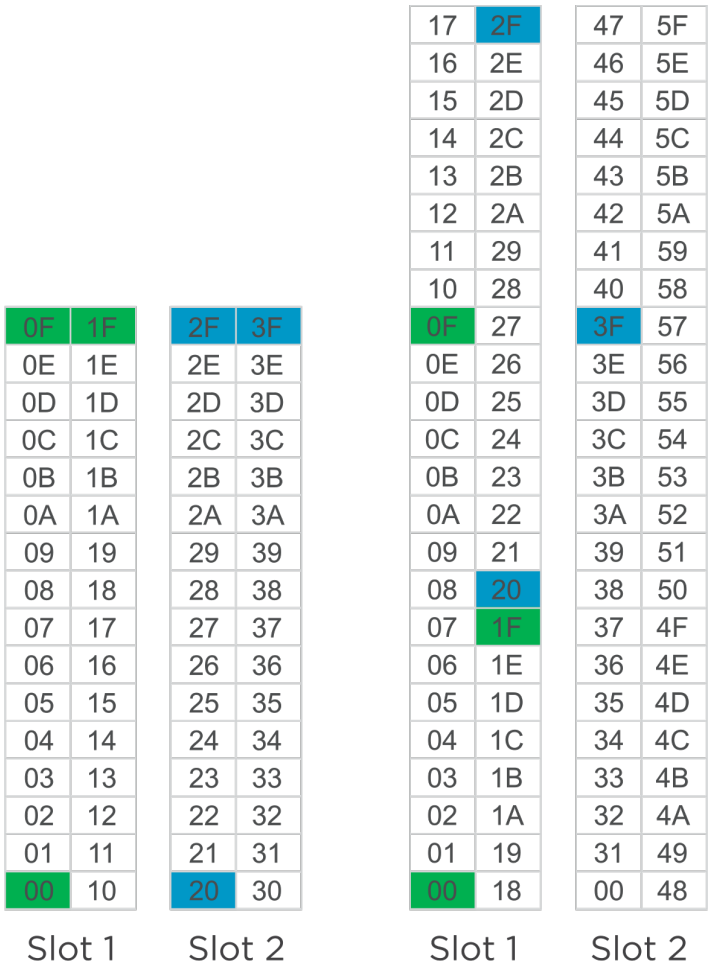

The IBM CHPID Mapping Tool has not been updated to account for 48-port blades. To aid in link address planning, there is an Excel workbook available from your Brocade representative.

3.8 Chassis Configuration Options Using 48-Port Cards

FC addresses can be assigned to port cards in any order. The only rule is that the FC addresses must be unique. Sequential addressing is typically used because it keeps documentation simple and is less susceptible to cabling and maintenance action errors.

With 48-port cards, an 8-slot director can contain as many as 384 ports, but since link addressing is limited to 256 addresses per switch by the FICON architecture, enterprises must consider how to make use of the 48-port card.

The following table shows the options for creating FICON logical switches in a chassis.

Table 4: Chassis Configuration Options with 48-Port Cards

Chassis Configuration	Discussion
8-slot chassis; one FICON logical switch	Creating a single logical switch on an 8-slot chassis is not recommended but serves as a useful discussion of how 48-port cards alter the chassis configuration. A logical switch can be created on each chassis using just the lower 16 ports on each column of ports per blade, effectively making it a 32-port card. Most customers who configure the chassis this way order some empty blades and move SFPs from the upper unused ports. Instead of leaving the upper 16 ports vacant, they can be put into another logical switch such as for mirroring or replication applications. An 8-slot chassis can also be ordered with just five cards, for a total of 240 ports, or with six cards, for a total of 288 ports, 32 of which would be left unused or used for other purposes such as mirroring.
4-slot director	The maximum port count is limited to 192 ports, but if FRUs were considered in the channel path planning, a new CHPID map is likely required.
8-slot director; two FICON logical switches per chassis	Creating two FICON logical switches per chassis allows for the greatest density. Instead of two paths per switch in each chassis, use one path per logical switch with two logical switches per chassis. Instead of 254 usable ports per chassis, there can be as many as 384 usable ports per chassis, resulting in a greater than 50% increase in port density for the same footprint. If a fully populated director is not necessary, this approach can still be used. The recommended layout is to start populating the first logical switch with port cards on the left half of the chassis and then populate the right side of the switch with port cards for the second logical switch. Configuring a chassis in this manner allows for symmetric growth. See Section 3.6.1 for more information about using two FICON logical switches per chassis.

3.8.1 Two FICON Logical Switches per Chassis

For FICON environments, when using one logical switch per chassis, the maximum number of ports per logical switch is 256 ports (254 usable port addresses). Since 256 is not evenly divisible by 48, either five port cards can be used, which yields a total of 240 ports, or six port cards can be used, which yields 288 ports. In addition to losing port count with just five port cards, using an odd number of port cards creates some planning challenges with redundant pathing. Using six port cards wastes 32 ports. In either case, two or three slots cannot be used for FICON unless there is a second FICON logical switch in the chassis.

Although partitioning a chassis into two symmetrical FICON logical switches results in a maximum logical switch size of 192 ports, splitting redundant paths such that each path is in its logical switch allows an additional 64 link addresses per path group. With two FICON logical switches per chassis, 128 additional link addresses are available.

In the FICON architecture, link addresses xxFE and xxFF are logical addresses (xxFE is reserved for the CUP, and xxFF is simply reserved). If ports are associated with those addresses, those ports must be disabled. Since there is no need to associate the link addresses xxFE and xxFF in a 192-port switch, those physical ports are also available, resulting in a slightly better than 50% port-density improvement with this approach.

Figure 8: Eight Logical Switches for the 8-Way Pathing Option to Fully Utilize All Ports

3.8.2 ISL Bandwidth Sharing with Logical Switches

There are a number of other uses for logical switches as well. Figure 9 illustrates all possible options in a single drawing (note, it is not a recommendation). The base switch allows multiple logical fabrics to share the same ISLs. When using a base switch for transport, the physical links between switches are referred to as xISLs, and the logical connections in a xISL are called Logical ISLs (LISL).

An xISL has the following limitations:

- An xISL is not visible through CUP
- System automation tools cannot be used to manage or monitor an xISL
- RMF or Disk Magic do not have access to the 74-7 records on xISL links
- Quality of Service (QoS) is not supported on an xISL
- Bandwidth planning and congestion management considerations are required
- xISL oversubscription in one logical switch impacts bandwidth in other logical switches
- Traffic Isolation zones have been depreciated

It is possible for some logical switches to use the base switch (xISL), while other logical switches are configured with their own dedicated ISLs (DISL).

Figure 9: Logical Switch Options

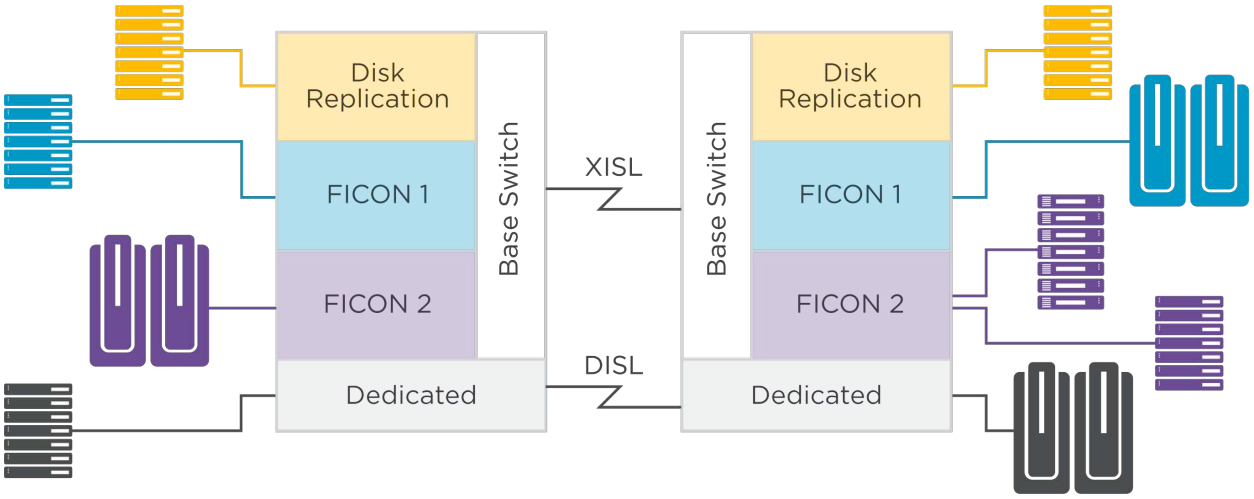

Chapter 4: Useful Resources

4.1 Customer Support Portal

All customers with a valid support contract either through an OEM or directly with Broadcom can register for access to the Customer Support Portal. In addition to traditional support-related links and documentation, there are links to the recommended education mentioned in this chapter.

https://support.broadcom.com/

4.2 Recommended Education

All education is free and can be found at the following site: https://www.broadcom.com/support/education/brocade/courses

4.2.1 SANnav

SANnav v2.0.0 Product Update Training (PUT-232)-Especially Module 10: SANnav FICON Support. This training is strongly recommended for all customers before upgrading from Brocade Network Advisor to SANnav.

Brocade Professional Mainframe Storage Network Architect (BPMSNA-300). This training is recommended for people who wish to be a subject matter expert (SME) with regard to FICON SAN fabrics.

4.2.2 Education for Administrators New to FICON

FICON Fundamentals (FICON-120) Introduction to Mainframe FICON (FICON-101)

4.3 IBM Migration Summary

The following paper has many links, detailed product support information, and other useful information associated with SAN migration.

SAN Health SAN Migration Tool

4.4 Brocade FICON Site

This site consolidates all links and mainframe material into a single location:
https://www.broadcom.com/info/brocade/mainframe-san

Revision History

MF-Migration-WP100; September 8, 2023

Initial release.