Redbooks
ibm.com/redbooks

IBM Storage Networking c-type FICON Implementation Guide

Aubrey Applewhaite

Mike Blair

Gary Fisher

Gavin O'Reilly

Lyle Ramsey

Fausto Vaninetti

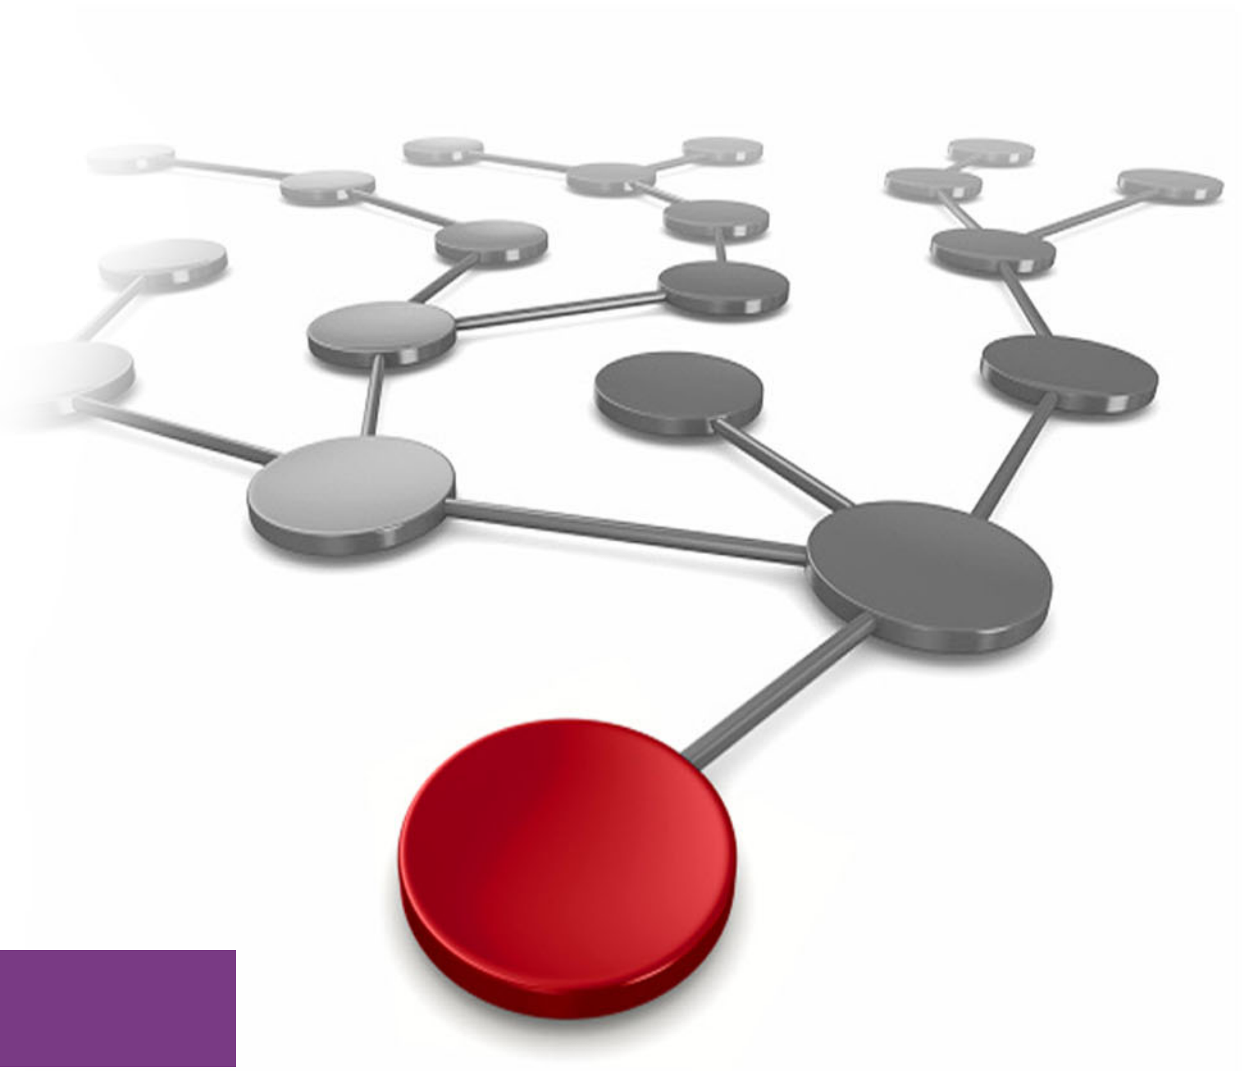

Storage

IBM

Redbooks

IBM Redbooks

IBM Storage Networking c-type FICON Implementation Guide

January 2022

SG24-8468-00

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

First Edition (January 2022)

This edition applies to IBM Storage Networking c-type Family switches and directors that are used with IBM Z platforms.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
C3®
DS8000®
Enterprise Storage Server®
FICON®
FlashCopy®
GDPS®
HyperSwap®
IBM®
IBM Cloud®
IBM FlashSystem®
IBM Z®
IBM z13®
Parallel Sysplex®
PowerPC®
Redbooks®
Redbooks (logo) ®
S/390®
System z®
Tivoli®
z/OS®
z/VM®
z13®
z15™

The following terms are trademarks of other companies:

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

Preface

The next-generation IBM® c-type Directors and switches for IBM Storage Networking provides high-speed Fibre Channel (FC) and IBM Fibre Connection (IBM FICON®) connectivity from the IBM Z® platform to the storage area network (SAN) core. It enables enterprises to rapidly deploy high-density virtualized servers with the dual benefit of higher bandwidth and consolidation.

This IBM Redbooks publication helps administrators understand how to implement or migrate to an IBM c-type SAN environment. It provides an overview of the key hardware and software products, and it explains how to install, configure, monitor, tune, and troubleshoot your SAN environment.

Authors

This book was produced by a team of specialists from around the world working at Cisco, Raleigh, North Carolina, and remotely.

Aubrey Applewhaite is an IBM Certified Consulting IT Specialist working for the Systems Lab Services team in the UK. He has worked for IBM since 1996 and has over 30 years experience in the IT industry. He has worked in many areas, including servers, operating systems (OSs), and technical support. He works as a Storage Consultant covering storage design, SAN design, and data migration. He specializes in IBM FlashSystem®, IBM SAN Volume Controller (SVC), and IBM b-type and c-type SAN hardware. He holds a Bachelor of Science degree in Sociology and Politics from Aston University.

Mike Blair is a senior technical leader at Cisco, after serving as a software engineer at IBM. He is the Lead Engineer for FICON and optical technologies for IBM Z, including testing and qualifications. He also is a Product Manager for IBM c-type SAN Switches (OEM of Cisco MDS Product Line). He acts as a IBM z/OS® Systems Programmer and Hardware Manager for the Cisco Mainframe lab and is an evangelist for Intersection of Cisco products with the IBM Mainframe platform. He holds a Bachelor of Science degree in Computer Engineering and lives in North Carolina.

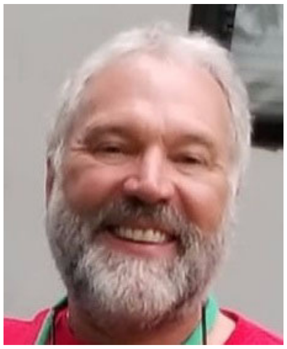

Gary Fisher has worked on many IBM software and hardware projects in computer connections for network and data transfer. Gary has received awards and co-authored patents, books, and articles, mostly for multi-system processes and automation for connectivity management. Gary received a Doctorate in Professional Studies in Computing from Pace University. Gary developed and taught several courses on mainframe networking, data communications, and storage area networking over distance. Gary is a mainframe connectivity consultant who is based in Poughkeepsie, NY, where he helps IBM customers worldwide manage interconnections between large and diverse computer installations.

Gavin O'Reilly is an IBM Storage Solutions Architect working in the IBM GTS Storage Service Line. He has 20 years of experience with IBM GTS and has worked in both storage delivery and solution roles. He works in the IBM Solution Integration Hub that is based in Dublin, Ireland. Gavin has extensive experience working with Cisco Fabric switching since 2008 with experience in mainframe FICON implementations and migration projects. He is an IBM Certified Storage subject matter expert (SME) across both midrange and enterprise storage products.

Lyle Ramsey is a veteran IT management executive and consultant, and oversees the IBM GTS Cisco SAN Architecture Global Strategy. He has more than two decades of experience leading large-scale information technology programs within private and government sectors and is a recognized SME and leader with extensive experience resolving complex issues. Lyle is a Cisco Multilayer Director Switch (MDS) and Data Center Network Manager (DCNM) SME who holds certifications from SNIA, Brocade, NetApp, EMC, Microsoft, and the Academy of Business. Lyle works closely with senior technical leaders across IBM GTS to contribute to global strategies and global best practices. Lyle is based in Phoenix, Arizona.

Fausto Vaninetti is a senior Technical Solution Architect at the European Datacenter Sales Organization for Cisco Systems. Fausto has been active in promoting SAN-related technologies since 2003, and he expanded his coverage to other domains like optical transmissions, compute, cloud, and machine learning. With his long IT experience, he has been the technical leader for large and complex data center projects with many international accounts. He contributes to product development strategies for both compute and storage networking. Fausto has authored several white papers for Cisco and is a speaker at major events. He serves as a member of the board of directors at SNIA EMEA.

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

Jon Tate
IBM Redbooks® (Retired)

Brent Anderson and Glenn Ferris
Cisco

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

IBM Storage Networking c-type family for mainframe IBM Fibre Connection environments

Welcome to the IBM Redbooks publication about the IBM Storage Networking c-type family for mainframe Fibre Connection (FICON) environments. The target audience of this book is architects, engineers, and administrators that are involved in the design, planning, deployment, and administration of FICON environments.

This document describes, explains, and shows how to deploy IBM c-type Directors in a mainframe FICON environment by following best practice processes and procedures. This book should be used by architects and administrators to support, maintain, and report on storage area network (SAN) fabrics and act as the guideline for a standard c-type mainframe deployment across companies to ensure service continuity.

The IBM c-type portfolio is based on network and storage intelligence. The switches allow for the configuration of scalable solutions that can help address the need for high performance and reliability in environments that range from small deployments to large, integrated enterprise SANs. The IBM family introduces various SAN capabilities. The family of products is designed for investment protection, flexibility, scalability, robustness, advanced diagnostics, and integration between low-cost environments and enterprise SANs.

Companies can use IBM c-type FC technology-based director and switches as resources to deliver high-performance connectivity across data center fabrics worldwide. This technology allows for scaling your SAN on demand and keeping the total cost of ownership (TCO) at a minimum.

The terms *director* and *switch* are used interchangeably. When comparing both, director generally means higher availability, larger port capacity, and more capabilities. IBM offers c-type switches and director, but unless there is a need to differentiate, within this book we simply refer to them as switches.

This book provides information about Generation 6 (GEN 6) 32 GbE c-type hardware products, software, and features; architectural design, installation, configuration, and security; and how to operationally maintain FICON environments.

Note: IBM c-type is 64 gigbit Ethernet (GbE) (GEN 7) FICON ready when using a combination of fabric-3 and supervisor-4 modules. For more information, see this white paper.

In addition, this book includes a hardware naming convention table (IBM and Cisco names) and introduces SAN technology features that are provided by the NX-OS operating system (OS).

Note: The term *FICON* represents the architecture that is defined by the International Committee for Information Technology Standards (INCITS) and published as ANSI standards. FICON is a fibre connected input/output (I/O) interface that is used to connect server systems to SANs and storage frames.

This chapter introduces the IBM Storage networking c-type range of hardware and software products in the portfolio, and it provides an overview of the hardware components and software features that are available for modern SAN data fabrics.

The following topics are covered in this chapter:

- IBM c-type hardware overview
- Enterprise SAN Directors
- IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 software licensing for NX-OS
- Extension switch model
- IBM c-type software

1.1 IBM c-type hardware overview

IBM c-type SAN products are IBM OEM products for the Cisco Multilayer Director Switch (MDS) director and switch product ranges.

The IBM Storage Networking c-type family provides storage connectivity for mission-critical applications, massive amounts of data, solid-state drives (SSDs), and cloud-based environments with a single, proven OS and a centralized management platform that enables evolutionary adoption and consistent SAN operations. Services-oriented SAN applications enable centralized solutions to meet customer needs, including data migration and acceleration of backup and replication performance between distant data centers.

The products are grouped into the following categories:

- Entry SAN switches
- Mid-range SAN switches
- Enterprise SAN Director switches
- Extension switches

For more information about IBM storage networking c-type, see IBM Storage Networking c-type family.

Table 1-1 provides a product matrix to correlate the Cisco products and models to the IBM product names and machine types and model numbers. Products that are withdrawn from marketing are not listed.

Table 1-1 IBM and Cisco product model cross-reference

Cisco product name	IBM product name	IBM machine type and model
Cisco MDS 9132T Switch	IBM Storage Networking SAN32C-6	8977 Model T32
Cisco MDS 9148T Switch	IBM Storage Networking SAN48C-6	8977 Model T48
Cisco MDS 9396T Switch	IBM Storage Networking SAN96C-6	8977 Model T96
Cisco MDS 9250i Multi-service Switch	IBM Storage Networking SAN50C-R	8977 Model R50
Cisco MDS 9706 Director	IBM Storage Networking SAN192C-6	8978 Model E04
Cisco MDS 9710 Director	IBM Storage Networking SAN384C-6	8978 Model E08
Cisco MDS 9718 Director	IBM Storage Networking SAN768C-6	8978 Model E16

> **Note:** The scope of this document is focused on switches and director that support the FICON protocol with the IBM Mainframe. The IBM Storage Networking SAN192C-6, IBM Storage Networking SAN384C-6, and IBM Storage Networking SAN50C-R switches are described.

1.2 Enterprise SAN Directors

IBM c-type Enterprise SAN Directors provide the data center networking infrastructure with enterprise solutions for the highest availability and scalability.

IBM offers the following enterprise SAN c-type Directors with support for FICON:

- IBM Storage Networking SAN192C-6
- IBM Storage Networking SAN384C-6

1.2.1 IBM Storage Networking SAN192C-6

IBM Storage Networking SAN192C-6 is a Director-class SAN switch that is designed for deployment in small to medium-sized storage networks that can support enterprise clouds and business transformation. It layers a comprehensive set of intelligent features into a high-performance, protocol-independent switch fabric.

IBM Storage Networking SAN192C-6 addresses the stringent requirements of large, virtualized data center storage environments. It delivers uncompromising availability, security, scalability, ease of management, and transparent integration of new technologies for flexible data center SAN solutions. It shares the OS and management interface with other IBM data center switches. By using the IBM Storage Networking SAN192C-6, you can transparently deploy unified fabrics with FC, FICON, and Fibre Channel over IP (FCIP) connectivity for low TCO.

For mission-critical enterprise storage networks that require secure, robust, and cost-effective business-continuance services, the FCIP extension module delivers outstanding SAN extension performance, reducing latency for disk and tape operations with FCIP acceleration features, including FCIP write acceleration and FCIP tape write and read acceleration and FICON tape acceleration.

Figure 1-1 shows the IBM Storage Networking SAN192C-6 switch.

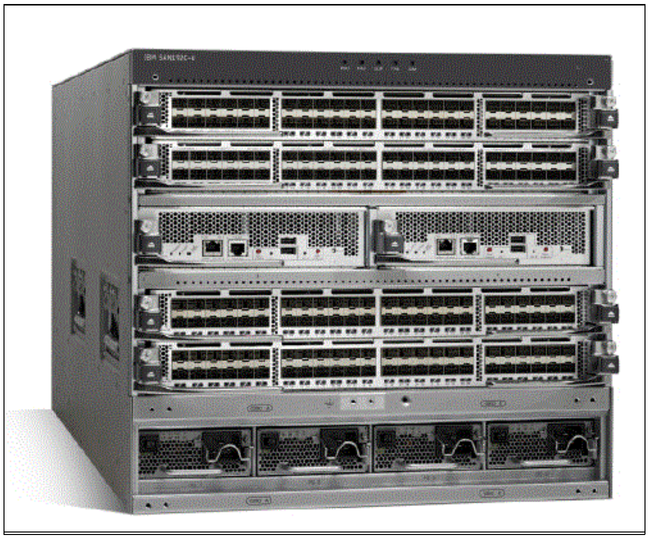

Figure 1-1 IBM Storage Networking SAN192C-6 switch

Product highlights

IBM Storage Networking SAN192C-6 offers several important features, which are described in this section.

Lower TCO with SAN consolidation

Organizations need efficient, cost-effective SANs to keep up with today's exponential data growth. IBM Storage Networking SAN192C-6 lets you easily consolidate data assets into fewer, larger, and more manageable SANs to reduce the hardware footprint and the associated capital and operational expenses. It offers industry-leading scalability with the following features:

- Up to 192 32-Gbps ports per chassis.
- Up to 12 Tbps front-panel, FC, line-rate, and non-blocking system-level switching capacity.
- Enables large and scalable deployment of SAN extension solutions with the IBM Storage Networking c-type family Directors 24/10-port SAN Extension Module.
- Exceptional capabilities with intelligent fabric services.
- Virtual storage area networks (VSANs) consolidate individual physical SAN islands while maintaining logical boundaries.
- Inter-VSAN Routing (IVR) share resources across VSANs.

Scalable expansion with outstanding investment protection

IBM Storage Networking SAN192C-6 is designed to make optimal use of valuable data center floor space. It is 15.6 inches tall (9RU) and allows up to four IBM Storage Networking SAN192C-6 Directors per standard 7-foot rack (42RU). A smaller footprint makes it an excellent candidate for deployment in smaller storage networks and pod-based converged data center infrastructure solutions for the cloud.

By using IBM Storage Networking c-type family Directors' switching modules, the IBM Storage Networking SAN192C-6 supports up to 192 ports in a 6-slot modular chassis, with up to 768 ports in a single rack. FC ports can be configured and auto-negotiated at 2/4/8-Gbps, 4/8/16-Gbps, or 8/16/32-Gbps speeds, depending on the optics and switching module selected.

IBM Storage Networking SAN192C-6 supports the same FC switching modules as the IBM Storage Networking SAN384C-6 and IBM Storage Networking SAN768C-6 Directors for a high degree of system commonality. Designed to grow with your storage environment, IBM Storage Networking SAN192C-6 provides smooth migration, common sparing, and outstanding investment protection.

The 24/10-Port SAN Extension Module is supported on IBM Storage Networking c-type family Directors. With 24 line-rate 2-, 4-, 8-, 10-, and 16-Gbps FC ports, and eight 1- and 10-GbE FCIP ports, this module enables large and scalable deployment of SAN extension solutions.

Enterprise-class availability

The IBM Storage Networking SAN192C-6 is designed for high availability (HA). In addition to meeting the basic requirements of non-disruptive software upgrades and redundancy of all critical hardware components, the IBM Storage Networking SAN192C-6 software architecture offers outstanding availability. It provides redundancy on all major hardware components, including the supervisors, fabric modules, and power supplies. The Supervisor Module automatically restarts failed processes, which makes the IBM Storage Networking SAN192C-6 exceptionally robust. In the rare event that a supervisor module is reset, complete synchronization between the active and standby supervisor modules helps ensure stateful failover with no disruption of traffic. Redundancy details are shown in Table 1-2.

Table 1-2 Redundancy details for IBM Storage Networking SAN192C-6

Item	Redundancy
Supervisor	1+1
Power supplies	Grid redundancy
Fabrics	N+1 redundancy

HA is implemented at the fabric level by using robust and high-performance Inter-Switch Links (ISLs). A port channel allows users to aggregate up to 16 physical links into one logical bundle. The bundle can consist of any speed-matched ports in the chassis, which helps ensure that the bundle can remain active if a port, ASIC, or module fails. ISLs in a port channel can have different lengths.

Business transformation with enterprise cloud deployment

Enterprise clouds provide organizations with elastic computing and network capabilities, which enable IT to scale up or down resources as needed in a quick and cost-efficient manner. IBM Storage Networking SAN192C-6 provides industry-leading scalability and the following features for enterprise cloud deployments:

- Pay-as-you-grow flexibility to meet the scalability needs in the cloud
- Robust security for multitenant cloud applications
- Predictable performance to meet stringent service-level agreements (SLAs)
- Resilient connectivity for an always-on cloud infrastructure
- Advanced traffic management capabilities, such as quality of service (QoS), to allocate network capabilities to cloud applications rapidly and cost-efficiently

Furthermore, Data Center Network Manager (DCNM) provides resource monitoring and capacity planning on a per-virtual machine (VM) basis. You can federate up to 10 DCNM servers to easily manage large clouds. Resource-use information can be delivered through Storage Management Initiative Specification (SMI-S)-based developer APIs to deliver IT as a service.

FCIP for remote SAN extension

FCIP for remote SAN extension simplifies data-protection and business-continuance strategies by enabling backup, remote replication, and other disaster recovery (DR) services over wide area network (WAN) distances by using open standards FCIP tunneling. The extension optimizes WAN resources for backup and replication by enabling hardware-based compression, hardware-based encryption, FCIP write acceleration, and tape read and write acceleration for both FCIP and FICON over IP. The SAN extension module supports four tunnels per interface and can scale up to 32 tunnels (four tunnels x 8 1/10 GbE ports).

Comprehensive solution for robust security

IBM Storage Networking SAN192C-6 offers an extensive security framework to protect highly sensitive data crossing today's enterprise storage networks. It employs intelligent packet inspection at the port level, including the application of access control lists (ACLs) for hardware enforcement of zones, VSANs, and advanced port-security features. It also uses Fibre Channel Security Protocol (FC-SP) and TrustSec Fibre Channel Link Encryption mechanisms to provide comprehensive security for storage networks.

IBM Storage Networking SAN192C-6 product specifications

The product specifications for IBM Storage Networking SAN192C-6 are described in this section. IBM Storage Networking SAN192C-6 supports FC connectivity for servers and storage. The IBM Storage Networking SAN192C-6 model supports all the FICON features and functions that are listed in this section, and requires NX-OS 8.1(1b) or later.

Table 1-3 summarizes the IBM Storage Networking SAN192C-6 product specifications.

Table 1-3 Product specifications for IBM Storage Networking SAN192C-6

Feature	Description
Product compatibility	IBM Storage Networking c-type.
Software compatibility	NX-OS Software Release 8.1(1b) or later.
OSs	For the most current and complete information, see IBM System Storage Interoperation Center (SSIC).

Feature	Description
Optional features	► 24/10-port SAN Extension Module (#AJL5) ► 48-Port 32-Gbps FC Switching Module (#AJL2) ► 48-Port 32-Gbps FC Switching Module Bundle (#AJL4) ► Fabric-1 Switching Module (#AJK9) ► Fabric-3 Switching Module (#AJN9) ► Sup-4 supervisor (#AJN6) ► Enterprise Package (#AJJ9) ► DCNM SAN Advanced Edition (#AJJA) ► Mainframe Package (#AJJB) ► SAN Insights (#AKJV) ► Small Form-factor Pluggables (SFPs) ► Fans
Indicators	► Power supply LED ► Fan LED ► Supervisor LED ► Fabric LED ► Line-card module LED
Chassis slot configuration	► Line-card slots: 4 ► Supervisor slots: 2 ► Crossbar switching fabric slots: 6 ► Fan trays: Three fan trays at the back of the chassis ► Power supply bays: 4
Performance and scalability	► Up to 12 Tbps front-panel FC switching bandwidth ► Supported FC port speeds: – 2/4/8-Gbps autosensing, optionally configurable – 4/8/16-Gbps autosensing, optionally configurable – 8/16/32-Gbps autosensing, optionally configurable ► Buffer credits: 48-port line-rate 32-Gbps FC module: – Default credits per port: 500 – With Enterprise license ► 8300 shared among a single port group of 16 ports ► 8191 maximum credits per port ► Ports per chassis: Up to 192 ports, which can be FC (4/8/16/32-Gbps) ► Ports per rack: Up to 768 FC (4/8/16/32-Gbps) ► Port channel: Up to 16 ports (the port channel can span any speed-matched port on any module in the chassis)

Feature	Description
Features and functions	
Fabric services	► Name server ► Registered State Change Notification (RSCN) ► Login services ► Fabric Configuration Server (FCS) ► Broadcast ► In-order delivery
Advanced functions	► VSAN ► IVR ► Port channel with multipath load-balancing ► QoS: flow-based and zone-based ► N_Port ID Virtualization (NPIV)
Diagnostic and troubleshooting tools	► Power-on self-test (POST) diagnostic tests ► Online diagnostic tests ► Internal port loopbacks ► Switched Port Analyzer (SPAN) and Remote Switched Port Analyzer (RSPAN) ► FC Traceroute ► FC Ping ► FC Debug ► IBM Fabric Analyzer ► Syslog ► Online system health ► Port-level statistics ► Real-Time Protocol (RTP) Debug
Network security	► VSANs ► ACLs ► Per-VSAN role-based access control (RBAC) ► FC zoning ► N-Port worldwide name (WWN) ► N-port FC-ID ► Fx-port WWN ► Fx-port WWN and interface index ► Fx-port domain ID and interface index ► Fx-port domain ID and port number ► FC-SP ► Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) switch-switch authentication ► DH-CHAP host-switch authentication ► Port security and fabric binding ► Management access ► SSHv2 implementing Advanced Encryption Standard (AES) ► Simple Network Management Protocol Version 3 (SNMPv3) implementing AES ► Secure File Transfer Protocol (SFTP) ► TrustSec Fibre Channel Link Encryption

Feature	Description
IBM FICON	► FC-SB-6 compliant ► Cascaded FICON fabrics ► Intermix of FICON and FC Fibre Channel Protocol (FCP) traffic ► FICON IBM Control Unit Port (CUP) management interface ► Exchange-based-routing ready ► FICON multihop
Serviceability	► Configuration file management ► Nondisruptive software upgrades for FC interfaces ► IBM Call Home ► Power-management LEDs ► Port beaconing ► System LEDs ► SNMP traps for alerts ► Network boot
Reliability and availability	► Online, nondisruptive software upgrades ► Stateful nondisruptive supervisor module failover ► Hot-swappable redundant supervisor modules ► Hot-swappable redundant fabric modules ► Hot-swappable 2N redundant power ► Hot-swappable fan trays with integrated temperature and power management ► Hot-swappable Enhanced SFP+ optics (8/16/32-Gbps FC and 10-GbE) ► Hot-swappable switching modules ► Stateful process restart ► Any module, any port configuration for port channels ► Fabric-based multipathing ► Per-VSAN fabric services ► Online diagnostic tests ► Port tracking ► Virtual Routing Redundancy Protocol (VRRP) for management

Feature	Description
Network management	► Access methods through the Supervisor-1 or Supervisor-4 Module: – Out-of-band 10/100/1000 Ethernet port – RS-232 serial console port – In-band IP over FC ► Access methods through the FC switching module In-band FICON CUP over FC ► Access protocols: – Command-line interface (CLI) for console and Ethernet ports – SNMPv3 that uses Ethernet port and in-band IP over FC access – FICON CUP ► Distributed Device Alias service ► Network security: – Per-VSAN RBAC that uses Remote Authentication Dial-In User Service (RADIUS)-based Terminal Access Controller Access Control System Plus (TACACS+)-based authentication, authorization, and accounting (AAA) functions – SFTP – SSHv2 implementing AES – SNMPv3 implementing AES – Management applications ► CLI ► DCNM GUI
Programming interface	► Scriptable CLI ► DCNM web services API ► NX-API

IBM Storage Networking SAN192C-6 physical specifications

Table 1-4 details the specific requirements for planning the installation of the devices in the data center rack. For more information, see IBM Storage Networking SAN192C-6, SAN384C-6, and SAN768C-6, Installation, Service, and User Guide.

Table 1-4 Physical specifications

Item	Description
Power and cooling	► Power supplies (3000 W AC). ► Input: 100 - 240 V AC nominal (±10% for full range), 16 A nominal; 50 - 60 Hz nominal (±3 Hz for full range). ► Output: 1451 W 50 V ±4%/28 A, 3.4 V ±4%/15 A (100 - 120 V AC input), 3051 W 50 V ±4%/60 A, and 3.4 V ±-4%/15 A (200 - 240 V AC input). ► Airflow: Front-to-back.
Power consumption (typical)	► Ports 96 - Watts 800. ► Port 192 - Watts 1490.

Item	Description
Environmental	► Temperature, ambient operating: 0 - 40°C (32 - 104°F). ► Temperature, ambient nonoperating and storage: -40 - 70°C (-40 -158°F). ► Relative humidity, ambient (noncondensing) operating: 10 - 90%. ► Relative humidity, ambient (noncondensing) nonoperating and storage: 10 - 95%. ► Altitude, operating: -60 - 2000 m (-197 - 6500 ft.).
Physical dimensions (H x W x D)	► Chassis dimensions (9RU): 39.62 x 43.9 x 81.3 cm (15.6 x 17.3 x 32 in.). ► Chassis depth including cable management and chassis doors is 96.52 cm (38 in.). ► Unit is rack-mountable in a standard 48.26 cm (19-inch) Electronic Industries Alliance (EIA) rack. Unit is also 2-post rack-mountable.
Weight	► Chassis only: 145 lb (65.8 kg). ► Fully configured: 325 lb. ► 48-port 16-Gbps FC line card: 17 lb (7.71 kg). ► 48-port 32-Gbps FC line card: 17.5 lb (7.94 kg). ► Power supply (3000W AC): 6 lb (2.7 kg). ► Fabric-1 module: 11 lb (5.0 kg). ► Fabric-3 module: 11 lb (5.0 kg) ► Supervisor-1 module: 7 lb (3.2 kg). ► Supervisor-4 module: 7 lb (3.2 kg) ► Fan tray: 8.5 lb (3.86 kg). ► Supervisor blank cover: 1.25 lb (0.57 kg). ► Line-card blank cover: 4.5 lb (2.04 kg).

Rack and cabinet requirements for the IBM Storage Networking SAN192C-6 chassis

The IBM Storage Networking SAN192C-6 chassis has the following rack and cabinet options:

- Standard perforated-doors cabinets
- Solid-walled cabinets with a roof fan module (bottom to top cooling)
- Standard open 4-post telco racks
- Standard open 2-post telco racks

Use a standard 19 inch, 4-post EIA cabinet or rack with mounting rails that conform to the imperial universal hole spacing, per section 1 of the NSI/EIA-310-D-1992 standard.

The depth of a 4-post rack or a cabinet must be 24 - 32 inches (61.0 - 81.3 cm) between the front and rear mounting vertical rails.

Ensure that the airflow and cooling are adequate and there is sufficient clearance around the air vents on the switch.

The rack must have sufficient vertical clearance for the chassis, 2 RU for the shelf brackets, and any needed clearance for the installation process.

The front and rear doors of enclosed racks must have at least 60% of an open area perforation pattern.

Additionally, you must consider the following site requirements for the rack:

- Power receptacles must be within reach of the power cords that are used with the switch.
- AC power supplies: Power cords for 3-kW AC power supplies are 8 - 12 feet (2.5 - 4.3 m).
- DC power supplies (ask IBM representatives for this option): Power cords for 3.0-kW DC power supplies are supplied and set by the customer.
- HVAC/HVDC power supplies (ask IBM representatives for this option): Power cords for 3.5-kW HVAC/HVDC power supplies are 14 feet (4.26 m) long.
- Where necessary, a seismic rating of Network Equipment Building Standards (NEBS) Zone 3 or Zone 4, per GR-63-CORE.

To correctly install the switch in a cabinet in a hot-aisle/cold-aisle environment, you should fit the cabinet with baffles to prevent exhaust air from recirculating into the chassis air intake. Work with your cabinet vendors to determine which of their cabinets meet the following requirements or see IBM Support for recommendations:

- The height of the rack or cabinet must accommodate the 9 RU (15.75 inches (40.0 cm)) height of the switch and its bottom support bracket. The bottom support bracket is part of the accessory kit for the switch.
- Minimum gross load rating of 2000 lb (907.2 kg) (static load rating) if supporting four switches.

1.2.2 IBM Storage Networking SAN384C-6

IBM Storage Networking SAN384C-6 is a Director-class SAN switch that is designed for deployment in large-scale storage networks to enable enterprise clouds and business transformation by adding enterprise connectivity options that support FICON.

IBM Storage Networking SAN384C-6 delivers a high-performing and reliable FICON infrastructure that is designed to support fast and scalable IBM Z servers.

Layering a comprehensive set of intelligent features onto a high-performance, protocol-independent switch fabric, IBM Storage Networking SAN384C-6 addresses the stringent requirements of large virtualized data center storage environments: HA, security, scalability, ease of management, and transparent integration of new technologies for flexible data center SAN solutions. Sharing the OS and management interface with other data center switches, IBM Storage Networking SAN384C-6 enables seamless deployment of fabrics with high-performance FC, FICON, and FCIP connectivity to achieve low TCO.

For mission-critical enterprise storage networks that require secure, robust, and cost-effective business-continuance services, the FCIP extension module is designed to deliver outstanding SAN extension performance, reducing latency for disk and tape operations with FCIP acceleration features, including FCIP write acceleration and FCIP tape write and read acceleration.

The IBM Storage Networking SAN384C-6 is shown in Figure 1-2 on page 13.

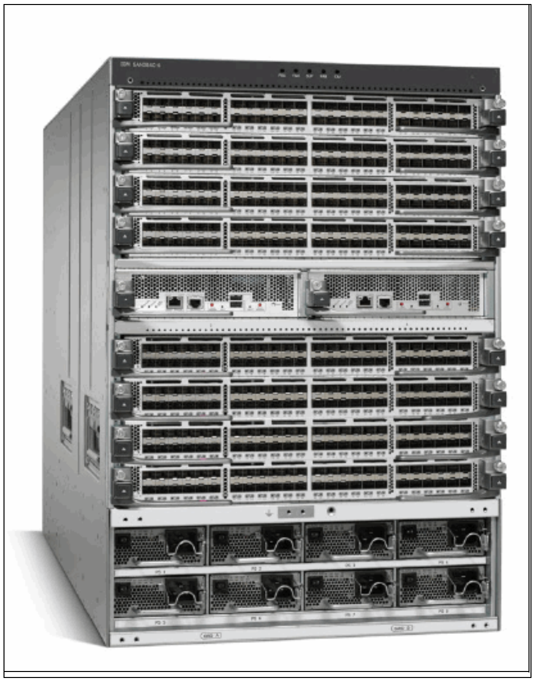

Figure 1-2 IBM Storage Networking SAN384C-6

Product highlights

IBM Storage Networking SAN384C-6 and its components offer the following main features:

- Outstanding SAN performance: The combination of the 32 Gbps FC switching modules and six Fabric-1 Crossbar switching modules enables up to 1.5 Tbps of front-panel FC throughput between modules in each direction for each of the eight IBM Storage Networking SAN384C-6 payload slots. With six Fabric-3 modules, 3 Tbps of front-panel FC throughput is possible. This per-slot bandwidth is double the bandwidth that is needed to support a 48-port 32 Gbps FC module at full line rate. Based on central arbitration and a crossbar fabric, the IBM Storage Networking SAN384C-6 architecture provides 32 Gbps line-rate, non-blocking, and predictable performance across all traffic conditions for every chassis port.
- HA: The IBM Storage Networking SAN384C-6 Director class switch enables redundancy on all major components, including the fabric card. It provides grid redundancy on power supply and 1+1 redundant supervisors. Users can include a fourth fabric-1 card to enable N+1 fabric redundancy at 768 Gbps slot bandwidth, like with Fabric-3 modules but scaled upwards. The IBM Storage Networking SAN384C-6 combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of major components for higher availability.
- Business continuity: The IBM Storage Networking SAN384C-6 Director enables large and scalable deployment of SAN extension solutions through the SAN Extension module.
- Outstanding scalability: The IBM Storage Networking SAN384C-6 Director provides up to 24 Tbps of FC backplane bandwidth with Fabric-1 modules and double with Fabric-3 modules. A single chassis delivers 384 4/8/16 Gbps, or 8/16/32 Gbps full line-rate autosensing FC ports. A single rack supports up to 1152 FC ports. The IBM Storage Networking c-type family Directors are designed to meet the requirements of even the largest data center storage environments.
- Deployment of SAN extension solutions: Enable large and scalable multi-site SANs with the 24/10-port SAN Extension Module.

- Intelligent network services: VSAN technology, ACLs for hardware-based intelligent frame processing, and fabric-wide QoS enable migration from SAN islands to enterprise-wide storage networks and include the following features:
  - Integrated hardware-based VSANs and IVR: Integration of VSANs into port-level hardware allows any port in a system or fabric to be partitioned to any VSAN. Integrated hardware-based IVR provides line-rate routing between any ports in a system or fabric without the need for external routing appliances.
  - Intelligent storage services: IBM Storage Networking SAN384C-6 operates with intelligent service capabilities on other IBM Storage Networking c-type family platforms to provide services, such as acceleration of storage applications for data replication and backup, and data migration to hosts and targets that are attached to the IBM Storage Networking SAN384C-6.
- Comprehensive security: The IBM Storage Networking c-type family supports a comprehensive security framework. It consists of RADIUS and TACACS+, FC-SP, SFTP, SSH Protocol, and SNMPv3 implementing VSANs, hardware-enforced zoning, ACLs, and per-VSAN RBAC.
- Unified SAN management: The IBM Storage Networking c-type family includes built-in storage network management with all features available through a CLI or DCNM, which is a centralized management tool that simplifies managing unified fabrics. DCNM supports the federation of up to 10 DCNM servers to manage up to 150,000 devices by using a single management window.
- Sophisticated diagnostic tests: The IBM Storage Networking SAN384C-6 provides intelligent diagnostic tests, protocol decoding, network analysis tools, and integrated Call Home capability for greater reliability, faster problem resolution, and reduced service costs.
- Multiprotocol intelligence: The multilayer architecture of the IBM Storage Networking SAN384C-6 enables a consistent feature set over a protocol-independent switch fabric. IBM Storage Networking SAN384C-6 transparently integrates FC and FICON.

Reducing TCO with SAN consolidation

With data growing exponentially, organizations need efficient, cost-effective, and large-scale SANs. You can scale while managing TCO with industry-leading port densities of up to 384 32-Gbps FC ports per chassis. You can deploy 1.5 Tbps24 Tbps24-Tbps front-panel FC line-rate nonblocking system-level switching.

You can deploy intelligent fabric services, VSANs for consolidating physical SAN islands while maintaining logical boundaries, and IVR for sharing resources across VSANs. You can consolidate your data into fewer, larger, and more manageable SANs, which reduce the hardware footprint and associated capital and operating expenses.

Enterprise-class availability

IBM Storage Networking SAN384C-6 is designed for HA. In addition to meeting the basic requirements of nondisruptive software upgrades and redundancy of all critical hardware components, the IBM Storage Networking SAN384C-6 software architecture offers outstanding availability. The supervisor modules automatically restart failed processes, which makes IBM Storage Networking SAN384C-6 exceptionally robust. In the rare event that a supervisor module is reset, complete synchronization between the active and standby supervisor modules helps ensure stateful failover with no disruption of traffic.

The IBM Storage Networking SAN384C-6 Director provides redundancy on all major active hardware components. Table 1-5 lists the redundancy that is available across all components.

Table 1-5 Redundancy details for IBM Storage Networking SAN384C-6

Item	Redundancy
Supervisors	1+1
Power supplies	Grid redundancy
Fabrics	N+1 redundancy

HA is implemented at the fabric level by using robust and high-performance ISLs. Port channel allows users to aggregate up to 16 physical links into one logical bundle. The bundle can consist of any speed-matched ports in the chassis, which helps ensure that the bundle can remain active if a port, ASIC, or module fails. ISLs in a port channel can have different lengths.

This capability is valuable in campus and metropolitan area network (MAN) environments because logical links can now be spread over multiple physical paths, which helps ensure uninterrupted connectivity even if one of the physical paths is disrupted. IBM Storage Networking SAN384C-6 provides outstanding HA, which helps ensure that solutions exceed the 99.999% uptime requirements of today’s most demanding environments.

Business transformation with enterprise cloud deployment

With industry-leading scalability and pay-as-you-grow flexibility, IBM Storage Networking SAN384C-6 enables you to quickly scale up or down enterprise clouds as needed. You also receive the following benefits:

- Robust security protects multi-tenancy cloud applications.
- Predictable high performance meets stringent SLAs.
- Resilient connectivity helps ensure an always-on cloud infrastructure.
- Advanced traffic management capabilities (such as QoS) quickly and cost-efficiently allocate elastic network capabilities to cloud applications.

Integrated mainframe support

IBM Storage Networking SAN384C-6 is mainframe-ready, with full support for IBM Z FICON and Linux environments. IBM Storage Networking SAN384C-6 supports transporting the FICON protocol in cascaded and non-cascaded fabrics, and an intermix of FICON and open systems FCP traffic on the same switch.

IBM Storage Networking SAN384C-6 product specifications

The product specifications for IBM Storage Networking SAN384C-6 are described in this section and summarized in Table 1-6. IBM Storage Networking SAN384C-6 supports FC connectivity for servers and storage. The IBM Storage Networking SAN384C-6 model supports all the FICON features and functions that are listed in this section, and requires NX-OS 8.1(1b) or later.

Table 1-6 Product specifications for IBM Storage Networking SAN384C-6

Feature	Description
Product compatibility	IBM Storage Networking c-type family.
OSs	For the most current and complete information, see IBM System Storage Interoperation Center (SSIC).

Feature	Description
Optional features	► 24/10-port SAN Extension Module (#AJL5) ► 48-Port 32-Gbps FC Switching Module (#AJL4) ► 48-Port 32-Gbps FC Switching Module Bundle (#AJL2) ► Supervisor-1 Module (#AJKE) ► Fabric-3 Switching Module (#AJNB) ► Supervisor-4 Module (#AJN6) ► Enterprise Package (#AJJ9) ► DCNM for SAN Advanced Edition (#AJJA) ► Mainframe Package (#AJJB) ► SAN Insights ► 3000 W AC power supply (#5960) ► SFPs ► Fans
Software compatibility	NX-OS Software Release 8.1(1b) or later is required for FICON support. The 48-port 32 Gbps FC switching module requires NX-OS Release 8.1(1b) or later.
Chassis slot configuration	► Line-card slots: 8 ► Supervisor slots: 2 ► Crossbar switching fabric slots: 6 ► Fan trays: Three fan trays at the back of the chassis ► Power supply bays: 8
Performance/Scalability	► Up to 24-Tbps front-panel FC switching bandwidth ► Supported FC port speeds: – 4/8/16 Gbps autosensing, optionally configurable – 8/16/32-Gbps autosensing, optionally configurable ► Buffer credits: 48-port line-rate 32-Gbps FC module: – Default credits per port: 500 – With Enterprise license: 8300 shared among a single port group of 16 ports ► Ports per chassis: Up to 384 4/8/10/16/32 Gbps FC ports ► Port channel: Up to 16 ports (The port channel can span any speed-matched port on any module in the chassis.)
Fabric services	► Name server ► RSCN ► Login services ► FCS ► Broadcast ► In-order delivery
Advanced functions	► VSAN ► IVR ► Port channel with multipath load-balancing ► QoS-flow-based, zone-based ► NPIV

Feature	Description
Diagnostic tests and troubleshooting tools	► POST diagnostic tests ► Online diagnostic tests ► Internal port loopbacks ► SPAN and Remote Switched Port Analyzer (RSPAN) ► FC Traceroute ► FC Ping ► FC Debug ► Fabric Analyzer ► Syslog ► Online system health ► Port-level statistics ► RTP Debug
Network security	► VSANs ► ACLs ► Per-VSAN RBAC ► FC zoning: – N_Port WWN – N_Port FC-ID – Fx_Port WWN – Fx_Port WWN and interface index – Fx_Port domain ID and interface index – Fx_Port domain ID and port number ► FC-SP: DH-CHAP switch-switch authentication ► DH-CHAP host-switch authentication ► Port security and fabric binding ► Management access: – SSHv2 implementing AES – SNMPv3 implementing AES – SFTP
FICON	► FC-SB-6 compliant ► Cascaded FICON fabrics ► Intermix of FICON and FC FCP traffic ► CUP management interface ► FICON multihop
Serviceability	► Configuration file management ► Nondisruptive software upgrades for FC interfaces ► Call Home ► Power-management LEDs ► Port beaconing ► System LED ► SNMP traps for alerts ► Network boot

Feature	Description
Reliability and availability	► Online, nondisruptive software upgrades ► Stateful nondisruptive supervisor module failover ► Hot-swappable redundant supervisor modules ► Hot-swappable redundant crossbar modules ► Hot-swappable 2N redundant power ► Hot-swappable fan trays with integrated temperature and power management ► Hot swappable Enhanced SFP (SFP+) optics (8/16/32 Gbps) ► Hot-swappable switching modules ► Stateful process restart ► Any module, any port configuration for port channels ► Fabric-based multipathing ► Per-VSAN fabric services ► Online diagnostic tests ► Port tracking ► Virtual Routing Redundancy Protocol (VRRP) for management
Network management	► Access methods through the Supervisor-1 or Supervisor-4 Module: – Out-of-band 10/100/1000 Ethernet port – RS-232 serial console port – In-band IP over FC ► Access methods through the FC switching module In-band FICON CUP over FC ► Access protocols – CLI for console and Ethernet ports – SNMPv3 that uses Ethernet port and in-band IP over FC access – FICON CUP ► Distributed Device Alias service ► Network security – Per-VSAN RBAC that uses RADIUS-based and TACACS+-based AAA functions – SFTP – SSHv2 implementing AES – SNMPv3 implementing AES – Management applications ► CLI ► DCNM GUI
Programming interface	► Scriptable CLI ► DCNM web services API ► NX-API

IBM Storage Networking SAN384C-6 physical specifications

Table 1-7 on page 19 details the specific requirements for planning the installation of the devices in the data center rack. For more information, see IBM Storage Networking SAN192C-6, SAN384C-6, and SAN768C-6 Installation, Service, and User Guide.

Table 1-7 Physical specifications and requirements

Item	Description
Power and cooling	► Power supplies (3000 W AC): – Input: 100 - 240 V AC nominal (±10% for full range); 16A nominal; 50 - 60 Hz nominal (±3 Hz for full range). – Output: 1451W 50 V ±4% 28A, 3.4 V120 V240 V240V AC input). ► Air flow: – The IBM Storage Networking SAN384C-6 provides 700 linear feet per minute (LFM) average system velocity, and 40 - 160 cubic feet per minute (CFM) total flow through each line-card slot depending on the line-card type and fan-speed setting. – With the IBM Storage Networking SAN384C-6 using front-to-back cold-aisle and hot-aisle air flow, you should maintain a minimum air space of 7 inches (17.78 cm) between walls, such as in a cabinet, and the chassis front and back air vents.
Power consumption (typical)	► Ports 192 - Watts 1600. ► Ports 288 - Watts 2100. ► Ports 384 - Watts 2740.
Environmental	► Temperature, ambient operating: 0 - 40°C (32 - 104°F). ► Temperature, ambient nonoperating and storage: -40 - 70°C (-40 -158°F). ► Relative humidity, ambient (noncondensing) operating: 10 - 90%. ► Relative humidity, ambient (noncondensing) nonoperating and storage: 10 - 95%. ► Altitude, operating: -60 - 2000 m (-197 - 6500 ft.).
Physical dimensions (H x W x D)	► Chassis dimensions (14 rack units [14RU]): 24.35 x 17.3 x 34.0 in. (61.9 x 43.9 x 86.4 cm). ► Chassis depth including cable management and chassis doors is 96.52 cm (38 in.). ► Unit is rack-mountable in a standard 48.26 cm (19-inch) EIA rack. Unit is also 2-post rack-mountable.

Item	Description
Weight	► Chassis (includes fans): 185.5 lb (84.2 kg). ► 48-port 16-Gbps FC line card: 17 lb (7.71 kg). ► 48-port 32-Gbps FC line card: 17.5 lb (7.94 kg). ► Power supply (3000W AC): 6 lb (2.7 kg). ► Fabric-1 module: 11 lb (5.0 kg). ► Fabric-3 module: 11 lb (5.0 kg ► Supervisor-1 module: 7 lb (3.2 kg). ► Supervisor-4 module: 7 lb (3.2 kg) ► Fan tray: 8.5 lb (3.86 kg). ► Supervisor blank cover: 1.25 lb (0.57 kg). ► Line-card blank cover: 4.5 lb (2.04 kg).

Rack and cabinet requirements for the IBM Storage Networking SAN384C-6 chassis

These are the rack and cabinet options:

- Standard perforated-doors cabinets
- Solid-walled cabinets with a roof fan module (bottom to top cooling)
- Standard open 4-post telco racks
- Standard open 2-post telco racks

Use a standard 19-inch, 4-post EIA cabinet or rack with mounting rails that conform to imperial universal hole spacing, per section 1 of the NSI/EIA-310-D-1992 standard.

The depth of a 4-post rack or a cabinet must be 24 - 32 inches (61.0 - 81.3 cm) between the front and rear mounting vertical rails.

Ensure that the airflow and cooling are adequate and there is sufficient clearance around the air vents on the switch.

The rack must have sufficient vertical clearance for the chassis along with 2 RU for the shelf brackets, and any necessary clearance for the installation process.

The front and rear doors of enclosed racks must have at least 60% of an open area perforation pattern.

Additionally, you must consider the following site requirements for the rack:

- The power receptacles must be within reach of the power cords that are used with the switch.
- AC power supplies: The power cords for 3-kW AC power supplies are 8 - 12 feet (2.5 - 4.3 m) long.
- DC power supplies (ask IBM representatives for this option): The power cords for 3.0-kW DC power supplies are supplied and set by the customer.
- HVAC/HVDC power supplies (ask IBM representatives for this option): The power cords for 3.5-kW HVAC/HVDC power supplies are 14 feet (4.26 m) long.
- Where necessary, a seismic rating of Network Equipment Building Standards (NEBS) Zone 3 or Zone 4, per GR-63-CORE.

To correctly install the switch in a cabinet in a hot-aisle/cold-aisle environment, you should fit the cabinet with baffles to prevent exhaust air from recirculating into the chassis air intake. Work with your cabinet vendors to determine which of their cabinets meet the following requirements or see IBM Support for recommendations:

- The height of the rack or cabinet must accommodate the 14-RU (24.5 inches (62.2 cm)) height of the switch and its bottom support bracket.
- Minimum gross load rating of 2000 lb (907.2 kg) (static load rating) if supporting three switches.

The rack must meet the following requirements:

- The minimum vertical rack space per chassis is 24.5 inches (62.2 cm) or 14 RU.
- The width between the mounting rails must be at least 17.75 inches (45.1 cm). For 4-post EIA racks, this is the distance between the two front rails and rear rails.

1.2.3 IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 supervisor modules

IBM provides two options for supervisor module configuration with the IBM Director switches.

- IBM Supervisor-1 Module
- IBM Supervisor-4 Module

The IBM Supervisor-1 Module has been the standard available for many years, and it provides port speed support up to 32 Gbps. The new IBM Supervisor-4 Module is the default option with a new IBM c-type Director Series switch, and it provides increased performance and functions with a port speed of 32 Gbps and future support for 64-Gbps-enabled port modules.

IBM Supervisor-4 Module

The IBM Supervisor-4 Module delivers the latest advanced switching technology with proven Cisco NX-OS software to power a new generation of scalable and intelligent multilayer switching solutions for SANs. The module is designed to integrate multi-protocol switching and routing, intelligent SAN services, and storage applications onto highly scalable SAN switching platforms. The IBM Supervisor-4 Module enables intelligent, resilient, scalable, and secure high-performance multilayer SAN switching solutions. The IBM c-type Enterprise Director family of storage networking solutions lowers the TCO for storage networking by combining robust and flexible hardware architecture, multiple layers of network and storage intelligence, and compatibility with all IBM c-type family switching modules.

This powerful combination helps organizations build highly available (HA), scalable storage networks with comprehensive security and unified management. The IBM Supervisor-4 Module is supported on the IBM Storage Networking SAN192C-6 and IBM SAN 384C-6 Series Multilayer Directors.

Figure 1-3 shows the IBM Supervisor-4 Module.

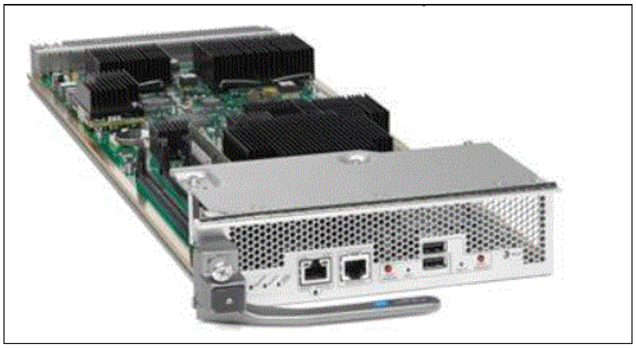

Figure 1-3 IBM Supervisor-4 Module

Industry-leading scalability

The IBM Supervisor-4 Module is designed to meet the requirements of the largest data center storage environments and combines industry-leading scalability and performance, intelligent SAN services, nondisruptive software upgrades, stateful process restart and failover, and fully redundant operation for a new standard in Director-class SAN switching.

Integrated performance

The combination of the IBM Supervisor-4 Module, IBM 48-Port 32-Gbps Fibre Channel Switching Module, and IBM Fabric-3 crossbar switching modules enables up to 3 Tbps of FC throughput between modules in each direction for each payload slot in the IBM c-type Series Director switches. This per-slot bandwidth is two times the bandwidth that is needed to support a 48-port 32-Gbps FC module at full line rate. The IBM c-type Series architecture, which is based on central arbitration and crossbar fabric, provides 64 Gbps line-rate, nonblocking, predictable performance across all traffic conditions for every port in the chassis.

High availability

The IBM Supervisor-4 Module and IBM c-type Series Multilayer Directors are designed for HA. In addition to meeting the basic requirement of nondisruptive software upgrades, the IBM c-type Series software architecture offers availability. The IBM Supervisor-4 Module can automatically restart failed processes, making it exceptionally robust. In the rare event that a supervisor module is reset, complete synchronization between the active and standby supervisor modules helps ensure stateful failover with no disruption of traffic**.**

Lower total cost of ownership

The IBM c-type Enterprise Director family provides advanced management tools for overall low TCO. It supports Cisco VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure, further decreasing TCO.

Comprehensive security framework

The IBM c-type family supports RADIUS and TACACS+; FC-SP 1, SFTP, SSH, and SNMPv3 implementing AES; and VSANs, hardware-enforced zoning, ACLs, and per-VSAN RBAC.

Unified SAN management

The IBM c-type family includes built-in storage network management, with all the features available through a CLI or DCNM, which is a centralized management tool that simplifies management of multiple switches and fabrics. Integration with third-party storage management platforms allows transparent interaction with existing management tools.

Intelligent network services

VSAN technology, ACLs for hardware-based intelligent frame processing, and fabric-wide QoS enable migration from SAN islands to enterprise-wide storage networks:

- Integrated hardware-based VSANs and IVR: Integration of VSANs into port-level hardware allows any port in a system or fabric to be partitioned to any VSAN. Integrated hardware-based IVR provides line-rate routing between any ports in a system or fabric without needing external routing appliances.
- Intelligent storage services: The IBM c-type Director Series operates with intelligent service capabilities on other IBM c-type family platforms to provide services such as acceleration of storage applications for data replication and backup and data migration to hosts and targets that are attached to the IBM c-type family.
- Smart Zoning: When the Smart Zoning feature is enabled, IBM c-type Series fabrics provision the hardware access control entries that are specified by the zone set more efficiently, avoiding the superfluous entries that allow servers (initiators) to talk to other servers or allow storage devices (targets) to talk to other storage devices. This feature makes larger zones with multiple initiators and multiple targets feasible without excessive consumption of hardware resources. Thus, smart zones can correspond to applications, application clusters, hypervisor clusters, or other data center entities, saving the time that administrators previously spent creating many small zones and enabling the automation of zoning tasks.

Advanced diagnostics and troubleshooting tools

The IBM c-type family integrates advanced diagnostics tools to perform analysis and debugging with port analyzing tools to identify problems in the traffic. Port-based and flow-based statistics enable sophisticated performance analysis and SLA accounting. The integrated Smart Call Home capability provides more reliability and enables faster problem resolution and reduced service costs.

Multiprotocol intelligence

The multilayer architecture of the IBM c-type Series enables a consistent feature set over a protocol-independent switch fabric. The IBM c-type Series transparently integrates FC, FCIP, and FICON.

- 2/4/8-Gbps, 4/8/16-Gbps, 8/16/32-Gbps, and 10-Gbps FC and 10-GbE: The IBM c-type Series supports both 2/4/8/16/32-Gbps and 10-Gbps ports on the IBM c-type 48-Port 32-Gbps Fibre Channel Switching Module. The IBM c-type Series also supports 10 GbE clocked optics carrying 10-Gbps FC traffic.
- FICON: The IBM c-type Director Series supports deployment in IBM Z FICON and Linux environments.

IBM Supervisor-1 Module

The IBM Supervisor-1 Module is designed specifically for the IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 chassis. The IBM Supervisor-1 Module delivers the latest advanced switching technology with NX-OS software to power a new generation of scalable and intelligent multilayer switching solutions for SANs. This supervisor module provides control and management functions for the switch and enables high-performance switching.

Designed to integrate multi-protocol switching and routing, intelligent SAN services, and storage applications onto highly scalable SAN switching platforms, the IBM Supervisor-1 Module enables intelligent, resilient, scalable, and secure high-performance multilayer SAN switching solutions when combined with the IBM Storage Networking c-type family switching modules.

Two IBM Supervisor-1 Module modules are required per system to use the reliability and availability features, such as Active-Active redundancy, Online nondisruptive software upgrades, hot-swappable modules, stateful process restart, and stateful nondisruptive supervisor failover.

Figure 1-4 shows the IBM Supervisor-1 Module.

Figure 1-4 IBM Supervisor-1 Module

This supervisor module supports the following features:

- Nondisruptive software upgrades
- Stateful process restart and failover
- Fully redundant operation
- Support for up to 384 FC ports in a single chassis
- Support for up to 24 Tbps of FC system bandwidth
- Multipathing based on Fabric Shortest Path First (FSPF)
- Ability to dynamically reroute traffic in the event of a switch failure
- Network management through the CLI and through DCNM
- Extensive security features, including RADIUS and TACACS+, FC-SP, SFTP, SSH, and SNMPv3 implementing AES; and VSANs, hardware-enforced zoning, ACLs, and per-VSAN RBAC
- Support for VSAN technology and IVR
- Network services such as ACLs and QoS
- Smart zoning

- POST and diagnostics
- SPAN and RSPAN

1.2.4 IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 crossbar fabric modules

The IBM Storage Networking SAN192C-6 supports up to six crossbar fabric modules. There is a crossbar fabric module that is designed specifically for the IBM Storage Networking SAN192C-6. The crossbar fabric modules are installed vertically at the back of the chassis behind the fan modules. Fabric slots 1 and 2 are behind fan module slot 1, fabric slots 3 and 4 are behind fan module slot 2, and fabric slots 5 and 6 are behind fan module slot 3.

Figure 1-5 shows a crossbar fabric module for the IBM Storage Networking SAN192C-6.

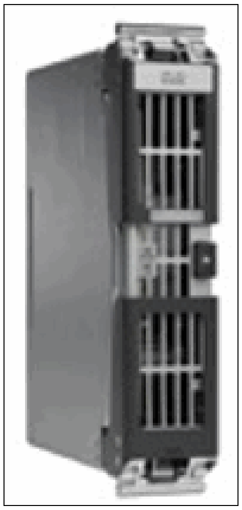

Figure 1-5 Crossbar fabric module for the IBM Storage Networking SAN192C-6

The IBM Storage Networking SAN384C-6 supports up to six crossbar fabric modules. There is a crossbar fabric module that is designed specifically for the IBM Storage Networking SAN384C-6. The crossbar fabric modules are installed vertically at the back of the chassis behind the fan modules. A minimum of three crossbar fabric modules are required to operate the switch. A fourth crossbar fabric module is required for N+1 protection.

Figure 1-6 shows a crossbar fabric module for the IBM Storage Networking SAN384C-6.

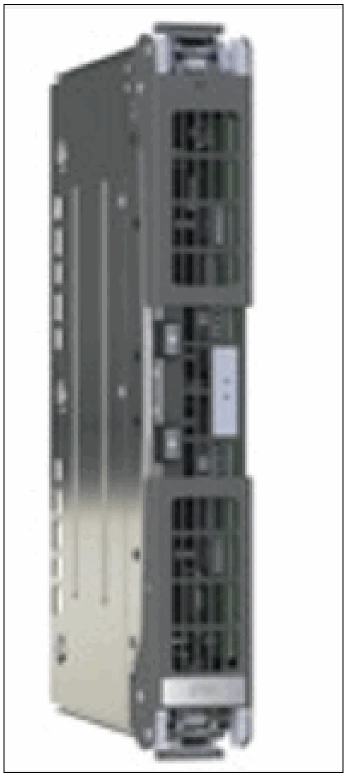

Figure 1-6 Crossbar fabric module for the IBM Storage Networking SAN384C-6

Each crossbar fabric module connects to four or eight switching modules and two supervisor modules. In addition, each crossbar fabric module supports four 55 Gbps fabric ports (F_Ports) that are connected to each switching module and one 55 Gbps F_Port that is connected to each supervisor module.

Because the crossbar fabric modules are behind the fan modules in the chassis, the LEDs on the crossbar fabric module are not easily visible from the back of the chassis. So, crossbar fabric status LEDs are provided on the fan modules too. Because each fan module covers two fabric modules, the status LEDs for two crossbar fabric modules are present on each fan module. If the fan module is removed, the status and locator LEDs on crossbar fabric modules will be visible.

When a fabric module must be found, the locator LED of the corresponding fan module must be activated, followed by the locator LED of the fabric module, by using the **`locator-led fan <fan module number>`** and **`locator-led xbar <xbar slot number>`**. For example, to find crossbar fabric module 4, the locator LED of fan module 2 must be activated followed by the locator LED of fabric module 4.

Each fabric-1 module provides 256 Gbps of FC bandwidth per line card slot (fabric-3 offers double that bandwidth). The maximum chassis bandwidth is 1.536 Tbps FC bandwidth per line card slot with six fabric modules installed.

Table 1-8 shows the switching capabilities per fabric.

Table 1-8 Switching capabilities per fabric

Number of fabric cards	Front panel FC bandwidth per slot
1	256 Gbps
2	512 Gbps
3	768 Gbps
4	1024 Gbps
5	1280 Gbps
6	1536 Gbps

Note: Fabric modules may be installed in any slot, but a best practice is one behind each fan tray.

1.2.5 IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 power supplies

The IBM c-type SAN switches and director support the following types of power supplies:

- 3000 W AC power supply (AC input and DC output)
- 3000 W DC power supply (DC input and DC output). Ask IBM representatives for this option.

Figure 1-7 shows the 3000 W AC power supply.

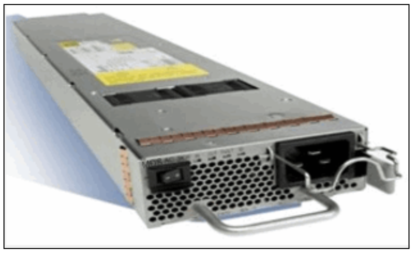

Figure 1-7 3000 W AC power supply

The IBM Storage Networking SAN384C-6 supports up to eight hot-swappable 3000 W AC power supplies (AC input). The IBM Storage Networking SAN192C-6 supports up to four hot-swappable 3000 W AC power supplies (AC input).

These AC power supplies hold 80Plus Platinum certification for maximum power efficiency.

The 3000 W AC power supply unit (PSU) may be connected to either 220 V or 110 V AC power sources. When connected to 220 V, each PSU has a maximum output capacity of 3000 W. When connected to 110 V, each PSU has a maximum output capacity of 1450 W.

Each power supply module monitors its output voltage and provides the status to the supervisor. In addition, the power supply modules provide information about local fans, power, shutdown control, and E2PROM to the supervisor.

A c-type SAN Director has a flexible power system providing different power modes. Any operational power supply provides power to the system power bus, which allows the power load of the system to be shared equally across all operational power supplies. Power supply output can be allocated to one of two pools. The available pool is available to start system components. The reserve pool is kept in reserve and not counted toward the available power.

The system can be configured in one of several modes that vary the size of the available and reserve power pools according to user requirements:

- Combined mode: This mode allocates the output power of all power supplies to available power for switch operations. This mode does not reserve any output power in case of power outages or power supply failures.
- Power supply redundancy mode (N+1): In this mode, one power supply's output is allocated to the reserve power pool, which provides the system with enough reserve power if a single power supply fails. The remaining power supplies are allocated to the available power pool. The reserve power supply must be at least as powerful as the most powerful power supply in the available pool to potentially replace the full power output of the failed power supply in the worst case. Because it is impossible to predict which power supply might fail, provision the system with power supplies of equal rating so that the output of any power supply that fails can be replaced by the remaining power supplies.

  For example, a system with four 3 kW power supplies in N+1 redundancy mode has a total of 12 kW. 9 kW are allocated to the available power pool, and 3 kW are reserved. If any of the power supplies fail, enough power is reserved that the remaining power supplies can still meet the 9 kW commitment.
- Input grid redundancy mode (grid redundancy): In this mode, half of the power supply's output is allocated to the reserve power pool and half to the available power pool, which provides the system with enough reserve power in the case of 50% of the power supplies failing, as when a power grid fails. The system logically allocates the left two columns of PSU bays to Grid A and sums the output power of operational PSUs. It does the same for the right two columns (Grid B) and uses the minimum of the two as the available power pool. To use maximum power, the sum of the power supply outputs of Grid A and Grid B PSU bays must be equal.

  For example, a system with four 3 kW PSUs in Grid A bays and three 3 kW PSUs in Grid B bays and in grid redundancy mode has 12 kW available from Grid A and 9 kW from Grid B. The minimum of the two grids is 9 kW, so 9 kW is allocated to the available power pool and 9 kW are reserved. If either grid fails, enough power is reserved that the remaining power supplies can still meet the 9 kW commitment. The output of the fourth PSU in Grid A is not considered in the calculations even though it provides power.
- Full redundancy mode: This mode supports both grid redundancy or N+1 redundancy. 50% of the power supply output is allocated to the reserve pool, and the other 50% of the power supply outputs are allocated to the available power pool. The reserved power may be used to back up either single power supply failures or a grid failure.

  For example, a system with six 3 kW power supplies in grid redundancy mode has a total of 18 kW. 9 kW are allocated to the available power pool and 9 kW are allocated to the reserve pool. If a grid failure occurs (half of the power supplies lose power), the full reserve power pool is available to meet the 9 kW commitment. Otherwise, as single power supplies fail, power is allocated to the available pool from the remaining reserve power pool until the reserve power pool is exhausted.

  Note: After a single power supply has failed in this mode, grid redundancy is no longer available**.**

Figure 1-8 shows how to connect power supplies in an IBM Storage Networking SAN384C-6 for grid redundancy.

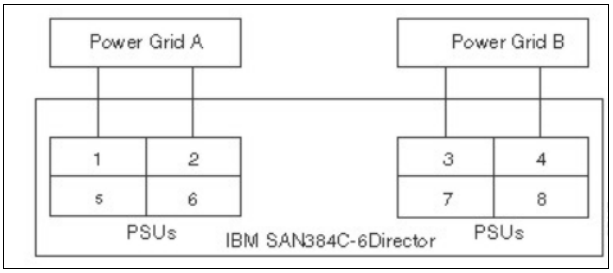

Figure 1-8 IBM Storage Networking SAN384C-6 Grid-PSU connections

Figure 1-9 shows how to connect power supplies in an IBM Storage Networking SAN192C-6 for grid redundancy.

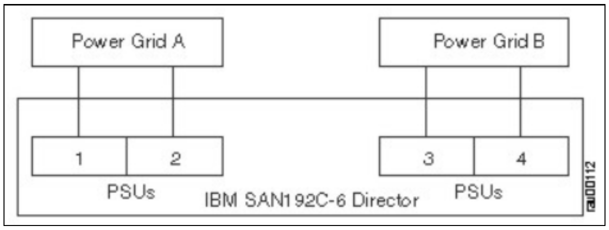

Figure 1-9 IBM Storage Networking SAN192C-6 Grid-PSU connections

Supported power cords and plugs

Each power supply has a separate power cord. Standard power cords or jumper power cords are available for connection to a power distribution unit with IEC 60320 C19 outlet receptacles.

1.2.6 IBM 48-Port 32-Gbps Fibre Channel Switching Module

The 48-Port 32-Gbps Fibre Channel Switching Module is designed specifically for the IBM c-type SAN Directors. The 48-Port 32-Gbps Fibre Channel Switching Module delivers predictable performance, scalability, and innovative features to support private and virtualized data centers. With industry-leading 768 line-rate 32-Gbps FC ports per director, the 32-Gbps 48-port Fibre Channel Switching Module meets the high-performance needs for flash memory and non-volatile memory express (NVMe) over FC (NVME-FC) workloads. It offers innovative services, including SSD awareness, on-board FC Analytics Engine, E-port and F-port diagnostics, integrated VSANs, IVR, and port channels. It delivers full-duplex aggregate performance of 1536 Gbps, making it suitable for high-speed 32-Gbps storage subsystems, 32-Gbps ISLs, high-performance virtualized servers, all-flash and NVMe arrays, and FICON environments.

Figure 1-10 shows the 48-Port 32-Gbps Fibre Channel Switching Module.

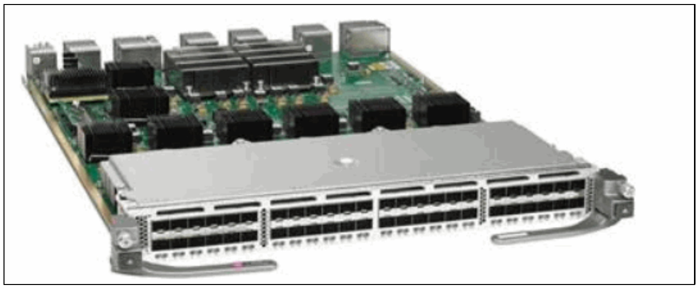

Figure 1-10 48-Port 32-Gbps Fibre Channel Switching Module

With 384 line-rate 32 Gbps FC ports per director, the 48-Port 32-Gbps Fibre Channel Switching Module meets the high-performance needs for flash memory and NVMe-FC workloads. The switching module is hot swappable and compatible with 4 Gbps, 8 Gbps, 16 Gbps, and 32 Gbps FC interfaces. This module also supports hot swappable Enhanced SFP+ transceivers.

Individual ports can be configured with 32 Gbps, 16 Gbps, and 8 Gbps SFP+ transceivers. Each port supports 500 buffer credits for exceptional extensibility without the need for extra licenses. With the Enterprise Package license, up to 8191 buffer credits can be allocated to an individual port, enabling full-link bandwidth over long distances with no degradation in link utilization.

Figure 1-11 shows the 48-Port 32-Gbps Fibre Channel Switching Module port group view.

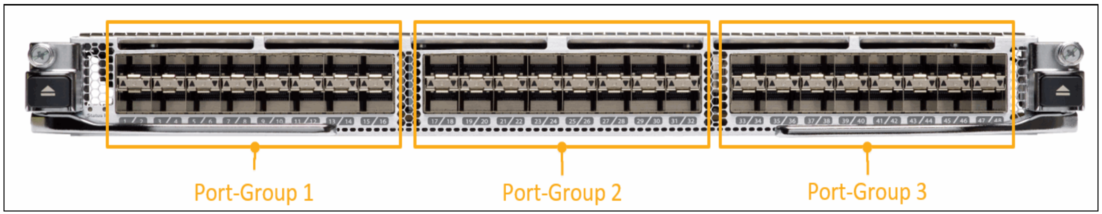

Figure 1-11 Port group view

The main features are:

- Forty-eight 32-Gbps FC line-rate ports.
- Common module for IBM Storage Networking c-type Director platform.
- 4/8/16/32-Gbps FC supported speeds in SFP+ format.
- Dedicated Analytics Network Processing Unit (NPU).
- Up to 1.5 Tbps front-panel bandwidth.
- Maximum of 8300 buffer credits per port group.
- 500 buffer-to-buffer (B2B) credits per port by default. Up to 8191 B2B credits per port (with enterprise license).
- Three port groups, 16 ports per port group, and 16 Ports per ASIC.

Note: SAN Analytics is not supported by the FICON protocol. Only FC SCSI and NVMe analytics are available on the switch port module.

Supported transceivers

The IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 support the FC SFP+ transceivers in either short wavelength (SWL) or long wavelength (LWL). The transceivers are field-replaceable and hot-swappable. You can use any combination of SFP+ transceivers that are supported by the switch. The only restrictions are that SWL transceivers must be paired with SWL transceivers; LWL transceivers must be paired with LWL transceivers; and the cable must not exceed the stipulated cable length for reliable communications.

SFP+ transceivers provide the uplink interfaces, laser transmit (Tx) and laser receive (Rx), and support 850 - 1610 nm nominal wavelengths, depending upon the transceiver.

Note: Use only Cisco transceivers in the IBM c-type SAN switches and director. Each transceiver is encoded with model information that enables the switch to verify that the transceiver meets the requirements for the switch.

SFP+ transceivers can be ordered separately or with the IBM c-type SAN switches and director.

Table 1-9 lists the FC receivers.

Table 1-9 FC transceivers

Description	Type
32 Gbps FC SW SFP+	SWL
32 Gbps FC LW SFP+	LWL
32 Gbps FC ELW SFP+	LWL
16 Gbps FC SW, SFP+	SWL
16 Gbps FC LW, SFP+	LWL
16 Gbps FC ELW, SFP+	LWL
8 Gbps FC SW, SFP+	SWL
8 Gbps FC LW, SFP+	LWL
8 Gbps FC ER SFP+	Extended Reach
8/16-Gbps FC CWDM SFP+	Coarse Wavelength-Division Multiplexing (CWDM)
8/16-Gbps FC DWDM SFP+	Dense Wavelength-Division Multiplexing (DWDM)
10GBASE-SR SFP Module	Short Reach
10GBASE-LR SFP Module	Long Reach
10GBASE-ER SFP Module	Extended Reach
10GBASE-DWDM SFP+	DWDM (40 different wavelengths are offered.)

For more information about a specific SFP+ transceiver, see SFP+ Transceiver Specifications.

1.2.7 IBM 24/10-Port SAN Extension Module

The capabilities of IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 can be extended with the 24/10-Port SAN Extension Module supported on IBM Storage Networking c-type Family Multilayer Directors. With 24 line-rate 2-, 4-, 8-, 10-, and 16-Gbps FC ports and eight 1- and 10-GbE FCIP ports, this module enables large and scalable deployment of SAN extension solutions. The SAN extension module has two independent service engines that can each be individually and incrementally enabled to scale as business requirements expand.

Figure 1-12 shows the 24/10-Port SAN Extension Module.

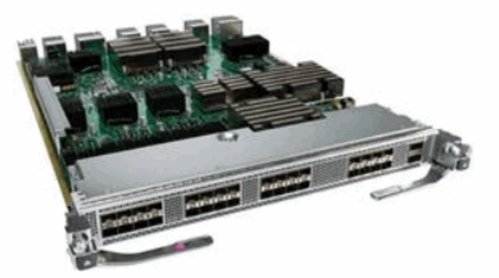

Figure 1-12 24/10-Port SAN Extension Module

The SAN extension module supports the full range of services that are available on other IBM Storage Networking c-type Family Fibre Channel Switching Modules, including VSAN, security, and traffic management services. The FCIP module uses IBM expertise and knowledge of IP networks to deliver outstanding SAN extension performance, reducing latency for disk and tape operations with FCIP acceleration features, including FCIP write acceleration and FCIP tape write and read acceleration. The switching module has two service engines on its system board.

Hardware-based encryption helps secure sensitive traffic with IP Security (IPsec), and hardware-based compression dramatically enhances performance for both high- and low-speed links, enabling immediate cost savings in expensive WAN infrastructure. Multiple FCIP interfaces within a single engine or across service engines can be grouped into a port channel of up to 16 links for HA and increased aggregate throughput.

Main features and benefits

The 24/10-Port SAN Extension Module is designed for mission-critical enterprise storage networks that require secure, robust, and cost-effective business-continuance services. The SAN extension module offers the following main features:

- FCIP for remote SAN extension:
  - Simplifies data-protection and business-continuance strategies by enabling backup, remote replication, and other DR services over WAN distances by using open standards FCIP tunneling.
  - Optimizes utilization of WAN resources for backup and replication by enabling hardware-based compression, hardware-based encryption, FCIP write acceleration, and tape read and write acceleration for both FCIP and FICON over IP. The SAN extension module supports four tunnels per interface and can scale up to 32 tunnels (four tunnels x eight 1/10 GbE ports).
  - Preserves IBM Storage Networking c-type Family Directors enhanced capabilities, including VSANs, advanced traffic management, and security, across remote connections.

- Integrated IP storage services in a high-density form factor: The module supports eight 1- and 10-GbE ports. Individual ports can be configured with hot-swappable SWL and LWL SFP connections.
- Integrated hardware-based VSANs and IVR: The module enables deployment of large-scale multi-site and heterogeneous SAN topologies. Integration into port-level hardware allows any port in a system or fabric to be partitioned into any VSAN. Integrated hardware-based IVR provides line-rate routing between any ports in a system or fabric without the need for external routing appliances.
- Intelligent network services: The module uses VSAN technology for hardware-enforced, isolated environments in a single physical fabric; ACLs for hardware-based intelligent frame processing; and advanced traffic management features, such as fabric-wide QoS, to facilitate migration from SAN islands to enterprise-wide storage networks.
- Sophisticated diagnostics: The module provides intelligent diagnostics, protocol decoding, and network analysis tools, in addition to integrated Call Home capability, for greater reliability, faster problem resolution, and reduced service costs.
- Comprehensive network security framework: The module supports RADIUS and TACACS+, FC-SP, SFTP, SSH, SNMPv3 implementing the AES, VSANs, hardware-enforced zoning, ACLs, and per-VSAN RBAC. RBAC provides separate control over management functions and access on a per-VSAN basis, enabling separation of duties among administrators on the same physical switch. GbE ports support IPsec authentication, data integrity, and hardware-assisted data encryption.
- IPv6 support: The module supports IPv6 as mandated by the US Department of Defense (DoD), Japan, and China. IPv6 support is provided for FCIP and for management traffic that is routed in band and out of band.

Integrated FCIP for remote SAN and mainframe channel extension

Data-distribution, data-protection, and business-continuance services are significant components of today’s information-centered businesses. The capability to efficiently replicate critical data on a global scale helps ensure a higher level of data protection for valuable corporate information, and also increases utilization of backup resources and lowers total cost of storage ownership. The 24/10-Port SAN Extension Module uses the open-standards FCIP protocol to extend the distance of current FC and FICON solutions, enabling interconnection of SAN islands over extended distances.

Advanced SAN extension features

The 24/10-Port SAN Extension Module supports hardware-based FCIP compression to increase the effective WAN bandwidth of SAN extension solutions. The module can deliver compression ratios in the range of 4:1 - 5:1 over a wide variety of data sources.

The SAN extension module supports AES 256 IPsec encryption for secure transmission of sensitive data over extended distances. Hardware enablement of IPsec helps ensure line-rate throughput. Together, hardware-based compression and hardware-based encryption provide a high-performance, highly secure SAN extension capability.

Additionally, the SAN extension module supports FCIP write acceleration, a feature that can significantly improve application performance when storage traffic is extended across long distances. When FCIP write acceleration is enabled, WAN throughput is optimized by reducing the latency of command acknowledgments.

VSANs

Suitable for efficient, secure SAN consolidation, ANSI T11-standard VSANs enable more efficient storage network utilization by creating hardware-based isolated environments with a single physical SAN fabric or switch. Each VSAN can be zoned as a typical SAN and maintained with its own fabric services for greater scalability and resilience. VSANs allow the cost of SAN infrastructure to be shared among more users, while helping ensure segmentation of traffic and retaining independent control of configuration on a VSAN-by-VSAN basis.

Integrated SAN routing

In another step toward deployment of efficient, cost-effective, and consolidated storage networks, the 24/10- Port SAN Extension Module supports IVR, the industry's first and most efficient routing function for FC. IVR allows selective transfer of data between specific initiators and targets on different VSANs while maintaining isolation of control traffic within each VSAN. With IVR, data can transit VSAN boundaries while maintaining control-plane isolation, thus maintaining fabric stability and availability. IVR eliminates the need for external routing appliances, greatly increasing routing scalability while delivering line-rate routing performance, simplifying management, and eliminating the challenges that are associated with maintaining separate systems. IVR reduces the total cost of SAN ownership.

Advanced traffic management

The advanced traffic management capabilities that are integrated into the 24/10-Port SAN Extension Module simplify deployment and optimization of large-scale fabrics:

- Virtual output queuing: Helps ensure line-rate performance on each port independent of traffic pattern by eliminating head-of-line blocking.
- Port channels: Allow users to aggregate up to 16 FCIP ISLs into a single logical bundle, providing optimized bandwidth utilization across all links. The bundle can consist of any speed-matched ports from any module in the chassis, helping ensure that the bundle can remain active even in the event of a module failure.
- FSPF-based multipathing: Provides the intelligence to load balance traffic across up to 16 equal-cost paths and, in the event of a switch failure, dynamically reroute traffic.
- QoS: Can be used to manage bandwidth and control latency to prioritize critical traffic.
- Shaper: Rate limits the WAN bandwidth according to the maximum bandwidth configured for the FCIP tunnel.

Comprehensive solution for robust network security

Addressing the need for fail-proof security in storage networks, the 24/10-Port SAN Extension Module offers an extensive security framework to protect highly sensitive data moving in today's enterprise networks. The module employs intelligent frame inspection at the port level, including the application of ACLs for hardware enforcement of zones, VSANs, and advanced port security features:

- Extended zoning capabilities restrict broadcasts to only the selected zones (broadcast zones).
- VSANs are used to achieve greater security and stability by providing complete isolation among devices that are connected to the same physical SAN.
- FC-SP provides switch-to-switch and host-to-switch Diffie-Hellman Challenge Handshake.
- DH-CHAP authentication supports RADIUS and TACACS+ to help ensure that only authorized devices can access protected storage networks.

1.3 IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 software licensing for NX-OS

The IBM c-type Director series comes with base software functions that allow the switch to operate at a basic level with functions that include Device Manager (DM) (embedded web server), VSAN, ACLs, fabric zoning, and trunking. This level of base function is not always sufficient to provide advanced features that are useful for large or complex SAN deployments for enterprise customers.

1.3.1 Licensing model

Any feature that is not included in a license package is bundled with the IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 switches and provided at no extra cost.

IBM provides four optional feature-based licenses for the Director-class switches. The features provide advanced functions, advanced management, analytics, and mainframe support. This book focuses on mainframe FICON supported features.

Note: The feature *SAN Insights* provides SAN Analytics, which is not supported by FICON.

1.3.2 Mainframe Package (#AJJB)

This package is a comprehensive collection of features that are required to use the IBM Storage Networking c-type Series switches in mainframe storage networks. The Mainframe Package enables the connectivity of IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 storage director to IBM Z servers with both FICON and FCP protocols. The Mainframe Package is licensed per director switch for all the ports on the director. For cascaded FICON environments, all director in the FICON VSAN must have a Mainframe Package.

The Mainframe Package includes a number of features:

- FICON cascaded Directors
- FICON Dynamic Routing (FIDR) and Cisco originator exchange ID (OXID)-based routing
- Enhanced ISL aggregation
- Lossless in-order delivery
- VSANs
- Dynamic port number assignment
- FICON CUP
- Fabric Binding
- Forward Error Correction (FEC) for 16-Gbps FICON channels, control units (CUs), and ISLs

1.3.3 Enterprise Package (#AJJ9)

The Enterprise Package enables a set of advanced traffic-engineering and security features that are recommended for large or complex SANs that require a level of function above and beyond what is provided with the base standard features.

Advanced features that are enabled with the Enterprise Package include:

- Advanced Traffic-Engineering Features:
  - IVR
  - QoS
  - Extended B2B credits
- Enhanced Network Security Features:
  - FC-SP
  - Port security
  - VSAN-based access control
  - IPsec
  - Digital certificates
  - Fabric binding for open systems FC
  - Cisco TrustSec

1.3.4 DCNM SAN Advanced Edition Package (#AJJA)

DCNM is an easy-to-use application that simplifies management across multiple switches and fabrics. Focused on supporting efficient operations and management, it provides a robust framework and rich feature set that meets the routing, switching, and storage administration needs of present and future virtualized data centers. DCNM streamlines provisioning of the unified fabric and proactively monitors the SAN components.

This feature enables the advanced functions that extend the DCNM features when the basic level of DCNM features does not provide the required level of management capability.

Table 1-10 provides a comparison list of “Standard” and “Licensed” functions that should be considered when determining whether the extra purchase cost is required. Typically, having the Advanced feature is advised for ease of management.

Table 1-10 DCNM standard and licensed features

Feature	DCNM unlicensed mode	DCNM with Advanced Feature licenses installed
FC, Fibre Channel over Ethernet (FCoE), FICON, and internet Small Computer Systems Interface (iSCSI) Topology View	Yes	Yes
Fabric, Device, and Summary Views	Yes	Yes
Port, Switch, and fabric-level configuration	Yes	Yes
Event and security management	Yes	Yes
Configuration analysis tools	Yes	Yes
Network diagnostic and troubleshooting tools	Yes	Yes
Real-time performance monitoring	Yes	Yes

Feature	DCNM unlicensed mode	DCNM with Advanced Feature licenses installed
One command multi-switch CLI access	Yes	Yes
DM	Yes	Yes
Template-based provisioning	Yes	Yes
Generic Online Diagnostics (GOLD)	Yes	Yes
Heterogeneous storage array discovery	-	Yes
Scale-out federation architecture	-	Yes
SAN Host Path Redundancy Analysis	-	Yes
Automatic fabric failover	-	Yes
VMware vCenter plug-in	-	Yes
Multiple fabrics management	-	Yes
Centralized management server with discovery	-	Yes
Continuous health and event monitoring	-	Yes
Historical performance monitoring and reporting	-	Yes
Event forwarding	-	Yes
DCNM proxy services	-	Yes
Configuration backup, archive, and compare	-	Yes
Roaming user profiles	-	Yes
VMpath analytics	-	Yes
Domain Dashboards	-	Yes
Capacity Manager	-	Yes
Event Snooze	-	Yes
Reporting	-	Yes

Base unlicensed feature

Here are the features for which no license is required:

- SAN and local area network (LAN) discovery.
- Event Registration (traps, syslogs, accounting, and threshold).
- Web Services (SOAP, XML, and API).
- Federation (up to 10 DCNM servers).

Restrictions on opening an unlicensed fabric

Here are a few restrictions regarding the opening of an unlicensed fabric:

- Opening a fabric from a remote SAN client requires Cisco DCNM Advanced License.
- If you are using a remote Cisco DCNM SAN client, you cannot open any unlicensed fabric. The fabric must be licensed or the Cisco DCNM Essential license does not work.
- If you are trying to open an unlicensed fabric from a SAN client running on the Cisco DCNM server, you can open only one unlicensed fabric at a time.
- If one instance is opened from a local SAN client, you cannot open another instance of an unlicensed fabric.

1.4 Extension switch model

You can deploy SAN extension solutions, distributed intelligent fabric services, and multi-protocol connectivity for open systems and mainframe environments. IBM offers a non-Director class solution for deployment in smaller size Mainframe environments with support for FICON: the IBM Storage Networking SAN50C-R.

1.4.1 IBM Storage Networking SAN50C-R

IBM Storage Networking SAN50C-R is an optimized platform for deploying high-performance SAN extension solutions, distributed intelligent fabric services, and cost-effective multi-protocol connectivity for both open systems and mainframe environments. With a compact form factor and advanced capabilities that normally are available only on Director-class switches, IBM Storage Networking SAN50C-R is an ideal solution for departmental and remote branch-office SANs, and large-scale SANs along with the IBM Storage Networking SAN384C-6 Director.

IBM Storage Networking SAN50C-R switch offers up to forty 16-Gbps FC ports and two 1/10-GbE IP storage services ports in a fixed two-rack-unit (2RU) form factor. The eight 10-GbE Fibre Channel over Ethernet (FCoE) ports are not used in FICON environments. The IBM Storage Networking SAN50C-R switch connects to existing native FC networks, protecting investments in storage networks.

The IBM SAN Extension over IP application package license is enabled as standard on the two fixed, 1/10-GbE IP storage services ports, enabling features such as FCIP and compression on the switch without the need for extra licenses. Also, by using the eight 10-GbE FCoE ports, the IBM Storage Networking SAN50C-R platform attaches to directly connected FCoE and FC storage devices, and supports multi-tiered unified network fabric connectivity directly over FCoE.

Figure 1-13 shows the IBM Storage Networking SAN50C-R.

Figure 1-13 IBM Storage Networking SAN50C-R

Product highlights

The IBM Storage Networking SAN50C-R switch provides unique multiservice and multi-protocol functions in a compact 2RU form factor:

- SAN consolidation with integrated multi-protocol support: The IBM Storage Networking SAN50C-R switch is available in a base configuration of 20 ports of 16-Gbps FC for high-performance SAN connectivity, 2 ports of 1/10-GbE for FCIP and iSCSI storage services, and eight ports of 10-GbE for FCoE connectivity.
- High-density FC switch with 16-Gbps connectivity: The IBM Storage Networking SAN50C-R switch scales up to 40 ports of 16-Gbps FC in a fixed configuration switch. The base configuration comes with 20 ports of 16-Gbps FC enabled for high-performance SAN connectivity. It can be upgraded onsite to enable additional 20 ports of 16-Gbps FC by adding the Port-On-Demand Activation license. Additionally, the IBM Storage Networking SAN50C-R cost-effectively scales up for FICON mainframe environments.
- Intelligent application services engine: The IBM Storage Networking SAN50C-R switch includes as standard a single application services engine that enables the included IBM SAN Extension over IP software solution package to run on the two fixed, 1/10-GbE storage services ports. The IBM SAN Extension over IP package provides an integrated, cost-effective, and reliable business-continuance solution that uses IP infrastructure by offering FCIP for remote SAN extension, along with various advanced features to optimize the performance and manageability of FCIP links.
- Hardware-based virtual fabric isolation with virtual VSANs and FC routing with IVR: VSANs and IVR enable deployment of large-scale multi-site and heterogeneous SAN topologies. Integration into port-level hardware allows any port in a system or in a fabric to be partitioned into any VSAN. Included in the optional IBM Storage Networking c-type Enterprise advanced software package, IVR provides line-rate routing between any of the ports in a system or in a fabric without the need for external routing appliances.
- Remote SAN extension with high-performance FCIP:
  - Simplifies data protection and business-continuance strategies by enabling backup, remote replication, and other DR services over WAN distances by using open-standards FCIP tunneling.
  - Optimizes the usage of WAN resources for backup and replication by enabling hardware-based compression, hardware-based encryption, FCIP write acceleration, and FCIP tape read and write acceleration. Virtual ISL connections are provided on the two 1/10-GbE ports through tunneling.
  - Preserves IBM Storage Networking c-type Family enhanced capabilities, including VSANs, IVR, advanced traffic management, and network security across remote connections.

- Cost-effective iSCSI connectivity to Ethernet-attached servers:
  - Extends the benefits of FC SAN-based storage to Ethernet-attached servers at a lower cost than is possible by using FC interconnect alone.
  - Increases storage usage and availability through the consolidation of IP and FC block storage.
  - Through transparent operation, it preserves the capability of existing storage management applications.
- Advanced FICON services: The IBM Storage Networking SAN50C-R supports FICON environments, including cascaded FICON fabrics, VSAN-enabled intermix of mainframe and open systems environments, and NPIV for mainframe Linux partitions. CUP support enables in-band management of IBM Storage Networking c-type switches from the mainframe management console.

  FICON tape acceleration reduces latency effects for FICON channel extension over FCIP for FICON tape read and write operations to mainframe physical or virtual tape. This feature is sometimes referred to as *tape pipelining*.
- Platform for intelligent fabric applications: The IBM Storage Networking SAN50C-R switch provides an open platform that delivers the intelligence and advanced features that are required to make multilayer intelligent SANs a reality, including hardware-enabled innovations to host or accelerate applications for data migration, storage backup, and data replication. Hosting or accelerating these applications in the network can dramatically improve scalability, availability, security, and manageability of the storage environment, resulting in increased utility and lower TCO.
- In-service software upgrades (ISSUs) for FC interfaces: The IBM Storage Networking SAN50C-R switch promotes high serviceability by enabling NX-OS software to be upgraded while the FC ports are carrying traffic.
- Intelligent network services: The IBM Storage Networking SAN50C-R switch uses VSAN technology for hardware-enforced, isolated environments within a single physical fabric, ACLs for hardware-based intelligent frame processing, and advanced traffic management features such as fabric-wide QoS to facilitate migration from SAN islands to enterprise-wide storage networks.
- High-performance ISLs: The IBM Storage Networking SAN50C-R switch supports up to 16 FC ISLs in a single port channel. Links can span any port on any module in a chassis for added scalability and resilience. Up to 253 B2B credits can be assigned to a single FC port to extend storage networks over long distances.
- Comprehensive network security framework: The IBM Storage Networking SAN50C-R switch supports RADIUS and TACACS+, FC-SP, SFTP, SSH Protocol, SNMPv3 implementing AES, VSANs, hardware-enforced zoning, ACLs, and per-VSAN RBAC. Additionally, the 10-GbE ports offer IPsec authentication, data integrity, and hardware-assisted data encryption for FCIP and iSCSI.
- IPv6 capable: The IBM Storage Networking SAN50C-R switch supports IPv6 as mandated by the US DoD, Japan, and China. IPv6 support is provided for FCIP, iSCSI, and management traffic routed inband and out of band.
- Sophisticated diagnostic tests: The IBM Storage Networking SAN50C-R switch provides intelligent diagnostic tests, protocol decoding, and network analysis tools, and integrated IBM Call Home capability for added reliability, faster problem resolution, and reduced service costs.

Architecture and key components

This section describes the architecture and key components of the IBM Storage Networking SAN50C-R.

VSANs

VSANs are ideal for efficient, secure SAN consolidation, enabling more efficient storage network usage by creating hardware-based isolated environments with a single physical SAN fabric or switch. Each VSAN can be zoned as a typical SAN and maintains its own fabric services for added scalability and resilience. VSANs allow the cost of a SAN infrastructure to be shared among more users while helping ensure complete segmentation of traffic and retaining independent control of configuration on a VSAN-by-VSAN basis.

IVR

In another step toward deploying efficient, cost-effective, and consolidated storage networks, the IBM Storage Networking SAN50C-R switch supports IVR, the industry's first routing function for FC. IVR allows selective transfer of data between specific initiators and targets on different VSANs while maintaining isolation of control traffic within each VSAN. With IVR, data can transit VSAN boundaries while maintaining control plane isolation, maintaining fabric stability and availability.

IVR is one of the feature enhancements that are provided with the enterprise advanced software package. It eliminates the need for external routing appliances, greatly increasing routing scalability while delivering line-rate routing performance, simplifying management, and eliminating the challenges that are associated with maintaining separate systems. Deploying IVR means lower total cost of SAN ownership

FCIP for remote SAN extension

Data distribution, data protection, and business continuance services are significant components of today's information-centric businesses. The capability to efficiently replicate critical data on a global scale helps ensure a higher level of data protection for valuable corporate information and increases the usage of backup resources and lowers total cost of storage ownership.

- Building on expertise and knowledge of IP networks, the IBM Storage Networking SAN50C-R switch uses open-standards FCIP to break the distance barrier of current FC solutions, enabling the interconnection of SAN islands over extended distances.
- The IBM Storage Networking SAN50C-R switch dramatically enhances hardware-based FCIP compression performance for both high-bandwidth and low-bandwidth links, providing immediate cost savings for an expensive WAN infrastructure. The IBM Storage Networking SAN50C-R achieves up to a 43:1 compression ratio, with typical ratios of 4:1 to 5:1 over a wide variety of data sources.
- The IBM Storage Networking SAN50C-R switch supports hardware-based IPsec encryption for secure transmission of sensitive data over extended distances. Hardware enablement of IPsec helps ensure high throughput. Used together, hardware-based compression and hardware-based encryption provide high-performance, highly secure SAN extension capabilities.

I/O Acceleration services

The IBM Storage Networking SAN50C-R switch supports I/O Acceleration (IOA) services, an advanced software package that can improve application performance when storage traffic is extended across long distances. When FC and FCIP write acceleration are enabled, WAN throughput is optimized through reduced latency for command acknowledgments. Similarly, the IBM Storage Networking SAN50C-R switch supports FC and FCIP tape write acceleration, which allows operation at nearly full throughput over WAN links for remote tape backup and restore operations. IOA can be deployed with disk data replication solutions to extend the distance between data centers or reduce the effects of latency. IOA can also be used to enable remote tape backup and restore operations without significant throughput degradation. Here are the main features of IOA:

- Extension of the acceleration service as a fabric service to any port in the fabric, regardless of where it is attached
- Fibre Channel Write Acceleration (FC-WA) and Fibre Channel tape acceleration (FC-TA)
- FCIP write acceleration (FCIP-WA) and FCIP tape acceleration (FCIP-TA)
- FC and FCIP compression
- HA by using port channels with acceleration over FC and FCIP
- Unified solution for disk and tape IOA over MANs and WANs
- Speed-independent acceleration that accelerates 2/4/8/16-Gbps FC links and consolidates traffic over 8/16 Gigabit ISLs

The IOA feature is not supported for the FICON protocol.

Mainframe support

IBM Storage Networking SAN50C-R is mainframe-ready and supports IBM Z FICON and Linux environments that are provided with the mainframe advanced software package. Qualified by IBM for attachment to all FICON-enabled devices in an IBM Z operating environment, IBM Storage Networking SAN50C-R switches support transport of the FICON protocol in both cascaded and non-cascaded fabrics, and an intermix of FICON and open-system FCP traffic on the same switch. VSANs simplify intermixing of SAN resources among IBM z/OS, mainframe Linux, and open-system environments, enabling increased SAN utilization and simplified SAN management.

VSAN-based intermix mode eliminates the uncertainty and instability that is often associated with zoning-based intermix techniques. VSANs also eliminate the possibility that a misconfiguration or component failure in one VSAN will affect operations in other VSANs. VSAN-based management access controls simplify partitioning of SAN management responsibilities between mainframe and open systems environments, enhancing security. FICON VSANs can be managed by using the standard DCNM, the CLI, or CUP-enabled management tools, including Resource Measurement Facility (RMF) and Dynamic Channel Path Management (DCM).

Advanced traffic management

The following advanced traffic-management capabilities are integrated as standard on the IBM Storage Networking SAN50C-R switch:

- Virtual output queue (VOQ): Helps ensure line-rate performance on each port, independent of traffic pattern, by eliminating head-of-line blocking.
- Port channels: Allow users to aggregate up to 16 physical ISLs into a single logical bundle, providing optimized bandwidth usage across all links. The bundle can consist of any speed-matched ports from any module in the chassis, helping ensure that the bundle can remain active even during a module failure.

- FSPF-based multipathing: Provides the intelligence to load balance across up to 16 equal-cost paths and, during a switch failure, dynamically reroute traffic.

The following extra advanced traffic-management capabilities are available on the IBM Storage Networking SAN50C-R switch with the optional enterprise advanced software package to simplify deployment and optimization of large-scale fabrics:

- Up to 253 B2B credits: Can be assigned to an individual port for optimal bandwidth usage across long distances.
- QoS: Can be used to manage bandwidth and control latency to prioritize critical traffic for specific applications.
- IVR: Eliminates the need for external routing appliances, greatly increasing routing scalability while delivering line-rate routing performance, simplifying management, and eliminating the challenges that are associated with maintaining separate systems.
- SCSI flow statistics: Collects logical unit number (LUN)-level SCSI flow statistics, including read, write, and error statistics, for any combination of initiators and targets.

Comprehensive solution for robust network security

To address the need for failure-proof security in storage networks, the IBM Storage Networking SAN50C-R switch includes as standard an extensive security framework to protect highly sensitive data crossing today's enterprise networks:

- When the Smart Zoning feature is enabled, IBM Storage Networking c-type family fabrics provision the hardware access control entries that are specified by the zone set more efficiently. Doing so avoids the superfluous entries that allow servers (initiators) to talk to other servers, or allow storage devices (targets) to talk to other storage devices.

  This feature makes larger zones with multiple initiators and multiple targets feasible without excessive consumption of hardware resources. Thus, smart zones can correspond to applications, application clusters, hypervisor clusters, or other data center entities, saving the time that administrators previously spent creating many small zones, and enabling the automation of zoning tasks.
- Intelligent packet inspection is provided at the port level, including the application of ACLs for hardware enforcement of zones, VSANs, and advanced port-security features.
- Extended zoning capabilities are provided to help ensure that LUNs can be accessed only by specific hosts (LUN zoning) to limit SCSI read commands for a certain zone (read-only zoning), and to restrict broadcasts to only selected zones (broadcast zones).

The following additional advanced security-management capabilities are available on the IBM Storage Networking SAN50C-R switch with the enterprise advanced software package to further help ensure the security of large-scale fabrics:

- Switch-to-switch and host-to-switch authentication helps eliminate disruptions that might occur because of unauthorized devices connecting to a large enterprise fabric.
- Port security locks down the mapping of an entity to a switch port to help ensure that SAN security is not compromised by the connection of unauthorized devices to a switch port.
- VSAN-based access control allows customers to define roles in which the scope of the roles is limited to certain VSANs.
- FC-SP provides switch-switch and host-switch DH-CHAP authentication that supports RADIUS and TACACS+ to help ensure that only authorized devices access protected storage networks.
- A comprehensive IPsec protocol suite delivers secure authentication, data integrity, and hardware-based encryption for both FCIP and iSCSI deployments.

- Digital certificates are issued by a trusted third party and used as electronic passports to prove the identity of certificate owners.
- Fabric binding for open systems helps ensure that the ISLs are enabled only between switches that are authorized in the fabric binding configuration.

Ease of management

To meet the needs of all users, the IBM Storage Networking SAN50C-R switch provides three principal modes of management: CLI, DCNM, and integration with third-party storage management tools.

The IBM Storage Networking SAN50C-R switch presents a consistent, logical CLI. Adhering to the syntax of the widely known I/O Subsystem (IOS) Software CLI, which is easy to learn and delivers broad management capabilities. The CLI is an efficient and direct interface that provides optimal capabilities to administrators in enterprise environments.

DCNM for SAN is an application that simplifies management across multiple switches and converged fabrics. It provides robust features to meet the routing, switching, and storage administration needs of present and future virtualized data centers, streamlines provisioning of the unified fabric, and proactively monitors SAN components. DCNM SAN can be used independently or with third-party management applications.

The solution is designed to scale to large enterprise deployments through a scale-out server architecture with automated failover capability. These capabilities provide a resilient management system that centralizes infrastructure and path monitoring across geographically dispersed data centers. DCNM SAN base management functions are available at no additional charge, but advanced features are unlocked by the DCNM SAN Advanced license. The DCNM SAN application can be installed on Linux and Microsoft Windows OSs and supports both PostgreSQL and Oracle databases.

Advanced software packages

The IBM Storage Networking SAN50C-R switch can be further enhanced through extra optional licensed software packages that offer advanced intelligence and functions.

The following software packages are available:

- Enterprise Package: This package includes a set of traffic engineering and advanced security features, such as extended-distance B2B credits, IVR, QoS, switch-to-switch and host-to-switch authentication, LUN zoning, and read-only zones, which are recommended for enterprise SANs.
- DCNM SAN Advanced: DCNM is an easy-to-use application that simplifies management across multiple switches and fabrics. Focused on supporting efficient operations and management of SSD-aware fabrics, it provides a robust framework and rich feature set that meets the routing, switching, and storage administration needs of present and future virtualized data centers. DCNM streamlines provisioning of the unified fabric and proactively monitors the SAN components.
- I/O Accelerator Services package: The IBM Storage Networking SAN50C-R switch supports IOA services, which are an advanced software package that can improve application performance when storage traffic is extended across long distances. When FC and FCIP write acceleration is enabled, WAN throughput is optimized through reduced latency for command acknowledgments.
- Mainframe Package: This package is a comprehensive collection of features that are required for using the IBM Storage Networking c-type Series switches in mainframe storage networks. These features include FICON protocol, FICON tape acceleration (read and write), CUP management, switch cascading, fabric binding, and intermixing.

Table 1-11 lists the base and license functions.

Table 1-11 Base and license functions

Description	Included or optional
DCNM SAN base version	Included
SAN Extension over IP	Included
DCNM SAN Advanced	Optional
Enterprise	Optional
On-demand Port Activation	Optional
Mainframe Package	Optional
I/O Acceleration Service	Optional

IBM Storage Networking SAN50C-R product specifications

This section describes the product specifications for the IBM Storage Networking SAN50C-R switch. The IBM Storage Networking SAN50C-R switch offers distributed intelligent fabric services, cost-effective and high-performing FC and FCIP connectivity for open systems, remote SAN extension, and fast DR.

Table 1-12 lists the specifications.

Table 1-12 Specifications for the IBM Storage Networking SAN50C-R switch

Feature	Description
Product compatibility	IBM Storage Networking c-type Family
Software compatibility	NX-OS Release 8.1(1b) or later
Cards, ports, and slots	Fixed configuration with 40 ports of 16-Gbps FC and 10 ports of 10 GbE
Fabric services	► Name server ► Internet Storage Name Server (iSNS) ► RSCN ► Login services ► FCS ► Public loop ► Broadcast ► In-order delivery
Advanced functions	► VSAN ► IVR ► Port channel with multipath load-balancing ► Flow-based and zone-based QoS ► FCIP tape read and write acceleration

Feature	Description
Diagnostic and troubleshooting tools	► POST diagnostic tests ► GOLD ► Internal port loopbacks ► SPAN and RSPAN ► FC traceroute ► FC ping ► FC debug ► Fabric Analyzer ► Syslog ► Online system health ► Port-level statistics ► RTP debug
Network security	► VSANs ► ACLs ► Per-VSAN RBAC ► FC zoning ► N-port worldwide name (WWN) ► N-port FC-ID ► Fx-port WWN ► Fx-port WWN and interface index ► Fx-port domain ID and interface index ► Fx-port domain ID and port number ► iSCSI zoning ► iSCSI name ► IP address ► FC-SP ► DH-CHAP switch-to-switch authentication ► DH-CHAP host-to-switch authentication ► Port security and fabric binding ► IPsec for FCIP and iSCSI ► Internet Key Exchange (IKE) v1 and IKEv2 ► Management access ► SSHv2 implementing AES ► SNMPv3 implementing AES ► SFTP
FICON	► FC-SB-6 compliant ► Cascaded FICON fabrics ► Intermix of FICON and FCP traffic ► CUP management interface
Serviceability	► Configuration file management ► ISSU for FC interfaces ► Call Home ► Power-management LEDs ► Port beaconing ► System LED ► SNMP traps for alerts ► Network boot

Feature	Description
Performance	► Port speed: 2/4/8-Gbps and 4/8/16-Gbps autosensing, optionally configurable ► Buffer credits: 64 per port (shared-mode ports) and up to 253 on an individual port (dedicated-mode ports with optional Enterprise Package license activated) ► Ports per chassis: Fourty ports of 16-Gbps FC, 8 ports of 10-GbE FCoE, and 2 ports of 1/10-GbE ► Ports per rack: Up to 1050 ► Port channel: Up to 16 physical links ► FCIP tunnels: Up to 6 per port
Reliability and availability	► ISSU ► Hot-swappable, 2+1 redundant power supplies ► Hot-swappable fan tray with integrated temperature and power management ► Hot-swappable SFP+ optics ► Passive backplane ► Stateful process restart ► Any port configuration for port channels ► Fabric-based multipathing ► Per-VSAN fabric services ► Port tracking ► VRRP for management and FCIP or iSCSI connections ► Online diagnostic tests
Network Management	► Access methods through the Supervisor-1 Module: – Out-of-band 10/100/1000 Ethernet port – RS-232 serial console port – In-band IP over FC ► Access methods through the FC switching module In-band FICON CUP over FC ► Access protocols: – CLI that uses console and Ethernet ports – SNMPv3 that uses Ethernet ports and in-band IP over FC access – FICON CUP ► Distributed Device Alias service ► Network security: – Per-VSAN RBAC that uses RADIUS-based and TACACS+-based AAA functions – SFTP – SSHv2 implementing AES – SNMPv3 implementing AES – Management applications ► CLI ► DCNM GUI
Programming interfaces	► Scriptable CLI ► DCNM web services API ► NX-API

IBM Storage Networking SAN50C-R physical specifications

Table 1-13 details the specific requirements for planning the installation of the devices in the data center rack.

Table 1-13 Physical and environmental specifications for IBM Storage Networking SAN50C-R

Feature	Description
Environmental	► Temperature, ambient operating: 0 - 40°C (32 - 104°F) ► Temperature, ambient nonoperating and storage: -40 - 70°C (-40 - 158°F) ► Relative humidity, ambient (noncondensing) operating: 10 - 90% ► Relative humidity, ambient (noncondensing) nonoperating and storage: 10 - 95% ► Altitude, operating: -60 to 2000 m (-197 to 6500 ft)
Physical dimensions	► Dimensions (H x W x D): 9.75 cm x 43.74 cm x 54.36 cm (3.84 in. x 17.22 in. x 21.4 in.), 2RUs. (All units rack-mountable in standard 19-inch EIA rack) ► Weight of fully configured chassis: 10.2 kg (22.4 lb)
Power and cooling	► Power supply: 300W AC ► Power cord: Notched C15 socket connector connecting to C16 plug on power supply ► AC input characteristics ► 100 - 240 V AC (10% range) ► 50 - 60 Hz (nominal) ► Airflow (front to back) ► 200 LFM through system fan assembly ► IBM suggests maintaining a minimum air space of 6.4 cm (2.5 in.) between walls and chassis air vents and a minimum horizontal separation of 15.2 cm (6 in.) between two chassis to prevent overheating.

Integrated Supervisor Module

The non-removable IBM Storage Networking SAN50C-R Integrated Supervisor Module provides the control and management functions of the IBM Storage Networking SAN50C-R switch, and it includes 40 integrated 16-Gbps FC switching ports and eight 10-Gbps Ethernet FCoE port modules.

The IBM Storage Networking SAN50C-R Integrated Supervisor Module has an IBM PowerPC® 8572E processor. It also has an internal CompactFlash card that provides 4 GB of storage for software images. The non-volatile random access memory (NVRAM) consists of a battery, a battery controller, and 512 K x16 static random access memory (SRAM). The SRAM stores event logs and the core dumps that must be stored after a power cycle occurs.

Fan modules

The IBM Storage Networking SAN50C-R switch has two fan trays that are installed vertically at the back of the chassis. Each fan module can be removed while the other fan module continues to move air through the chassis.

Power supplies

The IBM Storage Networking SAN50C-R switch has the capacity for up to three hot-swappable 300 W AC PSUs. Each PSU can provide information about itself to the supervisor. The two types of information that are available are status information (output voltage, fan state, and unit state) and part information (serial number and revision). When connected to a nominal 110 or 220 VAC input, each PSU provides 300 W of output power.

In the default configuration and with all three PSUs installed, the IBM Storage Networking SAN50C-R switch has N+1 PSU redundancy. The only power redundancy mode that is available is redundant; combined mode is not supported on this platform.

Typically, when FCOE ports are not used, grid redundancy with the IBM Storage Networking SAN50C-R switch is possible with only two PSUs.

Support for power redundancy with two online PSUs

The IBM Storage Networking SAN50C-R switch running NX-OS 8.1(1) or later supports power redundancy with two online PSUs. To enable power redundancy in this scenario, FCoE ports must be brought to the ADMIN DOWN state. The power supplies in an IBM Storage Networking SAN50C-R switch work in the power modes that are shown in Table 1-14 when the FCoE ports are in the ADMIN DOWN state.

Table 1-14 Power modes when the FCoE ports are in the ADMIN DOWN state

Three online PSUs	Two PSUs are connected to Grid A and one PSU is connected to Grid B; they work in N+1 redundant mode.
Two online PSUs	One PSU is connected to Grid A and another PSU is connected to Grid B; they work in N:N redundant mode.
One online PSU	The PSU is connected to any one grid; it works in non-redundant mode.

The power supplies in an IBM Storage Networking SAN50C-R switch work in the power modes that are shown in Table 1-15 when the FCoE ports are in the ADMIN UP state.

Table 1-15 Power modes when the FCoE ports are in the ADMIN UP state

Three online PSUs	These PSUs work in N+1 redundant mode.
Two online PSUs	These PSUs work in non-redundant mode.
One online PSU	This PSU works in non-redundant mode. In this case, FCoE ports automatically switch to an error-disabled state to save power for FC and IP ports.

Supported transceivers

The IBM Storage Networking SAN50C-R switch supports the following transceivers:

- 8 Gbps SW/LW LC Enhanced SFP+
- 10 GbE SR/LR/ER LC SFP+
- 16 Gbps SW/LW/ELW LC SFP+
- 4/8/16-Gbps FC LW SFP+ DWDM SM DDM 13 dB 40 km
- 4/8/16-Gbps FC LW SFP+ CWDM SM DDM 13 dB 40 km
- 4/8/16-Gbps FC/FICON LW SFP+ DWDM SM DDM 1550 nm 13 dB 40 km
- 2/4/8-Gbps FC LW SFP+ DWDM SM DDM 80 km

- 2/4/8-Gbps FC LW SFP+ CWDM SM DDM 23 dB 70 km
- 2/4/8-Gbps FC LW SFP+ SM DDM 80 km

Rack requirements

The rack-mount kit enables you to install the switch into racks of varying depths. You can use the rack-mount kit parts to position the switch with easy access to either the port connections end of the chassis or the end of the chassis with the fan and power supply modules.

General requirements for racks

The rack must be one of the following types:

- Standard 19-inch 4-post EIA rack, with mounting rails that conform to imperial universal hole spacing, per section 1, *Standard Open Rack Requirements,* of ANSI/EIA-310-D-1992.
- Standard two-post telco rack, with mounting rails that conform to imperial universal hole spacing per section 1, *Standard Open Rack Requirements,* of ANSI/EIA-310-D-1992.

Rack requirements for the IBM Storage Networking SAN50C-R chassis

- The minimum vertical rack space per chassis is 3.5 in. (8.8 cm).
- The width between the mounting rails must be at least 17.75 in. (45.1 cm). For 4-post EIA racks, this is the distance between the two front rails and rear rails.

1.5 IBM c-type software

This chapter introduces the IBM Storage networking c-type range software OS that runs on the switch hardware and the management tools.

These topics are covered:

- NX-OS
- DCNM

1.5.1 NX-OS

The NX-OS software is a data center OS that runs on the IBM c-type range of SAN switches. The software supports the modular switch hardware to provide the resiliency and serviceability at the core of its design. NX-OS has its roots from the proven Cisco MDS 9000 SAN-OS Software. NX-OS provides the level of continuous availability that enterprise SAN fabric solutions in the modern data center require for providing mission-critical solutions.

NX-OS offers reliability, innovation, and operational consistency across data center platforms. NX-OS runs on the Cisco Nexus Family of hardware-based network switches and Cisco MDS 9000 family storage switches and Cisco UCS 6000 Series Fabric Interconnects.

> **Note:** Initial configuration of NX-OS is described in Chapter 6, “Initial connectivity and setup” on page 183.

1.5.2 Data Center Network Manager

DCNM is the comprehensive management solution for all NX-OS network deployments spanning LAN fabrics, SAN fabrics, and IP Fabric for Media (IPFM) networking in the data center that is powered by Cisco. At deployment time, a specific operational mode must be selected. In this document, we describe only the SAN fabric mode and refer to it as *DCNM-SAN*.

Cisco DCNM provides management, control, automation, monitoring, visualization, and troubleshooting across IBM c-type switching solutions. DCNM provides an alternative to the CLI for switch configuration commands. It provides a powerful GUI for managing Cisco unified data center networks (FC, Ethernet, and FCoE).

Its value proposition becomes more apparent when fabric-level actions are required, as opposed to single device tasks. DCNM increases overall data center infrastructure uptime and reliability, thus improving business continuity. Statistically, the DCNM license is always associated to IBM c-type Director sales, but it has a reduced traction with stand-alone fabric switches due to the simplicity of the topology. In those situations where the fabric is just one switch, there is no strong need for a network-wide management tool, and some customers prefer more automation (for example, autozone) versus more management capabilities.

DM is both a stand-alone or embedded tool inside DCNM. DM provides a graphic representation of an IBM c-type family switch chassis, including the installed switching modules, the supervisor modules, the status of each port within each module, the power supplies, and the fan assemblies. Also, DM provides detailed information for verifying or troubleshooting a device-specific configuration and can be considered on feature-parity to the NX-OS CLI. With DM, you can perform switch-level configurations like:

- Configure zones for multiple VSANs.
- Manage ports, port channels, and trunking.
- Manage SNMPv3 security access to switches.
- Manage CLI security access to the switch.
- Manage alarms, events, and notifications.
- Save and copy configuration files and software images.
- View a hardware configuration.
- View a chassis, module, port status, and statistics.

Note: Initial configuration of DCNM is described in Chapter 6, “Initial connectivity and setup” on page 183.

2

IBM Storage Networking c-type features

This chapter provides insights into some standard and licensed features of the c-type storage area network (SAN) switches that are offered by IBM. The features that are described in this chapter are only a subset of all the features that are available.

The close partnership between IBM and Cisco continues to provide various information technology expertise that is leveraged to offer the best technical solutions and help you create and deploy updated platforms that provide your data center with flexible and secure connectivity. We are constantly learning and adapting to meet your evolving business needs. IBM is dedicated to building a secure foundation as technology continues to advance and provide extraordinary outcomes in the data center today from core to edge.

Globally over the years, more businesses have been deploying large-scale SANs due to exponential growth in the data center. IBM c-type SAN switches offer many features that will help clients lower total cost of ownership (TCO) coupled with the ability to consolidate SAN environments and deliver the tools for IT professionals to manage, report, diagnose, and monitor their SAN infrastructures. The ability to incorporate robust security into on-premises business deployments has transformed and crossed over to off-premise cloud environments.

The IBM c-type family portfolio has four SAN switches and three SAN Directors. Each of these components comes with hardware and software that is needed to implement a resilient, secure fabric for your storage and server devices. All IBM c-type hardware comes with 1 year of onsite 24x7 same-day maintenance with optional service options to upgrade. The IBM c-type platform offers a mainframe IBM Fibre Connection (FICON) intermix solution, which includes traffic isolation, management, quality of service (QoS), and scaling.

IBM mainframes have been the workhorse of the computing world for more than 6 decades and have gained the trust of many large organizations who entrust IBM Z servers with their most demanding mission-critical applications.

The following topics are covered in this chapter:

- IBM c-type features
- IBM c-type Mainframe Package features
- VSAN
- Port channels
- FICON and CUP
- Fibre Channel Protocol
- Fibre Channel over IP
- Trunking
- Fabric Shortest Path First
- Zoning
- Quality of service
- N_Port ID Virtualization and N_Port Virtualization
- Non-Volatile Memory Express over Fibre Channel
- Forward Error Correction
- In-service software upgrades
- Cisco SAN Analytics

2.1 IBM c-type features

In addition to the standard features that are provided at no cost in the NX-OS code, there is also the Mainframe Package and three other optional licensed feature packages that address different enterprise requirements:

- Mainframe Package:
  - Support for Fibre Channel (FC) / FICON intermix with virtual storage area networks (VSANs)
  - FICON with IBM Control Unit Port (CUP)
  - FICON switch cascading
  - Fabric binding
- Enterprise Package:
  - FC TrustSec and IP security (IPsec) encryption
  - Port security
  - VSAN-based access control
  - Host/switch authentication
  - QoS
  - Inter-VSAN Routing (IVR)
  - Extended buffer credits
- SAN Analytics package:
  - I/O flow-level traffic visibility
  - Frame headers inspection for analyses and correlation

- Data Center Network Manager (DCNM) SAN Advanced package:

  DCNM SAN Advanced (previously known as the FMS package) allows centralized discovery for those switches that are connected through Inter-Switch Links (ISLs), and roaming profiles, which eliminate the manual effort to maintain consistent topology maps and customized definitions across multiple client stations. It also adds advanced management and troubleshooting capabilities that include historical performance management, reporting, and more.

Important: Switch-based licenses are mapped to the serial number of an IBM c-type chassis, which is 11 digits.

Note: For more information about standard and licensed features, see Cisco MDS 9000 Series Licensing Guide, Release 8.x.

2.2 IBM c-type Mainframe Package features

In this section, we further break down the IBM c-type Mainframe Package features:

- Support for FC/FICON intermix with VSANs: Helps to enable workload isolation, division of management, and control of fault-domains, all within a single platform with superior reliability and scalability.
- Support for the CUP interface: This support within each FICON VSAN allows for management and performance data collection natively from the IBM Z server.
- Support for cascaded FICON topologies: Customers can scale to large FICON environments. They also can support FICON environments spanning across multiple data centers.

When the Mainframe Package is installed, IBM c-type devices can be used with the IBM z13®, IBM z14, and IBM z15™. The package includes support for the new FICON Express 16S feature and complete interoperability with Forward Error Correction (FEC) capabilities.

The IBM c-type Mainframe Package provides the following features:

- FICON VSANs: Like logical partitions (LPARs) on IBM Z servers, VSANs provide hardware-based partitioning of a single physical infrastructure into multiple logical FICON SANs. FICON VSANs provide isolation of traffic, segmentation of management, and management of fault domains. FICON VSANs can be used to separate production environments from test or development environments, FICON from Fibre Channel Protocol (FCP) applications, or even disk storage from tape storage. This separation can be achieved without compromising scalability, availability, manageability, and network security. IBM c-type FICON Directors support up to eight FICON VSANs, each with its own CUP device.
- Dynamic port number assignment: All FICON port numbers are virtualized in the IBM c-type switches, allowing any port number to be allocated on any port within the FICON VSAN by using the defined range of 0x00 - 0xFD. When multiple FICON VSANs are used for workload segmentation, ports for each FICON VSAN can be allocated on a per-port basis with no restrictions regarding line-card allocation or use of duplicate port numbers.

- FICON CUP: Implementation of CUP in the IBM c-type switch enables in-band management of the director from IBM Z servers. The CUP device also provides IBM Resource Management Facility (RMF) Type 74 Subtype 7 records to IBM Z, thus allowing host performance management software (such as IBM RMF) to create FICON Director activity reports in time synchronization with the rest of the system reports. The FICON VSAN also allocates special logical FICON port numbers for Fibre Channel over IP (FCIP) links and port channels so that performance of these special link types can be tracked at the IBM Z level.
- FICON Cascaded Directors: Director cascading supports a topology for FICON devices in which ISLs can be used between IBM Z and I/O devices. The fabric binding feature, which is required for cascaded FICON, allows only preauthorized director to participate in the FICON fabric, thus helping ensure high integrity for FICON fabrics. Although more complex FICON topologies will technically work, IBM typically supports two-hop ISL configurations except in certain specific channel extension applications. For more information about the topologies that IBM Z supports, see Cisco MDS 9000 FICON Solutions.

  The IBM Storage Networking SAN384C-6 and Storage Networking SAN192C-6 are fully interoperable with Dense Wavelength-Division Multiplexing (DWDM) solutions in addition to offering 8-Gbps, 16-Gbps extended long wavelength (LWL), and 32-Gbps extended LWL optics to facilitate metropolitan area applications.
- Enhanced ISL aggregation: The IBM c-type family and all earlier FICON-capable Directors allow exploitation of the no-cost port channel feature. Port channels are virtual interfaces that consist of multiple physical ISLs. Port channels have three valuable attributes:
  - The member links can span any available port that is on the installed line cards, with no port group or application-specific integrated circuit (ASIC) limitations. This situation makes sure that the port channel has a role for scaling bandwidth between switches and increases reliability of the interconnection.
  - The length of the member links of a port channel can be different.
  - Up to 16 physical members can be part of a single logical port channel.

  Given these three attributes, port channels are an excellent mechanism for interconnecting metropolitan area data centers with disparate length site-to-site links. This flexibility also provides high availability (HA) by reducing the size of failure domains.
- Lossless in-order delivery: When you use port channels for cascaded FICON VSANs, member links of the port channel can be nondisruptively disabled or enabled. In fact, new links can even be added to or removed from an active port channel without the IBM Z server experiencing a single error.
- The IBM c-type Mainframe Package can enable the connection of IBM Storage Networking c-type family storage director with the IBM Z server with both FICON and FCP. This support includes the IBM Storage Networking SAN384C-6, IBM Storage Networking SAN192C-6, and IBM Storage Networking SAN50C-R.

2.3 VSAN

VSAN technology provides security and stability, and helps enable workload isolation between device end points that are physically connected to the same physical hardware for secure sharing of physical infrastructure and enhanced FC/FICON intermix support.

VSANs can span switch modules and can vary in size. For example, one VSAN with multiple FC ports can span different line cards. Adding ports to a VSAN is a nondisruptive process. VSANs can contain up to 239 switches and have an independent address space that allows identical Fibre Channel IDs (FCIDs) to be used simultaneously in different VSANs. The maximum number of ports for a FICON VSAN is 253 (0xFE is CUP and 0xFF is reserved) due to FICON specific addressing rules that do not apply to FC.

With VSANs, you can specify and assign specific cascaded links. The VSAN feature allows you to maintain throughput and enable complete traffic isolation when using the same physical hardware throughout the FICON fabric, as shown in Figure 2-1.

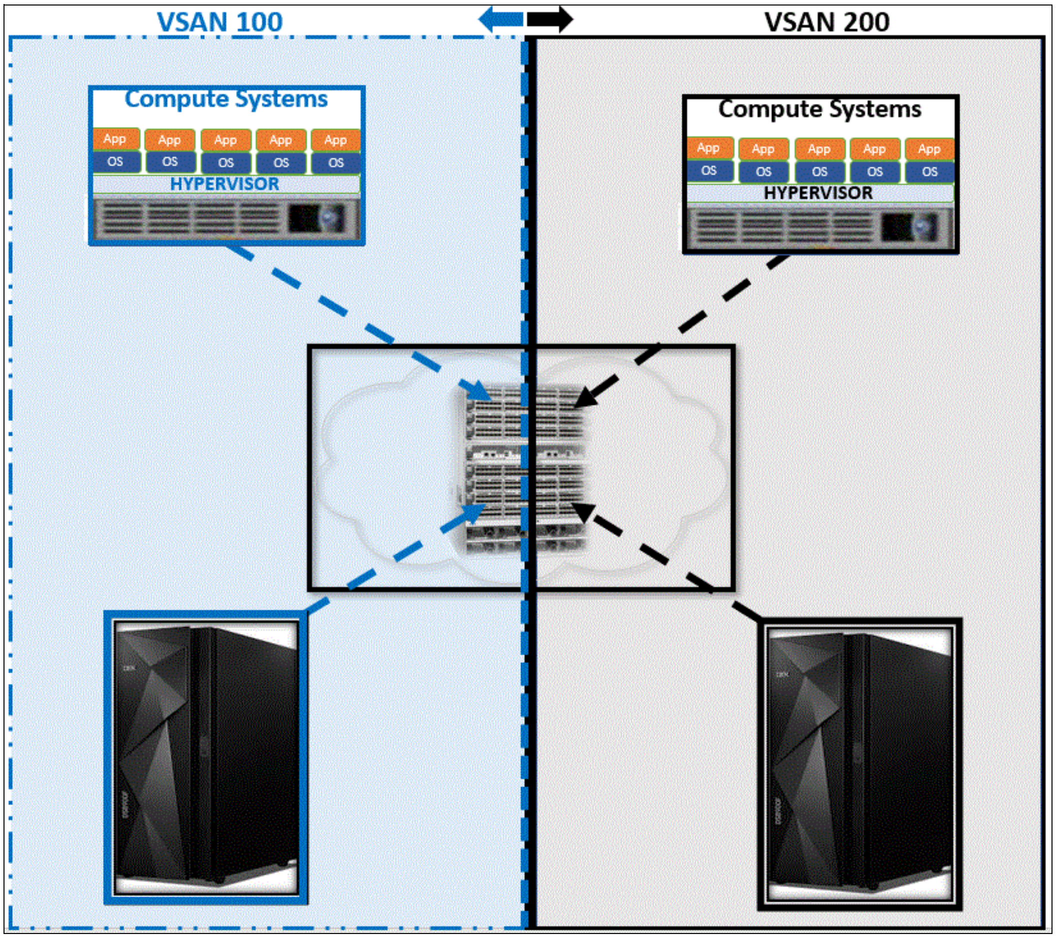

Figure 2-1 VSAN 100 and 200 Isolation

VSAN technology was pioneered by Cisco in 2003 and standardized by the International Committee for Information Technology Standards (INCITS) T11 in 2004. Since then, it has been part of the Cisco MDS 9000 offering, and it is also supported by all IBM c-type switches. VSAN technology has several unique points and differentiators:

- VSANs are included for free in all NX-OS releases.
- VSANs are available on all IBM c-type switches, whether they are director, fabric switches, or multiprotocol devices.
- VSANs are supported even if N_Port Virtualization (NPV) mode is configured on the switch.
- All VSANs are equivalent. They all can accommodate F-ports and E-ports with no limitations.
- The number of VSANs that was certified is 80 on director and 32 on switches, but more are technically possible, with a range of 4096.

Additional VSAN-related features include VSAN-based access control and IVR.

VSAN-based access control

This feature enables customers to define roles where the scope of the roles is limited to certain VSANs. For example, a network administrator role can be set up to allow configuration of all platform-specific capabilities, and VSAN administrator roles can be set up to allow only configuration and management of specific VSANs. VSAN-based access control reduces SAN disruptions by localizing the effects of user errors to the VSANs for which the user has administrative privileges. It adds a layer of security where only administrators can configure switches within specified VSANs.

Inter-VSAN Routing

In another step toward deployment of efficient, cost-effective, and consolidated storage networks, the IBM c-type supports IVR, the industry’s first and most efficient routing function for FC.

IVR allows selective transfer of data between specific initiators and targets on different VSANs while maintaining isolation of control traffic within each VSAN. With IVR, data can move across VSAN boundaries while maintaining control-plane isolation, thus maintaining fabric stability and availability.

IVR eliminates the need for external routing appliances, greatly increasing routing scalability while delivering line-rate routing performance, simplifying management, and eliminating the challenges that are associated with maintaining separate systems. IVR reduces the total cost of SAN ownership.

IVR allows you to do the following tasks:

- Access resources that are physically connected to an alternative VSAN that is not directly connected to your current VSAN.
- Integrate other vendor switches with a possible interoperation mode.
- Transport data traffic to different VSANs without having to form a single logical fabric.
- Create a disaster recovery (DR) solution in an efficient way when used with FCIP.
- Establish an interconnect between one or more VSANs.

With the IVR Zone wizard inside the DCNM management GUI, an administrator can simplify the process of setting up and configuring the IVR zones in a fabric. The wizard can help with checking all of the switches in the fabric for the code that is running on the switch.

Within the VSAN, IVR virtualizes the remote end devices in the native VSAN by using a virtual domain. As the end devices in disparate VSANs are configured, they modify the FFC header for all the back and forth communication. IVR uses an IVR SAN topology to determine how the traffic is routed between end points.

Note: FICON is not supported by IVR.

2.4 Port channels

The port channels term refers to the aggregation of multiple physical interfaces into one logical interface to provide HA and load0-balancing while allowing for higher aggregated bandwidth within a port channel. Port channels can consist of any speed-matched ports from any IBM c-type module in the switch chassis, which helps ensure that the aggregated links remain active even in the event of a module failure.

Any configuration changes that you apply to a port channel are applied to each member interface of that port channel.

Note: A port channel is *operationally up* when at least one of the member ports is up and that port's status is channeling. The port channel is *operationally down* when all member ports are operationally down.

2.5 FICON and CUP

FICON provides interface capabilities that enhance the IBM c-type family by supporting both mainframe and open systems SAN environments while in-bound management of the switch from FICON processors is provided by CUP support. The CUP function allows a z/OS system to communicate with the FICON Director through channel programs, which includes functions like blocking and unblocking ports, and continues to be an important in-band management tool that provides basic performance monitoring, powerful diagnostic capabilities, and error reporting functions. This enhanced visibility and capability allows IBM Z management of FICON SANs.

Over the last few years, there have been many new FICON CUP enhancements that were introduced by IBM, Cisco, and other vendors. The advanced technology enhancements that it provides are geared toward making it easier for administrators to get a better overall insight into FICON SAN fabrics because IBM Z has evolved and configurations have become more complex.

Note: IBM c-type FICON Directors support up to eight FICON VSANs, each with its own CUP device.

2.6 Fibre Channel Protocol

FC is a data transfer protocol running at up to 64 Gbps serially, providing lossless delivery of raw block data. FC is primarily used to connect servers, storage arrays, and tape libraries in open systems environments. FC networks form a switched fabric because the switches in a SAN can be configured to operate in unison as one or more large logical switches.

Currently, IBM c-type switches support many different speeds of connectivity up to 32 Gbps. In addition to this, 64 Gbps is expected to become available in the future. When available, the new linecards can be added to existing director chassis with no service disruption. As your data center grows or shrinks, these SAN switches can keep up with the dynamic pace of the changing environment.

The FCP ports on all the modules support an auto-sensing Small Form-factor Pluggable (SFP) that can determine the speed that needed to operate in that fabric.

2.7 Fibre Channel over IP

IBM c-type Directors can support FCP and FCIP ports on modules inside the chassis. These modules enable customers to use both FCP and FCIP on the same module and enable large scalable deployments of SAN fabrics. The FCIP module delivers outstanding performance while reducing the latency for disk and tape operations by using FCIP acceleration features like FCIP write acceleration and FCIP tape write and read acceleration.

> **Note:** FCIP is supported by IBM Storage Networking SAN50C-R and the IBM 24/10 Port SAN Extension Module.

In an expensive wide area network (WAN), the IBM c-type switches can use hardware-based compression, which enhances performance for both high and low speed links while keeping costs down. Hardware-based encryption can also be enabled to help secure sensitive traffic with IPsec.

Port channels of up to 16 links can be grouped to take advantage of HA and increased aggregate throughput.

2.8 Trunking

Trunking is a commonly used storage industry term, but unfortunately it has assumed different meanings. The NX-OS software and switches in the IBM c-type family adopt the meaning that is prevalent in the industry and implement trunking and port channels as follows:

- Port channels enable several physical links to be combined into one aggregated logical link.
- Trunking enables a link transmitting frames in the EISL format to carry (trunk) traffic from multiple VSANs. For example, when trunking is operational on an E-port, that E-port becomes a Trunking E-port (TE-port). A TE-port is specific to IBM c-type switches. An industry-standard E-port can link to other vendor switches and is referred to as a *non-trunking interface*.

2.9 Fabric Shortest Path First

Fabric Shortest Path First (FSPF) is a common path selection protocol that is used by FC fabrics and enabled by default. It automatically calculates the best path between any two switches in a SAN fabric and the most efficient way to route packets.

FSPF can dynamically establish the shortest and quickest path between any two switches supporting multiple paths and automatically determine an alternative path around a failed link. It provides a preferred route when two equal paths are available.

FSPF Link Cost tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost that is associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can be 1 - 65,535. For example, the default cost for 1 Gbps is 1000 and for 8 Gbps is 125.

2.10 Zoning

In this section, we describe four IBM c-type zoning features. Basic zoning and enhanced zoning are described together for convenience.

- Basic and enhanced zoning:
  - Basic zoning provides the basic approach of zoning initiators and targets.
  - The enhanced zoning feature is used to prevent overwrite of zoning configurations by concurrent operations through a fabric-wide lock of a zoning session.
- Smart zoning

  Smart zoning is used to reduce zoning time, reduce the size of the zone database, and optimize resource utilization on SAN switches by automatically creating multiple single-initiator single-target (SIST) pairings within a single large zone.
- Autozone

  The Autozone feature is used to eliminate the need for manual zoning in a single-switch SAN by automatically creating SIST zones.

> **Important:** Although there are several zoning options that can be leveraged based on use cases, we recommend enhanced zoning for standard deployments.

2.10.1 Basic and enhanced zoning

Zoning enables you to set up access control between SAN devices. If you have administrator privileges in your fabric, you can create zones to increase network security and prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field. This section defines various zoning concepts and provides details about zone set and management features.

Zoning has the following features:

- A zone consists of multiple zone members:
  - Members in a zone can access each other. Members in different zones cannot access each other.
  - If zoning is not activated, all devices are members of the default zone.

  - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
  - Zones can vary in size.
  - Devices can belong to more than one zone.
- A zone set can consist of one or more zones:
  - A zone set can be activated or deactivated as a single entity across all switches in the fabric.
  - Only one zone set can be activated at any time in any one VSAN.
  - A zone can be a member of more than one zone set.
- Zoning can be administered from any switch in the fabric:
  - Because zoning information is distributed to all switches in the fabric, zoning changes that are made on one switch are available in all switches.
  - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured nondisruptively.

  New zones and zone sets can be configured without interrupting traffic on unaffected ports or devices.
- Zone membership criteria are mainly based on port worldwide names (pWWNs) or FCIDs or relevant aliases:
  - pWWNs specify the WWN of a N_port that is attached to the switch as a member of the zone.
  - A fabric pWWN specifies the pWWN of the fabric port (F_Port) (switch port's WWN). This membership is also referred to as port-based zoning.
  - FCID specifies the FCID of an N_port that is attached to the switch as a member of the zone.
- The default zone membership includes all ports or pWWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy, which can be set to `permit`, which is the default setting, or `deny`.

Note: Enhanced zoning prevents multiple storage administrators from modifying zone sets at the same time. A best practice is to use enhanced zoning for all configured VSANs in a fabric.

2.10.2 Smart zoning

Smart zoning is an NX-OS feature that is used to reduce the administrative burden for configuring zones. It also drives other positive side effects like reducing the zone database size and the number of access control list (ACL) Ternary Content Addressable Memory (TCAM) entries. It allows multiple initiators and multiple targets to reside in the same zone but prevents initiator-initiator or target-target communication. Initiator-initiator communication is not recommended and can be problematic. Typically, target-target communication is not required, but there are exceptions, such as when using IBM SAN Volume Controller (SVC). If target-target communication is required, then use a basic zone and not a smart zone.

The alternative is to use either SIST zones or single-initiator multiple-target zones. In large environments, the creation of all these separate zones can be a significant operational overhead. Smart zoning combines the benefits of both these approaches.

2.10.3 Autozone

The Autozone feature eliminates the zoning configuration tasks on IBM c-type switches. After Autozone is enabled by a single command, any new device that is added to the FC fabric is automatically zoned without any human intervention. You do not require any zoning expertise to use Autozone. Due to its automated nature, it also eliminates any potential human errors that are associated with manual zoning configuration. This simple but powerful approach of Autozone helps bring agility to your business.

Autozone eliminates manual zoning configuration for selective use cases. Today, a user manually creates zones, adds initiators and targets to the zones, and adds each zone to a zone set. Autozone replaces these steps with a single command. The IBM c-type 32-Gbps fabric switches automatically detect and identify an end-device type as an initiator or a target. All end devices are zoned automatically by following the scheme of SIST zoning. Autozone is invoked only once. Any new devices are automatically detected as initiators or targets and zoned. You do not have to access the switch to modify the zoning configuration when a new or modified storage assignment is required. The final configuration that is made by Autozone is the same as that obtained by a manual-zoning configuration.

Here are more benefits of Autozone:

- Devices are automatically zoned when they log in to an IBM c-type switch fabric.
- No zoning expertise is required.
- Automated zoning eliminates human errors.
- Reduced effort and time in provisioning storage networks.
- Rapid deployments for single-switch SAN or Power on Demand (PoD) environments.

2.11 Quality of service

The QoS feature in IBM c-type NX-OS allows for data traffic to be classified into distinct levels for service differentiation (low, medium, or high priority).

QoS offers the following features:

- Provides a relative bandwidth guarantee to application traffic.
- Controls the latency that is experienced by application traffic.
- Prioritizes one application over another one (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

You can apply QoS to ensure that FC data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications, such as data warehousing.

The IBM c-type switches support QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high-priority frame. A high-priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic.
- Externally generated time-critical control traffic entering a switch from another vendor’s switch. High-priority frames originating from other vendor switches are marked as high priority as they enter an IBM c-type switch.

2.12 N_Port ID Virtualization and N_Port Virtualization

The following sections describe N_Port ID Virtualization (NPIV) and NPV.

N_Port ID Virtualization

NPIV allows an FC host connection or N_Port to be assigned multiple N_Port IDs or FCIDs over a single physical link. All FCIDs that are assigned can now be managed on an FC fabric as unique entities on the same physical host. NPIV is beneficial for connectivity between core and edge SAN fabrics.

NPIV can also be used on hosts. In a virtual machine (VM) environment where many VM operating systems (OSs) or applications are running on a physical server and using the same physical host bus adapter (HBA) to access the SAN fabric, each VM can now be managed independently for zoning, aliasing, and security.

N_Port Virtualization

NPV is an extension of NPIV. The NPV feature allows top-of-rack fabric switches to behave as NPIV-based HBAs to the core FC IBM c-type Director. The edge switch aggregates the locally connected host ports or N_Ports into one or more uplinks (pseudo-ISLs) to the core switches. NPV is primarily a switch-based technology that is designed to reduce switch management and overhead in larger SAN deployments. This IBM c-type software feature supports industry-standard NPIV, which allows multiple N_Port fabric logins concurrently on a single physical FC link.

NPV is a complementary feature that reduces the number of FC domain IDs in core-edge SANs. Fabric switches operating in the NPV mode do not join a fabric; they only pass traffic between core switch links and end devices as gateways, which eliminates the domain IDs for these switches. NPIV is used by edge switches in the NPV mode to log in to multiple end devices that share a link to the core switch without merging with the SAN fabric.

2.13 Non-Volatile Memory Express over Fibre Channel

The IBM c-type switches allow you to deploy Non-Volatile Memory Express (NVMe) over Fibre Channel (NVMe-FC). IBM c-type switches meet the stringent requirements of large virtualized storage landscapes to support more demanding storage performance-based workloads that can provide a higher insight into data analytics. The high-performance, optimization, and scalability is what allows NVMe-FC to interact with existing and next-generation NVM technologies. Host software may communicate with NVM over PCI Express (PCIe).

The benefits of using NVMe are that it can improve the performance of a centralized storage infrastructure over a SAN that is built by using IBM c-type switches. NVMe initiators can access NVMe targets over FC fabrics. In today’s data centers, FC continues to be the preferred protocol for connecting all-flash arrays due to its flexibility, scalability, availability, and high-performance plug-and-play architecture. Also, FC offers compatibility with earlier versions and a phased migration from SCSI to NVMe-based solutions.

FC-NVMe on IBM c-type switches has improved and achieved higher performance by using flash storage versus hard disk drives (HDDs) for various OSs. Multiprotocol flexibility supports both NVMe and SCSI in tandem over FC SAN fabrics so that businesses can leverage their existing infrastructure to deploy new FC-NVMe-capable end devices in a phased approach by sharing the existing FC SAN.

To sum up, NVMe can be enabled into your environment seamlessly and nondisruptively and require no hardware changes, which makes it an excellent candidate for SAN fabrics.

2.14 Forward Error Correction

For a storage switch, whose task is to interconnect end nodes and allow them to communicate, it is important to avoid introducing errors in the information exchange and even identify and correct eventual errors that are generated elsewhere, which makes communication more reliable.

For this reason, IBM c-type switches implement an advanced feature set to help ensure data integrity on all data paths. To help ensure the reliable transport of data, IBM c-type switches use several error-detection mechanisms to provide error-correction capabilities whenever possible:

- Error detection and correction on the supervisor memory.
- Cyclic redundancy check (CRC) for frame integrity on the ingress port.
- Internal CRC detection for frame integrity at the ingress port of the crossbar fabric module and the ingress port of the output line card, with automatic isolation of misbehaving components.
- Automatic dropping of corrupted frames.
- FEC on ISLs and F-ports.
- Syslog-based notifications to the administrator in case anomalies are detected.

Starting with 16 Gbps FC speeds, the robustness of FC networks has been further strengthened by introducing a new optional feature that is called FEC. The scope of FEC is not limited to identifying corrupted frames, but includes the possibility of correcting them in real time. Media such as loose transceivers (SFPs) or dirty cables might result in corrupted packets on ISLs and even links toward end nodes. FEC helps reduce or avoid data stream errors that would result in corrupted frames and lead to application performance degradation, as shown in Figure 2-2.

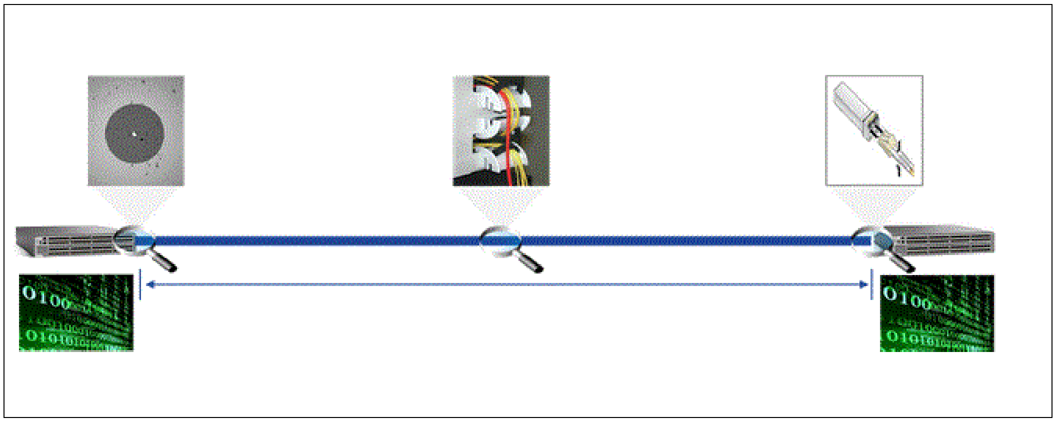

Figure 2-2 Link degradation examples

Standards bodies have made the usage of FEC mandatory on all 32 Gbps FC products and future higher bit rates.

FEC can be considered a first line of defense with the capability to correct wrong bits within frames up to a certain level. When the corruption is more severe, IBM c-type switches will not let those error frames flood the network and instead drop them before they hit the disk array, preventing a waste of resources (bandwidth and CPU on disk controllers) and avoiding possible data corruption inside the storage array.

Despite FEC being optional for 16 Gbps ports, it had a significant adoption in IBM Z environments, where it can be supported end-to-end, from host to control unit (CU), and not just across ISLs. For these FICON environments, transmitter training was required in combination with FEC.

When enabled, FEC on 16 Gbps ports would allow for recovery of up to 11 error bits in every 2112-bit transmission, thus enhancing the reliability of data transmissions. On 32 Gbps ports, the use of FEC is mandatory and its implementation has become even stronger. The mathematical algorithm that is implemented is Reed Solomon, which can correct 7 out of 514 symbols, with any symbol being 10 bits. This process translates into an error correction capability that is more than double what was possible at 16 Gbps.

Figure 2-3 shows an overview of different FEC algorithms at different FC speeds.

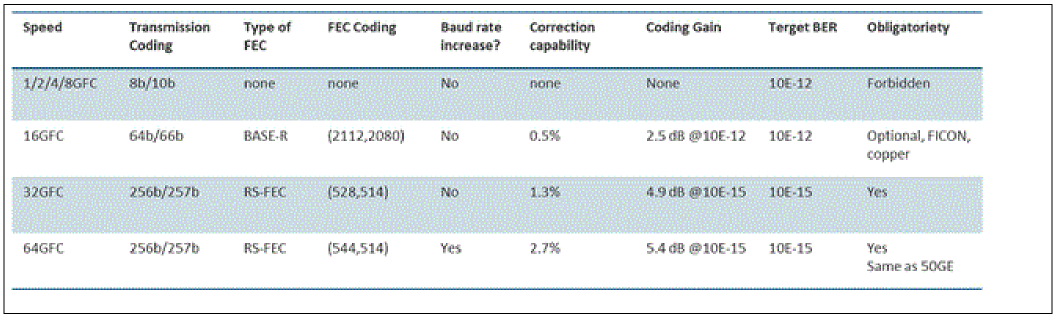

Speed	Transmission Coding	Type of FEC	FEC Coding	Baud rate increase?	Correction capability	Coding Gain	Terget BER	Obligatoriety
1/2/4/8GFC	8b/10b	none	none	No	none	None	10E-12	Forbidden
16GFC	64b/66b	BASE-R	(2112,2080)	No	0.5%	2.5 dB @10E-12	10E-12	Optional, FICON, copper
32GFC	256b/257b	RS-FEC	(528,514)	No	1.3%	4.9 dB @10E-15	10E-15	Yes
64GFC	256b/257b	RS-FEC	(544,514)	Yes	2.7%	5.4 dB @10E-15	10E-15	Yes Same as 50GE

Figure 2-3 FEC algorithms and speeds

Using FEC enables more statistics on the relevant ports. When everything is operating normally, the FEC corrected errors counter stay at zero, but show some value if any bit gets corrected. This way, even if links are working properly, the SAN administrator can easily see when FEC is correcting bits, and that is an indication that some action is needed at the next available maintenance window. In other words, FEC on IBM c-type switches allows the administrator to identify possible issues on the link even if applications are not yet negatively affected. This example is one of preventive maintenance.

Because this FEC capability is fully implemented in hardware and uses an in-band approach, there is no performance impact on throughput and no bit-rate change. The latency contribution is also maintained under 100 nanoseconds on every link, virtually eliminating any impact on application performance. All IBM c-type switches can detect and drop corrupted frames at the switch input, but FEC adds another layer of resiliency to help correct errors wherever feasible and reduce the number of packet drops. When enabled, this feature contributes to more resilient end-to-end frame delivery.

Note: IBM Storage Networking SAN50C-R does not support FEC.

The FEC capability can be enabled and monitored through the command-line interface (CLI) or the DCNM GUI.

2.15 In-service software upgrades

In-service software upgrades (ISSUs) are also known as nondisruptive upgrades. ISSUs allows you to upgrade the device software while the switch continues to forward traffic. An ISSU reduces or eliminates the downtime that is caused by software upgrades. ISSUs can be performed on IBM c-type Directors that have a single or dual supervisor. As a best practice, avoid changing the network while running an ISSU.

Similarly, nondisruptive software downgrades (in-service software downgrades (ISSDs) are also supported.

Not all NX OS releases are FICON certified, which means that you can run an ISSU only by jumping from one FICON certified release to another FICON certified release.

Important: For more information, see the following resources:

- Recommended Releases for Cisco MDS 9000 Series Switches
- Nondisruptive Upgrade Paths to Cisco MDS NX-OS Release 8.5(1)

2.16 Cisco SAN Analytics

The Cisco SAN Analytics solution offers end-to-end visibility into FC block storage traffic. The solution is natively available on the SAN due to its integrated-by-design architecture with the Cisco MDS 9000 switch family. Cisco SAN Analytics delivers deep visibility into I/O traffic between the compute and the storage infrastructure. This information is in addition to the already available visibility that is obtained from individual ports, switches, servers, VMs, and storage arrays.

Cisco SAN Analytics is an industry-first for storage networking devices. It can inspect FC, SCSI, and NVMe headers within FC frames, and it can correlate I/O flows and analyze them. It exposes initiator, target, LUN, and namespace identifiers in multiple views. A recent feature enhancement also provides virtual machine IDs (VMIDs). It scales up to 20,000 I/O flows for director line cards or switches, and for every flow it collects more than 70 metrics. It can be easily enabled on host ports, storage ports, or ISL ports.

Note: Cisco SAN Analytics is not available for FICON traffic. However, Cisco SAN Analytics can be enabled on ports within an FC VSAN when operating in the FC/FICON intermix mode.

3

IBM Storage Networking c-type security

This chapter provides information about securing your IBM Fibre Connection (FICON) network with best practices. The IBM c-type hardware and NX-OS software support advanced security features that provide security within a storage area network (SAN). These features protect your network against deliberate or unintentional disruptions from internal or external threats. Hardware and software authenticity is also a top concern in the modern age, where fake or maliciously modified products have negatively impacted some organizations. Security is a broad topic on its own, and this chapter does not pretend to be exhaustive, but provides a good starting point for the reader.

The following topics are covered in this chapter:

- Security at 360°
- Secure boot and anti-counterfeit technology
- Secure protocols
- Password policies
- Message of the day
- Authentication, authorization, and accounting, RADIUS, TACACS+, and LDAP
- Role-based access control and virtual storage area network-based authorization
- VSANs and zoning
- Fibre Channel Security Protocol
- Port security
- Fabric binding
- Hardcoding port definitions
- TrustSec
- IP Security

3.1 Security at 360°

The race to better security is clearly present across all IT sectors, and the mainframe world is no exception. IBM z15 sets a new and impressive standard in the protection of data with the IBM Data Privacy Passports technology that builds upon pervasive encryption to help clients protect and provision data and revoke access to that data at any time from any location, even for data that is not hosted by the z15.

The recent z15 system extends security, resiliency, and agility to hybrid clouds with encryption everywhere, instant recovery, and cloud native development.

This solution embeds data security policies and encryption with the data and enforces data privacy by policy across the entire enterprise, even when that data leaves your data center and moves to the cloud. This process happens on the same platform that enables you to use hybrid multicloud services, modernize z/OS applications in place, and integrate with Linux applications on and off premises.

Pervasive encryption enables customers to protect data and ensure privacy by encrypting data at the database, data set, or disk level. Using pervasive encryption does not mean that customers are required to change or adjust applications. Each application contains an internal encryption-decryption mechanism, allowing clients to apply cryptography without altering the application itself. These functions go a long way toward addressing data protection and privacy management challenges that often arise when an enterprise organization moves to using hybrid IT in a multicloud world.

Security and data protection is a significant operational challenge for any IT organization. The z15 meets this challenge by using many data-centric audit and protection mechanisms:

- Ability to track the location of all your data and the status of all applicable security mechanisms.
- Capability to build data protection and privacy into all applications and data platforms instead of relying on an assortment of third-party tools.
- Security of having data protection and privacy controls that are embedded into every layer of the computing stack.
- Reliability of having a consistent identity management process in place across your hybrid cloud environment.
- Predictability that is delivered by consistently deploying all computing platform elements.
- Flexibility to securely move data between infrastructure components and third parties.
- Comfort that is enjoyed by being able to meet new data privacy regulations and data sovereignty laws without fearing the risk and economic loss that is associated with data security and privacy failures.

The IBM Z family is also a key platform for all organizations with any hybrid multicloud transition strategy. IBM Cloud® Hyper Protect Services are cloud-specific security services that provide the following features:

- A complete set of encryption and key management services for a specific namespace.
- A database on-demand service with the ability to store data in an encrypted format but without the need for specialized skills.
- A secure Kubernetes cluster container service that can come in handy for packaging applications in a standardized, portable, and scalable way.

z15 also provides Encryption for Data in Flight (EDIF), which is payload encryption that happens directly on the mainframe for true end-to-end encryption. From a networking point of view, Fibre Channel over IP (FCIP) compression is minimally effective when data is encrypted by the mainframe, so enabling IPsec or TrustSec encryption between switches is not required.

The security features that are provided by the recent mainframe systems are matched by a multitude of security features that are available from IBM c-type networking devices. With the advent of optical and FCIP solutions that improve high availability (HA) and disaster recovery (DR), SANs often span outside a single data center, making security concerns even more important. Third-party IT hosting and colocation services add fuel to an already hot topic.

Network data security is about confidentiality, integrity, and authentication. The security features on IBM c-type devices can be classified into four main groups:

- Security at the device level, both hardware and software
- Security for device management access
- Security across devices at the fabric level
- Security for data in transit

Figure 3-1 represents the four groups of network security features on IBM c-type switches.

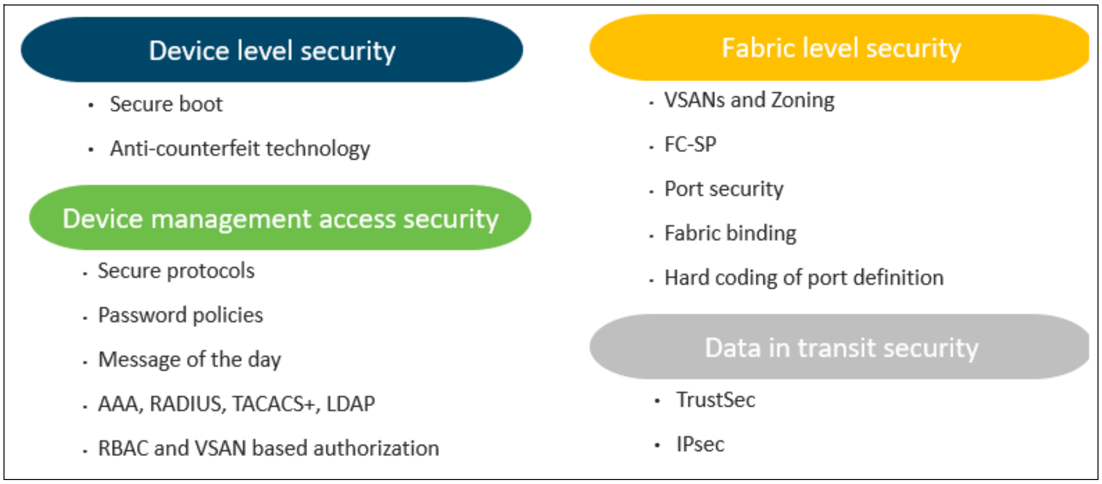

Device level security

- Secure boot
- Anti-counterfeit technology

Device management access security

- Secure protocols
- Password policies
- Message of the day
- AAA, RADIUS, TACACS+, LDAP
- RBAC and VSAN based authorization

Fabric level security

- VSANs and Zoning
- FC-SP
- Port security
- Fabric binding
- Hard coding of port definition

Data in transit security

- TrustSec
- IPsec

Figure 3-1 Network security features

Security must be part of any IT design, and it must be enforced by using effective security management strategies. In the modern digital age, with an expanded attack surface and new regulations now in place, security is not an optional item anymore. It *must* be there.

The following sections provide a high-level overview of the security features that are available on IBM c-type devices with some best practices with regard to implementing them. We do not push the definition of security to an extreme by including in that category all those features and best practices that help prevent *unintentional* service disruption, mostly from human error. However, to increase your security position, you must take more configuration steps.

3.2 Secure boot and anti-counterfeit technology

Secure boot support was introduced for all IBM c-type devices and switching modules with 32-Gbps Fibre Channel (FC) ports.

Secure boot ensures that the first code that is run on IBM c-type hardware platforms is authentic and unmodified. Secure boot anchors the microloader in immutable hardware, which establishes a root of trust and prevents network devices from running network software that was tampered with. It protects the boot code in the hardware, shows the image hashes, and provides the secure unique device identification (SUDI) certificate for the device. During the bootup process, if the authentication of the secure key fails, the line card module or the switch fail to boot up, which prevent the tampering of BIOS. Secure boot is enabled by default and no configuration is required. The feature is offered at no cost.

Anchoring the secure boot process in the hardware makes sure that the most robust security is achieved. In fact, a hardware modification is difficult, expensive, and not easy to conceal even if hackers have physical possession of the device. This approach is different from the rest of the industry and testifies to how important security is considered by IBM, making c-type switches the ideal network infrastructure for mainframe deployments.

Figure 3-2 shows the secure boot process.

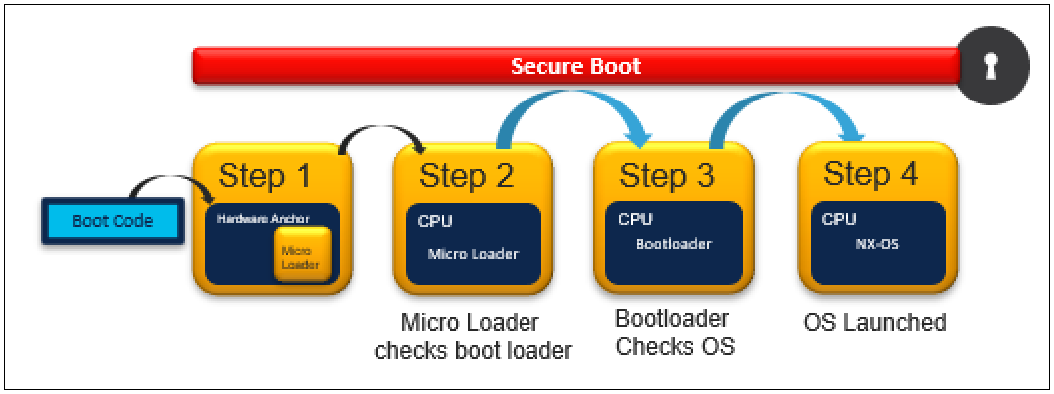

Figure 3-2 Secure boot process

The secure boot workflow is as follows:

- In the context of genuine hardware-anchored secure boot, the first instructions that run on a CPU are stored in immutable hardware.
- When the device boots up, the microloader verifies whether the next set of instructions is original by validating the digital signature on that set of instructions.
- The bootloader validates that the operating system (OS) is original by checking whether it is digitally signed in the correct way.
- The OS is launched if all the checks are passed. If any of the digital signature checks fail, the switch does not let that software boot, which ensures that malicious code does not run on the device.

Coupled to secure boot, anti-counterfeit measures were introduced on IBM c-type devices and switching modules with 32-Gbps ports. The anti-counterfeit measures ensure that IBM c-type hardware platforms with an NX-OS software image are genuine and unmodified, thus establishing a hardware-level root of trust and an immutable device identity for the system to build on.

The IBM c-type units incorporate an ACT2-enabled ASIC that embeds a corresponding SUDI X.509v3 certificate into the hardware. The SUDI certificate, the associated key pair, and the entire certificate chain are stored in the tamper-resistant trust anchor chip. The key pair is bound to a specific chip, and the private key is not exported. These features make cloning or spoofing of identity information impossible.

The SUDI is permanently programmed into the Trust Anchor Module (TAM) and logged during the closed, secured, and audited manufacturing processes. This programming provides strong supply chain security, which is important when the final products see contributions from multiple, component-level suppliers.

If an ACT2 authentication failure occurs, the following error message is displayed:

```
ACT2_AUTH_FAIL: ACT2 test has failed on module 9 with error: ACT2 authentication
failure
```

Note: Secure boot and anti-counterfeit technologies are available on 32 Gbps switching modules for IBM c-type Directors but not on IBM Storage Networking SAN50C-R.

3.3 Secure protocols

The IBM c-type family of switches offers strict and secure switch management options through switch access security, user authentication, and role-based access control (RBAC). Each switch can be accessed through the command-line interface (CLI) or SNMP, with Data Center Network Manager (DCNM) mostly relying on SNMP. As such, both access methods must be considered with regard to security. The FICON feature adds a third mechanism to manage the switch, and that is through the CUP protocol that is described in Chapter 2, “IBM Storage Networking c-type features” on page 53.

The following secure protocols are available on IBM c-type devices to secure management access by using the management port or in-band:

- SNMPv3 (SNMPv1 and SNMPv2c are supported but not recommended) provides built-in security for secure user authentication and data encryption. There is no need for certificates.
- SSHv2 (telnet is supported but not recommended) provides more controlled security by encrypting data, user ID, and passwords. By default, NX-OS software generates an RSA key by using 1024 bits and no X.509 certificates are required. SSH public key authentication can be used to achieve password free logins.
- The SCP, Secure File Transfer Protocol (SFTP), and HTTPS services (TFTP and HTTP are supported but not recommended) protocols offer secure ways to perform file transfers or exchange instructions. Bidirectional encryption and digital security certificates are used by HTTPS to protect against man-in-the-middle attacks, eavesdropping, and tampering. Typically, only the server is authenticated (by the client examining the server's certificate).

For more information about configuring secure protocols, see *Cisco MDS 9000 Series Security Configuration Guide, Release 8.x*.

3.4 Password policies

Passwords represent a basic but critical tool to secure access to devices. Some considerations apply. There is no default password for any IBM c-type switch. You must explicitly configure a strong password. The length of the password must be a minimum of eight characters, but the IBM recommendation is for 15. If a password is trivial (short or easy-to-decipher), the password configuration is rejected and the user account will not be created. Passwords for NX-OS devices are case-sensitive. As a best practice, *never* use the **`no password strength-check`**.

Passwords should not be trivial or easy to identify, but they should be easy to remember for the user who created them. Recently, specialized password generators have seen higher adoption. In this case, passwords are not easy to remember and might include printable characters that are not accepted by the devices under consideration.

A password should contain at least one alphabetic, one numeric, and a mix of capital and non-capital letters. It can also contain special characters if they are supported by the device.

When working with IBM c-type and DCNM, do not use any of these special characters in either usernames or passwords:

```
<SPACE> " & $ % ' ^ = < > ; : ` \ | / , .*
```

3.5 Message of the day

The message of the day (MOTD) feature is used to display some text on the terminal before the login prompt. The MOTD banner is designed to provide enough information, and it has a maximum of 40 lines of 254 characters each.

The MOTD banner can be used for security purposes. Assuming an intruder is trying to authenticate into the device, you do not really want to welcome them. Instead, you claim your exclusive rights on the device and command the non-authorized user to immediately log off. The MOTD banner also can be used for critical communication, such as informing users about known contingencies or planned maintenance windows. The MOTD should also contain the name, email address, and phone number of the device administrator so they can be easily reached if required.

The MOTD feature is available on both the IBM c-type NX-OS CLI and DCNM management tools.

The MOTD can be configured from the NX-OS CLI as a simple short message or an extended multi-line message. When a multi-line message is wanted, enter the command **`banner motd`** followed by the hash character (#) and then the carriage return (enter). Then, you can enter the text for each line. The same dash character indicates the end of input text.

Example 3-1 shows the output that a user would see when trying accessing the switch.

Example 3-1 MOTD banner message example

```
login as: admin
Pre-authentication banner message from server:
|
| ++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
| You have accessed an IBM c-type switch, property of ACME Corporation.
| Only authorized personnel have the right to access this switch.
| If you are not on the list of the authorized personnel,
| disconnect immediately or you will be prosecuted according to law.
|
| Planned maintenance activity:
| Saturday 27th Feb 2021 from 20.00 until 23.00
|
| If needed, contact Fausto Vaninetti
| email: fvaninet@examplecorp.com
| phone number: +01 323 745 745
| +++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++|
```

For more information about configuring the MOTD by using the NX-OS CLI for IBM c-type devices, see *Cisco MDS 9000 Series Fundamentals Configuration Guide, Released 8.x.*

The MOTD also can be configured for the DCNM web client splash window, this time in combination with an optional image. An administrator can do this task by selecting **Administration** → **DCNM Server** → **Customization**.

Figure 3-3 shows an example of the DCNM splash window with the MOTD included.

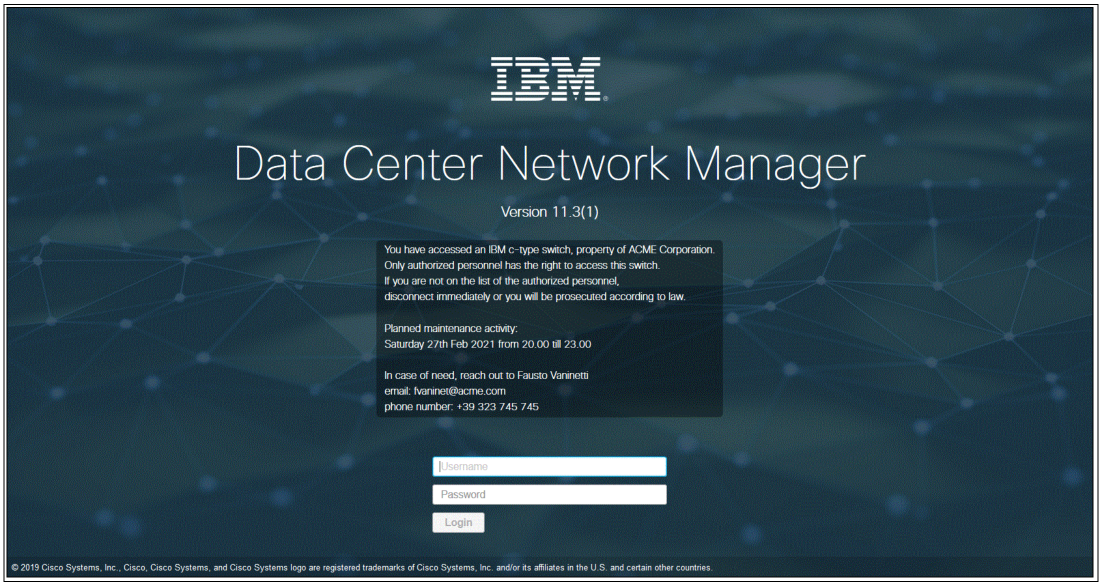

Data Center Network Manager

Version 11.3(1)

You have accessed an IBM c-type switch, property of ACME Corporation.
Only authorized personnel has the right to access this switch.
If you are not on the list of the authorized personnel,
disconnect immediately or you will be prosecuted according to law.

Planned maintenance activity:
Saturday 27th Feb 2021 from 20.00 till 23.00

In case of need, reach out to Fausto Vaninetti
email: fvaninet@acme.com
phone number: +39 323 745 745

© 2019 Cisco Systems, Inc., Cisco, Cisco Systems, and Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Figure 3-3 DCNM splash screen with the MOTD banner

3.6 Authentication, authorization, and accounting, RADIUS, TACACS+, and LDAP

In every organization, there is a need to regulate user access to devices and network management tools that are either CLI-based or GUI-based. The topic becomes even more relevant at scale, where a centralized access control solution appears to be the only practical way to cope with a multitude of devices, an evolving security landscape, and a dynamic assignment of users to roles. The problem is not new, and in fact network authentication, authorization, and accounting (AAA, pronounced triple-A) is a technology that has been in use for over 40 years.

AAA is a set of services for controlling access to IT infrastructure resources, enforcing policies, assessing usage, and providing the information that is necessary to auditing and eventually billing for services. These processes are considered important for effective network management and security, more so with large-scale network deployments.

In a large organization, you must make sure that only some users can log in to network devices and the relevant management tools. Due to differences in skills, roles, and responsibilities, different users have a different set of privileges and are able to perform a different set of tasks. Moreover, to implement an adequate governance approach, the management team needs to know *who did what and when*. Being able to answer this apparently simple question requires a proper implementation of the overall solution, but can be vital to identifying wrong behavior within the organization when a network operation is in jeopardy. Imagine a network outage due to human error. Someone might accidentally shut down a port and prevent proper communication across the FICON setup. A modern AAA solution allows for quick determination of which user made the mistake.

All IBM c-type network devices support local AAA services. Users log in, authenticate, and are authorized to perform some actions on specific resources. All activity is tracked for accounting purposes. The system maintains the username and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored information. Even better, all IBM c-type network devices typically support remote AAA services, which are provided through the Remote Authentication Dial-In User Service (RADIUS), Terminal Access Controller Access Control System Plus (TACACS+), or Lightweight Directory Access Protocol (LDAP) protocols. Remote AAA services offer the following advantages over local AAA services:

- User password lists for each switch in the fabric can be managed more easily.
- Remote AAA servers are deployed widely across enterprises and can be easily adopted.
- The accounting log for all switches in the fabric can be centrally managed.
- User role mapping for each switch in the fabric can be managed more easily.

For most network administrators, the genesis of AAA coincided with the development of the RADIUS protocol, which is based on UDP port 1812 and 1813. RADIUS became an internet standard through the Internet Engineering Task Force (IETF) in 1997, and it is still a widely accepted AAA protocol. Another commonly adopted AAA protocol is TACACS. It is described in RFC 1492, but it never became an internet standard. TACACS evolved into XTACACS, which added accounting capabilities, and then into TACACS+, which is the current version. TACACS+ runs over TCP port 49 and allows for encryption of all transmitted data, not just passwords, which overcomes the vulnerabilities that are found in RADIUS.

More recently, organizations have considered the widespread utilization of Microsoft Active Directory as the primary source of access control for all devices, services, and users, including network administrators. Leveraging the Microsoft version of LDAP with Kerberos authentication, network administrators can be redirected to a HA cluster of Active Directory servers for credential validation (authentication), assignment of privileges (authorization), and tracking of activity (accounting). For companies where Linux and open source software are preferred, the OpenLDAP client and its Directory Server are an alternative implementation, but others are available on the market. Secure LDAP is the typical protocol in use.

In essence, AAA represents a complete system for tracking user activities on an IP-based network and controlling their access to IT resources. AAA is often implemented as a dedicated and centralized server, sometime referred to as an *access control server*. Devices talk to the AAA directory server through a AAA daemon, following a classical client/server approach.

The general idea revolves around three fundamental security building blocks:

- *Authentication* is the process of determining whether someone, or something, is, in fact, who or what they are declaring to be. Authentication asks the question: *Who or what are you?*

  In computer networks (including the internet), authentication is commonly done by using usernames and passwords. Recently, for the most critical services, dual-factor, or two-factor, authentication (2FA) has become prevalent. The AAA directory server compares the user-provided authentication credentials with user credentials that are stored in a database. If there is a match, the user is recognized and permitted access to the IT resource. If the credentials do not match, authentication fails, and access is denied.

- *Authorization* is the process of specifying access rights to IT resources, which is related to security in general and to access control in particular. Authorization asks the question: *What are you allowed to do?*

  Authorization happens after authentication is successful. Authenticated users may be given different authorization levels that limit their access to IT resources. Authorization determination may be based on individual user-specific parameters, geographical location restrictions, date or time-of-day restrictions, membership to specific groups, and so on. In other words, authorization is the process of retrieving policies and mapping them to users, and determining what types of activities, resources, or services a user is permitted to access.

  The enforcement of the policy is also considered part of the authorization process. Within documentation, this policy is also referred to as the *role*. So, the user profile is the combination of username, password, and role. Practically, the role can be predefined on devices and management tools or manually configured by the administrator.

- *Accounting* refers to the record-keeping and tracking of user activities on an IP-based network for auditing, billing purposes, or trending or future analyses, which may include, but is not limited to, real-time accounting of the time that is spent accessing the network, the network services employed or accessed, and so on. Accounting asks the question: *What did you do?*

  The accounting feature tracks and maintains a log of every management session that is used to access the networking device and an associated timestamp. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally on devices or sent to remote AAA servers.

Sometimes, multiple AAA servers based on different protocols are used. Complex multi-domain systems, like those composed of computing, networking, and storage elements, might welcome the adoption of multiple authentication tools for the various technical domains and subject matter experts (SMEs). Each AAA server group is specific to one type of protocol or service or IT resource type.

Figure 3-4 represents the secure access protocols and the options for both local and remote centralized AAA services.

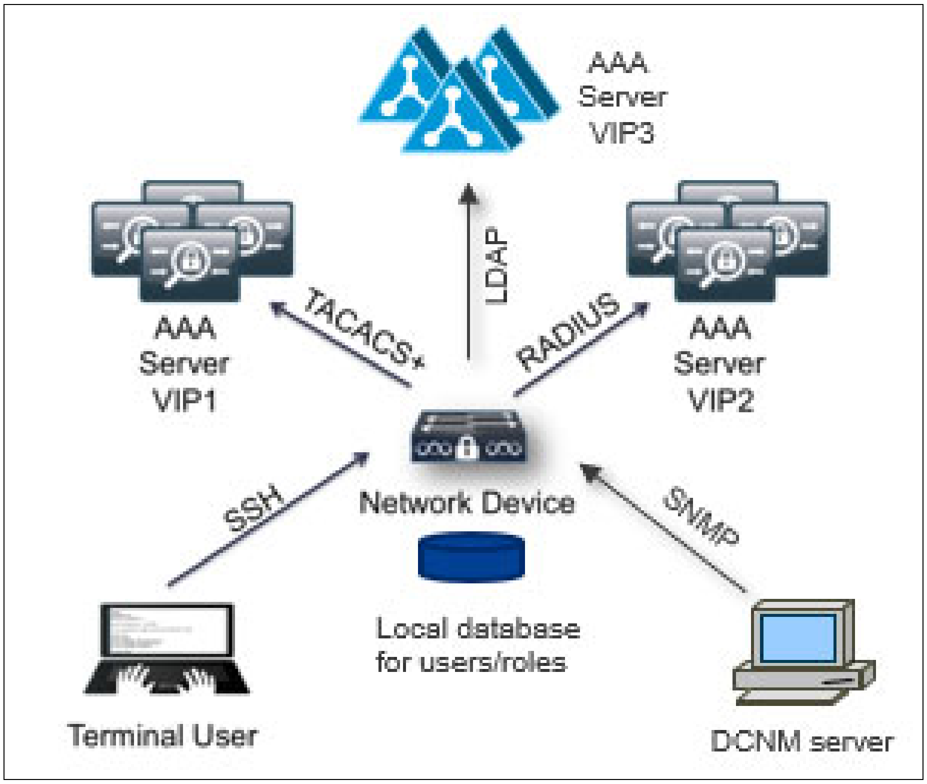

Figure 3-4 Secure access protocols

For step-by-step instructions about configuring remote centralized AAA services on IBM c-type devices, see Configuring Security Features on an External AAA Server.

3.7 Role-based access control and virtual storage area network-based authorization

All management access within the IBM c-type family is based on roles. Role-based authorization limits access to switch operations by assigning users to roles. Users are restricted to performing the management operations that are explicitly permitted based on the roles to which they belong.

By default, two roles exist in all switches:

- The *network-operator* has permission to view any portion of the configuration, but not the complete running and startup configuration. The operator can also apply licenses, but cannot make any configuration changes.
- The *network-admin* has permission to run all commands and make configuration changes.

These two default roles cannot be changed or deleted. Users belonging to the network-admin role are authorized to create and customize up to an extra 64 roles and add other users to those roles.

If a user has no role, they can be assigned a default role with minimum privileges or denied access, depending on switch configuration. The default role is network-operator.

Each role can be applied to multiple users and typically each user is assigned to a single role.

CLI and SNMP users sit in different local databases but share common roles. You can use SNMP to modify a role that was created by using CLI and vice versa. Each role in SNMP is the same as a role that is created or modified through the CLI. It is possible to limit the scope of authorization to specific virtual storage area networks (VSANs). In other words, custom roles can be restricted to one or more VSANs as required.

Figure 3-5 shows the VSAN-based authorization concept.

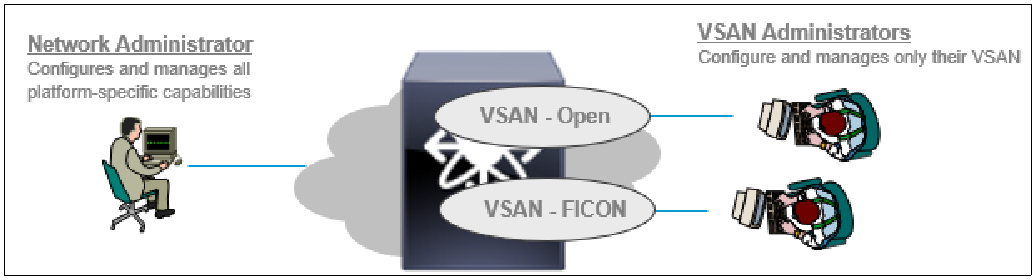
Network Administrator
Configures and manages all platform-specific capabilities
VSAN - Open
VSAN - FICON
VSAN Administrators
Configure and manages only their VSAN

Figure 3-5 VSAN-based authorization concept

One of the advantages of separating open systems traffic and FICON traffic into separate VSANs is that you can grant administrative access on a per VSAN basis, which means that if you have separate FICON and open systems administrative staff, complete administrative authority can be given to one VSAN while preventing access to a different VSAN. For example, it is possible to create a role that is called `ficon-admin` and allow it network-admin privileges when accessing the FICON VSAN only.

Example 3-2 lists the CLI commands to create, commit, and distribute a custom role.

Example 3-2 Commands to create a custom role

```
SAN384C-6# conf t
SAN384C-6(config)# role name ficon-admin
SAN384C-6(config-role)# description Custom role with true network-admin
privileges, FICON VSAN only
SAN384C-6(config-role)# rule 1 attribute-admin
SAN384C-6(config-role)# vsan policy deny
SAN384C-6(config-role-vsan)# permit vsan 30
SAN384C-6(config)# role commit
SAN384C-6(config)# role distribute
```

Role-based configurations must be committed before they take effect and it is best to distribute them to all switches in the fabric. This task is important for custom roles. To this end, the Cisco Fabric Services (CFS) infrastructure provides the necessary support to implement a single point of configuration for the entire fabric.

Example 3-3 shows the CLI output with details about the newly configured role.

Example 3-3 The show role command

```
SAN384C-6(config)# show role

<snip>

Role: ficon-admin
  Description: Custom role with true network-admin privileges, FICON VSAN only
  Vsan policy: deny
  Permitted VSANs: 30
  -------------------------------------------------
  Rule    Type    Command-type    Feature
  -------------------------------------------------
  1       permit  attribute-admin *
```

VSAN-based authorization is a powerful way of distributing administrative access to different groups. VSAN policy enforcement is done only for non-**show** commands. In other words, a VSAN-restricted user can see everything but can make changes only to allowed VSANs.

Configuring VSAN-based access control requires the ENTERPRISE_PKG license.

For more information about configuring RBAC- and VSAN-restricted users, see Configuring User Accounts and RBAC.

3.8 VSANs and zoning

VSANs are an effective tool for securing access to the fabric and preventing intentional or accidental wrongdoing. FC services are independent across VSANs, and no communication between any two end nodes occurs across VSANs, unless configured to do so by using Inter-VSAN Routing (IVR).

Note: IVR is not available for FICON VSANs.

There are two VSANs that are created on any IBM c-type switch by default: VSANs 1 and 4094. VSAN 4094 is referred to as the isolated VSAN. When a VSAN that has active ports is deleted, the ports are then moved to the isolated VSAN. Ports in the isolated VSAN cannot communicate with any other ports, including other ports that are in the isolated VSAN. Because of this behavior, moving all ports into the isolated VSAN at initial configuration is an effective way to secure ports in the fabric, despite it not being a common practice. Then, ports would require a manual configuration change to be placed in an active VSAN. By default, all ports are in the default VSAN, and that is VSAN 1. It is not a best practice to use VSAN 1 as your production VSAN.

Best practice: VSAN 1 should not be used for live traffic.

By creating a separate VSAN for your production traffic, you effectively isolate your production devices from any device that is later connected to the switch. Again, a manual configuration change is required to move a port from VSAN 1 to an active production VSAN.

In general, VSANs can be used to effectively separate traffic in the following scenarios:

- Different customers in storage provider data centers
- Production and test environments
- Low and high security requirements
- Backup traffic from user traffic
- Replication traffic from user traffic
- Virtualized applications from bare metal applications

Of course, VSANs are also used to separate FICON traffic from FC traffic.

The ease of use that is combined with the administrative and security benefits of VSAN technology explain why it is rare to find any customer that is not using it.

The use of VSANs does not preclude the use of zoning. The two features are complementary. The zoning process is per VSAN, which means that creating separate VSANs allows zoning granularity so that a misconfiguration of the zoning database in one VSAN does not cause a problem for any of the other VSAN.

In storage networking, zoning is the mechanism in FC fabrics that controls what ports are allowed to inter-communicate. Zoning is the partitioning of end nodes in an FC fabric into smaller subsets to restrict interference and add security by preventing unintentional communication. Although multiple devices are made available to a single device through a SAN, each system that is connected to the SAN should be allowed access only to a controlled subset of these devices. For example, we do not want an initiator to talk to another initiator because only initiator to target communication is useful and should be allowed. Target to target communication is also possible for data replication. In general, single initiator single target (SIST) zoning is recommended, which is also known as 1:1 zoning.

There can be only one active zone set per VSAN. Other zone sets can be configured but not active at a time. Several zones make up a zone set. IBM c-type switches support up to 16,000 zones and 20,000 zone members. Changes to the active zone set can be made non-disruptively. Zone members can be identified by using the following methods:

- Port worldwide name (pWWN): The worldwide name (WWN) of the attached device.
- Fabric port (F_Port) WWN: The WWN of the switch port.
- FCID of the attached device.
- FC alias: The alias name is in alphabetic characters and identifies a port ID or WWN. The alias can include multiple members.
- Device alias: The device alias name is like an FC alias but provides more scalability and works across VSANs.
- Domain and port: The domain ID is an integer 1 - 239. A port number of a non-Cisco switch is required to complete this configuration.
- IP address: The IP address of an attached device in 32 bytes in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- Internet Small Computer Systems Interface (iSCSI) Qualified Name (IQN): A unique identifier that is used in iSCSI to identify devices.
- Interface and sWWN: Based on a switch interface number and sWWN.

By default, all devices are initially placed into a default zone. When devices are moved to a user-defined zone, they are removed from this default zone. On an IBM c-type switch, the default zone is set to deny by default, which means that devices in this default zone cannot communicate with each other, which prevents accidental and unwanted communication. This behavior can be changed during the initial setup script or configured later, but it should stay as-is for better security in FC networks.

FICON requires the default zone to be set to permit because in FICON environments the devices that are allowed to communicate are explicitly defined in the Hardware Configuration Definition (HCD) file on the mainframe and security is derived from there.

Important summary: Default zoning is set to deny by default for open system VSANs. When you configure a FICON VSAN by using the CLI, you must change the zoning manually to permit. When using DCNM or Device Manager (DM) instead, the zoning is changed to permit automatically.

There is an alternative zoning approach that is called *smart zoning*, which is a valid alternative at scale. It provides a simpler operational environment while keeping the same 1:1 approach at an implementation level.

With smart zoning, the administrator creates large zones with both initiators and targets in them, grouping devices with some logical criteria of preference. Then, the “smart” capability applies 1:1 zoning to devices in agreement with best practices. This process is possible because when end nodes first register in an IBM c-type switch, they declare who they are (initiators versus targets), and this information is used by smart zoning to establish communication. This process is different than a *default zoning permit*, where any to any communication is allowed.

Just like VSANs, zoning can be considered a security feature. To make it better, IBM c-type devices support only hard zoning that is hardware-enforced, which means any frame entering a switch port goes through an inspection process and is allowed to reach only the end nodes that are configured in the ASIC TCAM table. This type of zoning is different from the so-called soft zoning, where there is no hardware enforcement and some control plane level of obfuscation of end node addresses. In this case, an intruder might have a frame reach a forbidden end node if they know the address. The hardware has no way to prevent that incident from happening.

Figure 3-6 on page 83 highlights the relationship between VSANs and zones.

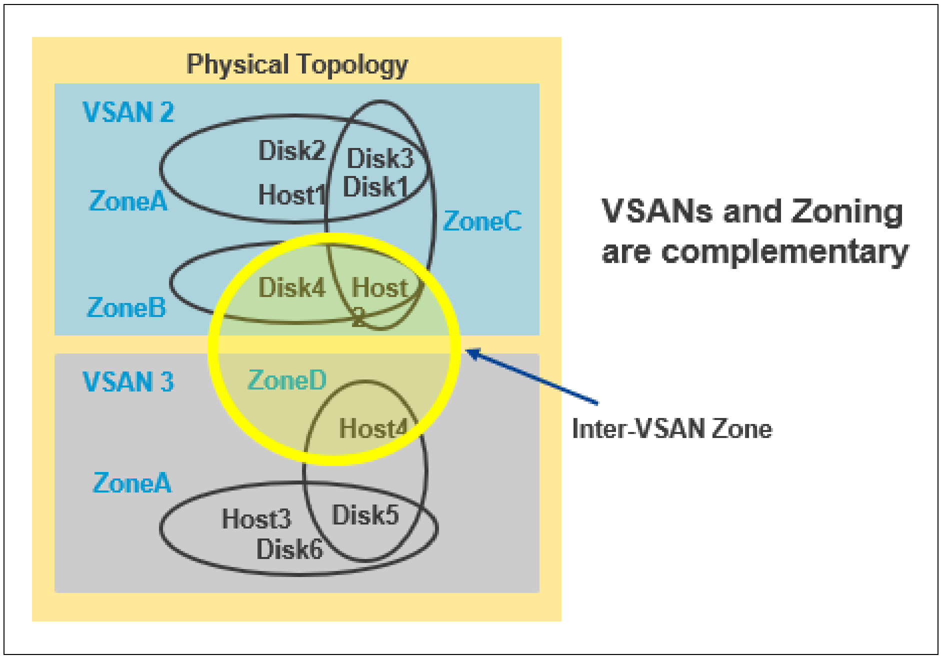

Figure 3-6 VSAN and zone relationship

For more information about zoning and a step-by-step guide about how to configure it, see Configuring and Managing Zones.

3.9 Fibre Channel Security Protocol

Although FC networks are often considered to be secure by themselves, there are always possibilities to make them better. Fibre Channel Security Protocol (FC-SP) capabilities provide switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. With FC-SP, switches, storage devices, and hosts can prove their identity by using a reliable and manageable authentication mechanism. With FC-SP, FC traffic can be secured on a frame-by-frame basis to prevent snooping and hijacking, even over untrusted links. A consistent set of policies and management actions are propagated through the fabric to provide a uniform level of security across the entire fabric. In simple terms, using the FC-SP authentication protocol helps to prevent either accidental or intentional fabric disruption by blocking unauthorized switches or devices from connecting to the fabric.

Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) is an FC-SP protocol that provides authentication between IBM c-type switches and other devices. DH-CHAP consists of the CHAP protocol that is combined with the Diffie-Hellman exchange. The FC-SP feature can be enabled with the NX-OS CLI command **`feature fcsp`**.

Configuring this feature requires the ENTERPRISE_PKG license.

DH-CHAP is a mandatory password-based, key-exchange authentication protocol. DH-CHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD-5 and SHA-1 algorithm-based authentication.

Configuring the DH-CHAP feature along with other IBM c-type features introduces some other considerations:

- port channel interfaces: If DH-CHAP is enabled for ports belonging to a port channel, DH-CHAP authentication is performed at the physical interface level, not at the port channel level.
- FCIP interfaces: The DH-CHAP protocol works with the FCIP interface as it would with a physical interface.
- Port security or fabric binding: Fabric binding policies are enforced based on identities that are authenticated by DH-CHAP.
- VSANs: DH-CHAP authentication is not done on per-VSAN basis.
- HA: DH-CHAP authentication works transparently with existing HA features.

We have only scratched the surface of the FC-SP capabilities. For more information about this advanced security feature, see Configuring FC-SP and DHCHAP.

3.10 Port security

In most cases, an FC device can attach to any SAN switch port and access SAN services that are based on VSAN and zone membership. Port security is a feature that was introduced to prevent unauthorized access to a switch port. This feature works against source ID (SID) spoofing because it allows access to the fabric based on device identity attributes. The port security feature requires the ENTERPRISE_PKG license.

When port security is enabled, all fabric login and initialization requests from unauthorized devices, including (Nx ports) and switches (xE ports), are rejected and the intrusion attempts are logged.

To enforce port security, you must configure the devices and switch port interfaces through which each device or switch is connected. You can use either the pWWN or the node worldwide name (nWWN) to specify the Nx port connection for each device. For switches, you use the switch worldwide name (sWWN) to specify the xE port connection. Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies is done on every activation and when the port tries to initialize. The port security feature requires all devices connecting to a switch to be part of the port security active database. The switch uses this active database to enforce authorization.

You can instruct the switch to automatically learn (auto-learn) the port security configurations. The auto-learn option allows any switch in the IBM c-type family to automatically learn about devices and switches that connect to it. Using this feature to implement port security saves tedious manual configuration for each port. Auto-learn is configured on a per-VSAN basis. If it is enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured port access. Learned entries on a port are cleaned up after that port is shut down.

By default, the port security feature is not activated. When you activate the port security feature, the auto-learn option is also automatically enabled. You can choose to activate the port security feature and disable auto-learn. In this case, you must manually configure the port security database by individually adding each port.

The steps to configure port security depend on which features you are using. For example, auto-learning works differently if you are using CFS distribution or not. For more information about configuring port security, see Configuring Port Security.

3.11 Fabric binding

The *fabric binding* feature ensures that Inter-Switch Links (ISLs) are enabled only between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis. Considering most FICON deployments use only one or two VSANs, this feature will not add much of an administrative burden.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

Fabric binding requires that you install either the MAINFRAME_PKG license or the ENTERPRISE_PKG license on your switch. You do not need both licenses.

Port security and fabric binding are two independent features that can be configured to complement each other. The main difference is that fabric binding works at the switch level and port security works at the interface level. From an administrative point of view, port security can benefit from the auto-learn feature. The CFS distribution is not available for fabric binding.

To enforce fabric binding, configure the sWWN to specify the xE port connection for each switch. Enforcement of fabric binding policies is done on every activation and when the port tries to come up. In a FICON VSAN, the fabric binding feature requires that all sWWNs are connected to a switch and that their persistent domain IDs are part of the fabric binding an FC VSAN, only the sWWN is required (the domain ID is optional).

Example 3-4 shows the usage of an sWWN and domain ID for a FICON VSAN.

Example 3-4 Using an sWWN and domain ID

```
switch(config)# fabric-binding database vsan 5
switch(config-fabric-binding)# swwn 21:00:05:30:23:11:11:11 domain 102
```

In a FICON environment, the purpose of the fabric binding feature is to ensure that ISLs in FICON cascaded topologies are enabled only for switches that are configured in the fabric binding database, which includes FC ISLs, FCIP ISLs, and port channels made up of these ISLs. Each FICON switch that is allowed to connect to the fabric must be added to the fabric binding database of every other FICON switch in the fabric. Activating fabric binding is a prerequisite for enabling FICON on a VSAN.

In a FICON cascaded topology the fabric binding database contains the sWWN and domain ID of all the switches that are authorized to join the fabric. The fabric binding authorization is enforced per VSAN because each VSAN is a logical fabric. In a FICON point-to-point switched topology, fabric binding is still required, but the fabric binding database is empty because defining the local sWWN and domain ID in the fabric binding database is not required.

There are two fabric binding databases:

- The *configuration database* contains all the manually configured WWNs and domain IDs of those switches that are authorized to join the fabric.
- The active database contains the entries that are being enforced in the fabric.

To start enforcing a newly created or modified configuration database, an activation sequence must be performed. The activation replaces the active database with the configured database. This activation fails if the configured database does not match the state of the fabric. For example, if a switch is in the fabric but not defined in the database, or if a switch is in the fabric but has a different domain ID than is defined in the database. Alternatively, the **force** option can be used to activate the new fabric binding configuration, which isolates the switch.

Attention: The **force** option must be used with discretion and care. In fact, it is easy to have a mistake in the configured fabric binding database, use the **force** option, and cause isolation to occur in the fabric.

The EFMD protocol makes sure that all switches in the fabric have identical fabric-binding databases when ISL links are started. The protocol does not distribute the database when it is changed on a single switch. The fabric binding database of each switch in the fabric or VSAN must be manually updated with the sWWN and domain ID of every other switch in the fabric.

When an ISL is initialized in a FICON VSAN, the following checks are performed:

- Is the peer sWWN present in the active fabric binding database?
- Does the domain ID of the peer switch match what is present in the active fabric binding database?
- Is the active fabric binding database of the peer switch identical to the active fabric binding database in the local switch? Again, this is the purpose of the EFMD protocol.

If during an ISL link negotiation the databases from the two switches do not match, the link does not allow the FICON traffic to flow. When switches are added to an existing fabric, all the switches must be configured to incorporate the new switches in their active databases.

Figure 3-7 on page 87 illustrates the concept of fabric binding.

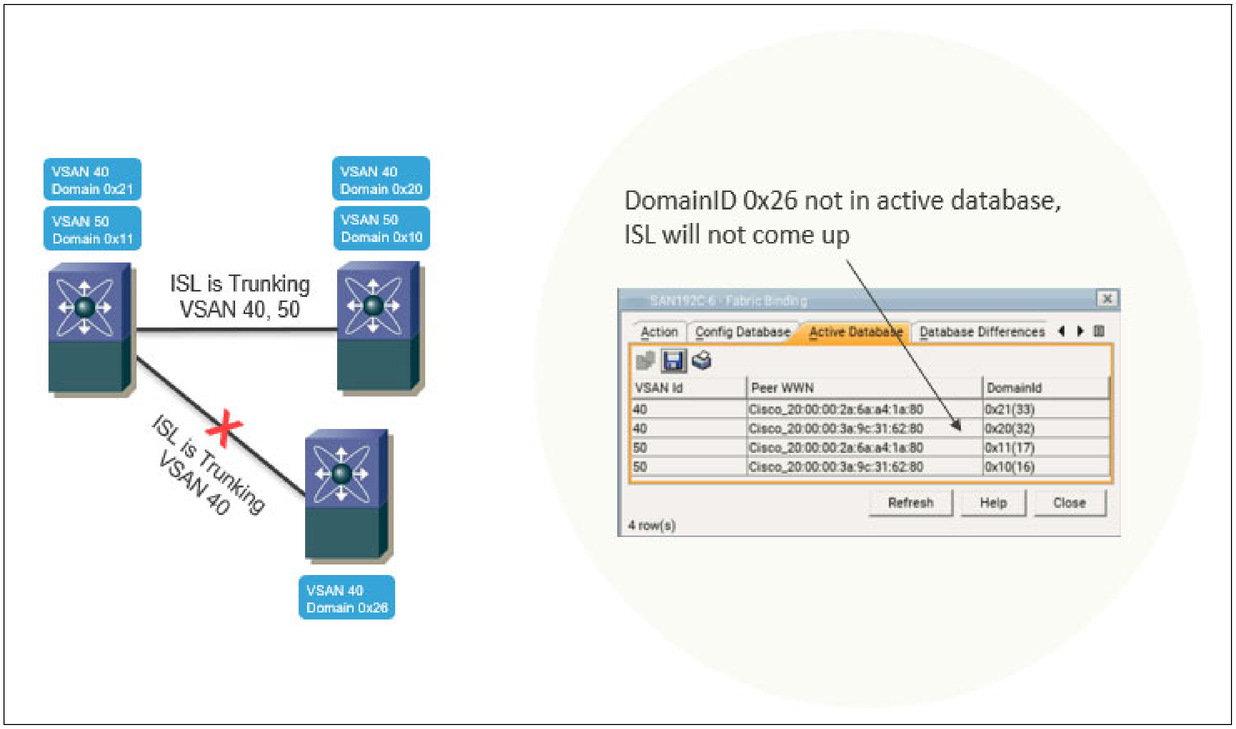

VSAN Id	Peer WWN	DomainId
40	Cisco_20:00:00:2a:6a:a4:1a:80	0x21(33)
40	Cisco_20:00:00:3a:9c:31:62:80	0x20(32)
50	Cisco_20:00:00:2a:6a:a4:1a:80	0x11(17)
50	Cisco_20:00:00:3a:9c:31:62:80	0x10(16)

4 row(s)

Figure 3-7 Fabric binding

Fabric binding configuration starts by enabling the feature by running the **`feature fabric-binding`** command. For more information about how to configure fabric binding, see Configuring Fabric Binding.

3.12 Hardcoding port definitions

Hardcoding E-port and F-port definitions can be considered a security best practice. This task requires some more configuration when compared to using the default setting (**`switchport mode auto`**), but it helps make sure that the topology is fully under control. In simple terms, configure the ports as E-ports when connected to other switches and as F-ports when connected to the mainframe and storage systems.

This approach is sometimes referred to as *port mode security*, and it is intended to protect edge ports from becoming ISL ports. It can be coupled with RBAC so that only certain users have the required privileges to change the port mode.

Best practice: Configure port mode security on your IBM c-type FICON switches.

3.13 TrustSec

Data integrity and confidentiality are undoubtedly the top priority for IBM customers, which is why IBM c-type switches have support for secure communication through both peer authentication and encryption for data in transit between peer switching devices. This support is often used with metro or long-distance links that by definition go outside the physical security perimeter of the data center and often use third-party communication networks. Organizations in the finance sector might need to comply with stringent national regulations and must turn on encryption inside the walls of their data center.

All current IBM c-type switches support peer authentication according to the FC-SP standard by using the DH-CHAP, but this process does not prevent unwanted activities such as traffic interception. To help ensure data integrity and privacy, data should also be encrypted.

For FC links, the capability to encrypt traffic on the wire is known as *TrustSec Fibre Channel Link Encryption*. This capability is an extension of the FC-SP standard and uses the existing FC-SP architecture. It enables either AES-Galois Counter Mode (AES-GCM) or AES-Galois Message Authentication Code (AES-GMAC). AES-GCM authenticates and encrypts frames with the 128-bit Advanced Encryption Standard (AES) algorithm, and AES-GMAC authenticates only the frames that are passed between the two peers. Encryption is performed at the line rate by encapsulating frames at switch egress. At switch ingress on the other side of the link, frames are decrypted and authenticated with integrity checks (a hop by hop encryption mechanism). Only E-ports and TE-ports that are configured between IBM c-type switches and their affiliates can support encryption.

> **Note:** IBM Storage Networking SAN50C-R supports peer authentication but not in-transit data encryption for native FC ports.

There are two primary use cases for TrustSec Fibre Channel Link Encryption:

- Customers are communicating outside the data center over native FC (for example, dark fiber or some type of wavelength division multiplexing).
- Encryption is performed within the data center for security-focused customers, such as defense, military, or financial institutions.

The beauty of TrustSec Fibre Channel Link Encryption is the simplicity of enabling it and the scale at which it can be enabled without affecting SAN performance. Both the NX-OS CLI and DCNM allow you to configure and provision this feature. To perform encryption between the switches, a security association (SA) must be established. An administrator must manually configure the SA before the encryption can take place. The SA includes parameters such as encryption keys and a salt (a 32-bit hexadecimal random number that is used during encryption and decryption). Up to 2000 SAs are supported per switch. Key management is not required because keys are configured and stored locally on the switches.

TrustSec Fibre Channel Link Encryption requires specific paths within the ASICs, so it is available only to a limited set of ports. For example, IBM c-type switches can support up to 12 encrypted FC ports on the 48-port 32 Gbps switching module for a total of 384 Gbps. This amount is much encrypted traffic on a single switching module, and is three times higher than the industry average. It is worth pointing out that enabling encryption on an FC port does not reduce the number of buffer-to-buffer (B2B) credits that are available for distance extension. That limitation, often found on other FC products, does not affect IBM c-type switches.

Enabling encryption on FC ports for IBM c-type switches does not impact performance, contrary to what might happen on alternative products. Due to an optimized switching architecture, ports continue transmitting at full line rate and latency is not increased in any perceivable way. Organizations are not required to find a compromise between SAN performance and SAN security. They can have both.

TrustSec Fibre Channel Link Encryption is competitively unique and a clear differentiator for high-security accounts. The TrustSec Fibre Channel Link Encryption feature requires the ENTERPRISE_PKG license.

Best practice: Configure TrustSec Fibre Channel Link Encryption with the affected ports administratively shut down.

For more information about how to configure TrustSec encryption, see Configuring Cisco TrustSec Fibre Channel Encryption.

3.14 IP Security

For securing FCIP links, the IP Security (IPsec) protocol framework of open standards, which is developed by the IETF, is used to provide data confidentiality, integrity, and peer authentication. Therefore, with IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing, making it an ideal feature when doing long-distance DR implementations. The IPsec feature is highly recommended on long-distance FCIP links because they are the most exposed to traffic interception. The encryption capability can be enabled natively on IBM c-type switches or alternatively on data center exit routers, depending on specific situations.

IPsec is composed of two protocols: one for key exchange and one for data flow encryption.

On IBM c-type switches, IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys that are used by IPsec. IKE provides authentication of the IPsec peers, negotiates IPsec SAs, and establishes IPsec keys. SAs are per direction, so a full duplex link has two of them. Conceptually, the SA provides all the parameters that are required to establish how data must be protected. Instead, the security policy indicates what traffic must be protected and inserted into the IPsec tunnel.

On IBM c-type switches, IPsec uses the Encapsulating Security Payload (ESP) protocol to achieve data confidentiality (encryption), integrity (hash), and peer authentication (signature or certificates). The 256-bit AES algorithm is typically used for encryption (this is referred to as the ESP-AES 256 transform set). Although there are two modes of operation that are allowed by the IPsec framework, the NX-OS implementation is tied to its specific security gateway use case, so supports only IPsec tunnel mode (and not IPsec transport mode, which is more suited for hosts). The IPsec tunnel mode encrypts and authenticates the entire IP packet, including its original header. It works by adding an outer IP header and the ESP header before the original IP packet and adding an ESP trailer and authentication header after it, essentially creating an envelope around the original IP packet. Because the outer IP header is in clear text, this approach remains compatible with NAT solutions.

Figure 3-8 Illustrates the IPsec ESP encapsulation mechanism.

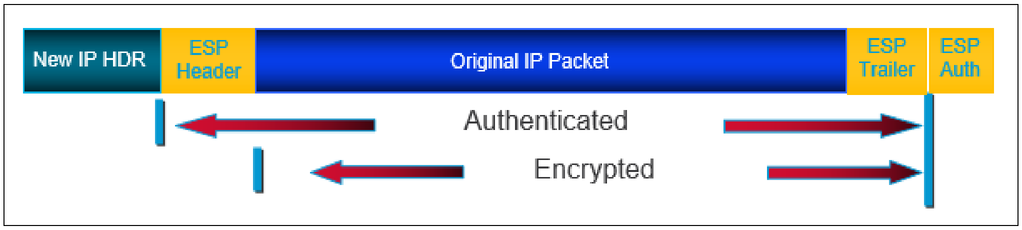

Figure 3-8 IPsec encapsulation mechanism

IPsec is applied to the physical interfaces of IBM c-type switches so that FCIP tunnels inherit this capability.

From a networking point of view, IPsec inserts an extra header into an existing TCP/IP packet. To avoid fragmentation and achieve higher performance, the encrypted packet must fit into the Ethernet interface maximum transmission unit (MTU). The maximum FC frame is 2148 bytes, and an extra 100 bytes are needed to accommodate IPsec encryption.

Best practice: Configure an MTU of 2300 bytes on long-distance links when using IPsec.

Some IP networks do not support jumbo frames and the Ethernet MTU value is set to 1500 bytes. In this case, an FC frame does not fit inside and fragmentation occurs, adding load on the transmitting and receiving switches and contributing to some extra latency. Even in this case, IPsec for FCIP links can be used, but performance is reduced.

IPsec is supported by all IBM c-type devices supporting FCIP, specifically the 24/10-port SAN extension module and the IBM Storage Networking SAN50C-R SAN extension switch.

Enabling IPsec on FCIP links requires the ENTERPRISE_PKG license.

For more information about how to configure IPsec encryption, see Configuring IP Security.

Figure 3-9 illustrates the combined use of TrustSec Fibre Channel Link Encryption and IPsec features in a design with three data centers.

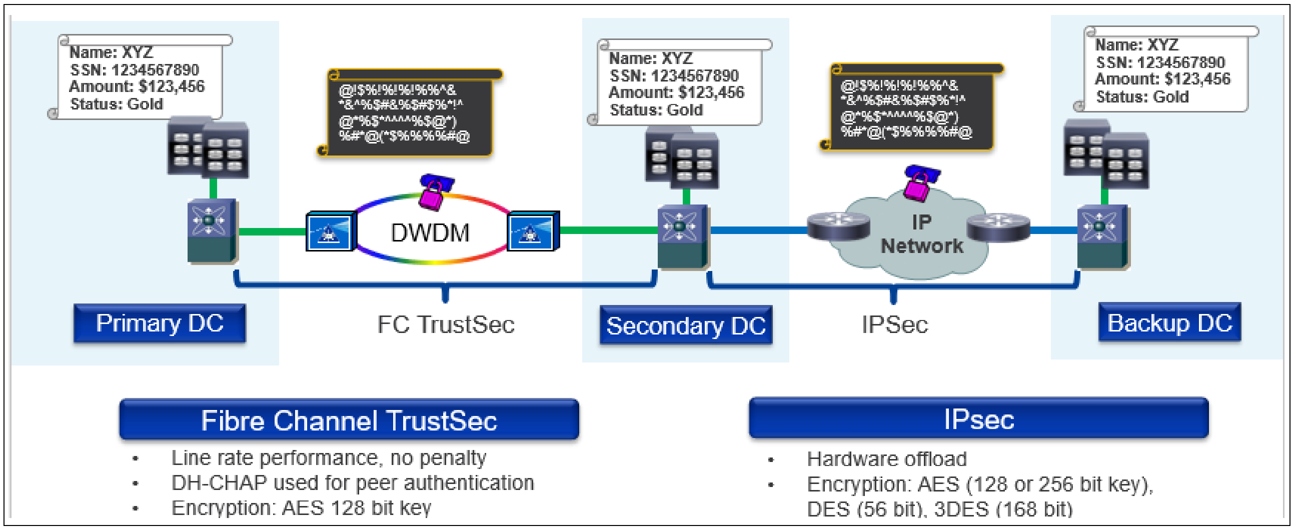

- Line rate performance, no penalty
- DH-CHAP used for peer authentication
- Encryption: AES 128 bit key

- Hardware offload
- Encryption: AES (128 or 256 bit key), DES (56 bit), 3DES (168 bit)

Figure 3-9 FC link encryption and IPsec features

4

Lab environment and topology

In this chapter, we provide a brief overview of the lab environment and topology that was used to test and document different aspects of IBM c-type switches in an IBM Z environment.

The following topics are covered in this chapter:

- Products that were used during testing
- Local DASD configuration
- Cascaded DASD configuration
- Cascaded VTS configuration
- Storage replication

4.1 Products that were used during testing

During the development of this publication, a lab was built to test and document different aspects of IBM c-type switches in an IBM Z environment. A combination of devices were used for the testing.

These devices included an IBM Z server, an IBM System Storage DS8870 server, third-party Direct Access Storage Device (DASD) systems, an IBM TS7760 Virtual Tape Server (VTS), and IBM c-type Fibre Channel Protocol (FCP)/ IBM Fibre Connection (FICON) switches and management software.

The test topologies are representative of the technologies that are required to implement specific functions, but they do not represent recommended deployments. They do not include redundant FICON/FCP fabrics and are used for demonstration and illustrative purposes only.

Figure 4-1 shows the overall test environment.

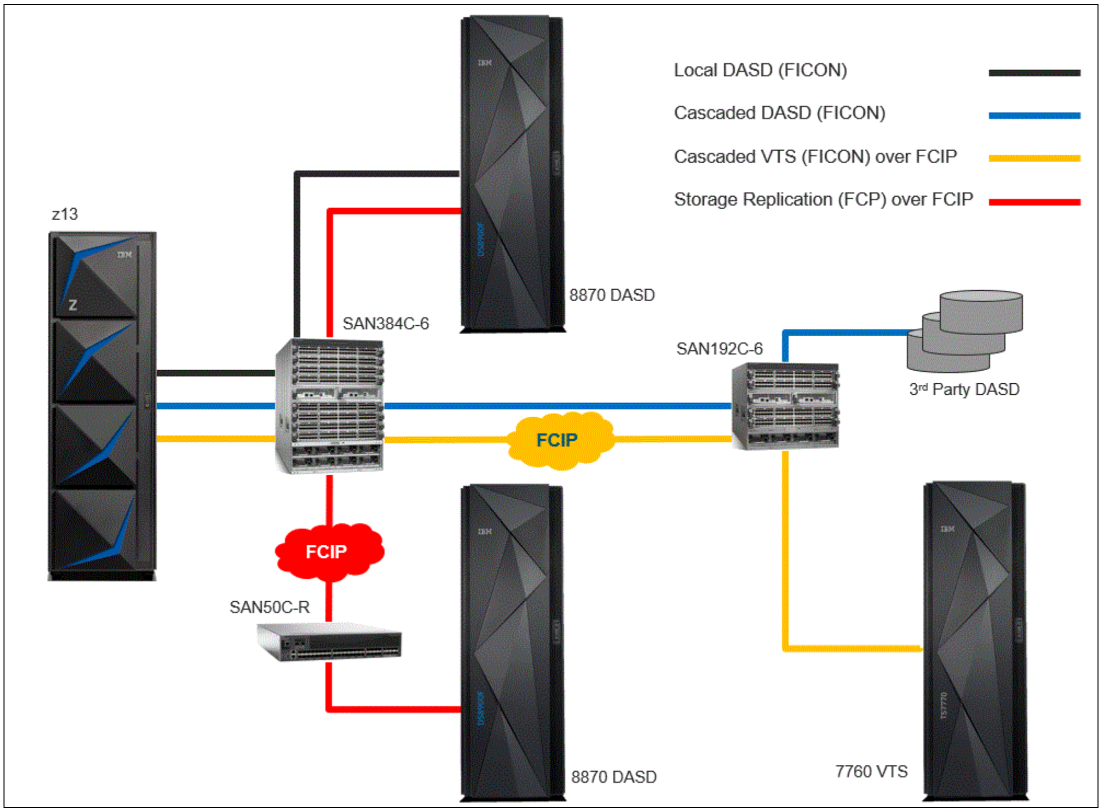

Figure 4-1 Test environment

4.2 Local DASD configuration

To represent a local FICON DASD configuration, the test environment includes an IBM Z server that is connected by using 16 Gb Fibre Channel (FC) connections to an IBM Storage Networking SAN384C-6 switch that is connected to an IBM DS8870 server, as shown in Figure 4-2 on page 93.

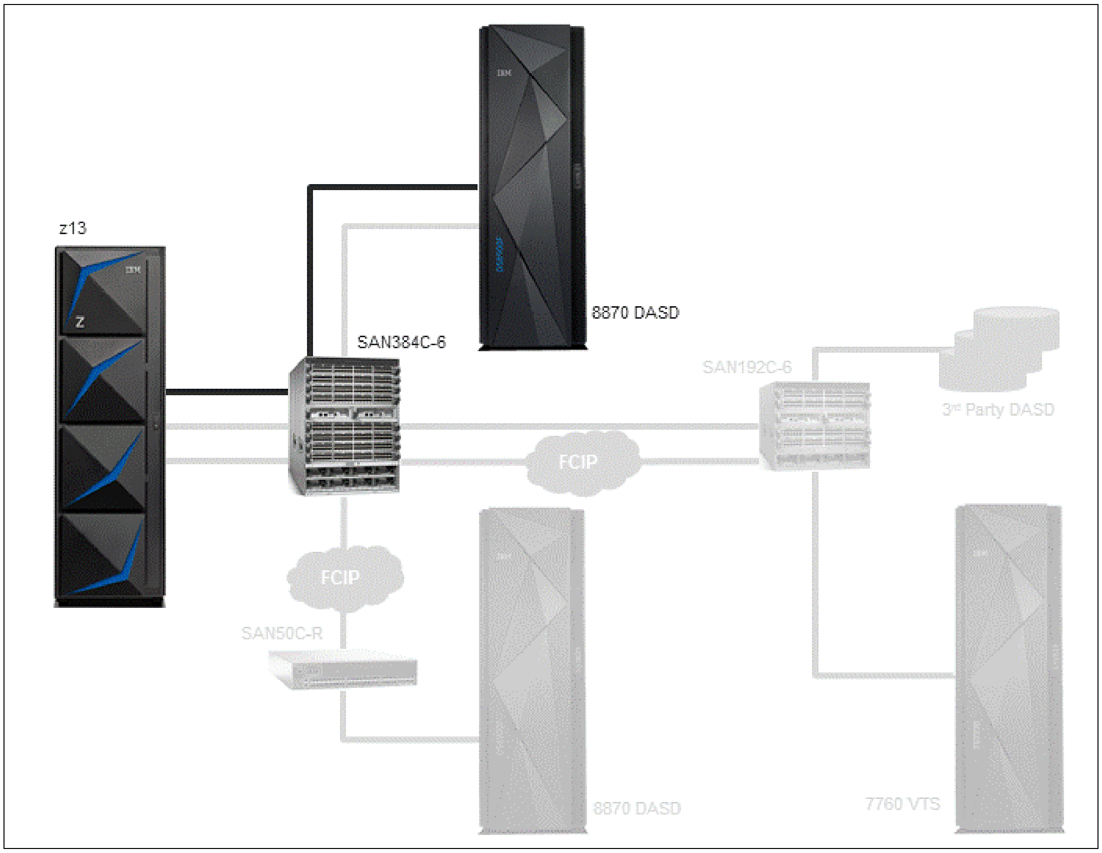

Figure 4-2 Local FICON topology

In the local FICON topology, the IBM Storage Networking SAN384C-6 switch, which is domain 0x20 in virtual storage area network (VSAN) 40, is connected to the IBM Z interfaces Channel Path ID (CHPID) 18 and CHPID 20 by using switch interfaces 0x00 and 0x30. The local DS8870 is connected to the switch interfaces 0x10 and 0x40, as shown in Figure 4-3.

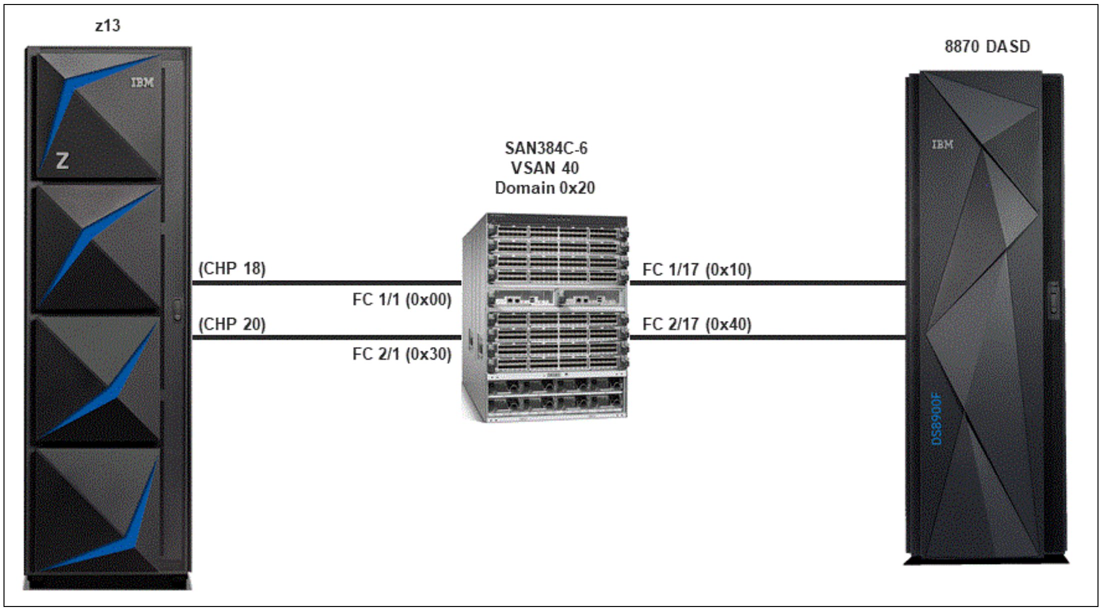

Figure 4-3 Local FICON topology connection details

4.3 Cascaded DASD configuration

To represent a cascaded FICON DASD configuration, the test environment includes an IBM Z server that is connected through 16 Gb FC connections to an IBM Storage Networking SAN384C-6 switch that is connected to an IBM Storage Networking SAN192C-6 switch, which is then connected to a third-party FICON DASD, as shown in Figure 4-4.

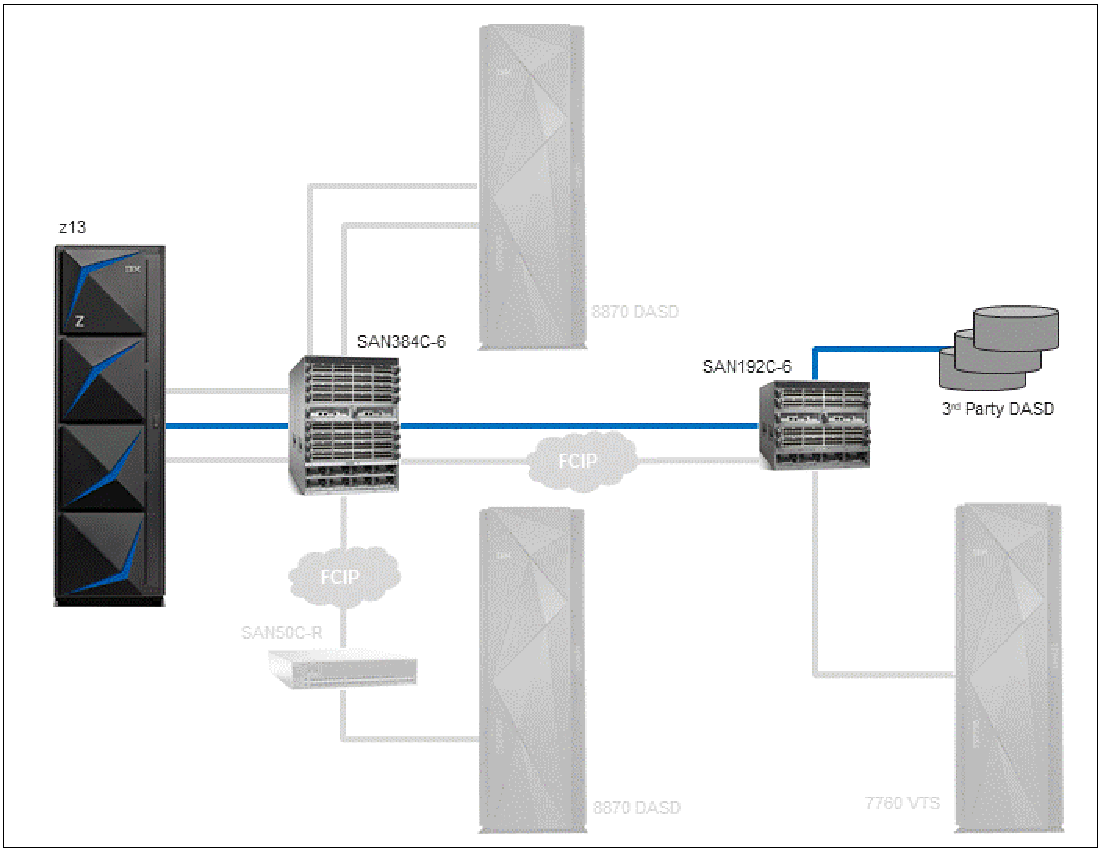

Figure 4-4 Cascaded FICON topology

In the cascaded FICON topology, there are two switches cascaded together through two physical interfaces, 0x2F and 0x5F, which are configured to form logical port channel 5 (0xF0). The function of the port channel is to provide a high availability (HA) connection between the two switches to prevent a single point of failure (SPOF) in the case of a link failure.

For this topology, VSAN 40 is used for the cascaded FICON traffic. The IBM Z interfaces are CHPID 20 and CHIPD 30, and they are connected to switch interfaces 0x22 and 0x42 respectively on the IBM Storage Networking SAN384C-6 switch. The remote DASD is connected to the IBM Storage Networking SAN192C-6 interfaces 0x0A and 0x4A, as shown in Figure 4-5 on page 95.

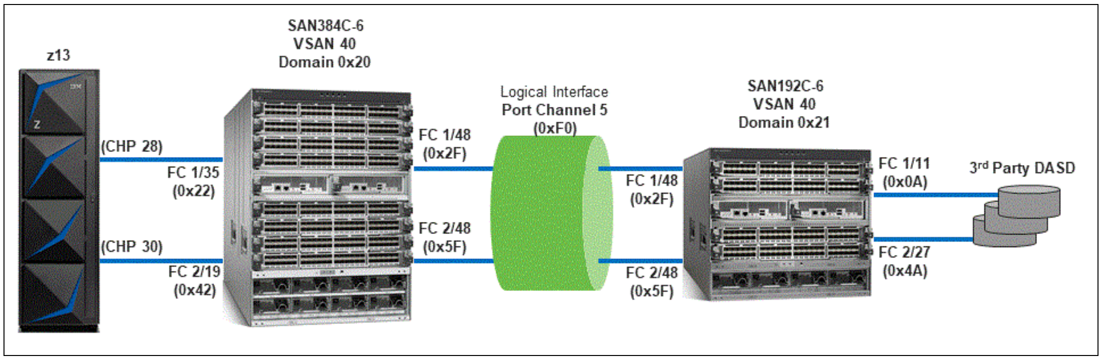

Figure 4-5 Cascaded FICON topology connection details

4.4 Cascaded VTS configuration

To represent a cascaded FICON VTS configuration, the test environment includes an IBM Z server that is connected by using 16 Gb FC connections to an IBM Storage Networking SAN384C-6 switch that is connected to an IBM Storage Networking SAN192C-6 switch over 10 GbE Fibre Channel over IP (FCIP) Inter-Switch Links (ISLs) to an IBM TS7760 VTS, as shown in Figure 4-6.

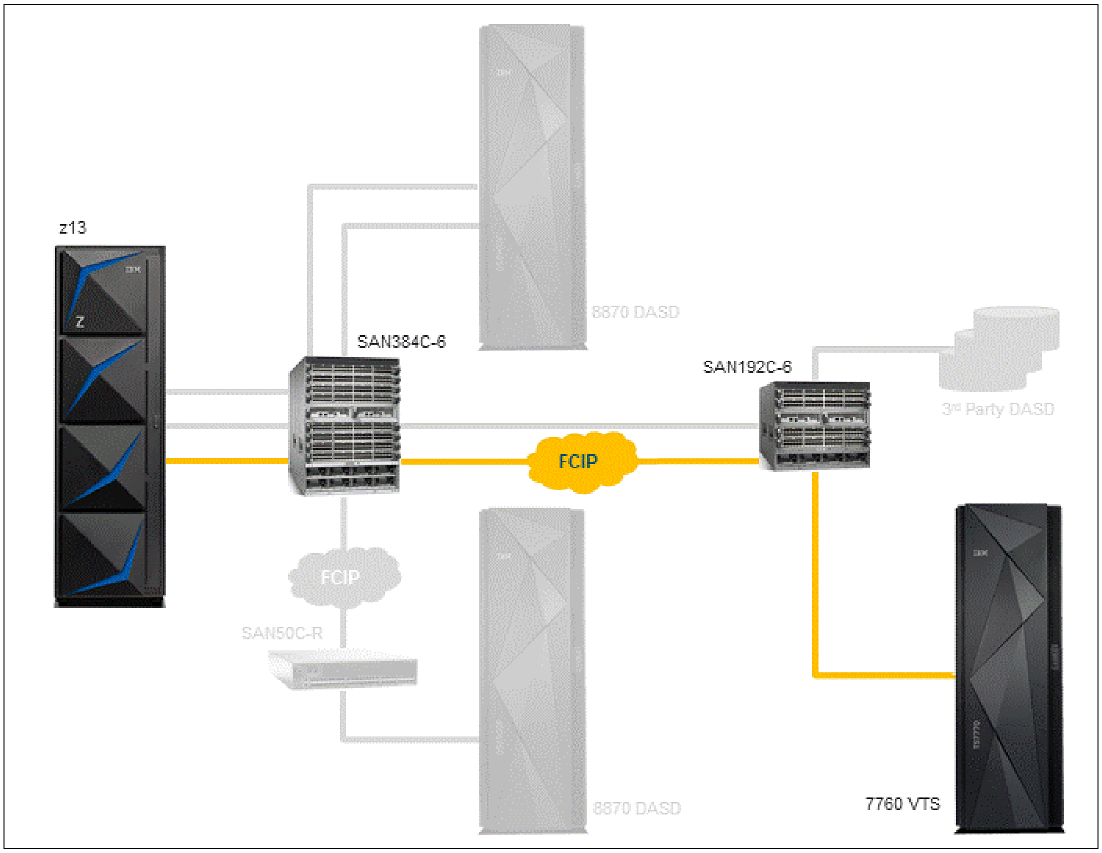

Figure 4-6 Cascaded VTS configuration

Three different FCIP topologies were used during the testing for this book. In all cases, the VTS traffic was configured to use VSAN 50 on both switches. The first topology connects the IBM Storage Networking SAN384C-6 switch to the IBM Storage Networking SAN192C-6 switch by using two 10 GbE interfaces with one 10 GbE FCIP tunnel over each interface. Logical FCIP interface 100 (0xF1) is mapped to physical interface IPs 7/1 on the IBM Storage Networking SAN384C-6 switch and physical interface IPs 6/1 on the IBM Storage Networking SAN192C-6 switch. Logical FCIP interface 110 (0xF2) is mapped to physical interface IPs 7/2 on the IBM Storage Networking SAN384C-6 switch and physical interface IPs 6/2 on the IBM Storage Networking SAN192C-6 switch.

Each interface and tunnel on the switches use different IP addresses and subnets, and they use the default FCIP TCP port of 3225. Additionally, the two FCIP interfaces are combined into a logical port channel (0xF5) for HA and redundancy.

The IBM Z interfaces CHPID 70 and CHPID 78 are connected to the IBM Storage Networking SAN384C ports 0x02 and 0x07, and the IBM TS7760 VTS is connected to the remote IBM Storage Networking SAN192C-6 switch on ports 0x01 and 0x05, as shown in Figure 4-7.

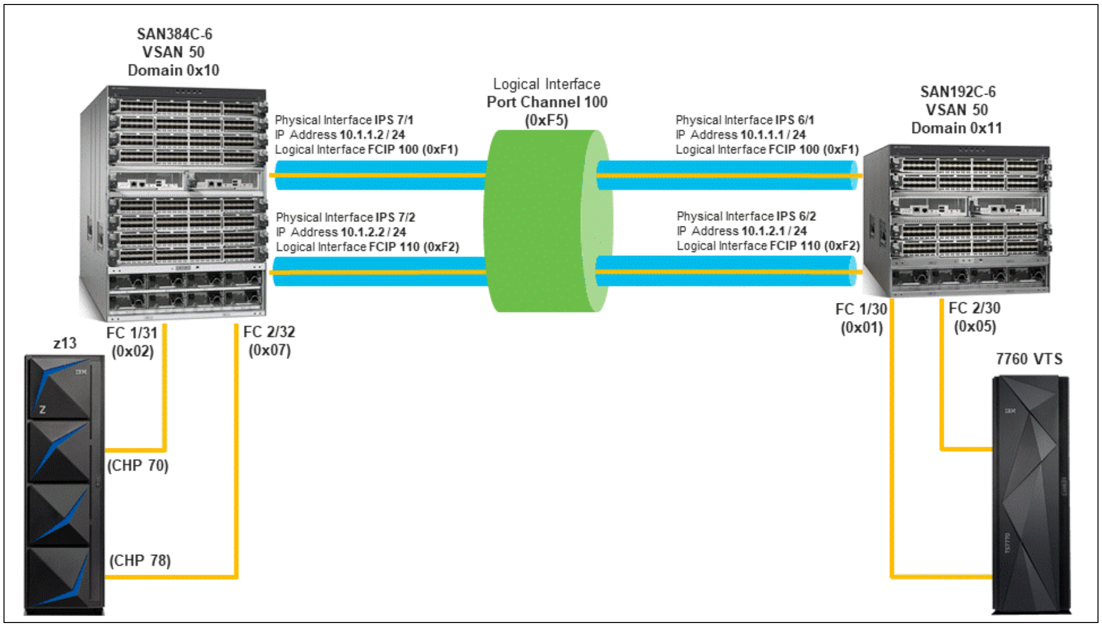

Figure 4-7 Cascaded VTS FCIP configuration with two physical interfaces

The second FCIP topology that is deployed depicts using two logical FCIP interfaces that are configured over a single physical interface. In this example, the physical interface is configured with logical subinterfaces with one 5 GbE FCIP tunnel that is defined for each subinterface. Each subinterface is mapped to a different VLAN and requires a connection to an Ethernet switch that understands VLAN tagging.

To make identification easier, the subinterface is named with the VLAN number as part of its name. So, VLAN 1000 is on subinterface IPS7/3.1000, and VLAN 1010 is on subinterface IPS7/3.1010. As in the previous example, each interface, in this case subinterface, has its own IP address and subnet and is associated to a single FCIP tunnel, as shown in Figure 4-8 on page 97.

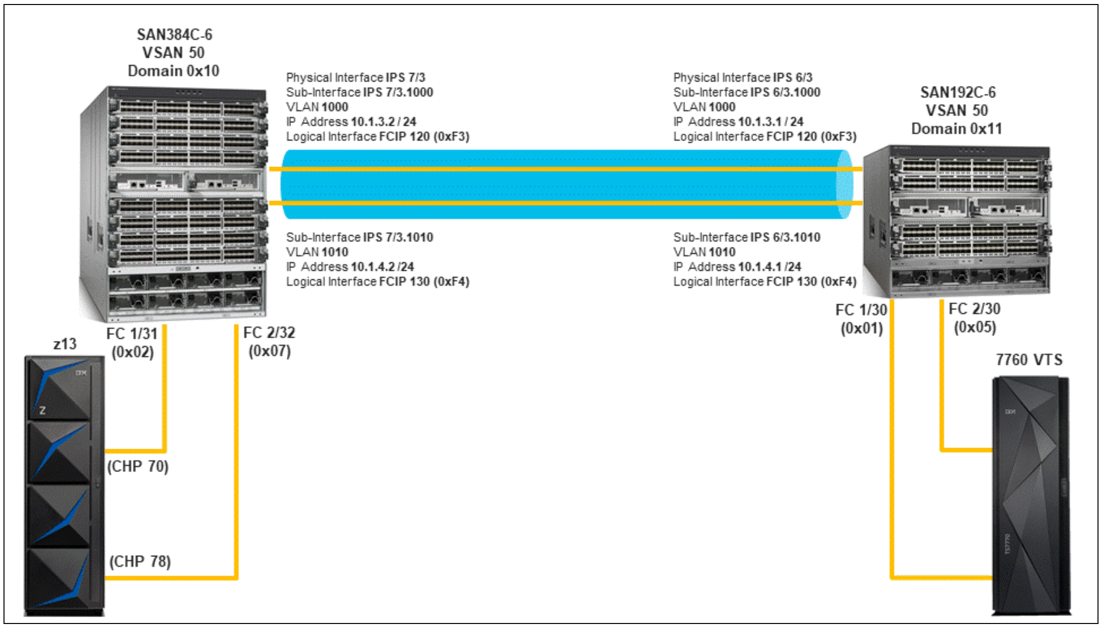
SAN384C-6
VSAN 50
Domain 0x10
Physical Interface IPS 7/3
Sub-Interface IPS 7/3.1000
VLAN 1000
IP Address 10.1.3.2 / 24
Logical Interface FCIP 120 (0xF3)
Physical Interface IPS 6/3
Sub-Interface IPS 6/3.1000
VLAN 1000
IP Address 10.1.3.1 / 24
Logical Interface FCIP 120 (0xF3)
SAN192C-6
VSAN 50
Domain 0x11
Sub-Interface IPS 7/3.1010
VLAN 1010
IP Address 10.1.4.2 /24
Logical Interface FCIP 130 (0xF4)
Sub-Interface IPS 6/3.1010
VLAN 1010
IP Address 10.1.4.1 /24
Logical Interface FCIP 130 (0xF4)
z13
FC 1/31
(0x02)
FC 2/32
(0x07)
FC 1/30
(0x01)
FC 2/30
(0x05)
7760 VTS
(CHP 70)
(CHP 78)

Figure 4-8 Cascaded VTS FCIP configuration with one physical interface and two subinterfaces

The third FCIP topology that is deployed shows using two logical FCIP interfaces that are configured over a single physical interface with no subinterfaces. In this example, each 5 GbE FCIP tunnel is defined to the same physical interface and uses the same IP address, but is differentiated in the configuration by using TCP ports 3225 and 3226, as shown in Figure 4-9.

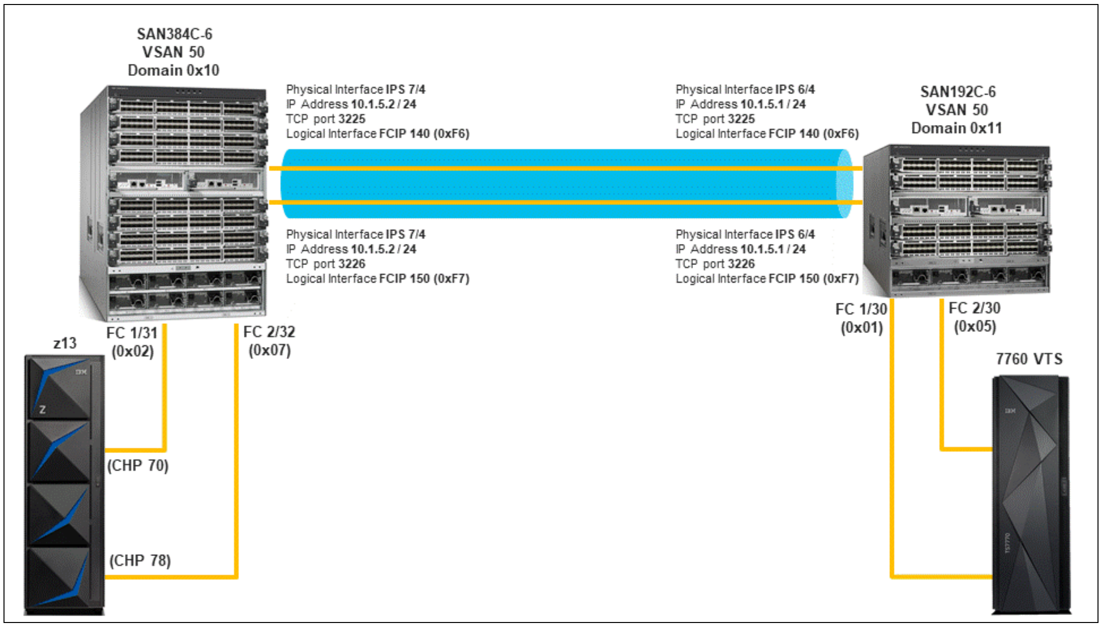

Figure 4-9 Cascaded VTS FCIP configuration with one physical interface and two FCIP interfaces with different TCP ports

4.5 Storage replication

To represent an FCP storage replication configuration, the test environment includes an IBM Z server that is connected by using 16 Gb FC connections to an IBM Storage Networking SAN384C-6 switch that is connected to an IBM DS8870 server. The IBM DS8870 server copies the data to a second remote IBM DS8870 server by using the IBM Storage Networking SAN384C-6 switch over a 500 Mb FCIP ISL to an IBM Storage Networking SAN50C-R switch, and then to the remote IBM DS8870 server, as shown in Figure 4-10.

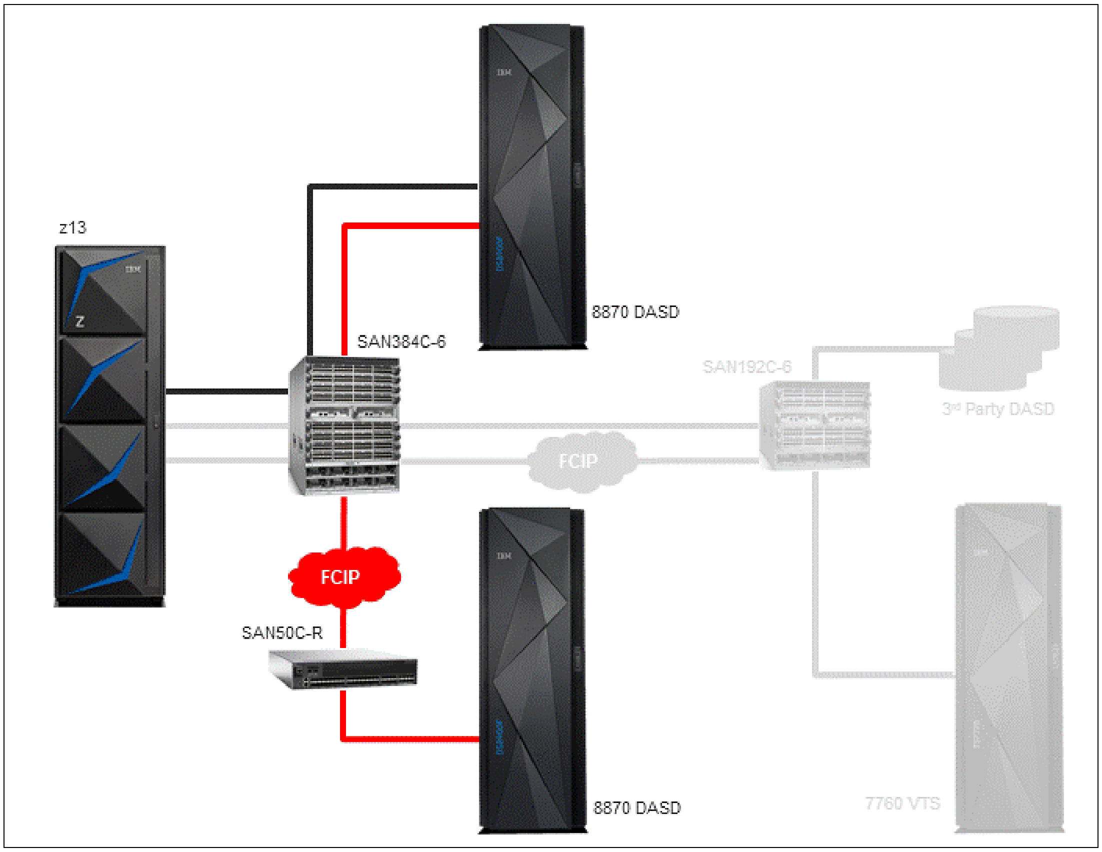

Figure 4-10 FCP storage replication

In this test setup, an open systems VSAN 100 is used on both switches. The local DASD is connected to the local IBM Storage Networking SAN384C-6 switch on ports FC1/15 and FC2/15, and the remote DASD is connected to the remote IBM Storage Networking SAN50C-R switch on ports FC1/17 and FC1/19. There is a single 1 GbE physical port on each switch with a 500 Mbps FCIP tunnel that is defined that connects from Raleigh, North Carolina to San Jose, California over a routed IP network, as shown in Figure 4-11.

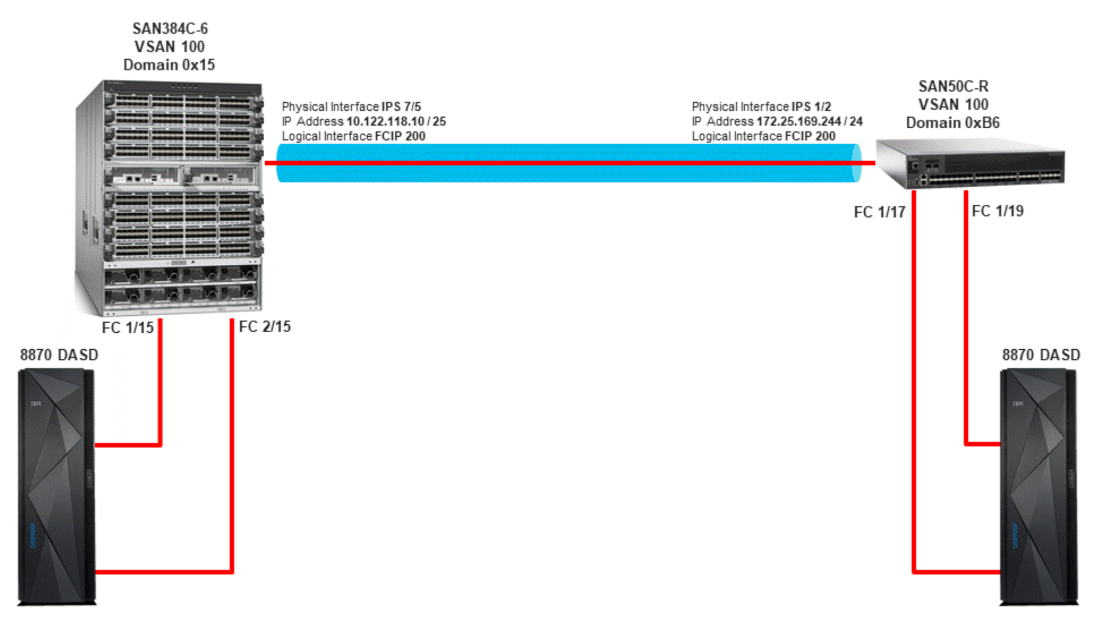
SAN384C-6
VSAN 100
Domain 0x15
Physical Interface IPS 7/5
IP Address 10.122.118.10 / 25
Logical Interface FCIP 200
Physical Interface IPS 1/2
IP Address 172.25.169.244 / 24
Logical Interface FCIP 200
SAN50C-R
VSAN 100
Domain 0xB6
FC 1/17
FC 1/19
FC 1/15
FC 2/15
8870 DASD
8870 DASD

Figure 4-11   FCP storage replication FCIP details

5

IBM Storage Networking c-type design considerations

This chapter describes the main design considerations for storage networks in an IBM Fibre Connection (FICON) environment, including topologies, physical implementation details, and a migration strategy from old deployments to modern ones. Some FICON fundamentals are also briefly presented to help you understand the main concepts and terminology and the peculiarities of FICON Directors and their impact on network design.

The following topics are covered in this chapter:

- IBM Z connectivity options
- FICON basics
- General FICON planning considerations
- IBM c-type considerations for FICON
- Topologies
- Inter-Switch Links and FICON routing options
- In order delivery of frames
- FCIP
- Power, cooling, racking, and cabling
- Migration strategy

5.1 IBM Z connectivity options

This section describes the origin of the IBM mainframe and its connectivity options.

5.1.1 The origin of the IBM mainframe

IBM mainframes are high-performance computers that are developed and sold primarily by IBM. They are characterized by large amounts of memory, processors, and attached storage devices, with the ability to process billions of calculations and transactions in real time. The mainframe is often most central to the storage and computing requirements of large enterprises. The term *mainframe* came into use as a reference to the initial shape of such devices. “Main” was because it was the largest, most central (or only) computer. The “frame” design looked similar to telephone switches and was notably different from racks and cabinets holding other components. Over time, the two words “main” and “frame” merged into a single one with the meaning of “big computer.”

An IBM mainframe computing system (also referred to as a central processor complex (CPC)) consists of a set of hardware products, including a processor unit (PU), and software products, with the primary software being an operating system (OS) such as IBM z/OS. The *central processor (CP)* is the functional hardware unit that interprets and processes program instructions. One or more CPs work together with other system hardware that can be shared between multiple CPs, such as I/O channels and storage. A single CP can process only one instruction at a time. To achieve more processing power, CPCs can contain multiple CPs, and an OS running on the CPC can use multiple CPs. When an OS has multiple CPs at its disposal, it can process instructions in parallel, which increases performance.

The mainframe is critical to commercial databases, transaction servers, and applications that require high resiliency, security, and agility. First introduced more than 50 years ago, mainframes are still omnipresent today. They handle massive amounts of heterogeneous processing tasks, reliably, securely and with great redundancy, and they offer compatibility with earlier programs and applications. Since 1998, there is support for Linux as an alternative to native mainframe OSs. The result was a unique combination of earlier and modern technologies. In fact, it is now commonplace to see mainframes that run COBOL applications on z/OS alongside Docker containers on Linux by using IBM z/VM®. Mainframes are in use at 92 of the world's top 100 banks, 23 of the 25 top airlines, all the world's top 10 insurers, and 71% of Fortune 500 companies.

5.1.2 Mainframe connectivity options

Mainframes offer many connectivity options, depending on what device they must communicate with and the required performance.

FICON is a storage area network (SAN) communication protocol that is used on IBM Z mainframe computers to exchange data with their storage arrays. The underlying transport network uses the same hardware as Fibre Channel Protocol (FCP) SANs, but there are some unique and critical differences. It is necessary to understand the differences, and their implications in a SAN design, before describing FICON design considerations and best practices. Currently, FICON is the most common method that is used to connect the mainframe to its auxiliary I/O devices, but Direct Attached Storage Devices (DASDs) were popular long before networked solutions became prevalent. Recently, a new implementation was introduced for mainframe to storage direct connection, and it is know as IBM zHyperLink. It complements, not replaces, FICON channels when there is a need for low-latency communication.

Mainframe server-to-server communications are most commonly implemented by using Open Systems Adapters (OSAs), coupling links, Shared Memory Communications over RMDA (SMC-R), and FICON Channel-to-Channel (FCTC) connections. Internal Coupling (IC) channel connections, HiperSockets, Shared Memory Communication - Direct Memory Access (SMC-D), and channel-to-channel (CTC) connections can be used for communications between logical partitions (LPARs) within an IBM Z platform. These technologies eliminate the need for any physical cabling or external networking connection between these virtual servers. It is a network in the box.

IBM Parallel Sysplex® technology represents a synergy between hardware and software and is composed of Parallel Sysplex capable servers, CFs, coupling links, Server Time Protocol (STP) and more. Parallel Sysplex technology is a highly advanced, clustered processing system. It supports high-performance, multisystem, read/write data sharing, which enables the aggregate capacity of multiple z/OS systems to be applied against common workloads.

Figure 5-1 shows the connectivity options with the IBM Z platform.

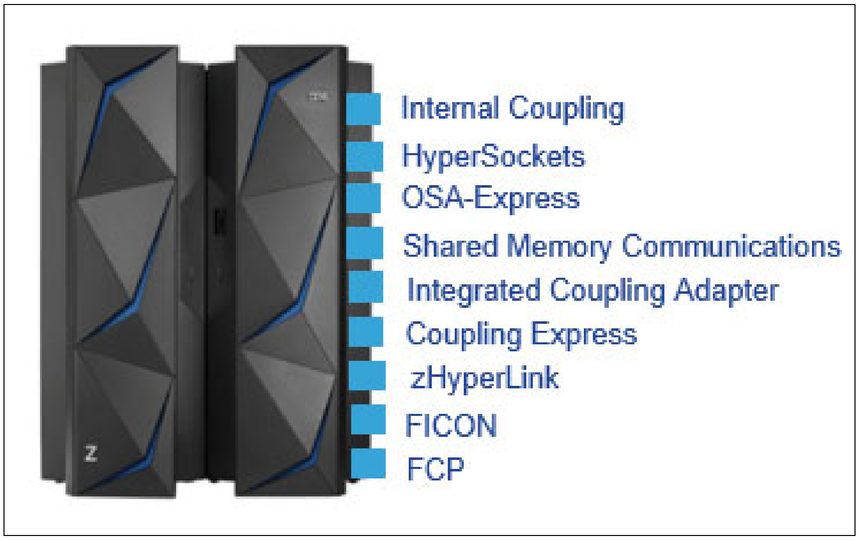

Figure 5-1 IBM Z connectivity options

For more information about connectivity options for the IBM Z platform, see *IBM Z Connectivity Handbook*, SG24-5444.

5.2 FICON basics

FICON channels enable full duplex data transfer, which means data travels in both directions simultaneously. Also, multiple concurrent I/Os can occur on a single FICON channel. The FICON channels currently support a speed of 16 Gbps. FICON channels experience minimal data rate droop at distances up to 100 km (62 miles) because of the specific data transfer method of this transmission protocol.

This section covers some fundamentals of the FICON protocol, and its origin, purpose, and terminology. It also explains why a FICON switch is different from a Fibre Channel (FC) switch.

5.2.1 Directors and switches

It is important to clarify the terms *director* and *switch*. In general, an FC *switch* is smaller in size than a director, has less port capacity, and has one or more single points of hardware failure. A *director* is a modular device, larger in size with more port capacity, and redundant components for no single point of hardware failure. However, the term switch is often used generically for both switches and director, and likewise the term director is sometimes used in a mainframe environment for both director and smaller switches.

A FICON switch (or director) supports I/O that contains Fibre Channel Single Byte (FC-SB)-6 payloads, supports the Fibre Channel Framing and Signaling (FC-FS) Extended Link Services (ELS) that FICON requires, and has support for the IBM Control Unit Port (CUP) function. All these items are described later.

5.2.2 The mainframe architecture: Channel subsystem and control units

The first mainframes used the System/360 (S/360) I/O architecture. The mainframe had special-purpose I/O processors that were called *channels* that allowed the main system processors to continue to perform work while the relatively slow I/O operations were performed concurrently. A *channel* is a physical element that acts as the interface between the main processor and the devices that it talks to. A group of I/O channels is referred to as the channel subsystem (CSS), and it provides a pipeline through which data is exchanged between servers or between a server and external devices. The CSS has evolved with the increased scalability of IBM mainframes, keeping pace with the augmented needs in terms of speed and throughput for data transfers. Today, multiple CSSs can be configured within the same mainframe.

On the other end of the communication link, a *control unit* (CU) was needed to interact and convert between the channel requests and the functions of devices, such as disk, tape, printers, and card readers. A CU provides the logical capabilities that are necessary to operate and control an I/O device, and it adapts the characteristics of each device so that it can respond to the standard form of control commands that are provided by the CSS. A CU may be housed separately, or it may be physically and logically integrated with the I/O device, the CSS, or within the server itself. Behind the CU are the I/O devices representing the communication target of the mainframe.

In z/OS environments, the combination of a channel, a CU, and an I/O device behind it is named a *channel path*, and it is statically defined. For switched environments, the channel path becomes a logical entity that is established through the switched fabric, and it must be statically defined too, as shown in Figure 5-2 on page 105.

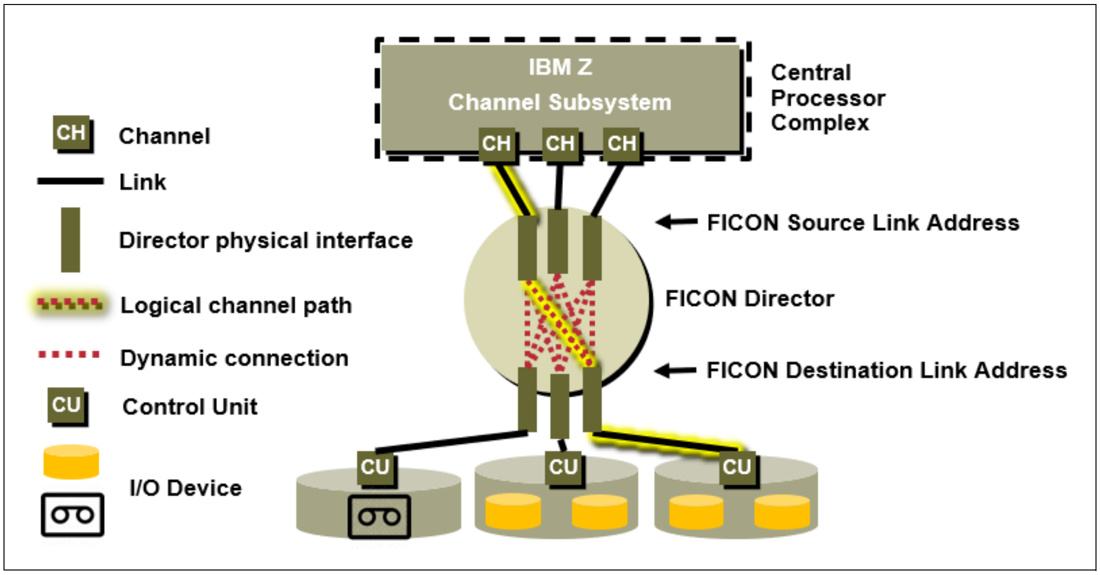

Figure 5-2 Logical channel path in a FICON switched fabric

5.2.3 Mainframe I/O and source-based routing

Historically, a mainframe I/O operation began with an instruction, the Start I/O (SIO), which indicated which device that it was directed to, such as 312, and how to reach it. Then, device 312 meant "use channel number 3 and send the device commands to CU 1, who will send them to device 2". Figure 5-3 shows how the complete path for the I/O operation was controlled by the program running on the mainframe.

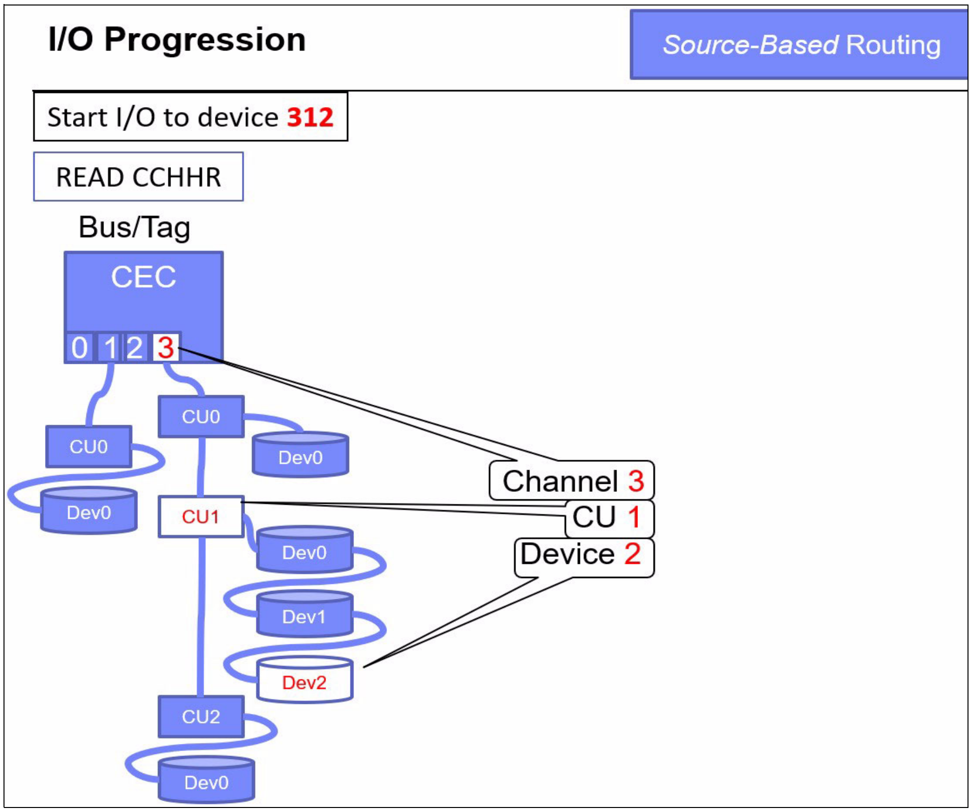

Figure 5-3 I/O in S/360: How a program controls the complete path for an I/O operation

We call this *source-based routing* because the communication path of the device is specified by the initiator of the I/O. All the channels, CUs, devices, and their connections are statically defined and loaded into the machine's memory by using a program that is called the Input/Output Configuration Program (IOCP).

This static approach was perfectly adequate in the period when storage devices were directly attached to the mainframe on a SCSI bus and the concept of storage networks had yet to come.

5.2.4 Extended Count Key Data

Unlike modern UNIX file system structures, back in 1964, data was stored in Count Key Data (CKD) format, which used storage space efficiently, which was critical because at the time storage space was a scarce and expensive resource. Every record had a length, or count, which was followed by a unique identifier, or key, and then the data.

In this way, white space such as blanks did not need to be stored, and the storage media, such as disk, could be filled to the limit with little space left unused.

As an improvement on the CKD method, Extended Count Key Data (ECKD) introduced support for nonsynchronous operation. In other words, the transfer of data between the channel and the CU is not synchronized with the transfer of data between the CU and the device, as shown in Figure 5-4.

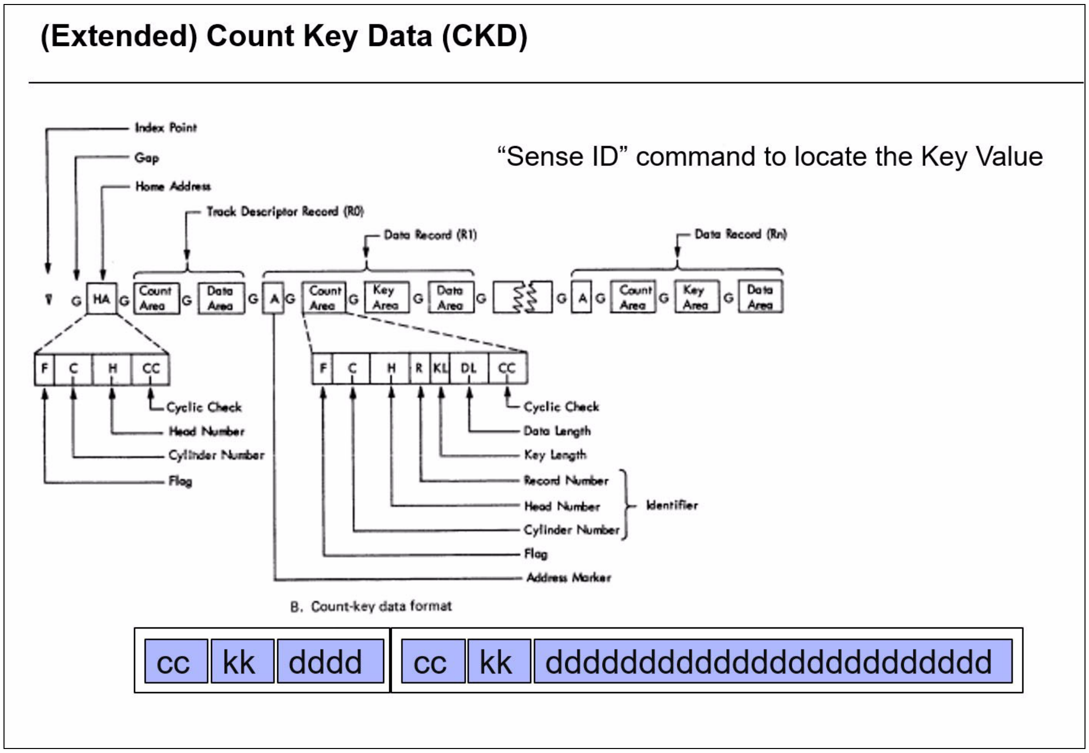
(Extended) Count Key Data (CKD)

"Sense ID" command to locate the Key Value

Figure 5-4 Extended Count Key Data format

To summarize, the ECKD approach can be described this way:

- All data blocks are variable length.
- The "Count" field contains the length and location (cylinder, track, and record number).
- The "Key" field, which is optional, describes the data.

- The “Data” field is the data of length “count”.
- I/O operations are conversational, with a series of requests and responses.

5.2.5 Channel programs

The SIO points to a channel program that is a set of instructions for the channel. Many instructions are sent by the channel to the CU, such as read/write. Other channel program instructions direct the running of the channel program, such as Transfer In Channel (TIC), which enable IF/THEN/ELSE programming constructs. A typical channel program to read data from a record, as shown in Figure 5-5, is:

- Sense (pointing to a data field with the key, such as X)
- TIC *-8 (if the sense did not find the record with key X, jump back to the previous instruction, where * 8 meaning “8 bytes before the current location”, knowing that channel commands are all 8 bytes)
- Read (pointing to a data location where the record's contents are to be placed)

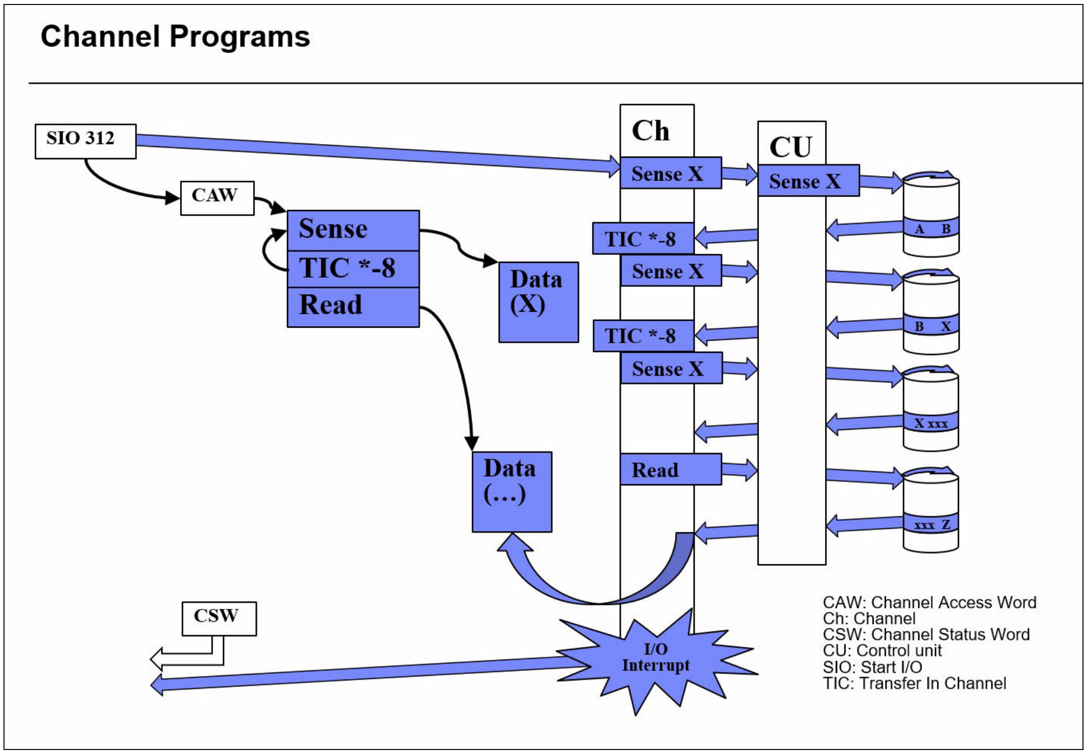

Figure 5-5 Conversational form of CKD I/O

The channel sends the sense instruction to the CU, which sends it to the device. When the device finds the next record ID under its magnetic read/write head, it sends the ID to the CU. If the CU sends a negative response back to the channel, it indicates that the requested ID, X, was not found in the current record. If the channel receives a negative response, it runs the next instruction, TIC *-8, sending the sense instruction back to the CU, which in turn sends the command to the device. The disk rotates and eventually the next record's ID is under the read head, and the process continues.

Eventually, the record with ID X might come under the read head. In that case, the TIC channel command does not jump back, but proceeds to the next channel command, Read, which the channel sends to the CU and then to the device. The device reads the records' data and sends it back to the CU and then the channel. The channel transfers the data to the location that the user specifies in the channel program. In this case, the channel program completes successfully, so the channel sends an I/O interrupt to the processor, which eventually returns to the program. The program continues with the instruction after the SIO. The program that issues the SIO has an indication of the success or failure of the channel program in the Channel Status Word (CSW). If successful, the requested data is in the indicated location.

As you can see, this method of I/O is a real-time dialog with many commands and responses back-and-forth between the channel and the CU, and between the CU and the device. The order of operations is critical because a read or write must follow immediately after the correct record ID is located.

5.2.6 Fixed-Block architecture

UNIX systems use a different method to store data that is called the Fixed-Block Architecture (FBA). Let us compare CKD to FBA to appreciate the differences.

FBA, as the name suggests, uses a consistent block size on I/O devices. The location of the block can be determined by the block's address and the rotation of the device, and the device can store the I/O request until the record falls under the head. On one hand, FBA is inefficient because it wastes disk space. The block size could be 512 bytes (4 K bytes in recent hardware), but that block might store only 1 byte of user data and the rest of the block's space is wasted. Also, it requires not one but multiple I/Os for large chunks of data.

However, FBA is simple and the channels (that is, host bus adapters (HBAs)) and devices are thus less expensive. An I/O operation can read/write a block of data in almost any order independent of the location of the devices' read/write head. For example, the writing of three blocks such as A, B, and C can be required by the program, but because of the rotational delay of the storage media, the order of write operations is B, then A, then C. Using FBA, the program, the HBA, and the device care much less about the order of I/O commands. Of course, there is no complete freedom about that order, but it is a lot less stringent than with the ECKD method.

To summarize, FBA (also known as Logical Block Addressing), can be described this way:

- All data blocks are the same size.
- Data blocks are sequentially numbered.
- The starting point is at Cylinder 0, Track 0, Record 0.
- The end point is device-dependent.
- 32-bit addressing with 512-byte sectors allows for 8 TB per device (more with 4 K byte sectors).
- The similar Cylinder, Head, and Sector (CHS) format is 24 bits.

Figure 5-6 on page 109 shows the FBA method.

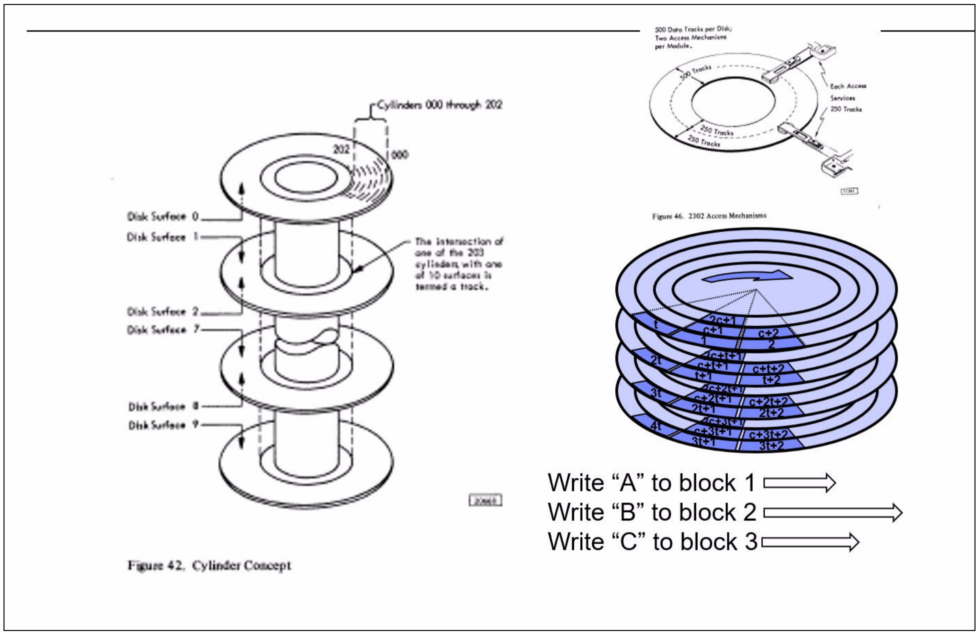
Cylinders 000 through 202
202
000
Disk Surface 0
Disk Surface 1
The intersection of one of the 203 cylinders with one of 10 surfaces is termed a track.
Disk Surface 2
Disk Surface 7
Disk Surface 8
Disk Surface 9
Figure 42. Cylinder Concept
500 Data Tracks per Disk; Two Access Mechanisms per Module.
500 Tracks
Each Access Services 250 Tracks
250 Tracks
250 Tracks
Figure 46. 2302 Access Mechanisms
t
2t
3t
4t
2c+1
c+1
c+2
Write “A” to block 1
Write “B” to block 2
Write “C” to block 3

Figure 5-6 The FBA method

5.2.7 Multipath I/O

An I/O processor might have more than one channel to access a CU and the devices behind it, and this implementation is called multipath I/O (MPIO). Many OSs (both mainframe native and UNIX variants) use MPIO to enhance performance and provide more reliable access to a device. Mainframes have a maximum limit of eight paths to a CU, and (if the device's CU supports it) at least two paths are expected as a minimum. Also, some CUs can determine which paths are to the same machine and OS and can respond on a different channel than from where the I/O command came, essentially multipathing in the other direction. There are several techniques for choosing which path should be selected, including the simple round-robin approach. The different paths might be different lengths or go through different hardware, and thus provide different response times.

5.2.8 Physical Channel ID and Channel Path ID

The CSS enables communication from the mainframe CPC to peripheral components through channel connections. The channels in the CSS permit transfer of data between the mainframe and storage devices (or other servers) under the control of a channel program. The CSS allows channel I/O operations to continue independently of other operations within the mainframe server itself, which allows other functions to resume after an I/O operation is initiated. The CSS also provides communication between LPARs within a physical server by using internal channels. A CSS can have up to 256 channel paths. A channel path is a single interface between a server and one or more CUs. Commands and data are sent across a channel path to perform I/O requests.

I/O adapters are often packaged in hardware units (called cards, blades, or modules), so they can be added, removed, relocated, or replaced as needed. Often, there is more than one channel adapter in the card for space and cost concerns. For example, ESCON cards often have eight channels, and FICON cards have as many as four, but most often come with two channels.

The Physical Channel ID (PCHID) is used to identify a specific channel interface. The PCHID is determined by the location in the mainframe where the I/O card under consideration is inserted (slot/port). It is typically three hexadecimal characters with a range of 000 - FFF (4095 elements).

The Channel Path ID (CHPID) is a logical value that is assigned to each channel path of the system that uniquely identifies that path. The CHPID number range is hexadecimal 00 - FF (256 elements) and must be unique within a CSS. With IBM Z servers, you can define more than 256 CHPIDs in the system by using multiple CSSs. CHPIDs are logical values and provide a level of abstraction that enables easier migration of programs from machine to machine. CHPIDs are not preassigned on IBM Z platforms. The administrator must assign CHPIDs to PCHIDs by using the appropriate tools.

5.2.9 IOCP

To implement source-based routing, the channels and the I/O subsystem (IOS) of the OS must have the configuration of channels, CUs, and devices. This task is done by defining all these elements and how they are attached.

Figure 5-7 shows a simple example of a definition statement from one channel, 3, to a CU, 1000. The 256 unit addresses on the CU provide device addresses in the range 1000 - 10FF. The `UNIT=2105` parameter defines the type of CU.

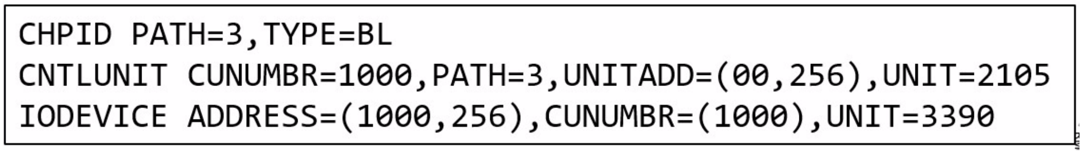

```
CHPID PATH=3,TYPE=BL
CNTLUNIT CUNUMBR=1000,PATH=3,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(1000,256),CUNUMBR=(1000),UNIT=3390
```

Figure 5-7 Definition statement example

The list of all definition statements is saved in a input/output definition file (IODF) that serves as the input to the IOCP. The IOCP creates a compiled and specially formatted version that is loaded in to memory and used by the channels. This compiled version is called the Input/Output Configuration Data Set (IOCDS). The IOCP also invokes the Multiple Virtual Storage Configuration Program (MVSCP) to create the version that is used by Multiple Virtual Storage (MVS).

> **Note:** MVS was one of the primary mainframe OSs on IBM S/390® computers and the precursor of z/OS. Older MVS releases are no longer supported by IBM, but z/OS supports running older 24-bit and 31-bit MVS applications and newer 64-bit applications.

There can be many IOCDSs, one of which is selected at Initial Machine Load (IML) or at Initial Program Load (IPL). Mainframes are resilient, so an IML from a powered-down machine is rare, but IPLs of individual LPARs are far more common. The IOCDSs are kept at the Service Element (SE) and are attached to the mainframe, which provides direct management and support services for the mainframe.

5.2.10 Hardware Configuration Definition

You must define an I/O configuration for the OS software and for the CSS hardware. The Hardware Configuration Definition (HCD) makes the definition of an IOCP input file much simpler with a hierarchical dialog format, help for the user, real-time error checking and ease of management of multiple IODFs. It is a best practice that the HCD is used to build and control the IBM Z ICODSs, as opposed to writing IOCP statements directly.

HCD can make dynamic I/O configuration changes for both hardware and software. An I/O configuration defines the hardware resources that are available to the OS and the connections between these resources. The resources include the channels, the ESCON/FICON Directors (switches), the CUs, and the devices.

To summarize, the HCD element of z/OS supplies an interactive user dialog box to generate the IODF and the IOCDS. The validation checking that HCD performs as data is entered helps eliminate errors before the new I/O configuration is implemented. The output of HCD is an IODF, which is used to define multiple hardware and software configurations to the z/OS OS.

When you activate an IODF, HCD defines the I/O configuration to the CSS and the OS. With the HCD activate function or the MVS activate operator command, you can change the current configuration without performing an IPL of the LPAR software or a power-on reset (POR) for the hardware. Making changes while the system is running is known as *dynamic reconfiguration*.

Figure 5-8 shows the relationships of the definition statements, processes, data sets (or files), and memory areas.

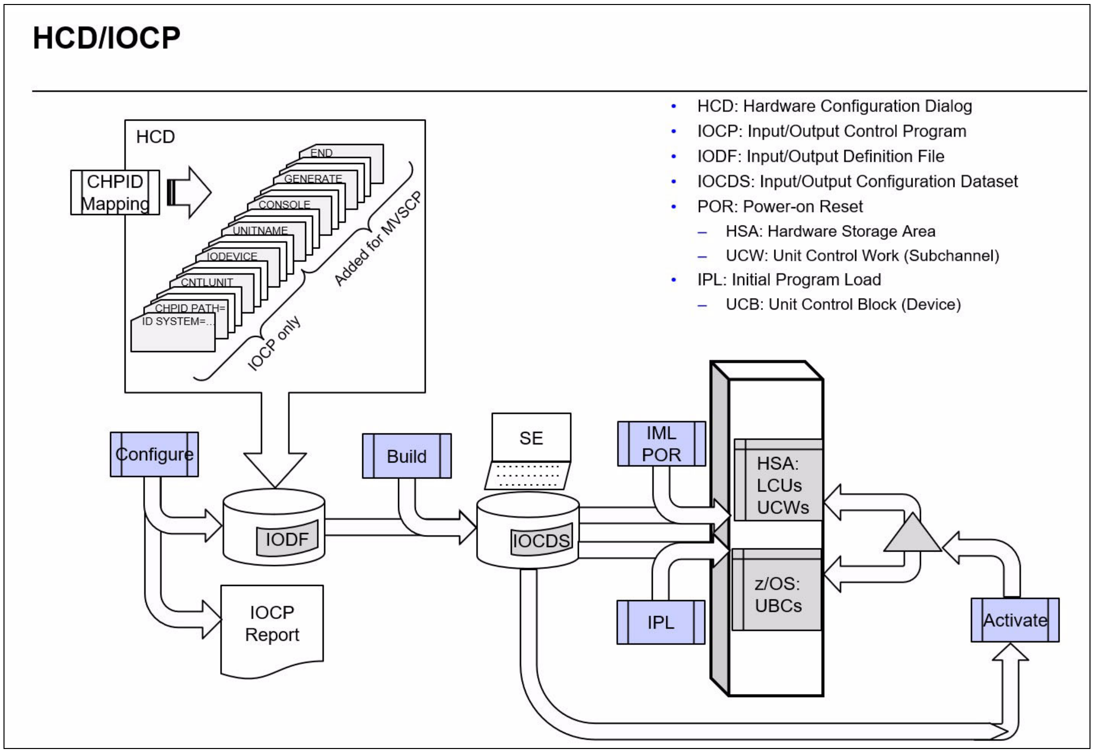
HCD/IOCP

- HCD: Hardware Configuration Dialog
- IOCP: Input/Output Control Program
- IODF: Input/Output Definition File
- IOCDS: Input/Output Configuration Dataset
- POR: Power-on Reset
  - HSA: Hardware Storage Area
  - UCW: Unit Control Work (Subchannel)
- IPL: Initial Program Load
  - UCB: Unit Control Block (Device)

Figure 5-8 HCD/IOCP relationship

Where:

- The Configure process of HCD creates the IOCP and MVSCP input statements and stores them in an IODF.
- The Build process of IOCP compiles the IODF and creates an IOCDS that is stored at the machine's SE.
- The IML or POR loads the specified IOCDS directly into the machine’s Hardware Storage Area (HSA) in a predigested format that can be used immediately by the CSS.
- The IPL loads the OS configuration in a similar manner.
- The Activate process dynamically updates the current running configuration by creating a “delta” that contains only the differences from the in-memory IOCDS and the new IOCDS, and then adds or removes those changes into memory.

5.2.11 CHPID Mapping Tool

On IBM Z, a logical CHPID number is assigned to a PCHID that corresponds to a hardware location (slot or port). This assignment is done by the user by using the CHPID Mapping Tool (CMT) or directly by using the configuration build process through HCD or IOCP. Of course, for internal channels (IC links and HiperSockets), there is no PCHID association and everything stays at the logical level.

Mainframe environments are designed for reliability. MPIO has redundant paths to the same device, which allows I/O to a device to continue if there is any failure of a single channel, the channel card, the fiber connections, the switch, the I/O adapter on the CU, or the card on the CU. If both MPIO channels to a CU are on the same channel card, a single failure of the card can cause the loss of both channels, as shown in Figure 5-9.

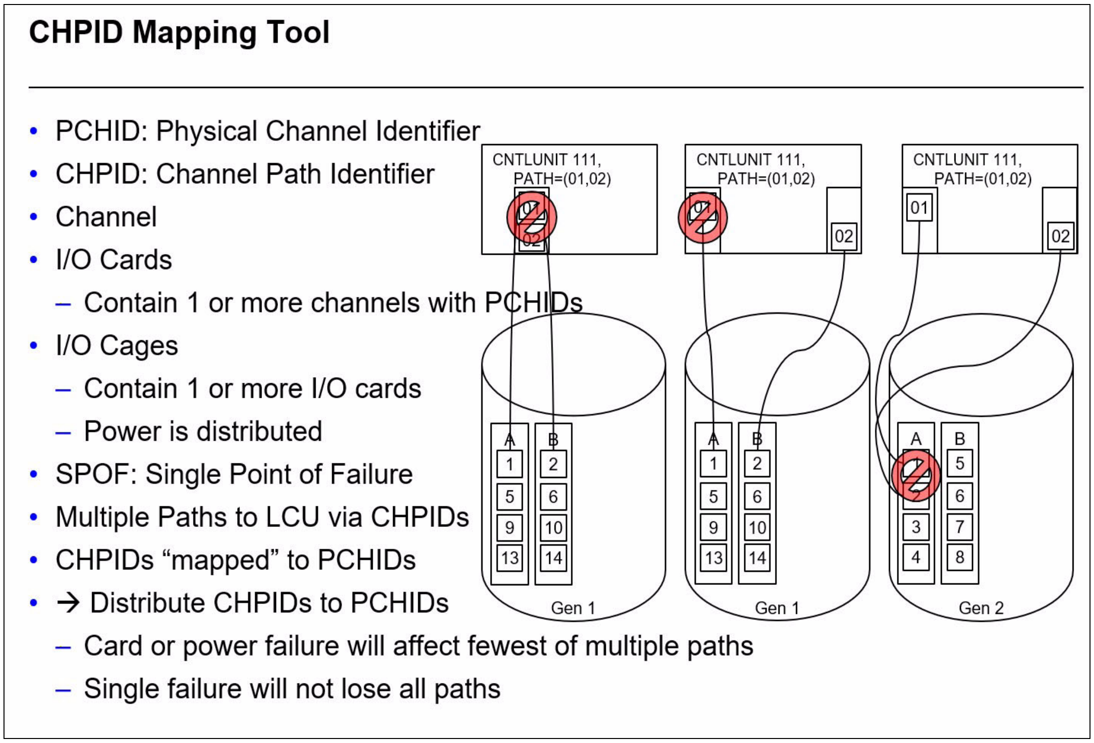

CHPID Mapping Tool

- PCHID: Physical Channel Identifier
- CHPID: Channel Path Identifier
- Channel
- I/O Cards
  - Contain 1 or more channels with PCHIDs
- I/O Cages
  - Contain 1 or more I/O cards
  - Power is distributed
- SPOF: Single Point of Failure
- Multiple Paths to LCU via CHPIDs
- CHPIDs “mapped” to PCHIDs
- → Distribute CHPIDs to PCHIDs
  - Card or power failure will affect fewest of multiple paths
  - Single failure will not lose all paths

Figure 5-9 CHPID Mapping Tool

It can become difficult for a data center planner to ensure that all the channels that access a single CU are spread out to different channel cards when a single mainframe can have many channels and CUs, mostly with MPIO. To address that complexity, IBM supplies the CMT, which takes a CFReport file and the IOCP source input file as input. A CFReport file comes with the mainframe when it is delivered and after any Miscellaneous Equipment Specification (MES) operation, which is the procedure that is used on a mainframe to install or remove hardware elements such as channel cards. The IOCP source file might be an initial attempt to define the I/O connectivity that is only concerned with logical definitions and not the location of the PCHIDs. The CMT examines both the available channels and on which cards they are located, and the MPIO paths that are defined. It generates a new CFReport and IOCP source file. HCD uses the new IOCP source file, with PCHIDs assigned and distributed to avoid placing multiple CHPIDs, in MPIO to the same CU on the same physical card.

The data center hardware planner is responsible for ensuring that the paths in the FICON Directors and the CUs that are used in MPIO are not all on the same card and that the different cards are not all in the same hardware element.

FICON Directors and CUs also have multiple adapters on the same card. They also have multiple and separate hardware control elements to provide redundancy in case of hardware or power failures. Mainframe environments should always have multiple SAN or FICON switches, and you should ensure that the MPIO paths are distributed and not all on the same switch.

Important: CMT does not examine the location of the ports on CUs or FICON Directors.

The data center hardware planner is responsible for ensuring that the paths in the FICON Directors and the CUs that are used in MPIO are not all on the same card and the different cards are not all in the same hardware element.

5.2.12 Modernizing mainframe I/O

Mainframes are compatible with earlier versions, where a new machine can run older programs with no changes. This mainframe feature has kept it as critical feature in so many data centers today. Over time, the mainframe's I/O structures were improved to satisfy user requirements, increase speed and capacity, and address changes due to rapidly increased hardware capabilities. All these changes were made with compatibility with earlier versions included. If the I/O methods could be rewritten with new specifications to take advantage of improved technology, it is likely the resulting methods would be much simpler, but all the programs that perform I/O would have to be rewritten and retested. Thus, some of the terms and constructs might appear to be more complex or difficult than you might think that they should be.

Here we present a brief history of the evolution of mainframe I/O to explain the changes and what was needed to accommodate compatibility with earlier versions:

- CKD was improved and extended (ECKD), which required changes at both the channel and the CU to support new I/O commands. However, old programs did not need to change when used with ECKD commands.
- The scope of the I/O constructs greatly increased:
  - CUs and devices were merged into one box with virtual CU and device images. The different logical CU images each had a CU address (CUADDR).
  - Multiple OS images can run concurrently on the same physical mainframe in separate LPARs. The different LPARs each have a Multiple Image Facility Identifier (MIF ID).

  - Multiple CSSs and multiple subchannel sets expanded the number and location of channels and devices while still providing compatibility with earlier versions.
- The device address became a subchannel that did not explicitly describe the path to the device, along with SIO becoming Start Subchannel (SSCH). This change removed an I/O queuing bottleneck by breaking the channel into separate virtual mini-channels, one per device, each with its own I/O queue. Thus, a subchannel is a view of a device even though it has “channel” in its name.
- A channel became a CHPID that did not explicitly describe the physical location of the hardware's I/O adapter. CHPIDs allow for easier migrations of a program from machine to machine and other advantages.

> **Note:** We use the term *channel* unless there is an important technical reason for the distinction.

- A PCHID enabled more than 256 physical channels on a mainframe and also enabled mainframes to run more LPARs on the same machine. A program's view of a channel is a CHPID, but CHPID 3 might be mapped to PCHID 127.
- Pipelining was introduced, where some channel commands could be transmitted without waiting for a response to the prior command. Many channel programs were reading or writing multiple blocks of data consecutively in one channel program, so why require waiting for a response after each one? This change increased the speed of I/O, especially over longer distances.
- Interconnection methods such as PCI Express (PCIe) and InfiniBand increased the speed and reach of mainframe I/O.

Despite all these changes, the essence of CKD's I/O flow remained and old channel programs, and the programs that invoked them, could run unmodified.

5.2.13 ESCON

In 1990, IBM introduced the first SAN with Enterprise I/O System Connection (ESCON). In addition to using fiber optics rather than copper cables, it also introduced the *dynamic* I/O switch, also known as the ESCON Director. It was called dynamic because the physical connections of channels (or HBAs) and CU adapters were static (unchanging), but the logical connections were made in the switch between channels and CUs during an I/O operation and then released so that they could be used to connect to different channels or CUs in the next I/O operation. The mainframe and its source-based routing required the addition of a *destination link address* for each CU and the *source link address* from each channel, as shown in Figure 5-10 on page 115, to tell the switch which ports are involved when making a new dynamic connection. The destination and source link addresses (each 1 byte) were added to the I/O operation.

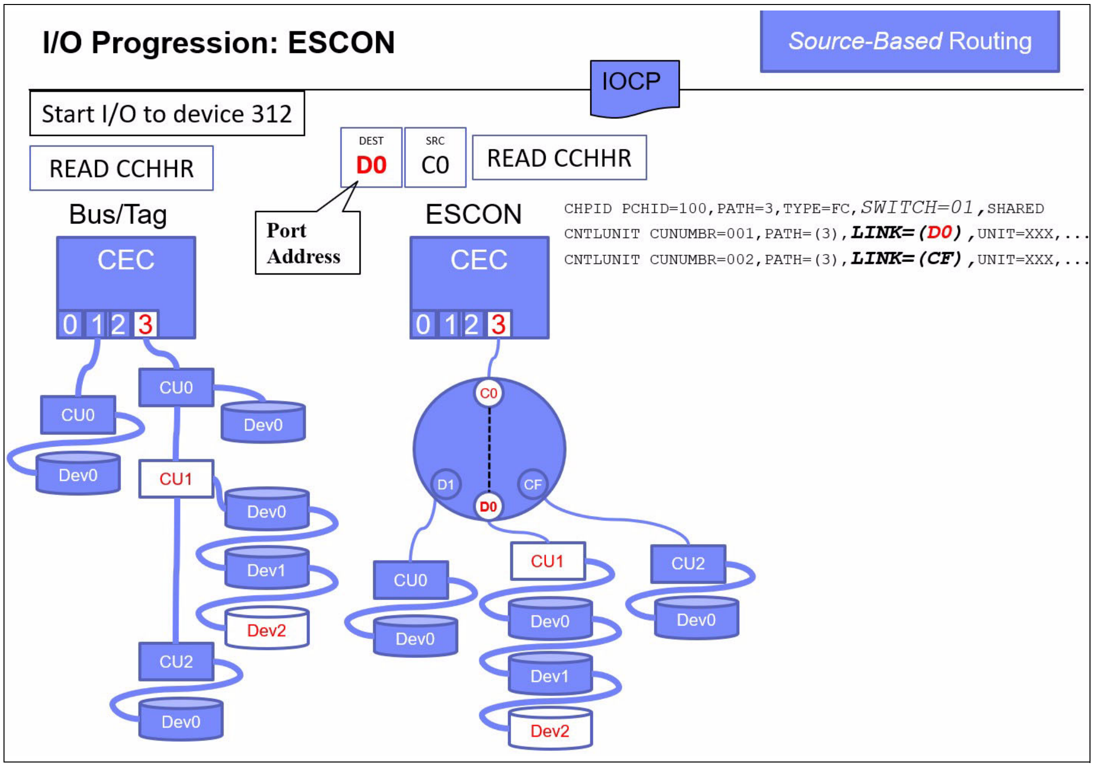

```
CHPID PCHID=100,PATH=3,TYPE=FC,SWITCH=01,SHARED
CNTLUNIT CUNUMBR=001,PATH=(3),LINK=(D0),UNIT=XXX,...
CNTLUNIT CUNUMBR=002,PATH=(3),LINK=(CF),UNIT=XXX,...
```

Figure 5-10 ESCON I/O and the SAN link address

The IOCP definition statements were extended to allow the user to specify these link addresses, as shown in Figure 5-11.

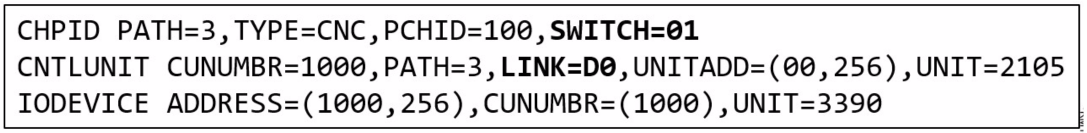

```
CHPID PATH=3,TYPE=CNC,PCHID=100,SWITCH=01
CNTLUNIT CUNUMBR=1000,PATH=3,LINK=D0,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(1000,256),CUNUMBR=(1000),UNIT=3390
```

Figure 5-11 IOCP definition statements: ESCON example

ESCON also introduced Link-Level and Device-Level functions. The Link-Level functions allowed end units (such as channels, ESCON Director ports, and CU adapters) to initialize, configure, validate, and monitor hardware ports. Only after the Link-Level functions had verified the hardware could any channel or CU send Device-Level functions, such as channel program commands and responses (status and data).

5.2.14 FCP

The FCP enabled migration from a SCSI bus I/O architecture to a SAN by using many of the inventions of ESCON. FCP allowed for direct connect arbitrated loop (a topology analogous to the SCSI bus), but most importantly, FCP enabled multiple SAN switches to be connected to each other to create a fabric of I/O connectivity. Because of these changes, the link address increased from 1 to 3 bytes:

- 1 byte (8 bits) for the domain ID (the unique switch identifier in the fabric)
- 1 byte (8 bits) for the link address (inherited from ESCON)
- 1 byte (8 bits) for the arbitrated loop port address (AL_PA)

In FCP, the SAN users do not specify the link address for I/O operations directly but often gain access to devices by using unique worldwide names (WWNs). The switch fabric used link-level functions to build and maintain a database containing the relationships of WWNs to domain IDs, link addresses, and AL_PAs. We call this *fabric-based routing* because the switch fabric determines the communication path to the device. The path is not specified by the initiator (an HBA or channel) or the target (typically a device adapter), which are primarily running SCSI FBA commands.

With FCP, a new switch entering a fabric can dynamically obtain a random domain ID from a principal switch or use one that is pre-specified. Internally, fabric-based routing must know the domain IDs, but not the external initiators and targets.

An FCP fabric might have different routes from an initiator to a target, with the best one chosen by the internal Fabric Shortest Path First (FSPF) algorithm. If one path should fail, other paths might exist, and I/O operations can be rerouted by using FSPF and continue, as shown in Figure 5-12.

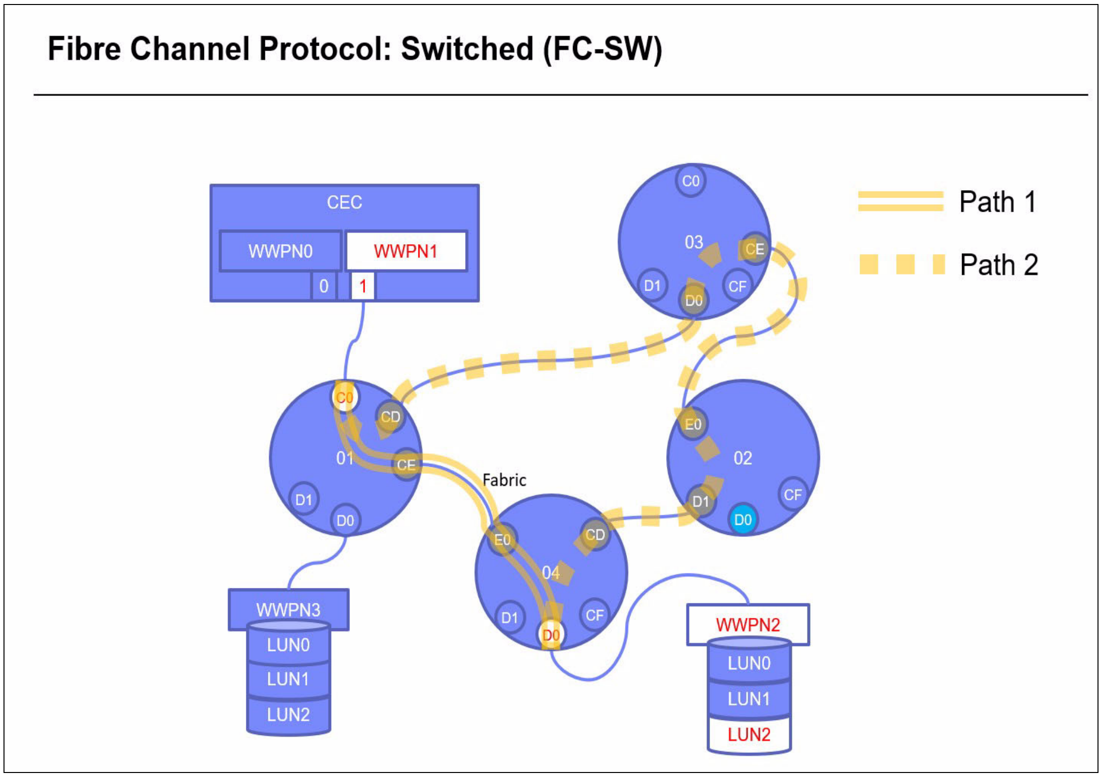

Figure 5-12 FCP with multiple paths through a fabric

SCSI had few bus addresses, usually 8, so commands were broadcast to all units on the bus. Broadcasts did not work well with large, scalable fabrics. To group sets of initiators and targets into small broadcast domains and isolate host-to-device sets from other hosts, FCP introduced *zones*. Consider a zone to be an access control list (ACL) or an allowlist, where only the members of the list are allowed to communicate. The zones are established in the switch fabric by SAN administrators. *Zone sets* or zoning configurations contain the zones that are used in the fabric. Aliases are used to simplify the zones by giving WWNs more human friendly identifiers. There can be hundreds of zones in a zone set, each one allowing its members to keep the simple addressing of a SCSI bus.

SAN switches initially had a limit of 256 ports in a single physical switch because of a 1-byte link address. Because an FCP user does not specify the link address, it was possible to modify it and use some high-order bits of the otherwise unused AL_PA field. This change was called 10-bit addressing, as shown in Figure 5-13.

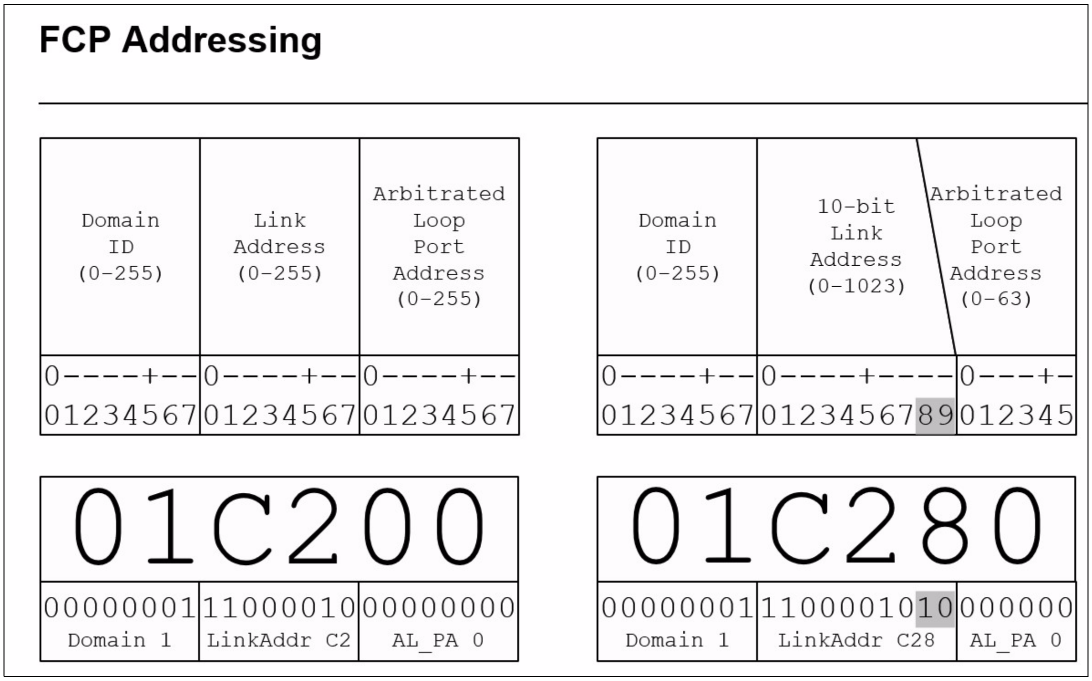

Figure 5-13 FCP link address and 10-bit addressing

SCSI has only eight addressable units on the same bus, meaning only 3 bits of the AL_PA field are required (the five high-order bits are unused). With 10-bit addressing, a single switch can have up to 1024 ports. The N_Port ID Virtualization (NPIV) technology, which is used sometime by FCP initiators and targets, relies on the AL_PA bits too.

5.2.15 FICON

Just as ESCON addressed constraints of SCSI and engendered FCP, FCP solved some constraints in ESCON, such as duplex communication, which improves both throughput and latency. Compared to the days of ESCON, a typical 4:1 reduction in the number of paths to each CU was possible with FICON and still provides adequate I/O bandwidth. Even with four times fewer channels, the FICON native configuration is capable of more I/O operation concurrency than an ESCON configuration.

FICON was the mainframe version of FCP. The FC specification for FICON is FC-SB. As the name indicates, only the link address field is used, which is 1 byte, as it was in ESCON.

Several restrictions were placed on the FCP protocol so that mainframes could provide compatibility with older ECKD I/O programs in an FCP-style SAN. ESCON used only 1 byte for a link address, so FICON kept the limit of 256 link addresses per switch domain, but FCP allowed some AL_PA bits to be used. FICON cannot use 10-bit addressing.

FICON uses the FC-layered architecture. FC logical layers are designed to allow multiple upper layer protocols (ULPs) to coexist in the same fabric. The combination of layers FC-0, FC-1, FC-2, and FC-3 define the functions of an FC port, and they are common for both FCP and FICON and all other ULPs. The FC-4 layer is where the unique mapping of ULPs such as SCSI, FICON, and other ULPs occur. Although the FCP protocol defines the mapping of SCSI to FC-0 – FC-3, the FC-SB protocol defines the mapping of Single-Byte Command Code Sets (SBCCSs) to FC-0 – FC-3. The same FC switch and port can support multiple ULPs, including FCP and FC-SB, so FICON devices can participate in an open-systems FC SAN. FICON and FC can co-exist on the same physical fabric. However, FCP and FICON devices cannot communicate with each other, as shown in Figure 5-14.

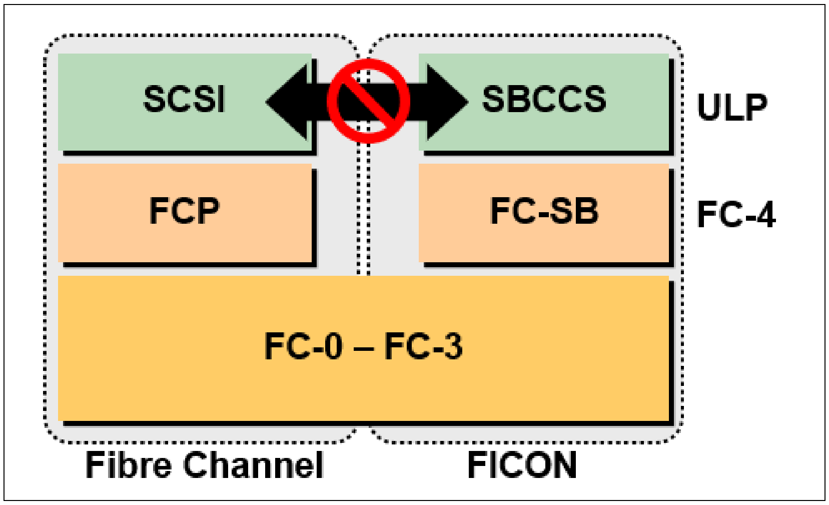

Figure 5-14 FC and FICON layered architecture

In FICON, each channel and CU is mapped to an N_Port. When the CU N_Port is attached to a fabric, the CU and its devices can be accessible to all channels that are attached that fabric. A CU can communicate simultaneously with more than one channel, much like a channel can communicate with more than one CU.

IBM originally announced the commercial availability of FICON channels for the S/390 9672 G5 processor in May of 1998, 2 years later than the relevant specification. Since then, FICON evolved from 1-Gbit to the 16-Gbit FICON Express16S channels, announced in January 2015 with the IBM z13. Each generation of FICON channels offered increased performance and capacity. Additionally, new topologies were allowed, like cascading and multihop cascading, and new routing mechanisms provided better load-balancing on Inter-Switch Links (ISLs). Major improvements also occurred with IBM High-Performance FICON for IBM System z® (zHPF) and transport mode. We revisit all these changes with a historical perspective.

FICON to ESCON bridge: FC-SB (1996)

Initially, FC-SB allowed mainframes to use FICON channels, which increased the link speed and enabled duplex communication, both of which improved I/O throughput. It used FICON Converter or bridge channels to ESCON switches and devices, enabling a migration from existing devices. The migration from ESCON to FICON took several years and was a significant topic back in the day.

Native FICON Channel to CU: FC-SB-2 (2000)

The next generation of FICON, FC-SB-2, introduced FICON ULPs, which meant that the mainframe could leverage FICON more fully. FC-SB-2 introduced:

- Attachments of FICON channels to native FICON devices such as DASD and tape, which increased base I/O rates from 17 MBps to 100 MBps on a 1-Gbps link.
- FCTC through direct connect or though FICON Directors. Unlike ESCON CTC, a single FCTC channel could both read and write, so only one was required per processor. To avoid a single point of failure (SPOF), at least 2 FCTCs were used.
- FICON Directors instead of bridging to ESCON Directors.
- An Information Unit (IU) was an ECKD command inside an FCP-style frame, thus allowing FICON channels to continue to use existing ECKD channel programs. IU pacing allowed the channel to send multiple channel commands to a CU without requiring a response for each command, thus supporting pipelining. A command response (CMR) could be indicated in any Channel Command Word (CCW), and the CU would have to immediately respond. The channel could speed up or slow down the IU rate (that is, the pace itself) based on the response time from CMRs during the channel program.
- Command Mode allows a FICON channel to perform link-level functions that FCP added to the functions that were created by ESCON.

Figure 5-15 shows the different scenarios for FICON Converter and FICON Native mode.

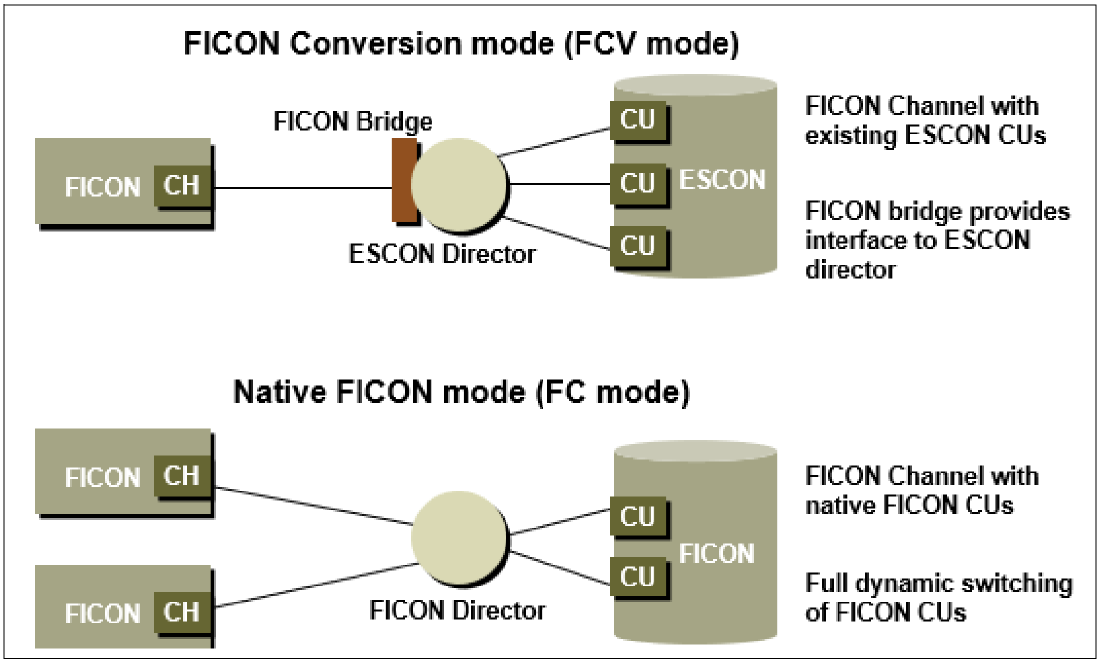
FICON Conversion mode (FCV mode)
FICON Bridge
FICON CH
ESCON Director
CU
CU
CU
ESCON
FICON Channel with existing ESCON CUs
FICON bridge provides interface to ESCON director
Native FICON mode (FC mode)
FICON CH
FICON CH
FICON Director
CU
CU
FICON
FICON Channel with native FICON CUs
Full dynamic switching of FICON CUs

Figure 5-15 FICON Converter and FICON Native mode

Cascaded FICON Directors: FC-SB-3 (2003)

The next generation, FC-SB-3, introduced FICON Cascading, which allowed two FICON Directors to connect to each other with ISLs, which was a significant improvement for several reasons:

- Persistent IU pacing improved performance over longer distances. This mechanism is an end to end flow control mechanism, as shown in Figure 5-16, which was introduced on top of the hop-by-hop buffer-to-buffer (B2B) credits mechanism that is used by the FC protocol.

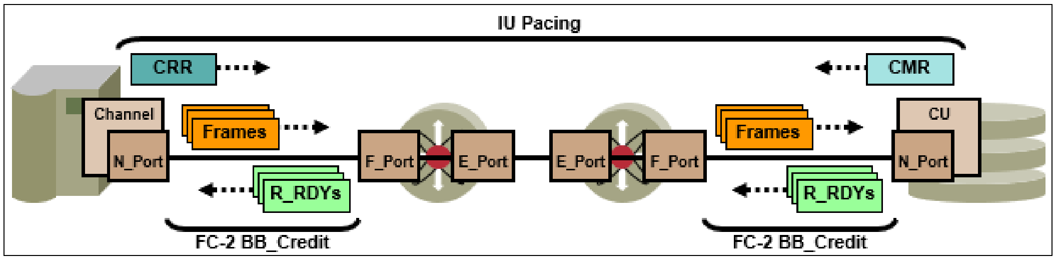

Figure 5-16 IU pacing flow control mechanism

- Cascading expands the number of CUs that a channel can access, and the number of channels that a CU adapter can access.
- Cascading increases the distance between a channel and a CU (or CTC for FCTC, or CU to CU for background data replication solutions like Peer-to-Peer Remote Copy (PPRC) or Symmetrix Remote Data Facility (SRDF))
- Most importantly, FICON Cascading allowed multiple data centers in a campus environment to be connected as one, which was useful for the first level of disaster recovery (DR) because the data centers could be on different power grids, different cooling systems, and so on. Also, the data could be replicated from one data center to another, either by dual write, eXtended Remote Copy (XRC), or PPRC or SRDF.
- FICON Cascading required 2-byte link addresses, where the first byte specified the switch's Domain ID, and the second byte specified the link address on the specified switch. The name of the FICON specification, FC-SB (where SB stands for single byte) remained, even though a 2-byte link address made it no longer technically correct.

Figure 5-17 on page 121 shows the FICON Link Addressing schema.

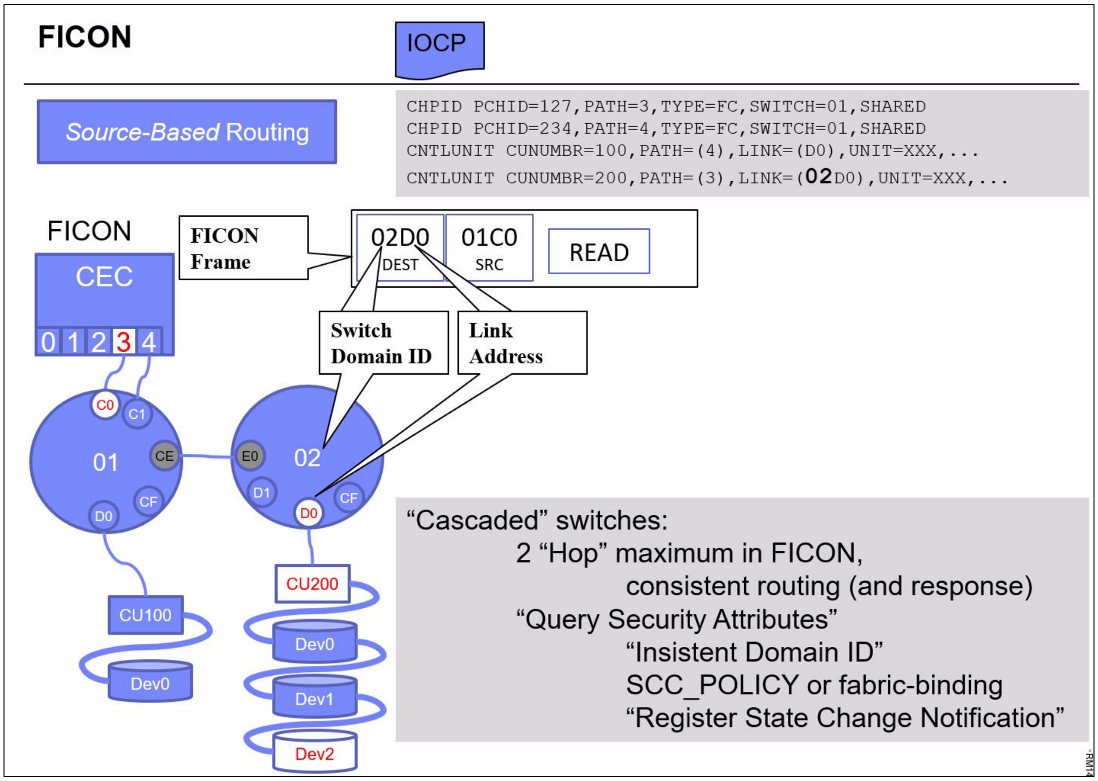

FICON

IOCP

```
CHPID PCHID=127,PATH=3,TYPE=FC,SWITCH=01,SHARED
CHPID PCHID=234,PATH=4,TYPE=FC,SWITCH=01,SHARED
CNTLUNIT CUNUMBR=100,PATH=(4),LINK=(D0),UNIT=XXX,...
CNTLUNIT CUNUMBR=200,PATH=(3),LINK=(02D0),UNIT=XXX,...
```

“Cascaded” switches:
2 “Hop” maximum in FICON,
consistent routing (and response)
“Query Security Attributes”
“Insistent Domain ID”
SCC_POLICY or fabric-binding
“Register State Change Notification”

Figure 5-17 FICON I/O and the extended link address

The channel specification also changed:

- If the channel accesses only locally attached CUs (that is, the CU adapter is attached to the same FICON Director as the channel), the destination domain ID is not required but may be specified. The channel can access only locally attached CUs.
- If the channel accesses any remotely attached CUs (that is, the CU adapter is attached to the cascaded FICON Director), the destination domain ID is required. The channel can now access locally attached or remotely attached CUs.
- All FICON channels learn the link address where they are attached to a switch and the switch's Domain ID dynamically at link-initialization time. The domain ID for the channel can optionally be specified in the **SWITCH** parameter of the CHPID statement in the IOCP, but the channel always uses the domain ID and link address it learns from the switch. The IOCP specification is only a comment, and because it can be incorrect, some customers choose not to specify it.
- The IOCP changed to allow a link address that is 2 bytes long (4 hex characters) to contain the domain ID as a prefix to the destination link address. Moreover, changes were made to the IOCP to support multiple switches, as shown in Figure 5-18.

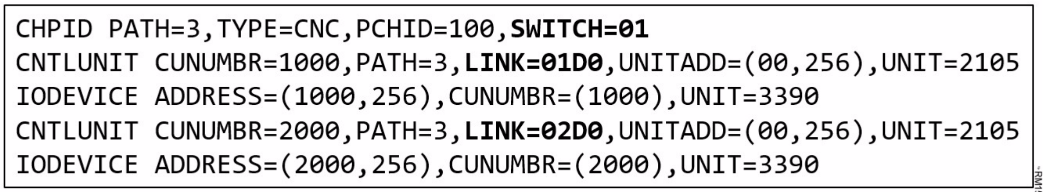

```
CHPID PATH=3,TYPE=CNC,PCHID=100,SWITCH=01
CNTLUNIT CUNUMBR=1000,PATH=3,LINK=01D0,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(1000,256),CUNUMBR=(1000),UNIT=3390
CNTLUNIT CUNUMBR=2000,PATH=3,LINK=02D0,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(2000,256),CUNUMBR=(2000),UNIT=3390
```

Figure 5-18 IOCP definition statements for FICON

For a mainframe to use cascaded FICON Directors, several new restrictions to FCP were added:

- Insistent domain ID (IDID): FICON, by using source-based routing, specifies a static switch domain ID in the IOCP. Thus, FICON requires that the domain IDs are specified and not changeable, which is called IDID in the standards. Cisco prefers to call IDID a static domain ID and the following text might use both terms interchangeably. Also, persistent FCID is not needed when the domain ID is static.
- Only 1 or 2 switches in the path: FICON is using source-based routing, IU pacing, and CMRs to manage and monitor the response time down each path to a device to choose the best path or avoid poorly performing paths. For this reason, FICON cannot have FSPF reroute I/O between different paths in a fabric, which limits a FICON fabric to a maximum of two switches in the path.
- Fabric-binding: When a FICON channel initializes at the link level, it can query the directly attached switch to verify that it has IDID, but the channel cannot query the second switch in an I/O path, which makes it difficult to ensure that there are a maximum of two switches in the path, and verifying that the second switch (even if it has the correct domain ID with IDID) is the correct switch and not an impostor. To address that issue, an allowlist with the WWNs of the switches is set and must match in each switch that comes into the FICON fabric. Cisco calls this *fabric-binding*. and it represents the digital implementation of the old saying “Trusting is good, not trusting is better”.
- Registered State Change Notification (RSCN): Switches can enter or leave a fabric for many reasons: new installations, hardware failures, and system reconfigurations. FICON requires that all switches in a FICON fabric send a notification to all FICON channels about switches coming in or leaving the fabric. Without such a notification, a FICON channel would be unaware of a new path that is made available to a CU and its devices, or the first path that is made available to a CU and new devices. This process is called RSCN.

A link-level function, Query Security Attributes (QSA), is used at link initialization by all FICON channels. If the response to QSA shows that the switch has IDID and fabric-binding, the channel can use RSCN to be informed of fabric changes, complete initialization, and then be available for I/O operations.

If the channel is attached to a switch that does not have one or more of the required restrictions, it will have the `FICON INCOMPLETE` status, as shown in Figure 5-19.

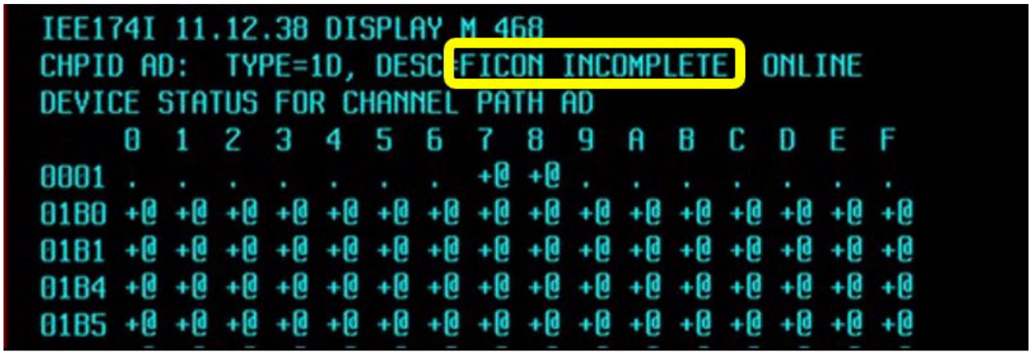

```
IEE174I 11.12.38 DISPLAY M 468
CHPID AD:   TYPE=1D, DESC:FICON INCOMPLETE  ONLINE
DEVICE STATUS FOR CHANNEL PATH AD
      0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
0001  .  .  .  .  .  .  .  +@ +@ .  .  .  .  .  .  .
01B0 +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@
01B1 +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@
01B4 +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@
01B5 +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@ +@
```

Figure 5-19 FICON INCOMPLETE channel status

After you set IDID and fabric-binding at the switch, you might have to toggle the channel at the HMC or SE to force it to go through the link-level initialization again.

Recall that a channel might not have any 2-byte link addresses specified because it accesses only locally attached CUs. In that case, the channel will not issue QSA or RSCN because it does I/O only to local ports. Thus, for some versions of mainframe hardware and OSs, a switch that does not have IDID or fabric-binding can be used by those single-byte channels. However, that is definitely *not* recommended for many reasons:

- Another channel might be attached to the same switch but use 2-byte link addresses and would be `FICON INCOMPLETE`, making debugging confusing.
- The channel learns of the switch's domain ID at initialization time. If the switch does not have IDID, it could potentially have the same domain ID as another switch, making performance analysis confusing.
- Future releases of hardware or software might change the requirements even on locally attached, single-byte link addresses.

To summarize, a FICON Director:

- Must respond to QSA by indicating that IDID is enabled, and that fabric-binding was used to create an allowlist of FICON Directors in the fabric.
- Must provide RSCN for fabric changes for ports that register for them.

zHPF: FC-SB-4 (2009)

The distances between a FICON channel and a CU are getting longer because of channel extension techniques such as Fibre Channel over IP (FCIP) or Dense Wavelength-Division Multiplexing (DWDM) systems. These extensions are useful and commonly used for data replication solutions and failover techniques, such as IBM HyperSwap® or AutoSwap.

Among the data replication techniques, some use only FCP, and others use FICON. For example, PPRC or SRDF use only FCP, while XRC uses FICON. IBM Geographically Distributed Parallel Sysplex® (IBM GDPS®) uses many of these data replication techniques.

ECKD, even with pipelining and persistent IU pacing, is still highly controversial. ECKD response time suffers with increased distances because the next channel command cannot be sent until the response from the CU is received, which is after one or more round trips.

To counteract this effect, FC-SB-4 introduced Transport Mode, where the previous approach was known as Command Mode. In Transport Mode, a channel sends the entire channel program to the CU, which is possible because the CUs are increasingly sophisticated and many channel programs follow a pattern, so now every possible channel program is addressed. In simple terms, Transport Mode is set up in the initial channel program to the CU, the channel program is sent (and possibly data, for a write operation), and the response is received (and possibly data, for a read operation). There is far less interaction than with ECKD and the response time improvement is significant.

This enhancement is called zHPF. zHPF is not enabled by default so that compatibility with earlier systems is possible. During link initialization, both the channel and the CU indicate whether they support zHPF or not. As you might expect with such a significant change to how mainframe I/O is performed, zHPF requires mainframe, CU hardware, OS, and CU microcode updates.

Using zHPF with the FICON channel, the z/OS OS, and the CU reduces the impact of the FICON channel. This reduction is achieved by protocol optimization and reducing the number of IUs that are processed, which results in more efficient usage of the fiber link. Currently, all mainframes and almost all disk arrays support zHPF.

System z Dynamic Auto Discovery: FC-SB-5 (2012)

With FC-SB-5, FICON is using FCP link-level functions to query the state of the fabric and can dynamically learn about CU adapters that are added or removed from the fabric. If a new CU is discovered, the mainframe's MVS OS can inform the operators, who can initiate a process to add the CU and device to the HSA and the MVSCP, like an HCD activate operation.

This enhancement is called System z Dynamic Auto Discovery (zDAC). zDAC is not enabled by default. zDAC brings plug-and-play capabilities to the mainframe and can speed up and simplify the process of adding storage devices.

FC-SB-5 also introduced bidirectional transport mode to make zHPF even faster.

FICON Dynamic Routing (2015)

FICON Dynamic Routing (FIDR) changed how FICON frames are sent across ISLs, which had an impact when multiple ISLs were grouped in a port channel.

In FCP and FICON, a read or write operation with much data is broken up because the largest payload in an FCP or FICON frame is 2112 bytes. The frames are chained together in a sequence so that they arrive at the program or device in the correct order. A channel program, either FICON or SCSI over FCP, must be sent in order (we have described ECKD, but the SCSI channel program for a single FBA block must also be in the correct order). Multiple sequences, such as a complete channel program, compose an exchange. With many channel programs from the same channel or a CU adapter that is accessed by many channels, the frames have identifiers of their sequences and exchanges, and can be interleaved.

The relationship of exchanges, sequences, and frames is shown in Figure 5-20.

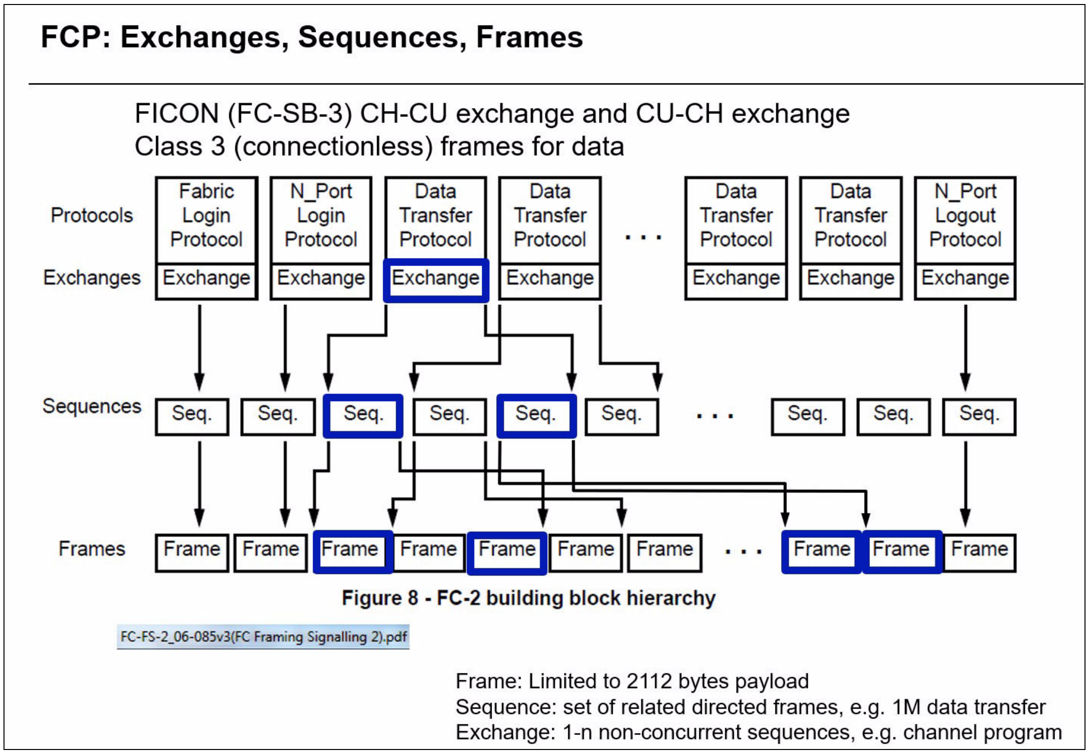

FCP: Exchanges, Sequences, Frames

FICON (FC-SB-3) CH-CU exchange and CU-CH exchange
Class 3 (connectionless) frames for data

Figure 8 - FC-2 building block hierarchy

FC-FS-2_06-085v3(FC Framing Signalling 2).pdf

Frame: Limited to 2112 bytes payload
Sequence: set of related directed frames, e.g. 1M data transfer
Exchange: 1-n non-concurrent sequences, e.g. channel program

Figure 5-20 Exchanges, sequences, and frames

Before FIDR, an ISL was chosen by using the source ID (SID)/destination ID (DID) for the I/O request from the channel while considering the current load on the ISLs. Typically, the ISL with the fewest paths that are allocated to it would get the new request. The disadvantage to SID/DID is that the choice is static: After the ISL is chosen, all future I/O between the channel and the CU is sent over the same ISL. There are many cases where the utilization of the ISLs is not evenly distributed. Another disadvantage is that separate ISL port channels are necessary when using multiple virtual storage area networks (VSANs) a VSAN's I/O might dominate an ISL and delay other VSAN's I/O.

With FIDR, the ISL is chosen by using SID, DID, and originator exchange ID (OXID). The advantage is that the ISLs are far more evenly used, especially over time. Also, the traffic from multiple VSANs, even FICON and FCP VSANs on the same physical switch, can use the same ISL port channel because the load is redistributed for each I/O.

However, FIDR does have some disadvantages:

- Diluted errors: Some components of an ISL can cause intermittent errors, such as when an optical component approaches its end of life. If the ISL paths from channels to devices are static, the error reports show up on the same path or paths. The single ISL can be identified when the path errors are analyzed. With FIDR, the errors show up as I/O failures to different devices from different channels, so identifying the offending ISL is less obvious.
- Slow drain: Some devices are called *slow-drain* devices because they “drain” frames away from the network at an insufficient rate. Other devices, such as tape, tend to have bursts of large or long-running I/O operations, which can cause a slow-drain condition because the I/Os are queued and take longer than normal before the resources (such as ISL buffers) are free for use by other I/O operations. Slow-drain conditions on an ISL slow down all other I/O that uses the same ISL. If the ISL paths are static, the OSs can detect that the path is slower than normal, and they will avoid using the path until it starts to respond more quickly. They will select other paths (if available). With FIDR, a slow-drain condition can span multiple ISLs at the same time, possibly all the ISLs in a port channel. They will affect different I/O operations and different paths, making it more difficult for the OS to detect and avoid the problematic links. Also, IU pacing and path selection are far less effective.

FICON Multi-hop (2017)

The maximum of two switches in a FICON path is restrictive. Mainframe environments require reliability, and there are some fabric configurations that provide more reliability than can be created when only two switches are in any path. Multiple 2-switch routes from a single channel do not break the 2-switch rule, as shown in Figure 5-21.

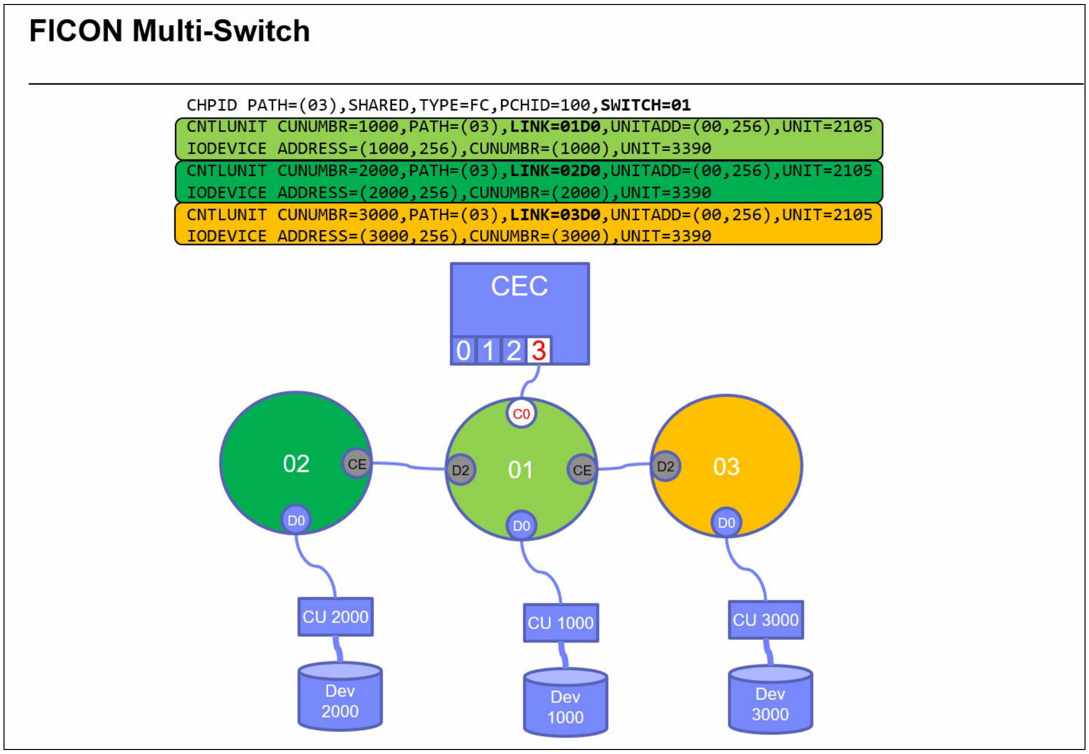

FICON Multi-Switch

```
CHPID PATH=(03),SHARED,TYPE=FC,PCHID=100,SWITCH=01
CNTLUNIT CUNUMBR=1000,PATH=(03),LINK=01D0,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(1000,256),CUNUMBR=(1000),UNIT=3390
CNTLUNIT CUNUMBR=2000,PATH=(03),LINK=02D0,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(2000,256),CUNUMBR=(2000),UNIT=3390
CNTLUNIT CUNUMBR=3000,PATH=(03),LINK=03D0,UNITADD=(00,256),UNIT=2105
IODEVICE ADDRESS=(3000,256),CUNUMBR=(3000),UNIT=3390
```

Figure 5-21 FICON multiple 2-switch paths from Switch 01

Some FICON configurations have a hop of no consequence, typically when using smaller switches for FCIP. In that situation, all the switches in the path must be in the fabric-binding database, and paths cannot be dynamically rerouted to other paths in a fabric, as shown in Figure 5-22 on page 127.

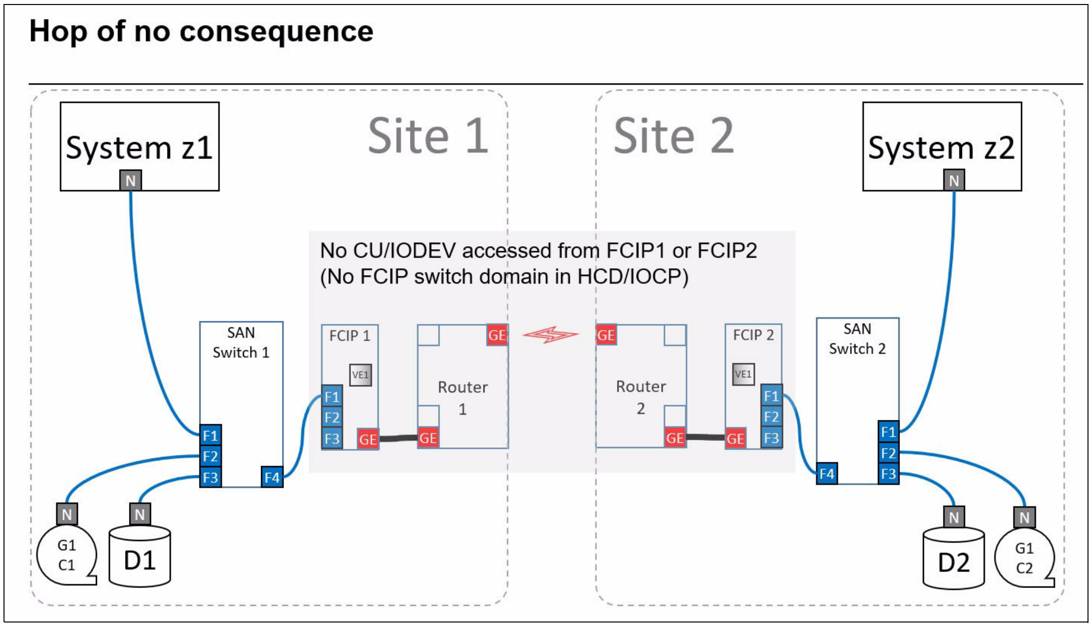

Figure 5-22 Hop of no consequence

FICON Multi-hop removed the 2-switch restriction for a set of specific fabric configurations that many IBM customers requested because they are useful for multi-site connectivity. These configurations are similar to the hop of no consequence. These configurations are shown in Figure 5-23.

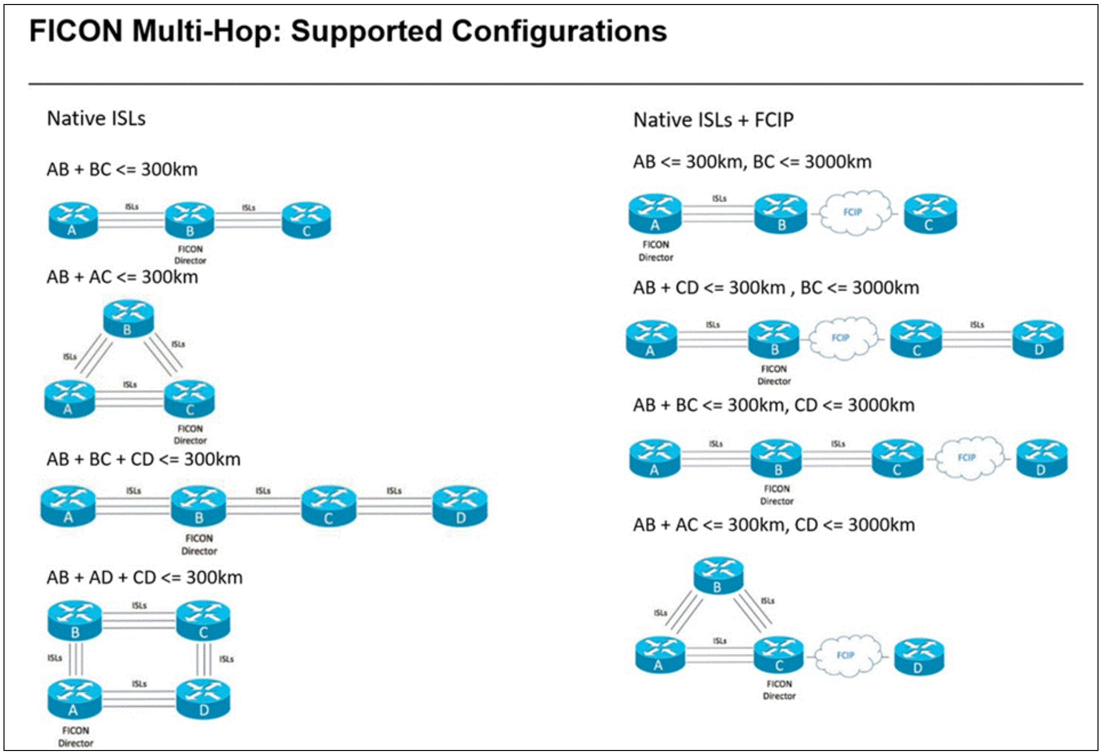
FICON Multi-Hop: Supported Configurations
Native ISLs
AB + BC <= 300km
AB + AC <= 300km
AB + BC + CD <= 300km
AB + AD + CD <= 300km
Native ISLs + FCIP
AB <= 300km, BC <= 3000km
AB + CD <= 300km , BC <= 3000km
AB + BC <= 300km, CD <= 3000km
AB + AC <= 300km, CD <= 3000km
ISLs
FCIP
FICON Director
A
B
C
D

Figure 5-23 FICON Multi-hop supported configurations

Extended Distance zHPF: FC-SB-6 (2016)

FC-SB-6 introduced extended distance transport mode, which improved zHPF by removing some of the restrictions that zHPF still required. This change significantly improved FICON response time over great distances.

FICON and the FC-SB standard have maintained compatibility with earlier versions. The characteristics of FICON are integrity, security, flexibility, availability, serviceability, transactions, efficiency, and reliability, which can be remembered by using the mnemonic IS FASTER.

5.2.16 FICON requirements

Table 5-1 summarizes the functions in FICON that can cause problems for IBM Z I/O programs and architectures, and what items must be modified or restricted to make FICON operational.

Table 5-1 IBM Z summary of functions

FCP feature	Problem for mainframe I/O	FICON modifications
Paths through the fabric (up to 239 switches) can be dynamically changed to adjust to link or switch failures.	inconsistent and variable I/O response times based on the path through the fabric.	Limit of two switches between end points (extended in some cases with multi-hop).
Domain IDs can be dynamically determined.	The host-based specification of the destination domain ID requires a consistent and static ID.	Static Domain ID.
Link addresses are dynamic and can change as ports or modules are added to a switch, or if an I/O entity is moved.	The host-based specification of the destination link address requires a consistent and static address.	Assign a port number.
If a switch has more than 256 physical ports, link addresses can use high-order bits from the AL_PA.	The host-based specification of the destination link address has a maximum of 8 bits.	The FCID last byte is 0.
Switch and port failures cause rerouted traffic patterns.	The channel does not know about a device path failure until the time of the I/O operation.	Use a Registered Link Incident Report (RLIR) and RSCN.
The first switch with the static domain ID will win.	The domain ID alone is not sufficient to guarantee that the second switch is correct (and not a malicious copy).	Use an allowlist of switches (that is, a fabric-binding database).
Any switch can connect to any other switch in the fabric.	The channel cannot verify that the second switch in the path has the same required characteristics of the attached switch (that is, a static domain ID and allowlist).	All switches must verify that any attached switch has the same restrictions (static domain ID, fabric-binding database, and Cisco Fabric Services (CFS) distribution).

FCP feature	Problem for mainframe I/O	FICON modifications
A route change in the network can introduce a path that might be faster or less congested than the old route. When a link change occurs in a port channel, the frames for the same exchange or the same flow can switch from one path to another faster path.	I/O on the new path might arrive before previously sent I/O commands to the same device.	Use In Order Delivery (IOD) (only when LIOD is not configurable). For example, with FCIP links, when the simpler LIOD feature is enabled and a port channel link change occurs, the frames crossing the port channel are treated as follows: ► Frames that use the old path are delivered before new frames are accepted. ► The new frames are delivered through the new path after the switch latency drop period has elapsed and all old frames are flushed. Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped.
Load-balancing attributes indicate the use of the source-destination ID (src-dst-id) or the OXID (src-dst-ox-id, which is the default) for load-balancing path selection.	An exchange (OXID) within an I/O might use a different route and arrive before a previously sent exchange.	The default load-balancing scheme is SID/DID. (SID/DID/OXID is used if FIDR is supported on all nodes within that specific FICON VSAN.)

5.2.17 IBM Control Unit Port for FICON

A mainframe user might want to access a FICON Director through an I/O function to control a FICON Director (such as enable or disable a port), or access information (such as errors or port utilization) from the director.

For a mainframe to access a device, it must have a definition for a channel that connects to a CU, which accesses the device. In a FICON environment, there are channels that are attached to a FICON Director that can be used for I/O. An internal (not physical) port is used as the destination link address for the virtual CU of the director. This port is called the director's CUP. ESCON port addresses were 1 byte long, so the range was hexadecimal 00 - FF. In the first ESCON Director, only 60 ports were physically implemented by using the range of hexadecimal C0 - FC. Port address FF was reserved for broadcast (but never implemented). Port address FE was chosen for the CUP. There also must be a device definition for the ESCON Director, but CNTLUNIT and IODEVICE are collectively called the CUP. FICON Directors also use port address FE for the CUP. ESCON Directors are unit type 9032, and FICON Directors are unit type 2032, but otherwise the definitions are the same.

Note: Port addresses FE and FF are reserved. FF cannot be used as the destination port address for any CU. FE can be used only as the destination port address for the CUP.

Figure 5-24 shows ESCON Director 08 with two channels (CHPID 07 attached to port C1 and CHPID 17 attached to port C2) and two physical CUs (1000 attached to port D0 and 2000 attached to D1). The figure shows the IOCP control statements that define the CHPIDs and CU 2000 and its devices. Next are the statements that define the ESCON Director as a CU and device. In this example, the same CHPID 07 is used to access the Direct Attached Storage Device (DASD) and the CUP. The destination link address of the CUP is FE (`LINK=FE`). There is only one device address that is needed for this CU (`UNITADD=(00,1)`) and device (`UNITADD=00`).

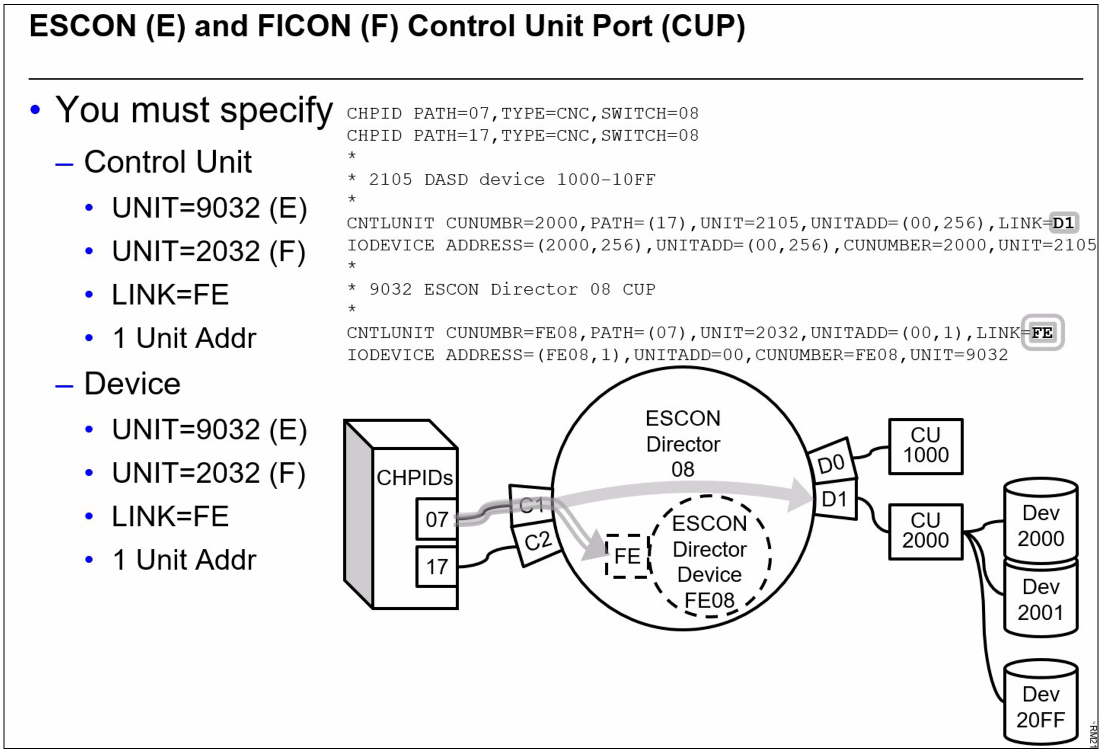

ESCON (E) and FICON (F) Control Unit Port (CUP)

- You must specify
  - Control Unit
    - UNIT=9032 (E)
    - UNIT=2032 (F)
    - LINK=FE
    - 1 Unit Addr
  - Device
    - UNIT=9032 (E)
    - UNIT=2032 (F)
    - LINK=FE
    - 1 Unit Addr

```
CHPID PATH=07,TYPE=CNC,SWITCH=08
CHPID PATH=17,TYPE=CNC,SWITCH=08
*
* 2105 DASD device 1000-10FF
*
CNTLUNIT CUNUMBR=2000,PATH=(17),UNIT=2105,UNITADD=(00,256),LINK=D1
IODEVICE ADDRESS=(2000,256),UNITADD=(00,256),CUNUMBER=2000,UNIT=2105
*
* 9032 ESCON Director 08 CUP
*
CNTLUNIT CUNUMBR=FE08,PATH=(07),UNIT=2032,UNITADD=(00,1),LINK=FE
IODEVICE ADDRESS=(FE08,1),UNITADD=00,CUNUMBER=FE08,UNIT=9032
```

Figure 5-24 Control Unit Port

There were programs such as ESCON Manager (for ESCON Directors) or IBM Tivoli® IOOPS (for FICON Directors) to use CUP for control and data access. These functions were incorporated into the Z/OS IOS. The System Management Facility (SMF) collects information from any defined CUP. Also, any defined CUP can send I/O alerts and error messages to the system console to warn about port failures, for example.

> **Caution:** FICON introduced cascading, but the relatively unsophisticated device driver for director was created before cascading when there could be only one director that was accessible by a channel. Cascading logically allows a single channel to access multiple CUPs, but CUP support was never extended to support that feature.

Figure 5-25 on page 131 shows the IOCP statements that allow for two-CUP access (at the top) and IOCP statements that are allowed as valid statements, but only one CUP can be accessed by the channel program (at the bottom). For multiple CUPs, you must use a different channel for each one.

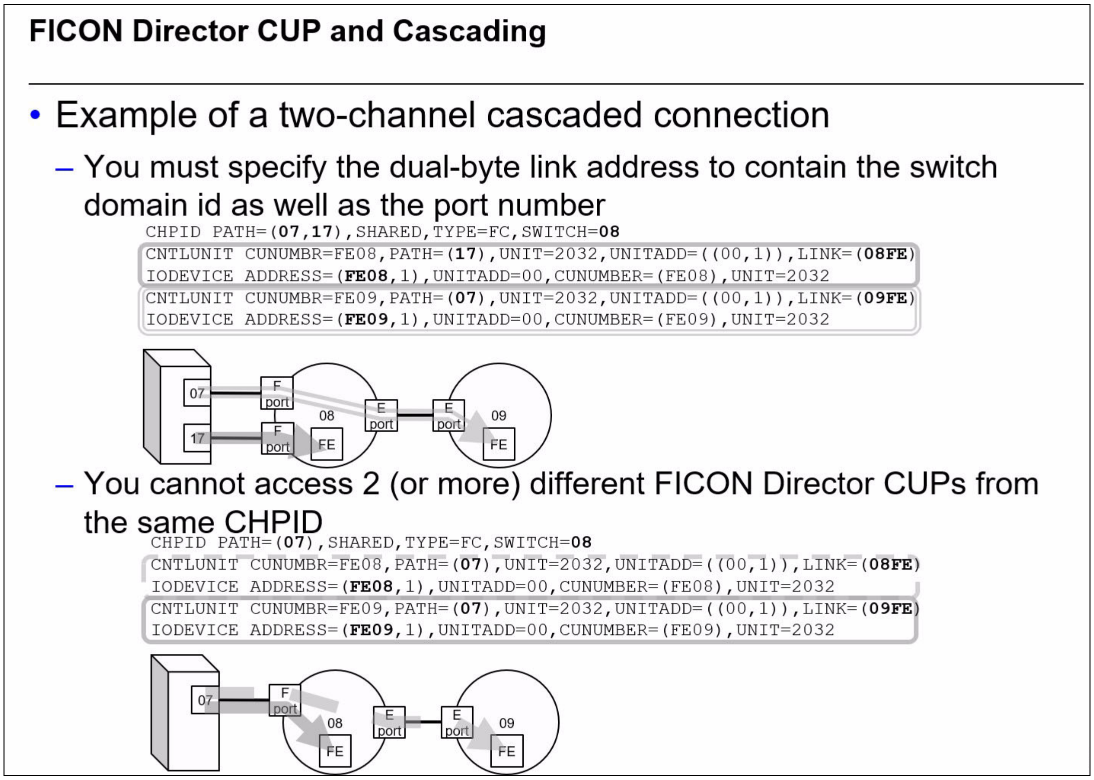

FICON Director CUP and Cascading

- Example of a two-channel cascaded connection
  - You must specify the dual-byte link address to contain the switch domain id as well as the port number

```
CHPID PATH=(07,17),SHARED,TYPE=FC,SWITCH=08
CNTLUNIT CUNUMBR=FE08,PATH=(17),UNIT=2032,UNITADD=((00,1)),LINK=(08FE)
IODEVICE ADDRESS=(FE08,1),UNITADD=00,CUNUMBER=(FE08),UNIT=2032
CNTLUNIT CUNUMBR=FE09,PATH=(07),UNIT=2032,UNITADD=((00,1)),LINK=(09FE)
IODEVICE ADDRESS=(FE09,1),UNITADD=00,CUNUMBER=(FE09),UNIT=2032
```

  - You cannot access 2 (or more) different FICON Director CUPs from the same CHPID

```
CHPID PATH=(07),SHARED,TYPE=FC,SWITCH=08
CNTLUNIT CUNUMBR=FE08,PATH=(07),UNIT=2032,UNITADD=((00,1)),LINK=(08FE)
IODEVICE ADDRESS=(FE08,1),UNITADD=00,CUNUMBER=(FE08),UNIT=2032
CNTLUNIT CUNUMBR=FE09,PATH=(07),UNIT=2032,UNITADD=((00,1)),LINK=(09FE)
IODEVICE ADDRESS=(FE09,1),UNITADD=00,CUNUMBER=(FE09),UNIT=2032
```

Figure 5-25 FICON Director CUP and cascading

In summary, the CUP feature on FICON Directors is a legacy of CUP on ESCON. Simply put, it is an in-band management capability. The CUP function allows z/OS to manage a FICON Director with a high level of control and security. Host communication includes control functions like blocking and unblocking ports, and performance monitoring and error reporting functions. FICON Directors have an internal N_Port that is reserved for CUP. The address of this port is defined (thee FE address), and by combining it with the DID of a switch, CUP can work well in single switch or cascaded topologies.

5.2.18 Resource Management Facility and FICON Director Activity Report

The Resource Measurement Facility (RMF) is a mainframe-centric view of the channel and device activity. CUP can give RMF insight into how fabric connectivity is affecting I/O operations on any CHPID. In fact, RMF can collect information from CUP and store it in 74 subtype 7 records. These records contain statistics for all switch ports, including details about B2B credits and port utilization, which are useful for performance and capacity management. This information is information to have because the average frame size of z/OS systems is 820 bytes, which is smaller than FC solutions, and that has implications on the number of buffer credits that are required to cover a distance. There are a few dozen CUP commands for monitoring and controlling switch functions.

For more information about the various reports beyond RMF that are available after FICON CUP is enabled and defined as a device, see 8.7, “FICON Director Activity Report” on page 426.

5.2.19 Microcode selection

FICON restricts FCP for security, timing, routing, and other considerations. IBM has worked with Cisco to qualify their devices. This qualification claims that the device has been tested in a mainframe environment, describes the uses and configurations that are supported for each device, and gives assurance to customers that the device will perform adequately and according to the protocols that are used in a mainframe environment. The rigorous testing for qualification extends to the specific level of firmware for each device. Vendors do not qualify each device type that they produce or each new release or modification level of the firmware. For more information about FICON releases for IBM c-type switches, see IBM Resource Link.

Note: Use only qualified levels of firmware in a FICON environment.

When using multiple FICON Directors in a data center, especially if they are cascaded, they should be on the same qualified level of firmware. When this setup is not possible, make sure to refer to release notes for details and verify that they are never more than two levels of difference. Generally, most levels of firmware are compatible, but they might support only the functions and features of the lower level.

5.2.20 Replication concepts and methods

Typically, the stored data in a mainframe environment is so mission-critical that one or more layers of replication and DR methods and plans are in place.

Note: The replication methods affect the FICON and FCP SAN configuration.

In the diagrams that follow, for brevity and readability, we use abbreviations like F=FICON, fc=FCP, M=Metro Mirror (MM), and so on.

A processor's channel can have a direct, point-to-point connection to a CU adapter, or a processor's channel can have a connection to a FICON Director for dynamic connections to one or more CU adapters, as shown in Figure 5-26.

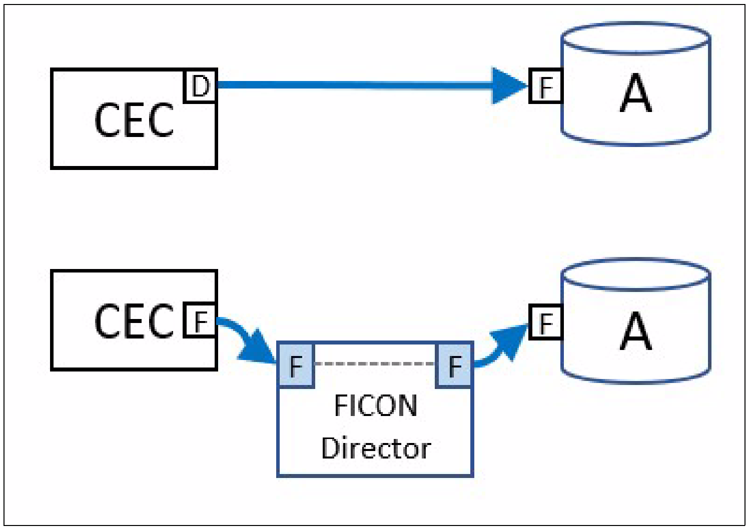

Figure 5-26 Storage Replication: Local I/O direct attach or through a FICON Director

MM is synchronous replication that is configured and managed by the CUs. Synchronous replication ensures that both copies of data are identical before the host's I/O request gets the response. IBM uses PPRC, EMC uses SRDF, and both use FCP connections between CUs either directly attached or through an FCP SAN, as shown in Figure 5-27.

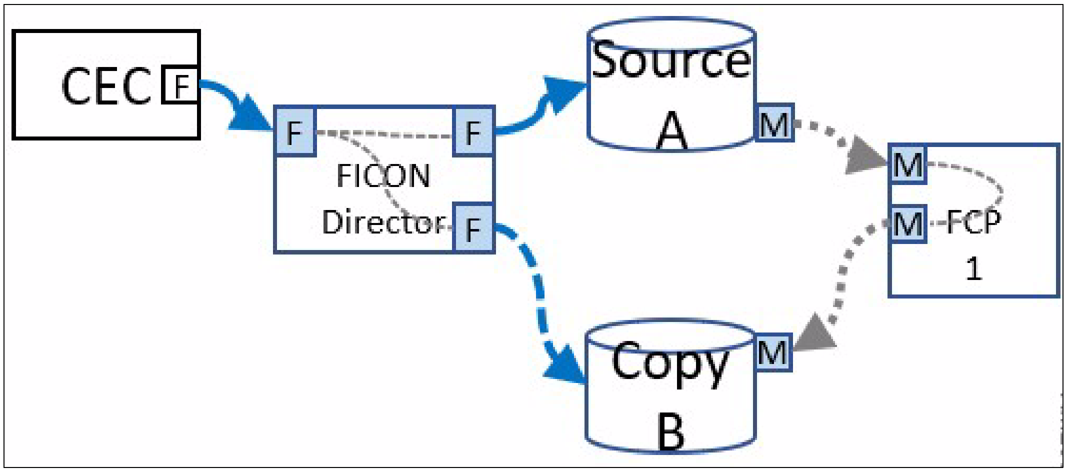

Figure 5-27 Storage replication: Metro Mirror

If there is a failure of the primary device (the A copy), HyperSwap or EMC AutoSwap will redirect I/O (dashed line) to the secondary device (the B copy) so that in most cases the application does not see any I/O issues. Because the I/O is not complete until both copies are complete, the A and B copies are relatively close to limit any increase in response time due to distance-induced latency. They might be in another data center in the same campus, but in different buildings with separate power and cooling sources.

Global Mirror (GM) is asynchronous replication that also is configured and managed by the CUs. Asynchronous replication can have a delay before new source data updates are reflected in the remote device. FCIP or DWDM systems extend the distance beyond the limit of an ISL, as shown in Figure 5-28.

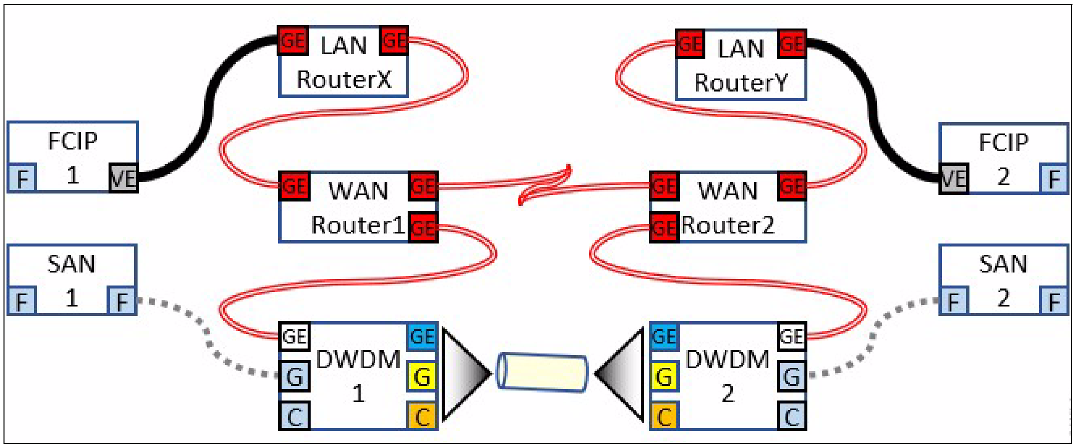

Figure 5-28 Storage replication: Global Mirror over FCIP

As shown, FCIP can be connected to local area network (LAN) switches or routers, then consolidated to wide area network (WAN) routers. WAN routers can either be attached to remote sites over a network (dotted network line) or over a DWDM (dashed network line) that is carrying other inter-site traffic. FCIP switches can also be attached directly to a DWDM system (dotted FCIP line).

IBM FlashCopy® (FC in the figures) creates a point-in-time copy of the data that is not synchronized with the source. FlashCopy also uses FCP, either direct-attach or through an FCP SAN switch, as shown in Figure 5-29. Point-in-time copies are useful to test DR procedures in a secondary site to not use the production GM copy.

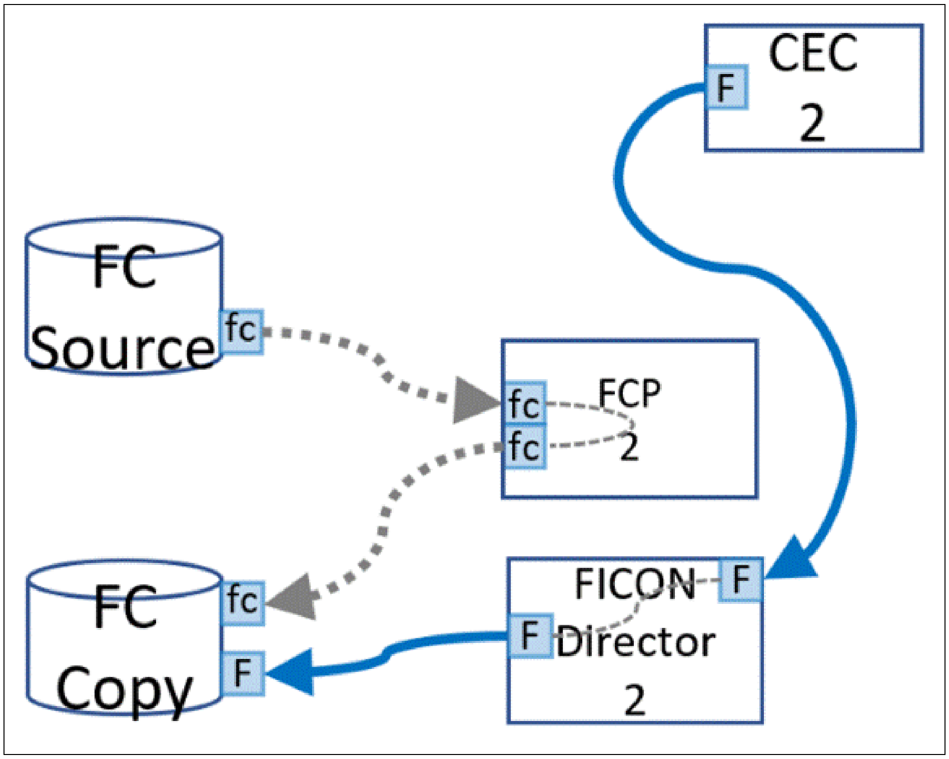

Figure 5-29 Storage replication: FlashCopy through an FCP SAN

XRC is another replication technique that uses an application that is called Sysplex Data Mover (SDM) in a remote site. Using FICON I/O, Figure 5-30 shows that SDM (1) monitors the devices in the main site (2) for updates and copies them back to the remote site. The updates are collected on Indexed Journal devices (3), and when the updates on devices that are associated in a consistency group are all copied, they are copied by using FlashCopy to the device that holds the remote copy (4).

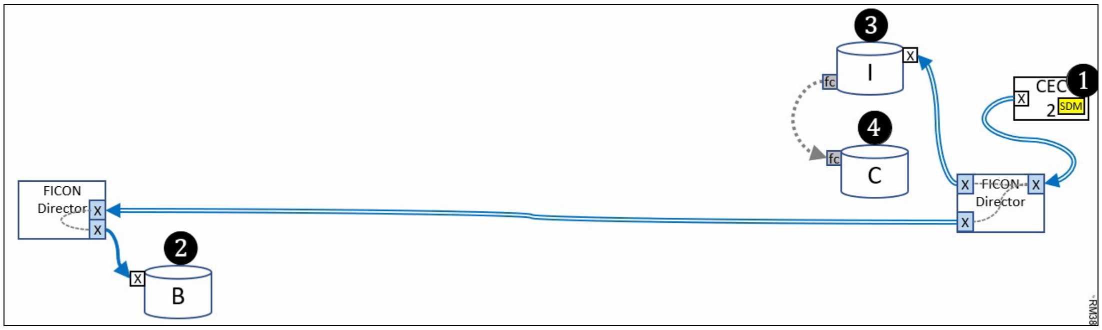

Figure 5-30 Storage replication: XRC

Global Copy (GC) sends fuzzy updates asynchronously to a remote device by using FCP without ensuring consistency. An application such as Logical Corruption Protection (LCP), with a FICON connection to the source and copy devices, as shown in Figure 5-31, can briefly switch the copy to use MM, which creates a consistent copy. Multiple copies can then be held and used as point-in-time checkpoints in case a fallback to a previous date and time is required.

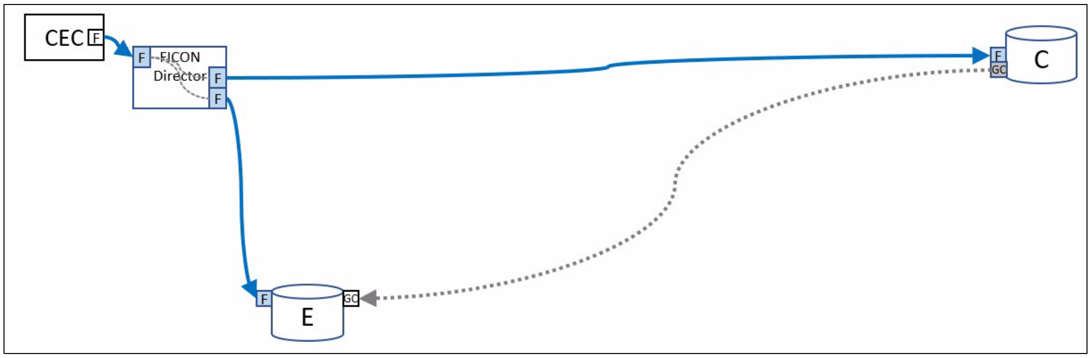

Figure 5-31 Storage replication: Global Copy

The combination of these techniques results in a 3-site Metro/Global Mirror (MGM) configuration, as shown in Figure 5-32.

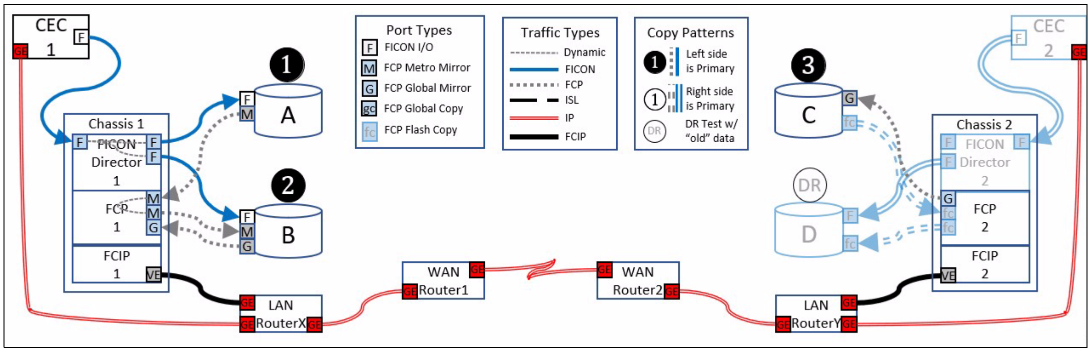

Figure 5-32 Storage replication: MGM

In a director-level device, the FICON VSAN and the FCP VSAN can all be placed on one physical chassis, and both can use the same FCIP connections to the remote site. FICON is also between sites to allow the local channel on CEC 1 to access the remote GM Copy, or to allow the SDM if using XRC. Network connections also are between sites and can be on the same LAN and WAN as FCIP, but should be on separate circuits to ensure that FCIP has dedicated network capacity. Furthermore, this approach shows only one item of each component (for example, channel, device, FICON Director, or WAN router), where at least two of each component is necessary for fault tolerance.

In more detail, Figure 5-32 shows:

- DASD groups:
  - Local I/O to A (FICON)
  - MM to B (FCP and synchronous)
  - GM to C (FCP, asynchronous, and over FCIP)
  - FlashCopy from C to D (FCP and synchronous) to create a point-in-time copy:
    - No further updates from C.
    - Right side can use D to test DR recovery at site 2 concurrently.

- Chassis contains:
  - FICON VSAN
  - FCP VSAN
  - FCIP module with:
    - GbE ports.
    - Virtual expansion ports (E_Ports) that are accessible by FCP or FICON VSANs.

These methods are summarized in Figure 5-33.

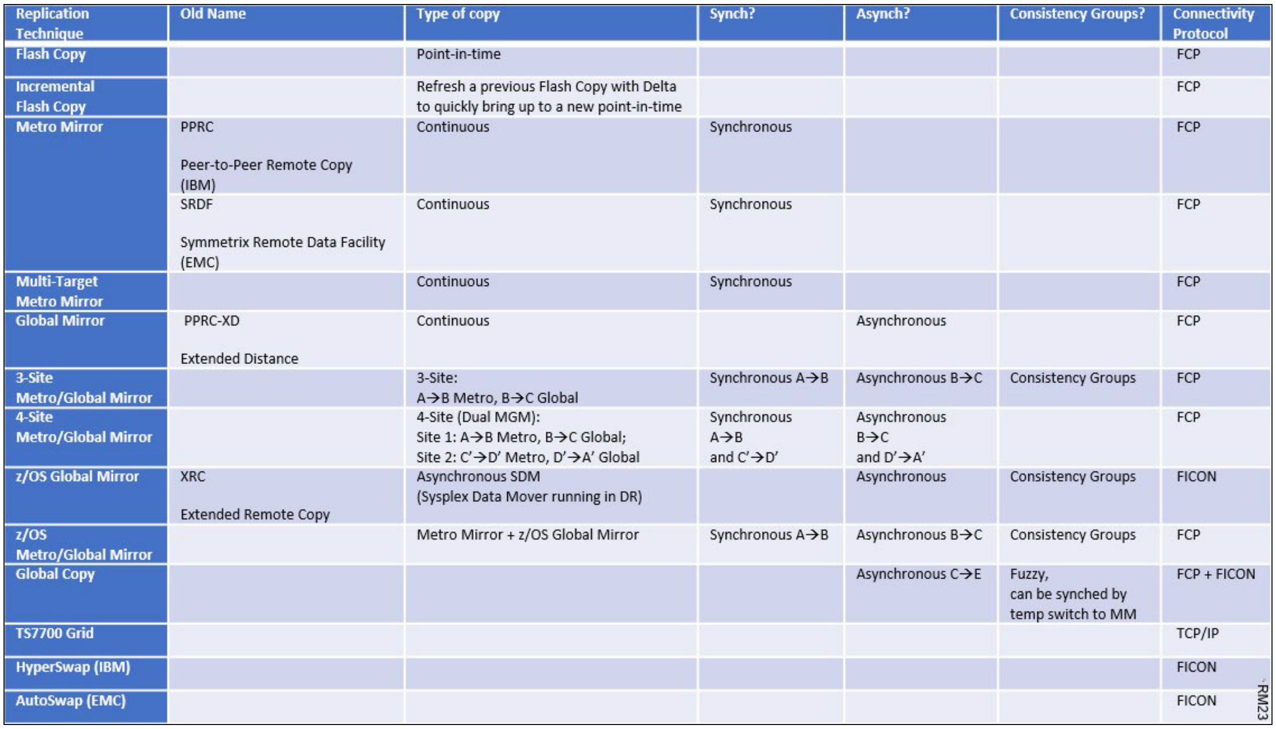

Replication Technique	Old Name	Type of copy	Synch?	Asynch?	Consistency Groups?	Connectivity Protocol
Flash Copy		Point-in-time				FCP
Incremental Flash Copy		Refresh a previous Flash Copy with Delta to quickly bring up to a new point-in-time				FCP
Metro Mirror	PPRC Peer-to-Peer Remote Copy (IBM)	Continuous	Synchronous			FCP
	SRDF Symmetrix Remote Data Facility (EMC)	Continuous	Synchronous			FCP
Multi-Target Metro Mirror		Continuous	Synchronous			FCP
Global Mirror	PPRC-XD Extended Distance	Continuous		Asynchronous		FCP
3-Site Metro/Global Mirror		3-Site: A→B Metro, B→C Global	Synchronous A→B	Asynchronous B→C	Consistency Groups	FCP

4-Site Metro/Global Mirror		4-Site (Dual MGM):
		Site 1: A→B Metro, B→C Global;
		Site 2: C'→D' Metro, D'→A' Global
			Synchronous
			A→B
			and C'→D'
				Asynchronous
				B→C
				and D'→A'
						FCP

z/OS Global Mirror	XRC Extended Remote Copy	Asynchronous SDM (Sysplex Data Mover running in DR)		Asynchronous	Consistency Groups	FICON
z/OS Metro/Global Mirror		Metro Mirror + z/OS Global Mirror	Synchronous A→B	Asynchronous B→C	Consistency Groups	FCP
Global Copy				Asynchronous C→E	Fuzzy, can be synched by temp switch to MM	FCP + FICON
TS7700 Grid						TCP/IP
HyperSwap (IBM)						FICON
AutoSwap (EMC)						FICON

Figure 5-33 Summary of replication techniques

5.3 General FICON planning considerations

Much like the FICON protocol itself has evolved over time, the same is true for requirements, expectations, and underlying hardware products. Network topologies also have diversified, becoming more flexible and scalable. Planning considerations for FICON deployments have evolved during this journey.

The following items should be considered during the installation planning phase for the IBM c-type family in a FICON environment:

- Which topology must be implemented?

> **Best practice:** Create a reasonably detailed picture of both the current and proposed topology during the planning phase. It is useful to describe the deployment to other individuals and facilitate disambiguation efforts (that is, verify that different terms or concepts that are used by different technical areas are describing the same or different things).

- Will this environment be a pure FICON environment, or a blend of FICON and open systems SAN traffic?

  > **Note:** Open systems and FICON can coexist in separate logical fabrics by using VSAN technology. A pure FICON environment can be built by using a single VSAN, with multiple paths in the fabric managed by z/OS. IBM c-type switches have native support for VSAN technology. It is always enabled, much like LPARs on modern mainframes.

- How many FICON channels do you need?
  - As always, it depends. If you have one IBM Storage system (such as Enterprise Storage Subsystem with up to 8 TB of capacity, two or four FICON paths to the storage array from a single IBM Z host is a good starting point. Measure the channel and switch port utilization to determine whether more paths are necessary. Mainframes can have as many as eight paths from an LPAR to a CU, but use at least two at a minimum. Add two more FCP paths to each storage array if you will be routing MM traffic through your fabric.

    > **Note:** A minimum of two FICON paths from an IBM Z host to the IBM Storage system is required. Also, there is a minimum of two FCP paths that are needed for any open systems traffic, which cannot share the FICON paths.

  - Add more FCP connections for concurrent open systems access to the IBM Storage system.
  - Other storage subsystems also have similar minimum pathing requirements.
  - For more than 8 TB of IBM Z capacity, use to six or eight FICON paths to the IBM Storage system from the IBM Z host.
  - With multiple IBM Z hosts, route their paths through one or more FICON Directors, and use the same rules as above based on the IBM Storage capacity for the total number of paths from the directors to the IBM Storage system.
  - Mainframe channels are defined as either FICON or FCP because they cannot be both. CU adapters are also defined as either FICON or FCP because they cannot be both. An ISL on a switch can be trunked to carry both FICON and FCP VSANs.
- How many FICON channels and CU ports will be connected to each FICON Director?
  - The number of channels and CU ports to be connected to each director depends on the number of FICON channels on the server or servers and CU ports on the devices, and the individual performance, availability, and growth requirements. It is possible to install and define more than eight paths to any physical CU (from the same IBM Z or S/390 processor image) when the physical CU has two or more logical CUs. A maximum of only eight channel paths may be defined to any one logical CU. This approach can be used for physical CUs that support greater than eight concurrent I/O transfers and that have a customer requirement for a high I/O rate, such as the IBM DS8000® or IBM Enterprise Storage Server®. For example, each of the eight FICON paths to an Enterprise Storage Server may be used to address eight different logical CUs in that Enterprise Storage Server.
  - Each FICON path can address up to 16,384 device addresses, and each FICON path can support multiple (up to 16 or more) concurrent I/O operations.

- How should the channels, CU ports, or ISLs be distributed among the switching modules in the directors?
  - Do not put multiple connections to the same end devices on the same switching module in your director (if possible).
  - Spread ISLs among multiple switching modules to increase reliability.
  - Each FICON switch or director configuration should satisfy all availability, performance, and growth requirements, so do not design a system that is 100% used from the beginning.
- When should ISL port channels be considered?
  - With port channels, you can group up to sixteen 32 Gbps ISLs to reach a cumulative data rate of up to 512 Gbps. port channels should be considered when a high volume of data traffic is anticipated between two FICON Directors.
  - VSAN traffic may be selectively managed by controlling VSAN access to some port channels.
- What distances can be accepted?

  The maximum distance between any host, CU, or FICON (or SAN) Director (or switch) is 10 km. With two cascaded switches between a channel and CU, the maximum distance is 30 km. However, the distance between a channel and a switch, or a switch and a CU, is contained within the same data center, so it is more likely to be within meters or tens of meters but nowhere near 10 km. IBM expects that distances greater than 10 km are provided by features on the FICON Directors, such as:

  - Extended Long Wave Small Form-factor Pluggables (ELW SFPs) are special, high-powered optics for extended distances that can support up to 25 km between switches, giving you a maximum of 10+25+10 = 45 km from mainframe to CU.
  - DWDM can provide reliable delivery of FC/FICON frames up to 5000 km. FICON qualifications reach only 300 km.
  - FCIP can provide reliable delivery of FC/FICON frames that are encapsulated in TCP segments at network distances (that is, terrestrially unlimited). FCIP links can be used to provide connectivity between cascaded switches.

  > **Note:** The speed of light latency of 5 microseconds per kilometer can, at sufficient distance, exceed the maximum allowable response time for I/O or applications. The response time is highly variable and depends on many factors within the application, and it cannot be stated as a standard limit.

- Should the switch IDs in the IOCP definitions and domain ID in the FICON Directors match?
  - There are two locations where the switch ID is specified in the IOCP:
    - `LINK=` on the `CNTLUNIT` statement
    - `SWITCH=` on the `CHPID` statement
  - For the `LINK=` parameter when identifying a CU adapter, the domain ID of the destination FICON Director must match that of the switch ID value that is defined at the FICON switch:
    - The domain ID is required in a cascaded environment.
    - The domain ID is optional in a noncascaded environment.
    - If the domain ID in `LINK=` is incorrect, the I/O will fail, or worse, go to the wrong CU.

  - For the **SWITCH=** parameter of a CHPID that is attached to a FICON switch or director, the domain ID of the destination FICON Director is not required to match that of the switch ID value that is defined at the FICON switch. However, as a best practice, specify the same value that is defined at the FICON switch to avoid any potential confusion or errors.

- Should zoning be used to separate specific channels and CU ports from other connected channels and CU ports?

  FICON on mainframes uses source-based routing that is defined in the IOCDS, unlike SCSI, which uses a broadcast to all entities on a shared bus. Thus, SCSI relies on FCP zones to limit the scope of the broadcast, but zones are not required for FICON. However, FICON ports must be allowed to connect to any other port on a FICON Director VSAN. There are two methods to provide FICON with the “any-to-any” connectivity it requires:

  - Define the FICON VSAN to access the default-zone, for example:

    ```
    zone default-zone permit vsan FICON_VSAN_Number
    ```

  - Define a single zone that contains all ports in the FICON VSAN, for example:

    ```
    zone name FICON_ZONE_Name vsan FICON_ZONE_Number
    member domain-id Did portnumber Pnum
    ```

    This single zone approach, despite theoretically possible, should never be used.

  If you have both FICON and FCP on the same physical switch and SAN network, as a best practice, separate VSANs for FICON and FCP so that their traffic can coexist on the same physical network but is isolated by different logical networks. It is possible, but not recommended, to create zones for FICON channels and CU ports within a VSAN that contain both FICON and FCP. The VSAN would require all the FICON configuration settings, which might not be correct for the FCP traffic in the same VSAN, so there is no advantage to that method.

  > **Important:** VSAN isolation, not zoning, is the recommended way to isolate FICON traffic.

- Should the out-of-band management port of the FICON Director or Directors and Data Center Network Manager (DCNM) server be connected to a separate LAN or to the corporate network?

  The directors and the DCNM server should be reachable on a separate LAN to isolate director management traffic from other IP traffic. When remote access is required to operate and maintain the FICON Director or Directors, connect the DCNM server to your corporate network through an IP router.

- What Small Form-factor Pluggable (SFP) transceivers should be used for FICON?

  The recommended pluggable transceivers to be used in a FICON environment are 16 GbE or 32 GbE LWL transceivers for single-mode fiber with LC connectors.

As part of system planning activities, you must decide where to locate the equipment, how it will be operated and managed, and what the business continuity requirements are for DR, tape vaulting, and so on. The types of software (OSs and application programs) that are intended to be used must support the features and devices on the system.

As a historical note, in the early days of FICON, you had more CU ports than CHPIDs. The ratio was 1:10 typically, which was the opposite of FC, where initiators were 10 times more than targets. Today, things are different after the introduction of logical images. A CHPID is only one adapter or port of a processor. A single processor has many CHPID ports. A CU has many HBAs or CU ports.

With today's large DASD CUs, it is common to run one CHPID per CU adapter, and a different CHPID for each logical CU image. Physically, that means more CHPIDs than CUs, but about 1:1 for CHPIDs to HBAs. Long ago with the original bus-and-tag approach, there were many CUs that were accessed through one channel. ESCON did not change that much. ESCON Multiple Image Facility (EMIF) allowed a CHPID to be shared by multiple LPARs and the ratio got higher. But, the arrival of logical images on the CU by using CUADDR (the flip side of EMIF) changed that drastically. With CUADDR, a single CU could compress 16 physical CUs into one box. Then, many adapters on a physical CU made the 1:1 ratio (CHPID to HBA) more normal.

5.4 IBM c-type considerations for FICON

This section covers the following topics:

- IBM c-type architecture and FICON optimizations
- FICON port numbering and addressing on IBM c-type switches
- Domain IDs, FCID allocation, and fabric binding
- How to implement more than one FICON VSAN
- FICON configuration files
- Code version selection in FC and FICON intermix environments

5.4.1 IBM c-type architecture and FICON optimizations

In this book, the terms *switch* and *director* are often used interchangeably, even if originally they were different. In general, a switch is smaller in size than a director and is a single point of hardware failure. A director is a modular device that is larger in size with no single point of hardware failure. Some times, the generic term of switch is used for indicating devices with the functional capability to switch frames from one port to another. Under this meaning, it differentiates between director-class switches and fabric switches.

Mission-critical directors represent the latest innovation in the architecture of modular devices. They provide a degree of redundancy for the critical components that cross-connect line cards among themselves. Architectures in use 10 years ago had only two such components in active-active mode, which led to a 50% drop in total throughput in a single failure. Now, a state-of-the-art mission-critical director can be deployed with N+1 redundancy for those elements so that an eventual single failure has no impact on the total throughput. The difference is shown in Figure 5-34 on page 141. IBM c-type mission-critical directors are the perfect match for IBM Z deployments.

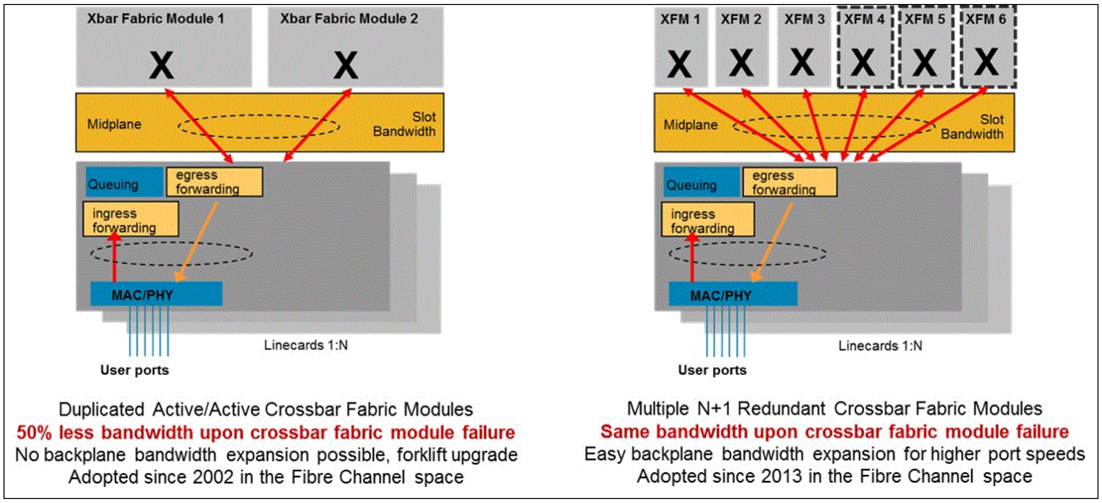
Xbar Fabric Module 1
Xbar Fabric Module 2
X
X
Midplane
Slot Bandwidth
Queuing
egress forwarding
ingress forwarding
MAC/PHY
Linecards 1:N
User ports
Duplicated Active/Active Crossbar Fabric Modules
50% less bandwidth upon crossbar fabric module failure
No backplane bandwidth expansion possible, forklift upgrade
Adopted since 2002 in the Fibre Channel space
XFM 1
XFM 2
XFM 3
XFM 4
XFM 5
XFM 6
Midplane
Slot Bandwidth
Queuing
egress forwarding
ingress forwarding
MAC/PHY
Linecards 1:N
User ports
Multiple N+1 Redundant Crossbar Fabric Modules
Same bandwidth upon crossbar fabric module failure
Easy backplane bandwidth expansion for higher port speeds
Adopted since 2013 in the Fibre Channel space

Figure 5-34 Architectural difference between a director and a mission-critical director

A FICON switch is used as a generic term for indicating an FC switch or director that supports the transfer of frames containing FC-SB-6 payloads, supports the FC-FS ELS that FICON requires, and has an internal logical N_Port that supports CUP.

Some models of the IBM c-type switches support FICON, FCP, and FCIP capabilities within a single, high availability (HA) platform. This combination simplifies the migration to shared mainframe and open systems storage networks, and it went through an extensive period of integration testing to meet the most stringent IBM FICON qualification requirements.

FICON is supported on the following IBM c-type Director-class and multiservice switches:

- IBM Storage Networking SAN192C-6 mission-critical director
  - IBM 48-Port 32-Gbps FC Switching Module (01FT644)
  - IBM SAN Director Supervisor Module 1 (01FT600)
  - IBM SAN Director Supervisor Module 4 (02JD753)
  - IBM 24/10 Port SAN Extension Module (01FT645)
- IBM Storage Networking SAN384C-6 mission-critical director
  - IBM 48-Port 32-Gbps FC Switching Module (01FT644)
  - IBM SAN Director Supervisor Module 1 (01FT600)
  - IBM SAN Director Supervisor Module 4 (02JD753)
  - IBM 24/10 Port SAN Extension Module (01FT645)
- IBM Storage Networking SAN50C-R multiservice switch

The minimum required NX OS Release is 8.1(1b) on IBM c-type devices. Although no hardware changes are required, you must have the MAINFRAME_PKG license to configure FICON parameters, or you can use the initial grace period of 120 days.

> **Note:** The *grace period* is the amount of time that an application can continue functioning without a license. In this case, the grace period is set to 120 days from the first occurrence of configuring any licensed feature without a license package. The grace period starts with the first check-out, and will be counted only for the days when that feature is enabled and configured (even if not used). If you remove configuration for this feature, the counter for the grace period stops incrementing.

FICON traffic may also be carried over 1- and 10-gigabit Ethernet (GbE) links between switches that have IP storage ports. There is no requirement for a specific SAN_EXT license to enable FCIP on IBM c-type switches.

When you enable the FICON feature within a VSAN on IBM c-type switches, the following considerations apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The switch load-balancing scheme is set to SID/DID by default. It can be changed to SID/DID/OXID for a more efficient approach if all the devices in the FICON VSAN support FIDR or exchange-based routing.
- The switch IPL configuration file is automatically created.

When the FICON feature on a VSAN is enabled, the switch IPL file is created automatically with a default configuration. The IPL file contains specific settings for FICON-enabled VSANs that are applied at restarts. The IPL file contains port configuration information about each FICON port regarding what other FICON ports are allowed to communicate with this port (prohibit function), whether this port is isolated from other FICON ports (block function), and the descriptive identifier of this FICON port (port name).

The IPL file also includes the port number mapping for port channels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration, in-order delivery guarantee, the static domain ID configuration, and the fabric binding configuration. This information is not stored in the startup-config or running-config of the switch as other configuration information is.

In general, changes to the active configuration (HCD settings) are saved to the IPL too. You can save up to 16 FICON configuration files on each FICON VSAN. The files are in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format, and they are saved in persistent storage so they persist after a reload of the switch. This IPL file works specifically with the CUP feature, but it can also be edited by using the NX OS command-line interface (CLI) or DCNM tool, as shown in Figure 5-35.

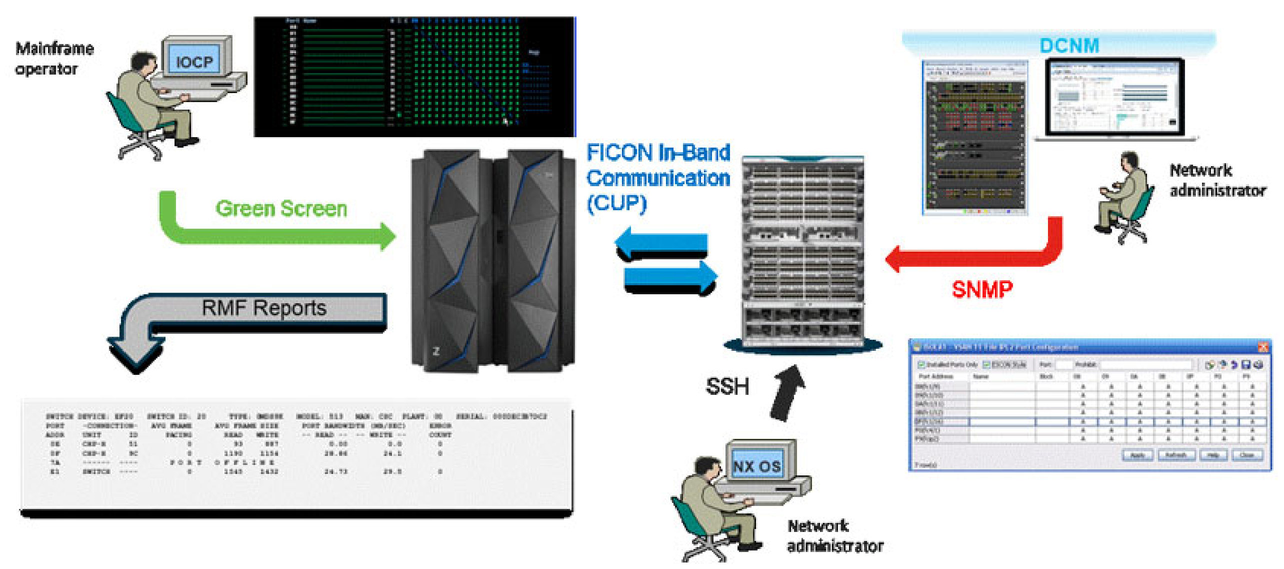

```
SWITCH DEVICE: EF20   SWITCH ID: 20    TYPE: [illegible]   MODEL: 513   MAN: CSC   PLANT: 00   SERIAL: [illegible]
PORT    -CONNECTION-   AVG FRAME     AVG FRAME SIZE    PORT BANDWIDTH (MB/SEC)     ERROR
ADDR    UNIT      ID      PACING      READ   WRITE    -- READ --  -- WRITE --     COUNT
 0E     CHP-H     51           0        93     887         0.00         0.0         0
 0F     CHP-H     9C           0      1190    1154        28.86        24.1         0
 7A     ------  ----        P O R T   O F F L I N E
 E1     SWITCH  ----           0      1543    1432        24.73        29.5         0
```

Figure 5-35 Management options for IBM c-type FICON Directors

5.4.2 FICON port numbering and addressing on IBM c-type switches

Despite having common FC-0, FC-1, and FC-2 protocol layers, FC and FICON have a different addressing method. One of the biggest differences is that FC is dynamic, and initiators and targets learn their fabric-assigned FC-2 FCID after a successful fabric login. Hosts learn their targets by joining the FC name server, and they learn logical unit numbers (LUNs) by querying the allowed targets.

Conversely, FICON addressing is static, with a limit of 256 link addresses per FICON Director or FICON VSAN. In z/OS environments, the IOCP is used to configure logical channel paths through the fabric, assigning link addresses, logical switch addresses, and Channel Path IDs (CHPIDs) by using the appropriate macros and parameters.

Port identification on networking devices can have physical or logical relevance. In IBM c-type switches, physical ports are identified based on the front panel location of the port and the specific slot in which the switching module resides. Considering a 48-port line card that is inserted in the upper slot of a modular chassis (slot 1 of an IBM Storage Networking SAN384C-6, for example), we refer to physical interface fc1/1 for the upper left port, and we refer to physical interface fc1/48 for the lower right port in the line card, as shown in Figure 5-36. The last port of an IBM Storage Networking SAN384C-6 Director, fully populated with 48-port line cards, is interface fc10/48 (the supervisors take slots 5 and 6).

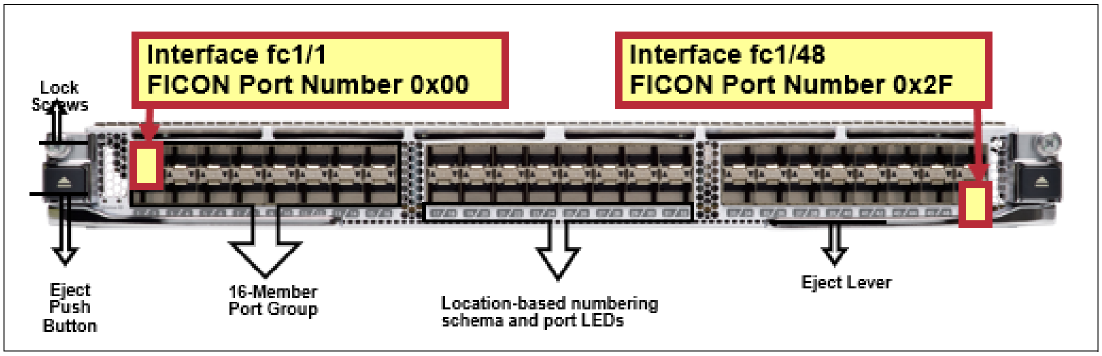

Figure 5-36 Interface numbering schema for a 48-port switching module in slot 1

Similarly, physical Ethernet ports that carry FCIP traffic are identified, for example, physical interface IPStorage 3/1. This type of port identification is not configurable, and it is unique per chassis. When referring to logical interfaces that have no clear match to a physical location, we refer to them with the logical interface type and a sequential integer number, like logical interface fcip 3 or logical interface port-channel 9. The port identification method is used for FCP and cabling activities but not for FICON.

Referring to the FICON feature, ports in IBM c-type switches are identified by a statically defined 8-bit value (256 combinations) known as the *FICON Port Number*. It is not the same as the numbering of FC port positions that start at the number one for each module. FICON Port Numbers are assigned in two ways:

- Automatically by NX-OS software based on switch type, the actual slot position of a module in the chassis, and the relative port position on the module.
- Manually forced by the administrator. Before assigning, changing, or releasing a port number, the port should be in the admin shut state.

Years ago, FICON Port Numbers could not be changed, and the first port in a switch always started at zero. On IBM c-type FICON Directors, the numbering schema was conceived to allow for a 48-port module to be present in any slot. The FICON Port Numbers are assigned to physical port positions in a chassis regardless of whether that position is filled, the port status is up or down, or the number of ports on the module (24 or 48). If a module has fewer physical ports than the number of FICON Port Numbers that is assigned to the slot, then the excess port numbers are unused.

For example, if a 24/10 SAN Extension module is inserted into slot 2 in an IBM Storage Networking SAN192C-6 Director that can potentially take a 48-port module, the port numbers 56 - 79 that are associated with positions 25 - 48 on that system board are considered uninstalled and are not used by a module that is installed into slot 3 in the same chassis. Because the IBM Storage Networking SAN384C-6 Director can host more than 256 ports, more than one FICON VSAN is required on those large chassis to use all ports.

On IBM Storage Networking SAN50C-R, there are up to 40 FC ports, and the FICON Port Numbers that are assigned by default are as follows:

- 0x00 - 0x27 (0 - 39) for ports FC 1/1-1/40.
- 0xF0 - 0xFD (240 - 253) are reserved for logical interfaces (FCIP and port channels).
- 0xFE (254) is reserved for CUP and cannot be assigned to any other interface.
- 0xFF (255) is a reserved port and cannot be assigned to any interface.

On IBM Storage Networking SAN192C-6, there are up to 192 FC ports, and the FICON Port Numbers that are assigned by default are as follows:

- 0x00 - 0x2F (0 - 47) for ports FC 1/1-1/48.
- 0x30 - 0x5F (48 - 95) for ports FC 2/1-2/48.
- 0x60 - 0x8F (96 - 143) for ports FC 5/1-5/48.
- 0x90 - 0xBF (144 - 191) for ports FC 6/1-6/48.
- 0xF0 - 0xFD (240 - 253) are reserved for logical interfaces (FCIP and port channels).
- 0xFE (254) is reserved for CUP and cannot be assigned to any other interface.
- 0xFF (255) is a reserved port and cannot be assigned to any interface.

On IBM Storage Networking SAN384C-6, there are up to 384 FC ports, and the FICON Port Numbers that are assigned by default are as follows:

- 0x00 - 0x2F (0 - 47) for ports FC 1/1-1/48.
- 0x30 - 0x5F (48 - 95) for ports FC 2/1-2/48.
- 0x60 - 0x8F (96 - 143) for ports FC 3/1-3/48.
- 0x90 - 0xBF (144 - 191) for ports FC 4/1-4/48.
- 0xC0 - 0xEF (192 - 239) for ports FC 7/1-7/48.
- 0xF0 - 0xFD (240-253) are reserved for logical interfaces (FCIP and port channels).
- 0xFE (254) is reserved for CUP and cannot be assigned to any other interface.
- 0xFF (255) is a reserved port and cannot be assigned to any interface.

Supervisor modules do not have FICON Port Number assignments, but FCIP interfaces and port channels, despite being logical and not physical interfaces, need a FICON Port Number that is different from any of the front panel ports that were described previously. Even CUP is a logical port that needs a FICON Port Number. FICON Port Numbers for the FCIP interfaces and port channels are allocated from the address space, which is reserved for logical ports (beyond the range of the maximum number of physical FICON ports). As an example, here is how to assign the FICON Port Number 234 to the interface port-channel 1:

```
ficon logical-port assign port-numbers 234
interface port-channel 1
ficon portnumber 234
```

FICON Port Numbers in the CLI of IBM c-type switches are expressed in decimal format. By default, FICON Port Numbers in hexadecimal format are the same as FICON Port Addresses, and they are used in the IOCDS definitions when defining the physical connections for z/OS, for example:

```
CNTLUNIT,CUNUMBR=07D0,PATH=(86,89),UNITADD=((00,016)),
LINK=(5022,5022),CUADD=1,UNIT=2105
```

To facilitate daily operations, Datacenter Network Manager and its embedded Device Manager (DM) can toggle the display between interface numbers and FICON Port Numbers. The default visualization with interface labels is shown in Figure 5-37.

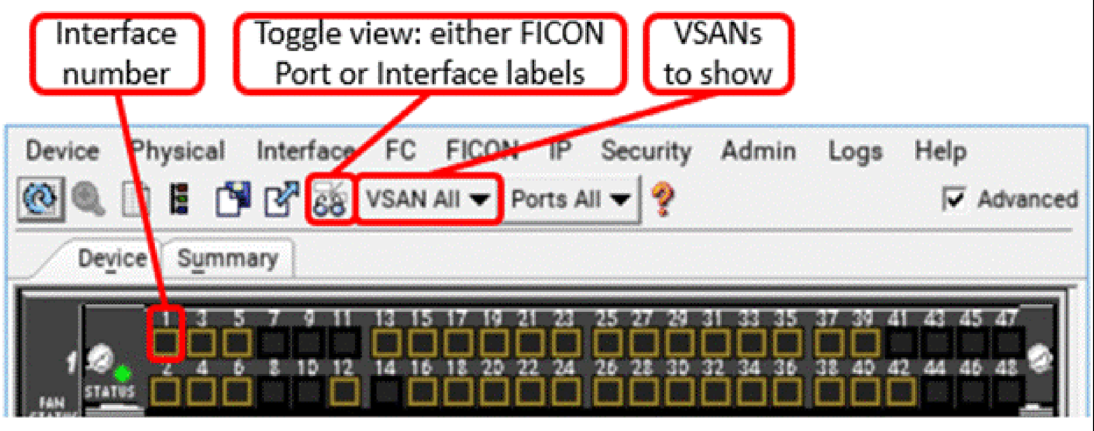

Figure 5-37 Default visualization from Device Manager with interface labels

After clicking the toggle icon, the corresponding FICON Port Number visualization opens, as shown in Figure 5-38.

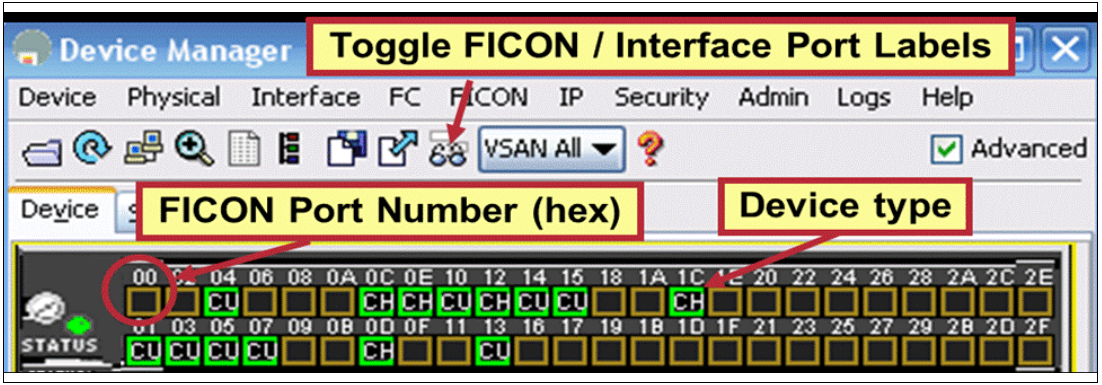
Device Manager
Toggle FICON / Interface Port Labels
Device Physical Interface FC FICON IP Security Admin Logs Help
VSAN All
Advanced
Device
FICON Port Number (hex)
Device type
00 02 04 06 08 0A 0C 0E 10 12 14 15 18 1A 1C 1E 20 22 24 26 28 2A 2C 2E
CU CH CH CU CH CU CU CH
01 03 05 07 09 0B 0D 0F 11 13 16 17 19 1B 1D 1F 21 23 25 27 29 2B 2D 2F
CU CU CU CU CH CU
STATUS

Figure 5-38 Device Manager visualization of FICON Port Numbers and Device Type

FICON Port Numbers are automatically assigned and cannot be changed. They represent the first and physical level of addressing in a FICON setup. A second and virtual level of addressing is introduced for IBM c-type FICON switches. For every FICON Port Number, there is one associated FICON Port Address.

The FICON Port Address is the value that is used for port identification in commands that are used against FICON ports (blocking and unblocking FICON ports, naming FICON ports, and prohibiting and allowing masking). The FICON Port Address is closely related to the link address as part of the definition of CUs within the host IOCP. FICON Port Addresses are used in the IOCDS definitions when defining the physical connections for z/OS, for example:

```
CNTLUNIT,CUNUMBR=07D0,PATH=(86,89),UNITADD=((00,016)),
LINK=(5022,5022),CUADD=1,UNIT=2105
```

By default, FICON Port Numbers are the same as FICON Port Addresses. More precisely, FICON Port Numbers in the CLI of IBM c-type switches are expressed in decimal format. When expressed in hexadecimal format, they are the same as FICON Port Addresses.

So, what is the reason for having both a FICON Port Number and a FICON Port Address for every FICON interface? There is a good reason: FICON Port Numbers cannot be changed, but you can swap the FICON Port Addresses by using the FICON Port Swap feature.

All traffic routing in FICON is accomplished based on the FICON Port Address (it is the value that forms the link address). So, what happens when a physical interface on a switch goes bad with the above architecture? To replace a single bad interface, most of the time you must replace the module it is on, which means that you could be affecting 24 or 48 ports, depending on the number of ports on the module. Thus, you would need a potentially large maintenance window to repair a single port issue.

This situation is what the FICON Port Swap function was created for. The FICON Port Swap function allows you to swap the FICON Port Addresses for two FICON Port Numbers. As an example, say that you start with two ports: 0x01, which is used for an active FICON connection, and 0x2F, which is a spare port and unused. Before the swap, each of these ports has the same FICON Port Number and FICON Port address, that is, FICON Port Address 0x01 is on physical FICON Port Number 0x01, and similarly for 0x2F. If we swap these two FICON Ports, we exchange the FICON Port Addresses for the FICON Port Numbers. This process means that FICON Port Number 0x01 will have FICON Port Address 0x2F, and FICON Port Number 0x2F will have FICON Port Address 0x01. To restore the I/O for FICON Port Address 0x01, move the fiber connection from FICON Port Number 0x01 to 0x2F and bring the port back online. The customer can be running I/O with their unchanged IOCP configuration on the mainframe without any major maintenance window.

Note: Port swapping is not supported for logical ports (port channels and FCIP links).

The above explanation is appropriate for FICON Directors with less than 255 usable ports and was used many years ago. But, the situation is more complex for larger directors, so a slightly different addressing technique is required. These many FICON ports must be split in to at least two FICON VSANs, but you are still constrained by the FICON Port Number limit of 256 values (8 bits) and assignment based on port location. Thus, FICON Port Numbers on IBM c-type FICON switches are now virtual, with the software allowing customers to assign the valid FICON Port Numbers to whatever physical interfaces fit their needs. For convenience, these virtual FICON Port Numbers are still associated with the same value for the FICON Port Addresses by default.

Of course, you cannot put two interfaces with the same FICON Port Number into the same FICON VSAN, but you have now an incredible level of flexibility, as shown in Figure 5-39 on page 147.

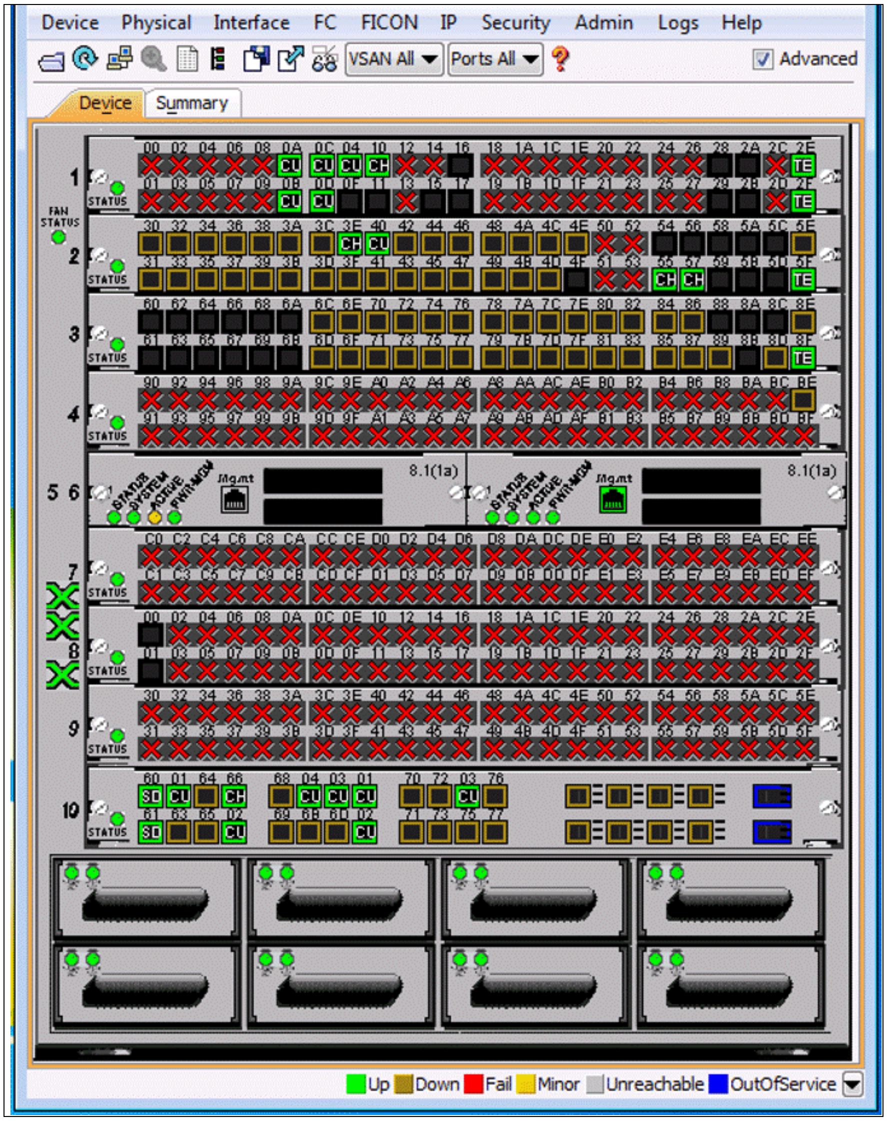

Figure 5-39 The flexibility that is achieved with virtual FICON Port Numbers

The FICON Director is displayed by DM when the toggle icon is clicked. The FICON Port Addresses are showing above each physical interface. On the first 48-port module (module 1), the FICON Port Numbers and FICON Port Addresses are configured as 0x00 - 0x2F. The FICON Port Numbers (and the associated FICON Port Addresses) continue sequentially through module 7, where the final port has FICON Port Number of 0xEF. Module 8 and Module 9 are reusing FICON Port Numbers 0x00 - 0x5F. The final module shows the ultimate flexibility of the virtual FICON Port Numbers because they can be placed and varied to nonsequential sequences. Now, two FICON Port Number 0x00s (the first port in module 1 and the first port in module 8) must be in two different FICON VSANs.

With this new ability to have virtual FICON Port Numbers, the FICON Port Swap function is obsolete but still available. With the now virtual FICON Port Number, if there is the same physical port failure as described above (active FICON Port Number 0x01 fails and must be moved to the spare FICON Port Number 0x2F), you can bring down both ports and exchange their FICON Port Numbers by using either the CLI or DM. When this task is done, you retain the 1:1 mapping between the FICON Port Number and the FICON Port Address, which keeps things as simple as possible. For this reason, it is a best practice to not use the FICON Port Swap function and instead reassign or exchange the actual FICON Port Numbers when the need arises.

5.4.3 Domain IDs, FCID allocation, and fabric binding

FICON requires a predictable and static FCID allocation scheme. When FICON is enabled, the FCID that is allocated to a device (CH or CU) is based on the FICON Port Address of the physical port to which it is attached. The FICON Port Address forms the middle byte of the fabric-assigned FCID address. In fact, FCIDs are 3 bytes. The first byte is the static domain ID of the switch, in hexadecimal format, which matches the **SWITCH** parameter on the CHPID macro in the IOCDS. The second byte of the FCID is the FICON Port Number (equivalent to the FICON Port Address). The last byte of the FCID (the Arbitrated Loop Address or Arbitrated Loop (Station) Address (ALA)) defaults to 0 for FICON VSANs. FICON requires the last byte of the fabric address to be the same for all allocated FCIDs. The value of the last byte can be changed if required, but only when the FICON VSAN is in the offline state.

Figure 5-40 shows the static FCID allocation for switched and cascaded FICON topologies.

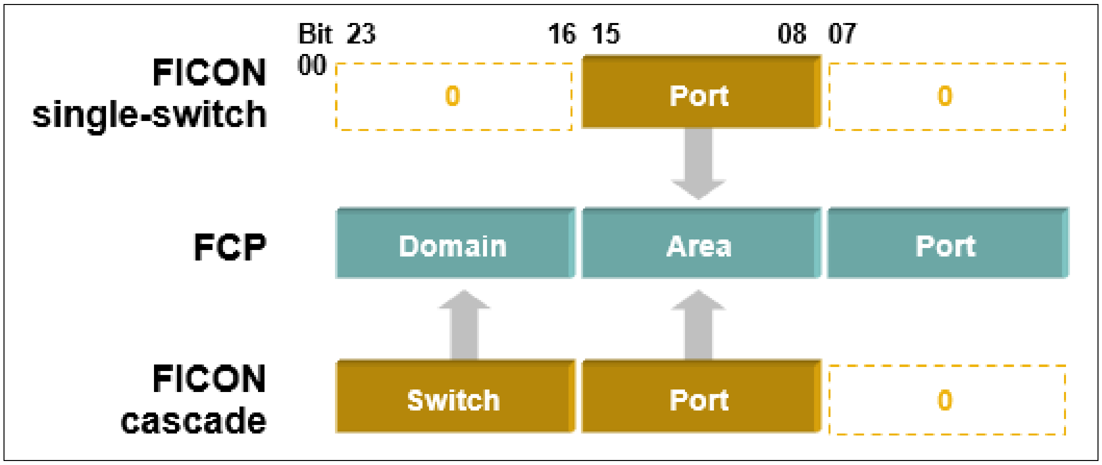

Figure 5-40 Static FCID allocation for switched and cascaded FICON topologies

It is important to clarify a possible confusion in terminology: On storage networking devices, the term FICON Port Address or FICON Port Number is used to indicate a port, but mainframe experts tend to use the term *link address*. Those concepts are in strict correlation. FICON Port Addresses are equivalent or partly match the link addresses that are used by mainframes. In a simple switched topology, the link address is 1 byte and equivalent to the FICON Port Address (0x00 - 0xFF). In a cascaded environment, the link address is 2 bytes and equivalent to the domain ID plus the FICON Port Address. This link address is what is coded in HCD or IOCP in the **LINK=** parameter for a **CNTLUNIT** statement.

The HCD or the IOCP input statements specify the destination link address that is used to access CU adapters in the path to a device. Mainframe channels learn of their link address from the FICON Director during their link initialization (their link address is not predefined in HCD or IOCP). Both the source link address (for example, the CHs) and the destination link address (for example, the CUs) are in the FICON I/O frame that is sent to the FICON Director, which uses these addresses to make the dynamic connection between the ports, as shown in the FC-SB header of Figure 5-41.

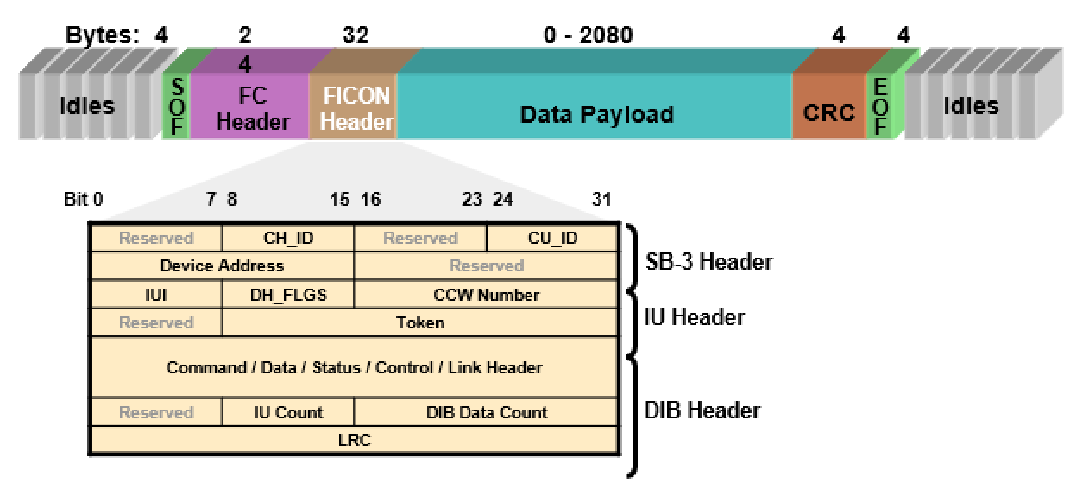

Reserved	CH_ID	Reserved	CU_ID
Device Address		Reserved	
IUI	DH_FLGS	CCW Number	
Reserved	Token		
Command / Data / Status / Control / Link Header			
Reserved	IU Count	DIB Data Count	
LRC			

Figure 5-41 The FC-SB frame format

Thus, a FICON Port Number is assigned, and it must remain the same even if the FICON Director restarts for some reason. The Port Number cannot be dynamic or randomly assigned by the switch.

Sometimes, there is confusion regarding the FCID and the FICON Port Number. To understand the differences, read the following points:

- FICON uses source-based routing because it explicitly identifies the switch ID and FICON Port Number on the destination switch based on the information in the I/O configuration that is specified in HCD/IOCP. Thus, the FICON Port Number is static for FICON I/O.
- FC uses fabric-based routing because the fabric determines the path and provides the requester (that is, the N_Port) with the destination FCID. As a result, the FCID can change if the N_Ports (that is, server HBAs or device I/O adapters) move. This change will not affect any N_Port configurations.

Despite IBM c-type switches having a dynamic FCID allocation scheme, when FICON is enabled on a VSAN, all the ports are changed to static FCIDs.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. Fabric binding is configured by using a set of switch worldwide names (sWWNs) and a persistent (static) domain ID and binds the fabric at the switch level. Enforcement of fabric binding policies is done on every activation and when the port tries to start. However, enforcement of fabric binding at the time of activation happens only if the VSAN is a FICON VSAN. A user-specified fabric binding list contains a list of sWWNs within a fabric. If a sWWN attempts to join the fabric and that sWWN is not in the user-specified list or the sWWN is using a domain ID that differs from the one that is specified in the allowed list, the ISL between the switch and the fabric is automatically isolated so that VSAN and the switch is denied entry into the fabric.

The persistent domain ID must be specified along with the sWWN. Domain ID authorization is required in FICON VSANs where the domains are statically configured, and the end devices reject a domain ID change in all switches in the fabric. The fabric authorization database is a list of the WWNs and associated information, such as Domain IDs of the switches that are authorized to join the fabric.

5.4.4 How to implement more than one FICON VSAN

There are default FICON Port Numbers per module, or specific FICON Port Numbers can be assigned to the ports by using the CLI or DCNM. All interfaces in a FICON VSAN must receive a unique FICON Port Number. Also, some FICON Port Numbers are not associated with a physical interface: the FICON CUP, FCIP interfaces, and port channel interfaces are all logical. To use port channels, FCIP, or FICON CUP, some FICON Port Numbers are reserved and cannot be applied to a physical interface. These numbers are typically the highest one, such as 240 - 253 (decimal).

An interface can be assigned to a FICON VSAN. By using a different CLI command or DCNM action, an interface is assigned a FICON Port Number. There is no single command or action that assigns an interface to a FICON VSAN and also assigns the FICON Port Number. It is up to the administrator to know the relationship among the following items:

- Interfaces (physical or logical)
- Assignment to each FICON VSAN
- Assignment of FICON Port Numbers

To offer a practical implementation guide when more than one FICON VSAN is configured, we show three examples about managing FICON VSANs and assigning FICON Port Numbers. For brevity, we show only the required configuration steps by using the CLI, but everything can be done with DCNM DM too.

Example 1: One FICON VSAN

Imagine that you want to create a FICON VSAN with all 48 physical interfaces on the module inserted into slot 1 of the IBM c-type modular chassis. You start by manually enabling FICON on the director, and then create VSAN 100 with domain 50 and make it a FICON VSAN. You use the NX-OS CLI and not the FICON setup wizard or DCNM DM.

Complete the following steps:

1. At the switch prompt, enter the following string:

```
Switchname#configure terminal
Switchname(config)#feature ficon
Switchname(config)#feature fabric-binding
Switchname(config)#vsan database
Switchname(config-vsan-db)#vsan 100
Switchname(config-vsan-db)#exit
Switchname(config)#in-order-guarantee vsan 100
Switchname(config)#fcdomain domain 50 static vsan 100
Switchname(config)#fabric-binding database vsan 100
Switchname(config)#fabric-binding activate vsan 100 force
Switchname(config)#ficon vsan 100
Switchname(config-ficon)#active equals saved
Switchname(config-ficon)#exit
```

2. VSAN membership is managed in the CLI through the **vsan database** subcommands. The interfaces should be shut (disabled) before moving them in case there is any I/O activity.

   ```
   Switchname# config t
   Switchname(config)# interface fc1/1-48
   Switchname(config-if)# shut
   Switchname(config-if)# vsan database
   Switchname(config-vsan-db)# vsan 100 interface fc1/1-48
   ```

   It is possible to view the resulting interface assignments to the FICON VSAN by running the following command:

   ```
   Switchname# show vsan 100 membership
   vsan 100 interfaces:
       fc1/1            fc1/2            fc1/3            fc1/4
       fc1/5            fc1/6            fc1/7            fc1/8
       fc1/9            fc1/10           fc1/11           fc1/12
       fc1/13           fc1/14           fc1/15           fc1/16
       fc1/17           fc1/18           fc1/19           fc1/20
       fc1/21           fc1/22           fc1/23           fc1/24
       fc1/25           fc1/26           fc1/27           fc1/28
       fc1/29           fc1/30           fc1/31           fc1/32
       fc1/33           fc1/34           fc1/35           fc1/36
       fc1/37           fc1/38           fc1/39           fc1/40
       fc1/41           fc1/42           fc1/43           fc1/44
       fc1/45           fc1/46           fc1/47           fc1/48
   ```

3. Assign the FICON Port Numbers. The FICON Port Numbers for physical interfaces are assigned for the entire module. The FICON Port Numbers for logical interfaces are assigned for the entire VSAN.

   ```
   Switchname# config t
   Switchname(config)# ficon slot 1 assign port-numbers 0-47 force
   Switchname(config)# ficon logical-port assign port-numbers 240-253
   ```

   We can view the FICON Port Number assignments by running the following command:

   ```
   Switchname# show ficon port-numbers assign
   ficon slot 1 assign port-numbers 0-47
   ficon logical-port assign port-numbers 240-253
   ```

Example 2: Two FICON VSANs

Imagine that you want to split this switching module into two FICON VSANs: VSAN 100 with domain 50 that uses ports 1 - 24 (the left half of the module) and VSAN 101 with domain 51 that uses ports 25 - 48 (the right half).

Complete the following steps:

1. Building on the previous configuration, you create the new FICON VSAN and move to it to the ports on the right half of the module by using the following string:

   ```
   Switchname#configure terminal
   Switchname(config)#vsan database
   Switchname(config-vsan-db)#vsan 101
   Switchname(config-vsan-db)#exit
   Switchname(config)#in-order-guarantee vsan 101
   Switchname(config)#fcdomain domain 51 static vsan 101
   Switchname(config)#fabric-binding database vsan 101
   Switchname(config)#fabric-binding activate vsan 101 force
   Switchname(config)#ficon vsan 101
   ```

```
Switchname(config-ficon)#active equals saved
Switchname(config-ficon)#exit
Switchname(config)# interface fc1/25-48
Switchname(config-if)# shut
Switchname(config-if)# vsan database
Switchname(config-vsan-db)# vsan 101 interface fc1/25-48
```

2. To view interface assignments to the new FICON VSAN, run the following command:

```
Switchname# show vsan 101 membership
vsan 101 interfaces:
    fc1/25    fc1/26    fc1/27    fc1/28
    fc1/29    fc1/30    fc1/31    fc1/32
    fc1/33    fc1/34    fc1/35    fc1/36
    fc1/37    fc1/38    fc1/39    fc1/40
    fc1/41    fc1/42    fc1/43    fc1/44
    fc1/45    fc1/46    fc1/47    fc1/48
```

3. Assign the FICON Port Numbers. The FICON Port Numbers are unique to each FICON VSAN, so the same FICON Port Number can be repeated if the ports are in different FICON VSANs. When you use the NX-OS CLI, this task is something the administrator must manage. The two ranges of FICON Port Numbers for a module are comma-separated.

```
Switchname# config t
Switchname(config)# ficon slot 1 assign port-numbers 0-23,0-23 force
```

You can now view the FICON Port Number assignments:

```
Switchname# show ficon port-numbers assign
ficon slot 1 assign port-numbers 0-23, 0-23
ficon logical-port assign port-numbers 240-253
```

Example 3: Port channels as FICON ISLs and FICON port numbers

A port channel is a group of ports that work as a single entity. Making a port channel of physical ISLs between two physical FICON Directors:

- Combines the capacity of all the ISLs.
- Provides redundant paths that fail over without loss of I/O.
- Makes them sharable by multiple FICON VSANs through VSAN trunking.

Like all ports in a FICON VSAN, a port channel must have a FICON Port Number. Recall that a port channel is logical, so its FICON Port Number must come from the range of reserved Port Numbers. To determine the FICON Port Number for the ISL Port-Channel, use the following string:

```
Switchname# show ficon port-numbers assign
ficon slot 1 assign port-numbers 0-23, 0-23
ficon logical-port assign port-numbers 240-249
Switchname# show ficon first-available port-number
Port number 240(0xf0) is available
```

Now, you can define a FICON port channel on the first two physical interfaces of the module and assign the first available FICON Port Number to it by using the following string:

```
Switchname# show port-channel usage
no port-channel number used
Switchname# config t
Switchname(config)# interface port-channel 9
Switchname(config-if)# channel mode active
Switchname(config-if)# ficon portnumber 0xf0
Switchname(config-if)# switchport mode e
Switchname(config-if)# switchport trunk allowed vsan 1
Switchname(config-if)# switchport trunk allowed vsan add 100-101
Switchname(config-if)# interface fc1/1-2
Switchname(config-if)# shut
Switchname(config-if)# switchport mode e
Switchname(config-if)# switchport trunk mode on
Switchname(config-if)# channel-group 9 force
```

5.4.5 FICON configuration files

FICON switches require a FICON configuration file to operate. You can save up to 16 FICON configuration files on each FICON-enabled VSAN and in persistent storage. The file format follows IBM requirements. These files can be read and written by mainframes by using the in-band CUP protocol, or they can be modified by using the NX-OS CLI or DCNM GUI.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, that is, an IPL. This file is created with a default configuration immediately after FICON is enabled in a VSAN. Multiple FICON configuration files with the same name can exist in the same switch if they belong to different FICON VSANs. When FICON is disabled on a VSAN, all its FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration options for each implemented FICON Port Address:

- Block / Unblock flag
- Prohibit / Allow mask
- Port Address name

You cannot prohibit or allow an ISL, a port channel, or an FCIP interface. If an interface is configured in E or TE mode and you try to prohibit that port, the prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode or TE mode. You cannot block or prohibit the CUP port (0XFE).

You cannot directly assign a FICON Port Number to a physical interface. You assign a range of FICON Port Numbers to the module slot by using the following command:

```
ficon slot 1 assign port-numbers 0-47
```

However, it is possible to assign FICON Port Numbers to logical interfaces such as FCIP interfaces and port channels:

```
interface port-channel 9
ficon portnumber 0xf0
```

5.4.6 Code version selection in FC and FICON intermix environments

Intermixing FICON and FC is a popular option in mainframe environments. Support of Linux on IBM Z is a major driver of this trend. The performance, scalability, and fabric-level virtualization capability that is provided by IBM c-type FICON switches make them ideally suited for this kind of deployment. Taking advantage of the high number of FC and FICON VSANs that are supported on IBM c-type switches, some organizations have moved beyond simple intermix and realized a mainframe-centric private cloud, where virtualization on servers, applications, switches, and storage devices is center stage.

IBM c-type networking devices require NX-OS 8.1.1b or later for FICON environments and NX-OS 8.3.2 version or later for FC. They can interoperate with Cisco Multilayer Director Switch (MDS) 9000 devices, even if they use older and different NX-OS versions. A complete interoperability matrix can be found at Interoperability Matrix for Cisco Nexus and MDS 9000 Products.

Best practice: In general, it is recommended to have the same NX-OS release on all devices that are part of the same fabric.

The recommended NX-OS releases for both FC and FICON protocols are documented and regularly updated at Recommended Releases for Cisco MDS 9000 Series Switches.

Recommended releases are based on field-proven evidence of stability and lack of significant issues. Because it takes some time before a statistically significant installation base is proven to operate with no issues, the recommended NX-OS release is almost always not the latest version. At times, customers that are highly concerned by security-related aspects might decide to deploy the most recent NX-OS version, even if it is not the recommended one. In general, FICON releases go through the longest and most extensive testing and qualification efforts and meet the highest-quality standards.

FICON capabilities enhance certain models of IBM c-type switches by supporting both open systems and mainframe storage network environments. When FICON is configured on devices, only those NX OS releases that are FICON qualified should be considered for code upgrades. In simple terms, FICON deployments tend to be more prescriptive in terms of what NX-OS versions are allowed.

Here is an example: NX-OS Release 8.1(1b) and Release 8.4(1a) are IBM-qualified FICON releases for IBM c-type devices. From the hardware point of view, NX-OS Release 8.4(1a) introduces FICON support on the IBM Storage Networking SAN192C-6 Crossbar Fabric-3 Switching Module (DS-X9706-FAB3), IBM Storage Networking SAN384C-6 Crossbar Fabric-3 Switching Module (DS-X9710-FAB3), and Supervisor-4 Module (DS-X97-SF4-K9). Several new capabilities also have come to fruition with the new software version. Customers that use NX-OS 8.1(1b) release might decide that they need to upgrade their switching infrastructure to NX-OS 8.4(1a) release. Is any intermediate NX-OS release required to perform the upgrade? Is there any impact on traffic during the upgrade?

In this example, no intermediate NX-OS release is required, and no intermediate NX-OS release should be considered (the upgrade from NX-OS Release 8.1(1b) to NX-OS Release 8.4(1a) should be direct). There is no intermediate FICON qualified release between the two. The reverse path is used for fallbacks.

Note: Always check the NX OS release notes documents for details.

Figure 5-42 shows the correct upgrade path for FICON on qualified IBM c-type devices.

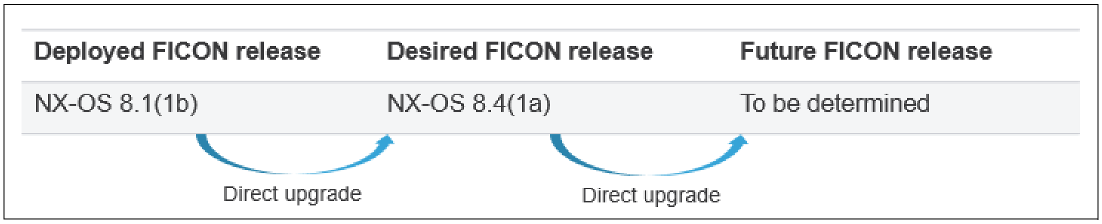

Deployed FICON release	Desired FICON release	Future FICON release
NX-OS 8.1(1b)	NX-OS 8.4(1a)	To be determined

Figure 5-42 Upgrade path for FICON qualified releases

FC and FICON ports can be upgraded nondisruptively, that is, without affecting traffic in any way. However, traffic on 1/10 Gbps Ethernet ports (IP storage ports) is momentarily disrupted during an upgrade, which includes IP storage ports on the IBM Storage Networking SAN50C-R and IBM 24/10 Port SAN Extension Module. Nodes that are members of VSANs traversing an FCIP ISL are impacted.

When operating in a FC/FICON intermix environment, you can mix FCP and FICON traffic on a single VSAN and use zoning to separate the two types of traffic, but it is not a best practice. In fact, one of the advantages of using IBM c-type switches is the intrinsic ability to separate FICON and FCP traffic into separate VSANs, which is considered a best practice. Using separate VSANs provides several benefits:

- Isolation is improved.
- VSAN-based roles for administrative access can be created.
- In-order delivery can be set per VSAN.
- Load-balancing behavior can be set per VSAN.
- Default zoning behavior can be set per VSAN.
- Persistent FCIDs can be set per VSAN.
- Domain ID allocation (static or dynamic) behavior can be set per VSAN.
- FC timers can be set per VSAN.

When operating in a FC/FICON intermix environment, it is important to evaluate the implications. Often, the intermix deployment leads to a footprint and cost reduction, both in terms of purchasing and operational costs. Some large and highly sensitive corporations have adopted this intermix approach. However, it imposes some limitations on the NX-OS versions that should be used. The NX-OS version works at the switch level, not the VSAN level, which means that you cannot have one NX-OS version for FICON VSAN and a different version for FC VSAN on the same chassis. You may have different NX-OS versions on different chassis. When performing NX OS upgrades in a FC/FICON intermix environment, the stricter FICON rules apply, and only FICON qualified releases should be deployed. Typically, there are two FICON qualified NX-OS releases per annum.

5.5 Topologies

This section covers the following topics:

- Protocols and network topologies
- Resiliency and redundancy in Fibre Channel networks
- FICON topologies

5.5.1 Protocols and network topologies

This section covers high-level network design considerations for different protocols. A typical network is composed of end devices that must communicate, and the switches and cabling that connects the devices together. Topologies are usually described in terms of how the switches are interconnected. Network topologies differ largely depending on transport protocol and use case. In general, Ethernet topologies show the largest variety. Ethernet networks can be found in the data center, campus, or across a WAN. Due to the different requirements, the resulting topology differs.

Figure 5-43 provides a quick glimpse at the diversity for Ethernet networks depending on where they are deployed.

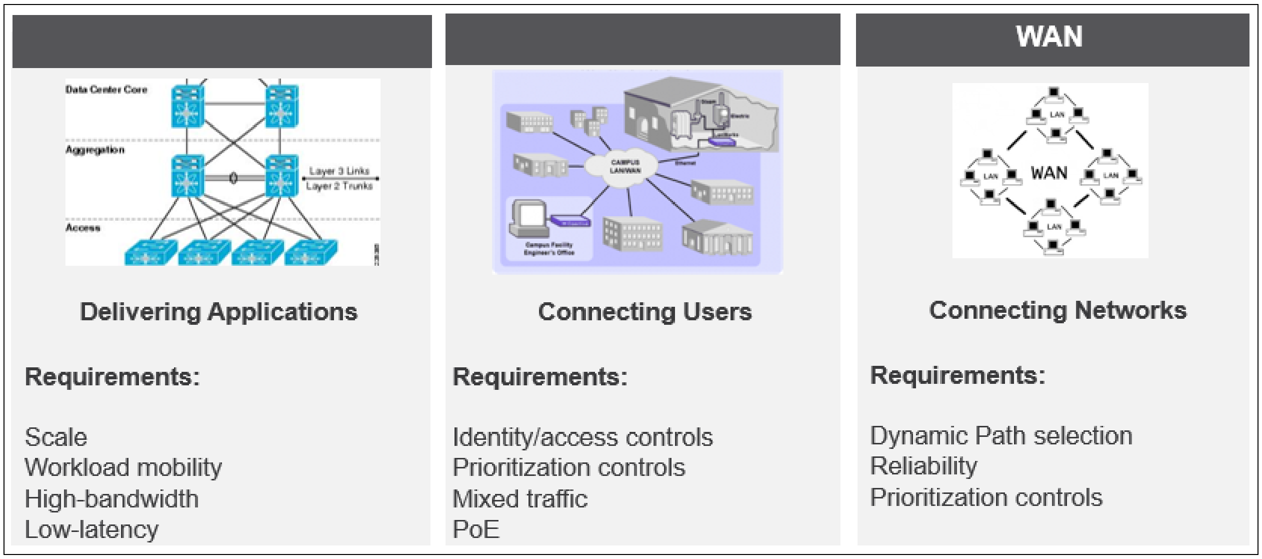

Delivering Applications

Requirements:

Scale
Workload mobility
High-bandwidth
Low-latency

Connecting Users

Requirements:

Identity/access controls
Prioritization controls
Mixed traffic
PoE

Connecting Networks

Requirements:

Dynamic Path selection
Reliability
Prioritization controls

Figure 5-43 Ethernet networks in the campus, data center, and WAN

In general, Ethernet network topologies are based on the combination of some basic models. Within data centers, the most common approach is the spine-leaf topology, which gradually replaced the N-tiered topology of the past decade.

Figure 5-44 represents various possible Ethernet topologies, including the spine-leaf approach.

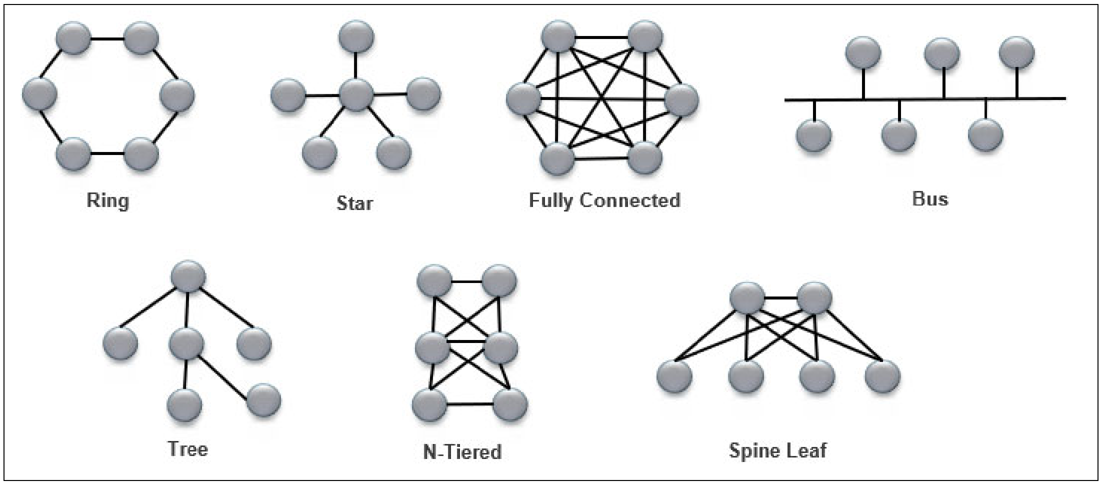

Figure 5-44 Variety of Ethernet network topologies

FC networks are built and optimized for storage. Being single-purpose, FC networks are suitable for specific optimizations that make them ideal for connecting hosts to their data. Traffic patterns are defined and traffic always flows from initiators (servers) to targets (storage) and vice-versa. FC networks also are within the perimeter of data centers and occasionally must cross a WAN to reach another data center. The scale of FC networks is also limited compared to their Ethernet counterparts, and having 10,000 ports in a single fabric is unique.

Based on the above considerations, it is easy to guess that FC topologies tend to be like each other and follow a determined schema. The goal is simple: You want scalable, reliable, and easy to operate networks with a lower capital investment.

The most typical topology is the *core-edge* design, where initiators connect to the edge and targets to the core. An appropriate quantity of ISLs are used between the core and edge switches to ensure that there is enough bandwidth to serve the workloads. If possible, with all flash arrays, the aggregate bandwidth reaching the target devices should be less or equal to the aggregate bandwidth on ISLs.

The core-edge topology is like the leaf-spine topology for Ethernet, but the traffic patterns are different. The core-edge topology is the recommended SAN topology to optimize performance, management, and scalability. With the servers and storage on different switches, this topology provides ease of operations and flexible expansion. Good performance is possible because initiator to target traffic always traverses a single hop from the edge to the core, making the solution deterministic.

In the past, the *edge-core-edge* topology was used to satisfy even larger port counts for a single fabric. Targets were connected to a set of edge switches, initiators were connected to another set of edge switches, and core switches established the communication paths. This topology, despite being scalable, comes with some drawbacks:

- The number of ports that are used on ISLs becomes significant, and the ratio of useful end ports to total ports goes down.
- End-to-end latency is higher because of the extra hop, and congestion on ISLs might occur.
- Troubleshooting becomes more complex when something negative happens, and a slow-drain phenomenon might occur.

The need for edge-core-edge topologies was reduced high port count switches and directors became available. For smaller size networks, the collapsed core-edge topology is often used. In this case, the network is a single switch in the form of a modular device where both initiators and targets are connected. There is no ISL.

Other network topologies exist. For example, a fully connected MeSH topology is sometimes adopted, but can lead to location-dependent, asymmetric performance levels. A tree topology also is sometimes used, and it is the result of a poor initial planning and inorganic growth over multiple years. In the end, the choice of the best topology depends on many actors, such as the workloads to be served, fixed versus modular platforms, cost, scalability, and oversubscription level.

Figure 5-45 shows the most common FC network topologies.

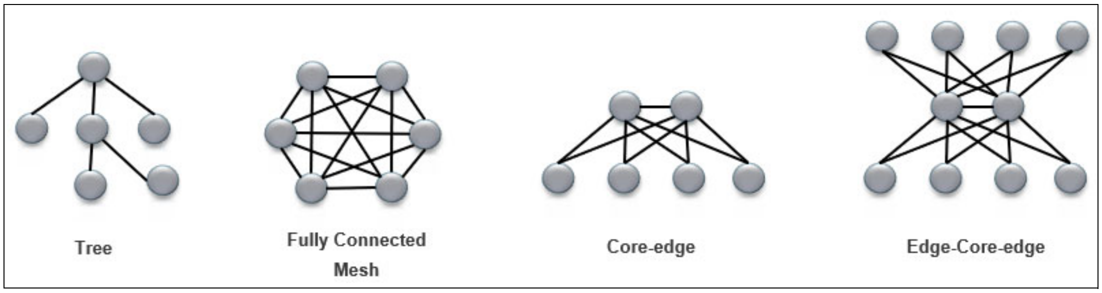

Figure 5-45 Fibre Channel common topologies

FICON topologies are even simpler. The reduced scale of FICON deployments and the need for extreme levels of availability and predictability leads to a defined set of topologies. Moreover, FICON environments require full qualification, and only qualified topologies can be used in production deployments.

5.5.2 Resiliency and redundancy in Fibre Channel networks

Two important aspects of any FC and FICON topology are resiliency and redundancy. The main objective for storage architects is to remove any SPOFs. *Resiliency* refers to the self-healing capability of the network, that is, its ability to automatically recover from a failure and still continue to function. *Redundancy* refers to the duplication of components, network elements, and even an entire fabric to eliminate any SPOFs in the network. Redundancy is the key to HA and enterprise-class installations.

For organizations that want to achieve business continuance under both foreseeable and unforeseeable circumstances, the elimination of any SPOFs should be a top priority, which is why a share-nothing approach should be used at the highest level of fabric design: The whole network should be redundant, with two separate fabrics and no network equipment in common. The use of logical segmentation (VSAN technology) on top of a single physical fabric can protect from human errors and other software-related issues, but it cannot provide the same degree of availability as two physically separated infrastructures.

Figure 5-46 contrasts and compares a physically redundant network with a logically redundant network.

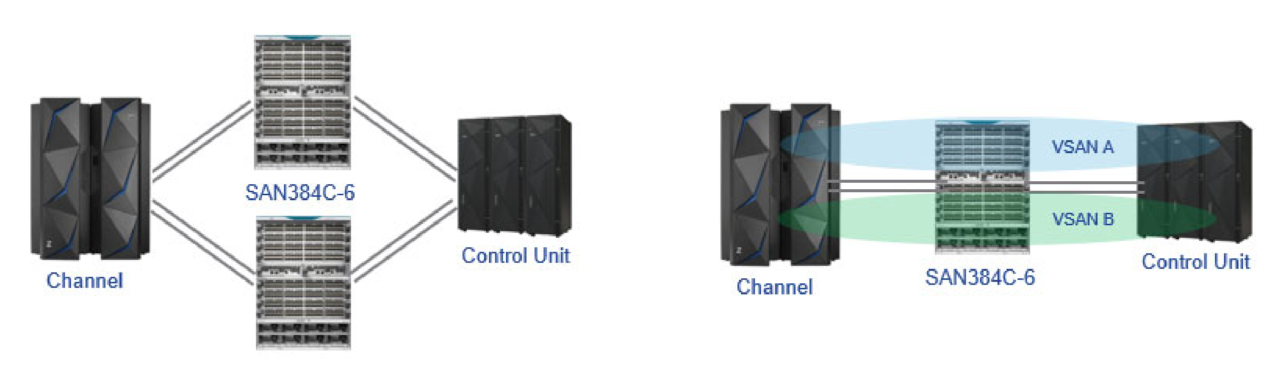

Figure 5-46 Physical versus logical redundancy

Servers and storage devices should be connected to both physical fabrics. Data traffic should flow across both networks transparently in either active-active or active-passive mode, depending on the settings that are applied to the multipath I/O (MPIO) solution.

MPIO is responsible for helping ensure that if one path on a host fails, an alternative path is readily available. Ideally, the two fabrics should be identical, but during migrations, differences in the way the fabric networks are designed and in the products that are used to build them are common. Generally, these two fabrics, which are identified as SAN A and SAN B, are in the same location. However, to provide greater robustness at the facility level, they are sometime kept in separate data center rooms.

Enterprises also might rely on secondary data centers to achieve business continuance or DR, depending on the distance between data centers and the recovery point objective (RPO). Using two fabrics locally and two separate locations within the territory provides an excellent approach for achieving complete redundancy.

Figure 5-47 exemplifies a typical redundant deployment in a dual-site FICON environment.

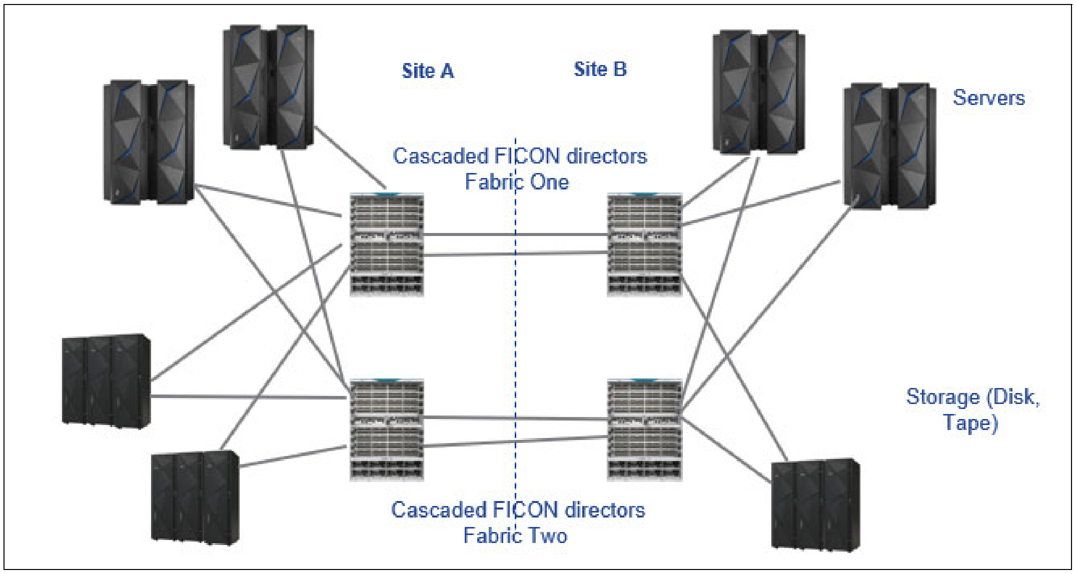

Figure 5-47 Redundant FICON deployment

In addition to redundant fabrics operating in parallel, HA within each individual fabric is required. It is a best practice to use redundant physical links between switches. A minimum of two ports, on two different line cards, should be used for ISLs, and they should be part of a logical bundle in an FC port channel. With this setup, the storage network administrator can take down one of the member links for diagnostic purposes without disrupting the traffic on the remaining FC port channel members.

This level of redundancy on ISLs also prevents fabric segmentation even if a link shuts down under a failure condition. Active mode should be preferred as a setting for port channels, and is the default on IBM c-type switches with NX-OS 8.4.1 and later so that recovery occurs automatically without the explicitly enabling and disabling the port channel member ports at either end of the link.

SAN extension line cards should be redundant within a single mission-critical director, and traffic should be shared among them. An IBM c-type mission-critical director can be filled with SAN extension line cards with no limitation on the number that can be accommodated. Members of the same FC port channel should be placed on different line cards, on different ASICs, or in different port groups whenever possible. Creating a logical bundle across ports that are served by the same ASIC has no positive effect on network availability and should not be considered as a best practice.

To improve operational ease and achieve greater availability, configuration and cabling should be consistent across the fabric. For example, do not configure ISLs at the upper left ports in one chassis and on the lower right ports in another chassis: mirrored configurations are recommended.

Here is a list of best practices for proper and reliable SAN design to help ensure application availability on FC networks:

- Avoid a SPOF by using share-nothing redundant fabrics.
- Use MPIO-based failover for server-to-storage connectivity by using redundant fabrics.
- Use redundancy features that are built in to individual fabrics.
- Use mirrored cabling and configurations for ease of operation and troubleshooting.

- Use a core-edge topology with separate storage and server tiers for independent expansion.
- Core switches should be mission-critical directors whenever economically viable.
- Spread port channel members across line cards and use the active-mode setting.
- Use the highest performing and most resilient switch in the fabric (typically a core mission-critical director) as the principal switch. Point to it for zoning changes, and use it as the seed for management tools.
- Always use static domain IDs.
- Enable and configure optional features that help prevent user misconfiguration by using checking, alerting, and monitoring functions.

IBM c-type fabrics have resiliency features that are built in that are derived from the NX-OS OS, which is the software that runs on all IBM c-type switches. The self-healing capabilities of NX-OS can quickly overcome most failures and repair the network. For example, when a link between switches fails, the FC port channel technology immediately moves all traffic flowing through that member link to the surviving member links. If an entire FC port channel fails (very unlikely), the FSPF process immediately recalculates the distribution of all traffic flows. All these functions require a second route to be available by using redundancy that is built in to the fabric design.

NX-OS also includes other capabilities that help make networks that are built by using IBM c-type devices resilient and HA. For example, processes can be gracefully shut down and restarted. VSANs isolate traffic flows at the hardware level to the point that a misconfiguration in the zoning database in a VSAN does not affect the others. When an individual port is administratively shut down, the process occurs gracefully, with buffers cleared and no packets lost.

A HA storage network must be paired with a HA storage infrastructure. Data must be available when it is accessed. Several approaches can be used for this purpose, including Redundant Array of Independent Disks (RAID) implementations, multiple copies of data spread across a clustered system, data replication over distance, and tape backup.

For a deeper understanding of resiliency and redundancy in FC networks, see *Design a Reliable and Highly Available Fibre Channel SAN*.

5.5.3 FICON topologies

FICON and zHPF deployments always are based on a limited number of switches per fabric (generally directors for higher reliability), so the possible topologies are limited too. Of course, redundant fabrics are the rule. LPARs on the fabric (with VSANs) are supported, even for FICON environments. A maximum of eight FICON VSANs is allowed, but FC VSANs have a maximum of 80.

A FICON channel in FICON native (FC) mode uses the FC communication infrastructure that is supported by IBM Z to transfer channel programs (CCWs) and data through its FICON/FICON Express adapter to another FICON adapter node, such as a storage device, printer, or another server (CTC).

A FICON native (FC) mode channel can operate in one of four topologies:

- Simple point-to-point (a mainframe host FICON channel is directly connected to a FICON capable CU). This deployment approach has a limited scale and does not include any IBM c-type switches.
- Switched point-to-point (through a single FC FICON capable switch to a FICON-capable CU). This deployment is widely adopted in single data center scenarios.
- Cascaded FICON (through two FC FICON capable switches to a FICON capable CU). This topology may include FCIP links.
- Cascaded FICON Multihop (through up to four FICON switches). This topology can include a single FCIP hop and some variance in the way it is built.

When in a cascaded configuration, up to 16 ISLs between the two adjacent FICON Directors may be grouped and become a FICON port channel.

Simple point-to-point topology

For small environments, the mainframe can be directly connected to its storage. This approach was the only one that was available before FICON was introduced. In a point-to-point connection, the FC link is between the processor’s FICON channel card (N_Port) and the FICON adapter (N_Port) in the CU. Despite of the obvious limitations, this simple point-to-point topology is not going away. In recent years, it received further validation. The IBM zHyperLink Express is a new mainframe I/O channel link option that was introduced with z14 and is based on a PCIe 3.0 bus. Essentially it is a short-distance (up to 150 meters) direct point-to-point connection between the mainframe and its IBM DS8880 storage, with the specific intent to reduce I/O latency and improve I/O throughput. This new connectivity option uses a 24-fiber ribbon cable with Multi-fiber Termination Push-on (MTP) connectors. zHyperLink is complementing and not replacing FICON connectivity.

A channel path that consists of a single link interconnecting a FICON channel in FICON native (FC) mode to one or more FICON CU images (logical CUs) forms a point-to-point configuration. A point-to-point configuration is permitted between a channel and CU only when a single CU is defined on the channel path or when multiple CU images (logical CUs) share an N_Port in the CU. The channel N_Port and the CU N_Port are responsible for managing the access to the link among the logical images. A maximum of one link can be attached to the channel in a point-to-point configuration. The maximum number of CU images that is supported by the FICON architecture over the FC link to CU is 256, so the maximum number of devices that can be addressed over a channel path that is configured point-to-point is equal to 256 times 256, or 65,536.

Figure 5-48 represents a simple point-to-point topology:

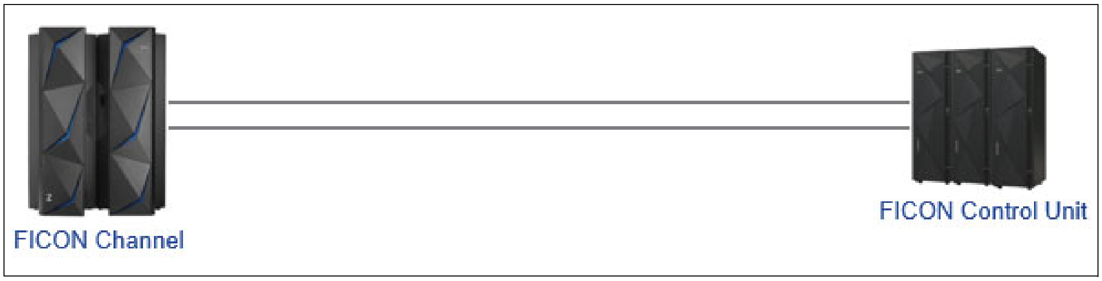

Figure 5-48 Simple point-to-point topology

The FICON channel determines whether the link that it is connected to is in a point-to-point or switched topology by logging in to the fabric, fabric login (FLOGI ELS), and checking the accept response to the fabric login (ACC ELS). The FLOGI - ACC (accept) response indicates whether the channel N_Port is connected to another N_Port (point-to-point) or a fabric port (F_Port).

Switched point-to-point topology

The simpler switched topology has a single FICON Director per fabric, and both mainframe and storage are directly connected to it. This switched topology was the only one that was supported when FC-SB-2 ULP was used until 2003. In a switched point-to-point connection, there is one FC link between the FICON channel card (N_Port) and the FICON Director (F_Port), and another FC link between the FICON Director (F_Port) and FICON adapter in the CU (N_Port).

Figure 5-49 represents a switched topology.

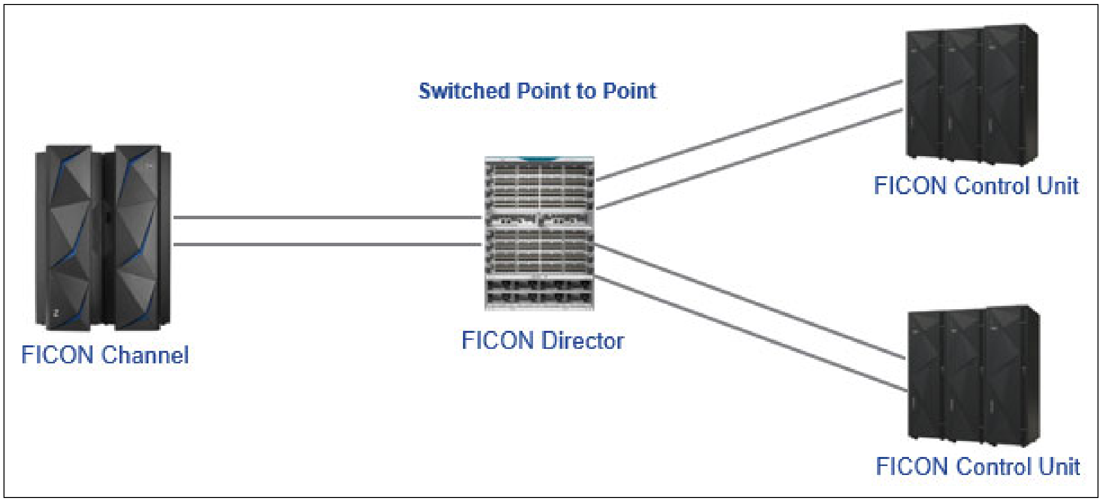

Figure 5-49 Switched topology

Multiple channel images and multiple CU images can share the resources of the FC link and the FC switch so that multiplexed I/O operations can be performed. Channels and CU links can be attached to the FC switch in any combination, depending on the configuration requirements and available resources in the FC switch. Sharing a CU through an FC switch means that communication from several channels to the CU can take place either over one switch to CU link (when a CU has only one link to the FC switch) or over multiple link interfaces (when a CU has more than one link to the FC switch). Only one FC link is attached to the FICON channel in a FICON switched point-to-point configuration, but from the switch the FICON channel can communicate with (address) several FICON CUs on different switch ports.

Cascaded FICON topology

All other supported switched topologies are like each other and generically called *cascaded topologies*. Cascaded FICON refers to an implementation where storage fabrics are linked through connections between pairs of FICON Directors or switches.

In a cascaded FICON topology, at least three FC links are involved: One is between the FICON channel on the mainframe and the local FICON Director; the second is between the FICON Directors; and the third is from the remote FICON director and the CU.

The connections between switching devices are called ISLs, and the corresponding ports on a director are called E_Ports. ISLs are flexible and can support processor-to-processor, processor-to-disk or tape subsystem, and subsystem-to-subsystem logical switched connections.

Cascaded FICON topologies provide several benefits:

- They facilitate the design and implementation of robust DR and business continuity solutions, such as GDPS.
- They reduce the infrastructure costs and complexities that are associated with these implementations.
- They introduce greater flexibility in the FICON architecture, a more effective utilization of fiber links, and higher data availability in the enterprise.

Figure 5-50 represents a cascaded FICON topology.

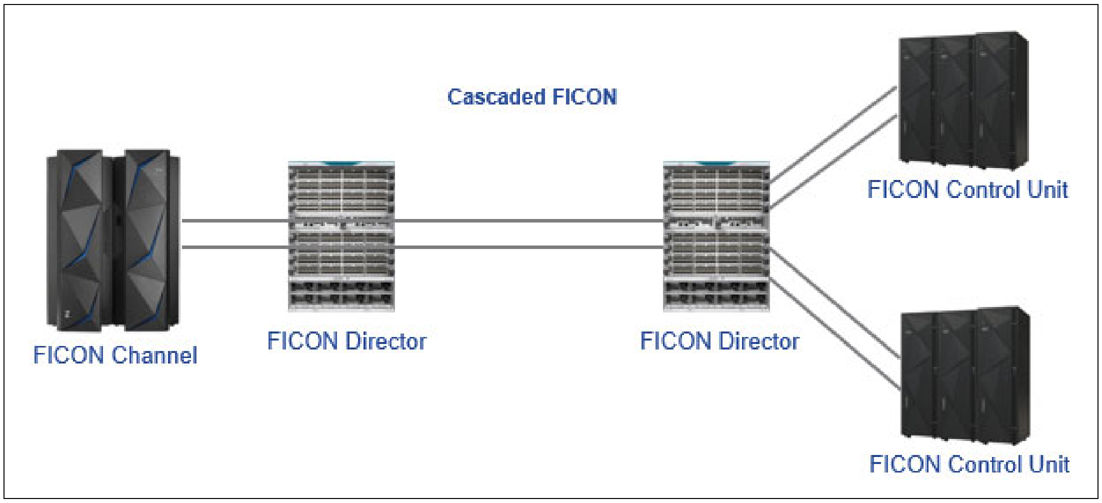

Figure 5-50 Cascaded FICON topology

The FICON channel in FICON native (FC) mode supports multiple concurrent I/O connections. Each of the concurrent I/O operations can be to the same FICON CU (but to different devices) or to a different FICON CU.

For cascaded connections, the HCD defines the relationship between channels and a director (switch ID) and specific switch ports. However, HCD does not define the ISL connections, and the management of the traffic over ISLs is controlled exclusively by the directors. In fact, during initialization, the directors identify their peers and create a routing table so that frames are forwarded to the correct director, which means extra ISL bandwidth can be added to a topology without any modification to the HCD definitions.

In the basic implementation, one FICON Director is connected to another one through an ISL in each fabric. A variation of this dual-device deployment uses FCIP for the ISL. Another possibility keeps the FCIP function on a dedicated pair of switches, and the director connects to them through an ISL. Even though this topology now includes four switching devices, it is still a single-hop FICON cascade deployment.

These topologies can be supported when port channels are used instead of individual ISLs.

The most recently supported cascaded topology is FICON Multihop. As the name implies, FICON Multihop allows support for cascading up to four switches, or three hops, which overcome the previous restriction of two cascaded switches or a single hop. A three-switch, dual-hop topology is also supported.

Note: FICON Multihop is supported only by using traditional static routing methods. FIDR is not supported.

Figure 5-51 represents a cascaded FICON Multihop topology.

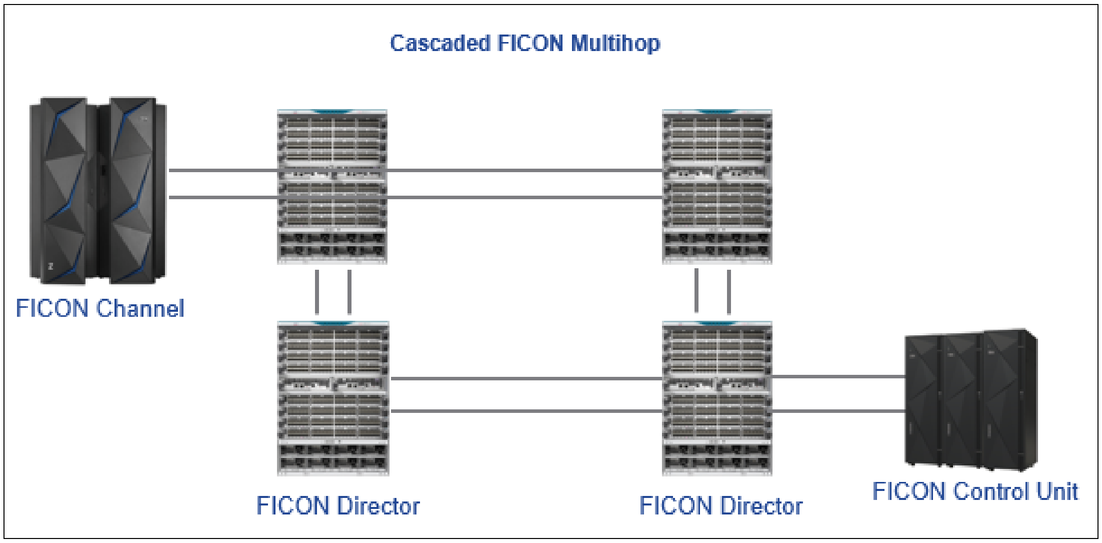

Figure 5-51 Cascaded FICON Multihop topology

FICON Multihop environments provide two benefits to users:

- The previous limit of a single hop led to some unnatural FICON switch configurations for multisite deployments. FICON Multihop allows both configuration and switch consolidation, which simplifies management.
- Availability increases. For example, using the square topology in a 4-site deployment achieves higher SAN availability in the event of a total connectivity break between two switches. The square FICON Multihop configuration allows for traffic to be routed around the failure to ensure that access is maintained to all devices and hosts that are connected to the SAN fabric.

Figure 5-52 on page 165 shows some alternative FICON Multihop topologies.

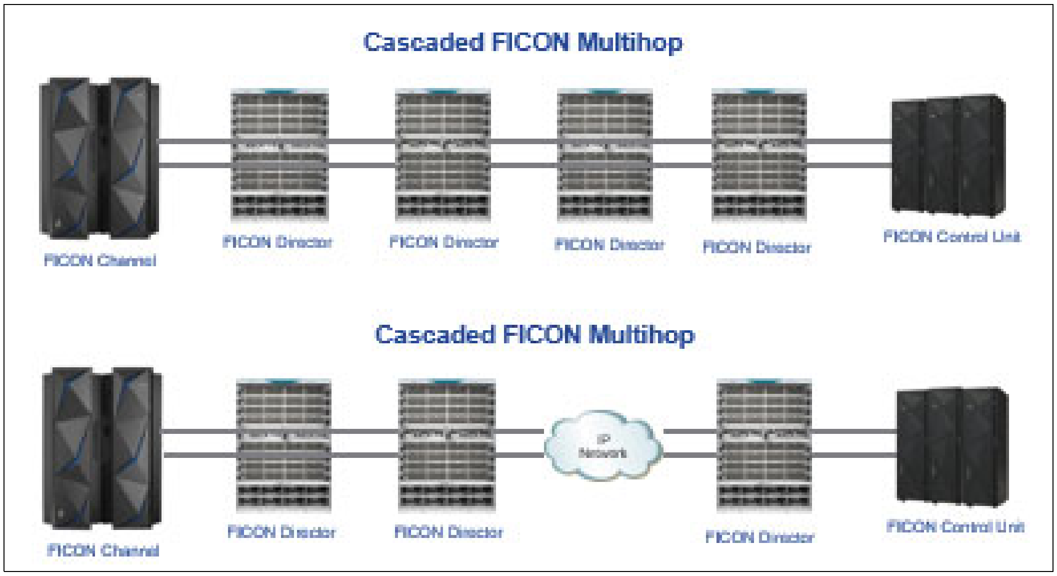

Figure 5-52 Alternative FICON Multihop topologies

There are some design implications that you should follow when implementing a FICON Multihop environment:

- Regarding bandwidth planning and allocation for failure scenarios, extra bandwidth should be provisioned on all routes to compensate for the eventual loss of connectivity between two switches and the subsequent rerouting of traffic over the nonfailed ISL paths.
- Regarding performance and the latency impacts of traffic that is rerouted to longer paths, there is little that can be done to mitigate these items, so there is a tradeoff for better availability.

FICON Multihop imposes some hardware requirements on all elements of the topology:

- The mainframe must be a z13 or later.
- The storage system must be an IBM DS8870 or later, or equivalent third-party storage system.
- The director must be an IBM c-type SAN192C-6, SAN384C-6, or SAN50C-R.
- The eventual DWDM infrastructure must be explicitly approved for this use.
- Software releases must meet a minimum version.

FICON Multihop deployments may use both native ISLs or FCIP extension networks. When native ISLs are used, the usual distance limitation of 10 km is imposed, unless colored SFPs or transponder-based DWDM equipment is adopted. When deploying long-distance FC ISLs, you might need to increase the number of buffer credits beyond the default values, which is possible by using the extended B2B credits feature with the enterprise package license. The number of buffer credits depends on distance, speed, and average frame size. For FICON traffic, consider using an average frame size of about 1 KB instead of 2 KB.

With FCIP, longer distances are possible. Only one FCIP hop is allowed per FICON Multihop configuration. Concurrently, you also must ensure that we do not exceed the FICON timeout limitations, even in the worst-case scenarios of a path loss. For both native ISL and FCIP configurations, the longest distance a FICON packet can traverse is 300 km.

Although the maximum distances for FICON and FCP protocols have been successfully tested up to 300 km, including FC ISLs and FCIP links, IBM requires an Extended Distance Request for Price Quotation (RPQ) for IBM Z and for LinuxONE to ensure that distances that are greater than 100 km adhere to the bounds of the qualification.

All these topologies are valid for the different mainframe generations, from z13 to the z15.

Direct connections between IBM Z FICON adapters and IBM c-type 32 Gbps optics can operate only at speeds of 16 Gbps and 8 Gbps because there are no 32 Gbps IBM Z FICON adapters at the time of writing. Links running at 32 Gbps are tested over ISLs between IBM c-type Directors. Targets can run at 32 Gbps too.

For more information about using FICON Multihop, its requirements, and supported configurations, see *FICON Multihop Requirements and Configurations*.

Intermixing FC/FICON topologies

All indicated topologies also can be used in a FC/FICON intermix solution. IBM c-type switches have additional flexibility when used this way, as shown in Figure 5-53.

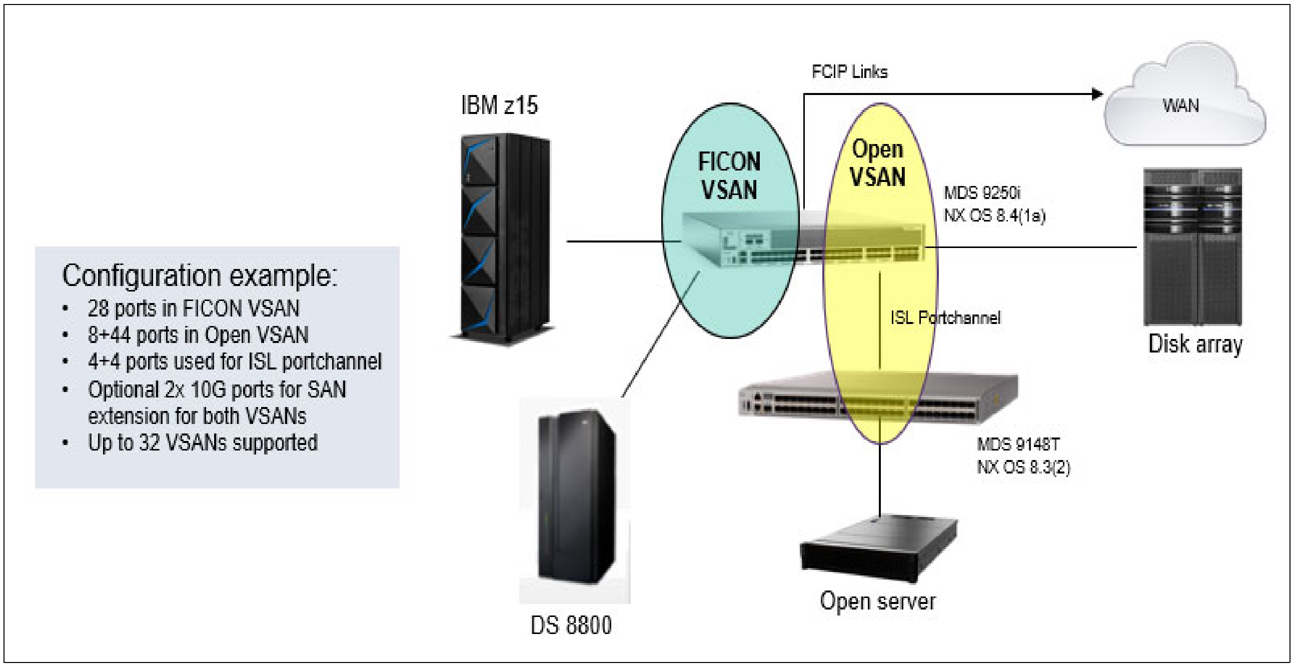

Figure 5-53 Port expansion in a FC/FICON intermix deployment scenario

The IBM Storage Networking SAN50C-R switch is qualified for FICON and offers up to 40 FC ports. It is typically used for SAN extension, but it also can be considered for local switching inside a data center.

Mainframe environments can be sized at deployment phase, but this situation is not always true for FC environments, so it is good to have some more port expansion flexibility. In our mixed FC/FICON deployment with IBM Storage Networking SAN50C-R, we could initially allocate 28 ports to FICON and eight ports to FC, leaving four ports unused. One VSAN is configured for FICON and one for FC. If we need more ports for the FC VSAN, we can create an ISL port channel between this device and another switch, like the IBM Storage Networking SAN48C-6. The new switch has no port in the FICON VSAN, and it is not FICON qualified. However, the proposed topology is valid and supported and also can be used for NVMe/FC traffic on the FC VSAN.

NX-OS releases can be different on the two switches. A positive side-effect of this solution is that it enables 32 GbE traffic on most ports of the FC VSAN while offering 2 GbE FICON compatibility to tape libraries.

5.6 Inter-Switch Links and FICON routing options

When more than one FICON Director is used in the same fabric, connections are required between pairs of director. These director-to-director connections are known as ISLs, and typically they are more than one on each segment.

Figure 5-54 shows an example of a single-hop, cascaded FICON topology. It has three ISLs between the FICON Directors.

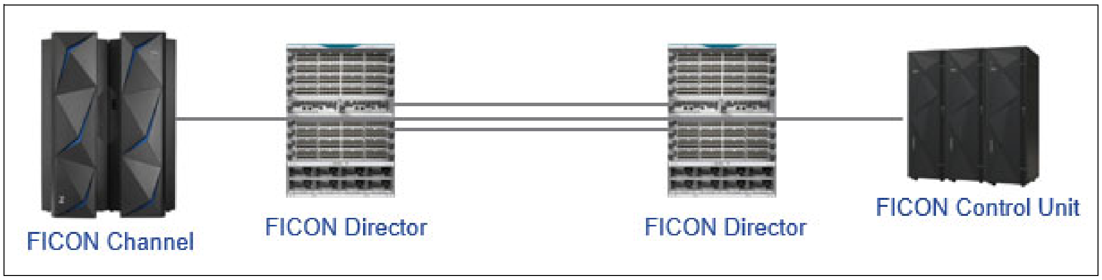

Figure 5-54 ISLs in a cascaded FICON topology

The end-to-end path from mainframe to disk includes three segments:

- The first segment starts at the FICON channel N_Port (node port) and stops at the F_Port on the FICON Director.
- The second segment connects an E_Port on the local director to an E_Port on the remote director. The link between two switches, and therefore between two E_Ports, is called an ISL, and we have three of them in our topology.
- The third segment connects the F_Port on the remote director to an N_Port on the CU subsystem.

Because we have multiple ISLs, how is the traffic distributed among them? To provide an answer, we must explain how frames are routed in an FC network.

5.6.1 Fabric Shortest Path First

Data moves through a FICON SAN fabric between storage devices and a mainframe and from switch to switch along one or more paths that make up a route. Routing policies determine the path for each frame of data. Before the FICON SAN fabric can begin routing traffic, it must discover the route that a frame should take to reach the intended destination. Route tables internal to the switches are listings that indicate the next hop to which a frame should be sent to reach a destination.

The assignment of traffic between directors over the ISLs is controlled by the director. The HCD defines the relationship between channels and directors and the specific switch port. But, the HCD does not define the ISL connections, and the distribution of traffic over ISLs is controlled exclusively by the director.

When FICON was first introduced, the only mechanism that was available for traffic distribution was the FSPF protocol. It is not our intention to go too deep into how FSPF works. Suffice to say, the FSPF protocol is the standardized routing protocol for FC (and FICON) SAN fabrics, and it is the foundation for all subsequent ISL routing mechanisms that have been introduced.

The FSPF protocol is a link-state-path selection protocol that directs traffic along the shortest path between the source and destination. The metric that is used to determine what path is the shortest is not the distance, but its administrative cost. FSPF has a notorious counterpart on IP networks that is called OSPF that follows the same school of thought. The FSPF protocol goes beyond determining the shortest route for traffic: It also detects link failures; updates the routing table; provides fixed routing paths within a fabric; and maintains the correct ordering of the frames.

After FSPF is established, it programs the hardware routing tables for all active ports on the switch. After a path is assigned to an ISL, that assignment is persistent. However, every time that a new ISL is added to the fabric, the ISL traffic assignments change. For this reason, this technique is not attractive to mainframe administrators because they cannot prescribe how the paths to a subsystem are mapped to the ISLs. In the worst case, all the paths to a subsystem might be mapped to the same ISL and overload it.

FSPF tracks the state of the links on all switches in the fabric and associates a cost with each link. The protocol computes paths from a switch to all the other switches in the fabric by adding the cost of all links that are traversed by the path, and chooses the path that minimizes the costs. This collection of the link states, including costs, of all the switches in the fabric constitutes the *topology database* or *link state database*.

The topology database is replicated and present in every switching device in the FICON SAN fabric. Each switching device uses information in this database to compute paths to its peers by using a process that is known as *path selection*. The FSPF protocol provides the mechanisms to create and maintain this replicated topology database. When the FICON SAN fabric is first initialized, the topology database is created in all operational switches. If a new switching device is added to the fabric or the state of an ISL changes, the topology database is updated in all the fabric’s switching devices to reflect the new configuration.

A Link State Record (LSR) describes the connectivity of a switch within the topology database. The topology database contains one LSR for each switch in the FICON SAN fabric. Each LSR consists of an LSR header and one or more link descriptors. Each link descriptor describes an ISL that is associated with that switch. A link descriptor identifies an ISL by the Domain_ID and output port index of the “owning” switch and the Domain_ID and input port index of the “neighbor” switch. This combination uniquely identifies an ISL within the fabric. LSRs are transmitted during fabric configuration to synchronize the topology databases in the attached switches. They are also transmitted when the state of a link changes and on a periodic basis to refresh the topology database.

Associated with each ISL is a value that is known as the link cost, which reflects the cost of routing frames through that ISL. The link cost is inversely proportional to the speed of the link: Higher-speed links are more wanted transit paths, so they have a lower cost. The topology database has entries for all ISLs in the fabric, which enables a switch to compute its least cost path to every other switching device in the FICON SAN fabrics from the information that is contained in its copy of the database.

As typical with other routing protocols, hello messages are used to establish bidirectional communication over an ISL. Hello messages are transmitted on a periodic basis on each ISL even after two-way communication is established to detect a switch or an ISL failure.

Hello messages are like heartbeats between the switching devices. The time interval between successive hellos (in seconds) is defined by a timer that is called the hello interval. This timer is configured on every switch and communicated as part of the hello messages. If a switch fails to receive a hello message within the expected timeframe, it assumes that something went wrong and removes the associated ISL from its topology database. When the ISL is restored, the switching devices must reestablish two-way communication and synchronize their topology databases before the ISL may be again used for routing frames.

After two-way communication is established through the hello protocol, the switches synchronize their topology databases. During the initial topology database synchronization, each switch sends its entire topology database to its neighbor switches. When it receives an acknowledgment, topology database synchronization is complete, and the switches are said to be “adjacent” on that ISL. The ISL is now in the “full state” and may be used for frame delivery.

Although the entire topology database is exchanged during this initial synchronization process, only updates are exchanged for the database maintenance phase to reflect eventual topology changes in the FICON SAN. This process makes the protocol more efficient and faster. The topology database must be updated whenever the state of any ISL changes (ISL failure or addition).

After the topology database is created and a switch has information about available paths, it can compute its routing table and select the paths that will be used to forward frames. FC uses a least-cost approach to determine the paths that are used for routing frames. Each ISL is assigned a link-cost metric that reflects the cost of using that link. The default cost metric is based on the speed of the ISL, but it can be administratively changed for traffic engineering purposes. When multiple paths are available to a destination but one has the lowest cost, the routing decision is easy, and that path selected by FSPF. The only way to influence this path selection is to administratively force a different cost on some of the paths.

But what happens when there are multiple paths with the same cost? Which one is selected? When there are multiple paths with the same cost to the same destination (which is typical for many cascaded FICON SAN architectures), a switching device must decide how to use these paths. The switch might select one path only (not ideal) or it might attempt to balance the traffic among all the available equal-cost paths and avoid congestion of ISLs, which is known as *equal-cost multipathing*. If a path fails, the switch may select an alternative ISL for frame delivery, which is where the different types of ISL routing options come into play. Over the past several years, multiple techniques beyond simple FSPF were introduced for routing traffic on FICON ISLs. These techniques fall under two categories: FICON static routing and FIDR.

5.6.2 Static routing

When static routing is configured, during the fabric initialization phase the switches identify their peers and create a routing table so that frames for remote subsystems may be forwarded to the correct director. In a first static routing approach, the choice of ISL route was based on the incoming port and the destination domain. The choice happened during fabric login and was a simple round-robin mechanism that is based on a first come, first served basis. Even ports that do not require remote connectivity are assigned an ISL under this model.

As a result, the data workload was not balanced across the different ISLs, which led to some ISLs being overloaded and others underutilized. In other words, this static routing approach did not leverage all the available ISL bandwidth between two switches. Early implementations demonstrated large imbalances of ISL utilization levels (above 50%), particularly when there were few CH/CU pairings. For large z/OS MM (previously known as PPRC) synchronous replication technology implementations and asynchronous GM implementations, this situation resulted in many suspends and poor replication performance. A better approach was needed.

IBM c-type SID/DID routing optimizes routing path selection and utilization based on a hash of the SID and DID of the path source and destination ports. Therefore, the ingress ports that are passing traffic locally only and not using the ISLs are not assigned to an ISL. The effect of this enhanced approach is a better workload balancing across ISLs with the guarantee that exchanges between a pair of devices would always stay in order. However, the ISL utilization level could still be unequal, depending on traffic patterns. Moreover, the routing table could change each time that the switch is initialized, leading to unpredictable and nonrepeatable results.

Despite not being perfect, this static routing approach has been used successfully for many years. It also works well with a slow-drain device. Static routing has the advantage of limiting the impact of a slow-drain device or a congested ISL to a small set of ports that are mapped to that specific ISL. If congestion occurs in an ISL, the IBM Z channel path selection algorithm detects the congestion because of the increasing initial CMR time in the in-band FICON measurement data. The CSS steers the I/O traffic away from congested paths and toward better performing paths by using the CMR time as a guide. More host recovery actions are also available for slow-drain devices.

Figure 5-55 illustrates the static SID/DID routing approach.

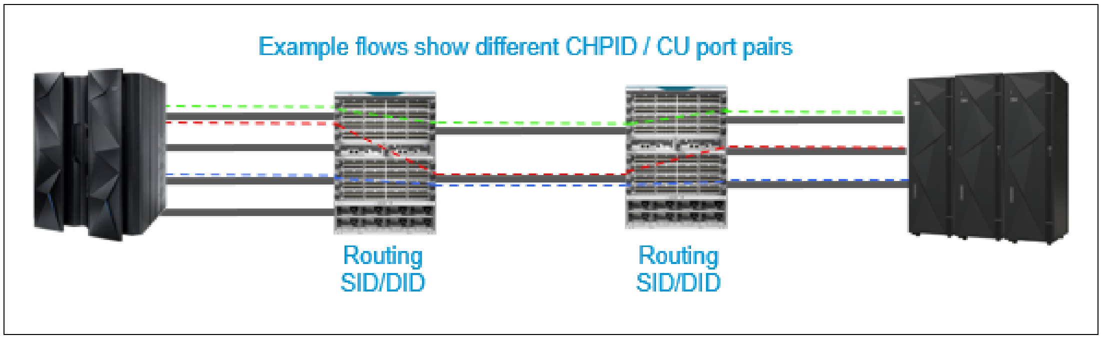

Figure 5-55   SIS/DID routing approach

5.6.3 Dynamic routing and the FIDR feature

With the z13, IBM announced that FICON channels are no longer restricted to using static routing policies in the SAN. The IBM Z servers started supporting dynamic routing in the SAN by using the FIDR feature. Despite being relatively new, this feature is leveraging the load-balancing approach that was used successfully for many years in the FC world. Users with cascaded FICON Director architectures can now rely on a routing policy where there are no fixed routes from the source ports to the destination ports. FIDR supports the dynamic routing policies that are provided by the FICON Directors in use.

Routes across ISLs are assigned to I/O operations dynamically on a per exchange (per I/O) basis. Loading of ISLs is applied at the time of the data flow, which provides the most effective mechanism for balancing data workload that traverses the available ISLs. Thus, every exchange can take a different path through the fabric. There are no longer fixed routes from the source ports to the destination ports. ISL resources are used equally, and the ISLs can run at higher utilization rates without incurring queuing delays.

High workload spikes resulting from peak period usage or link failures can also be dealt with more easily with FIDR. FIDR improves utilization of all available paths, thus reducing possible congestion on the paths. Every time that there is a change in the network that changes the available paths, the traffic can be redistributed across the available paths.

One example of such a dynamic routing policy is IBM c-type OXID routing. With FIDR, the routing assignments are based on the SID/DID and the FC OXID. Essentially, FIDR enables ISL routes to be dynamically changed based on the FC OXID parameter, which is unique for each I/O operation. With FIDR, an ISL is assigned at I/O request time, so different I/Os from the same source port going to the same destination port may be assigned different ISLs.

Figure 5-56 illustrates the SID/DID/OXID routing approach.

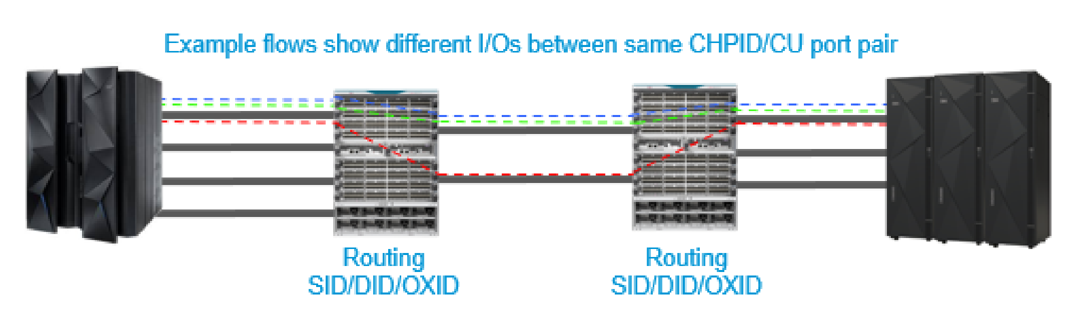
Example flows show different I/Os between same CHPID/CU port pair
Routing SID/DID/OXID
Routing SID/DID/OXID

Figure 5-56 SID/DID/OXID routing approach

The adoption of FIDR is reflected on both the NX-OS CLI and DCNM.

Figure 5-57 shows FICON VSANs with FIDR enabled as seen by DM. In fact, the LoadBalancing column indicates the SrcID/DestId/OxId schema.

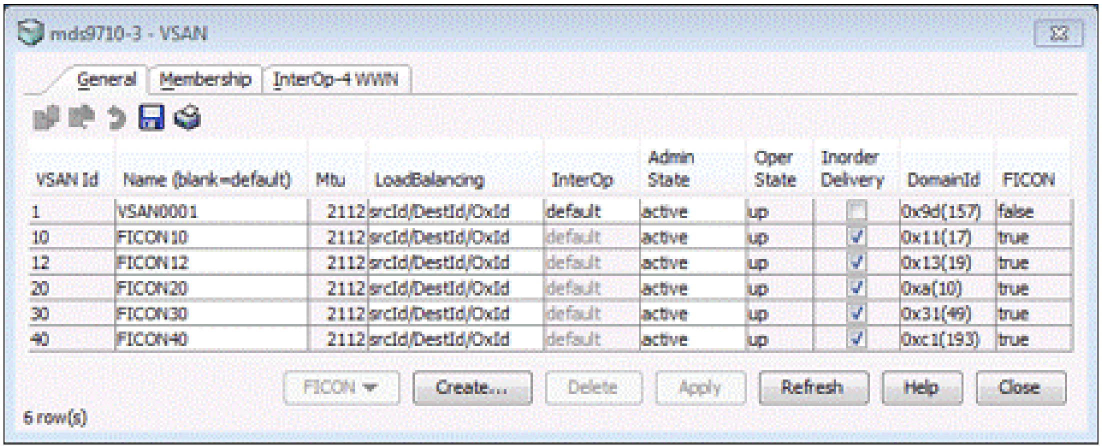

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	☐	0x9d(157)	false
10	FICON10	2112	srcId/DestId/OxId	default	active	up	☑	0x11(17)	true
12	FICON12	2112	srcId/DestId/OxId	default	active	up	☑	0x13(19)	true
20	FICON20	2112	srcId/DestId/OxId	default	active	up	☑	0xa(10)	true
30	FICON30	2112	srcId/DestId/OxId	default	active	up	☑	0x31(49)	true
40	FICON40	2112	srcId/DestId/OxId	default	active	up	☑	0xc1(193)	true

6 row(s)

Figure 5-57 FIDR on Device Manager

The characteristics of FIDR provide advantages in both performance and management for cascaded FICON configurations, depending on the specific circumstances and configuration:

- Support sharing of ISL links between FICON and FCP and consequent bandwidth consolidation.
- I/O traffic is better balanced between all available ISLs, and there is a better utilization of the available bandwidth on ISLs.
- Easier to manage with a predictable and repeatable I/O performance.

For many years IBM has recommended segmenting the traffic types by keeping FCP traffic (such as PPRC/MM) and FICON traffic on their own dedicated ISLs or group of ISLs. With FIDR, it is now possible to share ISLs among previously segmented traffic, which leads to cost savings for hardware: fewer ISLs means fewer FICON Director ports or DWDM links. Cost savings on the bandwidth between data centers might be even greater. On the downside, there are two behaviors that might occur in a FICON SAN with FIDR enabled: dilution of error threshold counts, and the impact of slow-drain devices.

Dilution of error threshold counts

It is possible for errors to occur in a FICON SAN that cause FC frames to get lost. One example scenario is when there are more than 11 bit errors in a 2112-bit block of data and Forward Error Correction (FEC) code cannot correct this large of a burst of errors. For IBM Z and the FICON protocol, an error is detected by the host OS. This error typically is an interface control check (IFCC) or a missing interrupt. The exact error type depends on exactly where in the I/O operation the error occurs. z/OS counts these errors, which are tracked to a specific FICON channel and CU link. With static routes, the error count covers the specific ISL that is in the path between the FICON channel and the CU. With FIDR mechanisms, the intermittent errors that are caused by a faulty ISL occur on many channel to CU links and are unknown to the OS or the IBM Z processor itself.

Therefore, the error threshold counters can become diluted, that is, they are spread across the different OS counters, which make it entirely possible that either the thresholds are not reached in the period that is needed to recognize the faulty link, or that the host thresholds are reached by all CUs that cross the ISL in question, which results in all the channel paths being fenced by the host OS. To prevent this behavior, the user should use the FICON SAN switching devices capabilities to set tighter error thresholds internal to the switch and fence/decommission faulty ISLs before the OS's recovery processes are invoked.

Slow-drain devices

When slow-drain devices are present, FICON SANs are likely to lack buffer credits at some point in the architecture. This lack of buffers can result in FICON switching device port buffer credit starvation, which in extreme cases can result in congested or even choked ISLs. Frames that are "stuck" for long periods (typically for periods >500 milliseconds) might be dropped by the FICON SAN fabric, which results in errors being detected. The most common type of error in this situation is an IBM C3® discard. When a slow-drain device event occurs and corrective action is not taken in short order, ISL traffic can become congested as the effect of the slow-drain device propagates back into the FICON SAN. With FIDR policies being implemented, a slow-drain device might cause the B2B credit problem to manifest itself on all ISLs that can access the slow-drain device. This congestion spreads and can potentially impact all traffic that must cross the shared pool of ISLs. With static routing policies, the congestion and its impact are limited to the one ISL that accesses the slow drain device.

Possible causes of slow-drain devices include:

- An insufficient number of buffer credits that is configured for links that access devices at a long distance.
- Disparate link speeds between the FICON channel and CU links.
- Significant differences in cable lengths.
- Congestion at the CU host adapters that is caused when the link traffic (across all ISLs) exceeds the capacity of the host adapter, which can occur when too many MM ISLs share the same host adapter as FICON production traffic ISLs.

IBM Z and z/OS have capabilities that mitigate the effect of slow-drain devices, such as channel path selection. The algorithms steer the I/O workload away from the paths that are congested by the slow-drain device toward the FICON channels in a separate, redundant FICON SAN. Best practices for IBM Z I/O configurations require at least two separate and redundant FICON SANs. Many users use four, and the largest configurations often use eight. For MM traffic, best practices call for the IBM Z user to use FIDR in the fabric for predictable and repeatable performance, resilience against workload spikes and ISL failures, and optimal performance. If a slow-drain device situation occurs in a FICON SAN fabric with MM traffic, it impacts the synchronous write performance of the FICON traffic because the write operations do not complete until the data is synchronously copied to the secondary CU. Because FICON traffic is subject to the slow-drain device scenarios today, using FIDR does not introduce a new challenge to the user and their FICON workloads.

Note: FIDR is not supported by FICON Multihop topologies. Use static routing instead.

Using FIDR maintains a per exchange IOD. Exchanges can be delivered out of order, which is described as loose IOD as opposed to strict IOD.

Note: Not all tape units can support FIDR.

More considerations

In any cascaded FICON architecture, it is a best practice to perform a bandwidth sizing study. Such a study can help you determine the bandwidth that is required for the cascaded links and the number of ISLs. Anticipated storage needs, type of supported traffic, replication method, need for GDPS or HyperSwap and other considerations make this sizing study complicated. As a best practice, use design tool for bandwidth sizing.

For proper operation, dynamic routing must be supported at all endpoints, the channel and connected devices, and FICON switches. IBM Health Checker for z/OS can identify any inconsistencies in the dynamic routing support within the SAN. When dynamic routing is enabled in the SAN, IBM Health Checker for z/OS verifies that the processor and attached DASD, tape, and non-IBM devices that are defined as type CTC support dynamic routing and identifies those endpoints that do not. z/OS uses the CUP device to gather information from the switch (such as topology and performance statistics for RMF). As part of the information that is returned from the CUP, there is an indication about whether dynamic routing is enabled for the SAN fabric or not.

Figure 5-58 is an example of the output where all devices and hosts support FIDR, and no inconsistencies are detected.

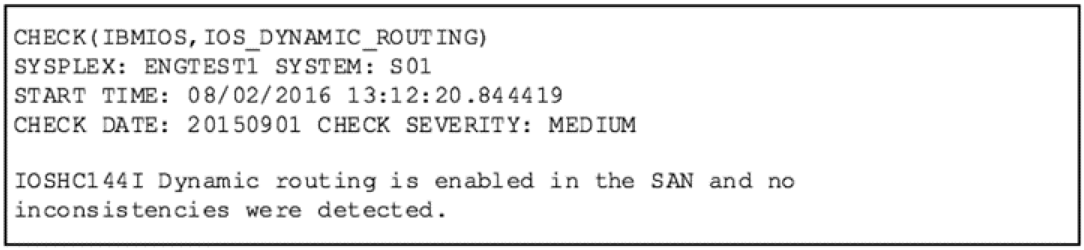

```
CHECK(IBMIOS,IOS_DYNAMIC_ROUTING)
SYSPLEX: ENGTEST1 SYSTEM: S01
START TIME: 08/02/2016 13:12:20.844419
CHECK DATE: 20150901 CHECK SEVERITY: MEDIUM

IOSHC144I Dynamic routing is enabled in the SAN and no
inconsistencies were detected.
```

Figure 5-58   IBM Health Checker for z/OS with no inconsistencies detected

Figure 5-59 is an example of the output where not all devices support FIDR. Because FIDR inconsistencies were detected, FIDR should not be enabled for this fabric.

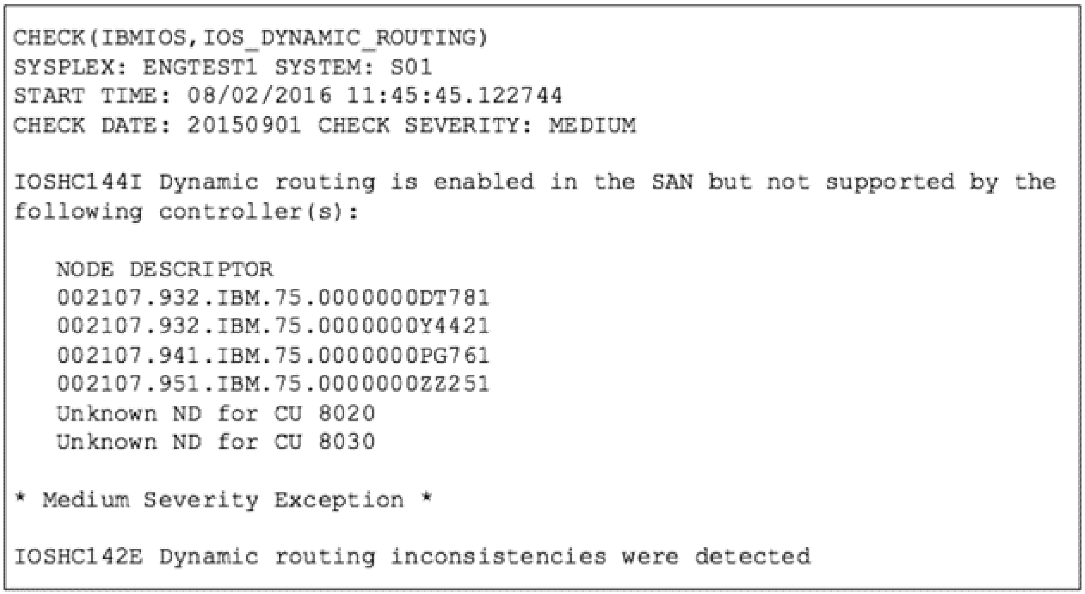

```
CHECK(IBMIOS,IOS_DYNAMIC_ROUTING)
SYSPLEX: ENGTEST1 SYSTEM: S01
START TIME: 08/02/2016 11:45:45.122744
CHECK DATE: 20150901 CHECK SEVERITY: MEDIUM

IOSHC144I Dynamic routing is enabled in the SAN but not supported by the
following controller(s):

   NODE DESCRIPTOR
   002107.932.IBM.75.0000000DT781
   002107.932.IBM.75.0000000Y4421
   002107.941.IBM.75.0000000PG761
   002107.951.IBM.75.0000000ZZ251
   Unknown ND for CU 8020
   Unknown ND for CU 8030

* Medium Severity Exception *

IOSHC142E Dynamic routing inconsistencies were detected
```

Figure 5-59   IBM Health Checker for z/OS with inconsistencies detected

In summary, FIDR provides significant technical improvements over the older FICON static routing policies, but it also introduces new concerns. Customers may decide what they consider best. Given that the main behavior of concern, that is, the impact of slow-drain devices, can be managed by the user by using good discipline, the IBM Health Checker for z/OS, and the IBM c-type FICON SAN management tools, the benefits of implementing FIDR should far outweigh the potential issues.

For more information about host and CU requirements for using FIDR, see *FICON Dynamic Routing (FIDR): Technology and Performance Implications*.

5.6.4 Speed definition: Static versus auto values

The interface speed can be configured as auto or set to a specific value. When left as auto, the two interfaces on the same link negotiate the highest possible speed. Although the auto setting normally works well, there are situations where it does not, for example, when DWDM equipment sits between two FICON network devices, the auto-negotiation protocol might encounter issues.

As a best practice, set the port speed to the required value in a FICON environment.

5.7 In order delivery of frames

Under normal operating conditions, the IBM c-type switches send all packets in the same order that they were sent by the originator. This intrinsic IOD of frames is part of the architecture of these devices. However, things can go differently at the fabric level. When ISL links are added or removed from the fabric, some frames might reach the intended destination out of order. In fact, the traffic is rehashed and redistributed on the available links when the failure is seen on each side, which can result in out of order frames when there is a combination of long and short ISLs, either within a port channel or within an Equal Cost Multipath (ECMP) scenario. Out of order frames can cause errors, particularly on the host side, and are not acceptable in general.

5.7.1 In Order Delivery

To correct that fabric-level behavior, a feature that is called IOD can be enabled on a per VSAN basis. The IOD feature was introduced to ensure IOD of frames in cases of fabric reconvergence events. With IOD enabled, frames are either delivered in order or not delivered at all. IOD works on ECMP links, port channels, and also their FCIP equivalent links. IOD is optional for FC VSANs, but mandatory for FICON VSANs. IOD quiesces traffic on all equal cost ISLs or members of a port channel for a little while (500 ms) to make sure that the fabric becomes stable again before more frames are delivered and in order again. During that short freeze timeframe, frames are dropped. In mainframe environments, this situation leads to many IFCCs, usually 1 per device that has an active data flow on it, meaning 100s - 1000s. So, IOD prevents delivery of out of order frames, but leads to frame drops.

5.7.2 Lossless IOD

As an enhancement to IOD, IBM c-type switches now support LIOD. This enhancement was initially introduced for FICON environments to achieve IOD with no frame drops, but now it is also supported on FC topologies.

LIOD works only on port channels and not in ECMP scenarios. The port channel is drained by all packets in flux when a chaser frame is sent to drain the queue on each ISL that is still up. While this task is happening, frames are queued on the interface but nothing is dropped. When the last chaser reply is received from the alternate side, the port channel is rehashed with the member interfaces that are present and traffic is released for normal flow. Traffic is halted for about 2x the round-trip time (RTT) of the port channel, which for most implementations is measured in µsec and does not affect response time in a serious way. Even a 10 km link has a hit of about 100 µsec for a single flow, and then the response time would return to normal.

With LIOD, if a port within a port channel is administratively taken down or brought up, there will be no drops at all and be truly lossless. If there is a surprise cut in the fiber, there will be a few frame drops, but 100 times less than during the 500 ms freeze time of IOD. LIOD takes effect by default when IOD is enabled on a VSAN.

LIOD works much like IOD, but it is different in the way it operates, which results in it being much faster. LIOD sends a command (chaser frame) to a peer switch to flush the queues instead of waiting for the 500 msec timer. LIOD works only for FC port channels. It does not work for FCIP port channels. FICON implementations with IBM c-type benefit from LIOD on FC port channels and achieve IOD with virtually no frame drops.

5.7.3 ECMP, port channels, and LIOD

In a more complex scenario, consider that you want two different equal-cost port channels, PC-A and PC-B, between a pair of switches. You enable IOD on one VSAN. You have both ECMP and port channels. I/O flows are directed to either PC-A or PC-B according to the ECMP load-balancing setting in use, that is, SID/DID or SIS/DID/OXID. If one member link inside PC-A is administratively brought down, the traffic that is hashed to PC-B is not affected. For the traffic that is hashed to PC-A, LIOD behavior manifests as described in 5.7.2, "Lossless IOD" on page 175. During the short period in which LIOD acts, there is no change to the way that ECMP hashing operates. Traffic still is directed to both PC-A and PC-B.

To enable IOD/LIOD on a VSAN, run the following command:

```
switch(config)# in-order-guarantee vsan 101
```

For more information about IOD and LIOD, see In-Order Delivery.

5.8 FCIP

The IBM c-type switches transparently integrate FCP, FICON, and FCIP in one system. The FICON implementation on IBM c-type Directors and IBM Storage Networking SAN50C-R multiprotocol switches supports IP tunneling to efficiently consolidate SANs over WAN distances. IP tunnels enable a globally accessible storage infrastructure. Using the FICON over FCIP capability enables cost-effective access to remotely located mainframe resources. With the IBM c-type platform, IBM storage replication services can be extended over metropolitan to global distances by using the existing IP infrastructure and further simplify business continuance strategies.

To facilitate FICON traffic over an IP network, the participating IBM c-type Directors must each have a 24/10 SAN Extension switching module with two or more FCIP ports in use to provide a redundant path capability. No license is required to activate the FCIP ports. Alternatively, you can use the IBM Storage Networking SAN50C-R switch for FCIP connectivity. The IBM Storage Networking SAN50C-R switch is interoperable with the 24/10 SAN Extension switching module.

The implementation of FCIP on IBM c-type switches is advanced. The TCP/IP protocol stack was enhanced to offer better performance and higher resiliency over unstable long-distance WANs. Moreover, data compression can be enabled to minimize WAN bandwidth usage. Data security is also possible with IPsec technology.

When configuring IBM z15 Fibre Channel Endpoint Security (FCES) connections across an FCIP tunnel, it is a best practice to turn off compression because the FCES encrypted data is not compressible.

Topics such as FCIP capacity planning and tuning are influenced by factors that are unique to each customer's configuration and are not served well with general rules. IBM, Cisco, and others offer professional services to assist with these complex topics. IBM Z and IBM LinuxONE Lab Services help clients build and deploy solutions on IBM Z and LinuxONE infrastructures. For more information, see IBM IT Infrastructure.

5.9 Power, cooling, racking, and cabling

One important item to pay attention to during the design phase of any new FICON network is the way that it will be hosted in the data centers. Power, cooling, racking, and cabling considerations will play a role. Product data sheets can help you find the correct specifications, but some general considerations might be useful.

IBM c-type Director-class switches always come with front-to-back air flow, which cannot be changed. With the port side of the chassis acting as the air intake, the installed SFPs operate at the minimum possible temperature, which boosts their reliability. The IBM Storage Networking SAN50C-R switch follows the same approach and offers port-side intake for air flow. Other IBM c-type switches, which are not qualified for FICON, come with both airflow direction options, which you can select when ordering.

Figure 5-60 illustrates the air flow direction and slot numbering for an IBM Storage Networking SAN384C-6 mission-critical director.

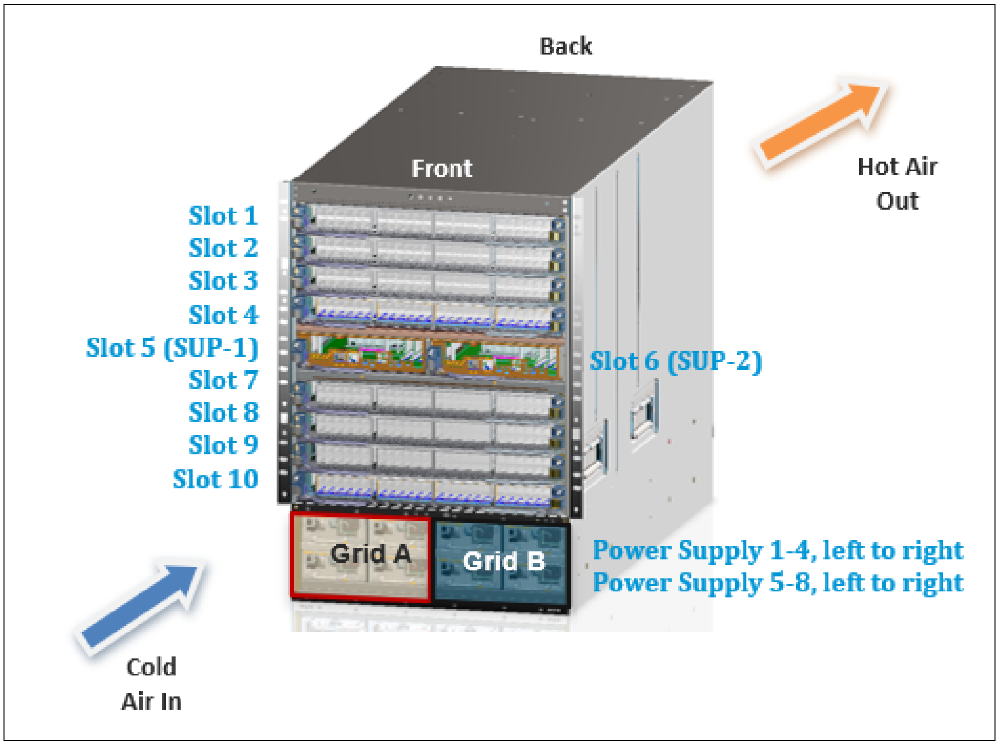

Figure 5-60 IBM Storage Networking SAN384C-6 air flow direction and slot numbering

IBM c-type networking devices use third-party certified power supplies. The 80Plus Platinum certification ensures the best energy efficiency in the industry (>94% top efficiency) and meets the stringent requirements of organizations undergoing IT green initiatives.

Figure 5-61 shows the official certification from the 80Plus organization for the IBM c-type 3 KW AC power supplies for a mission-critical director.

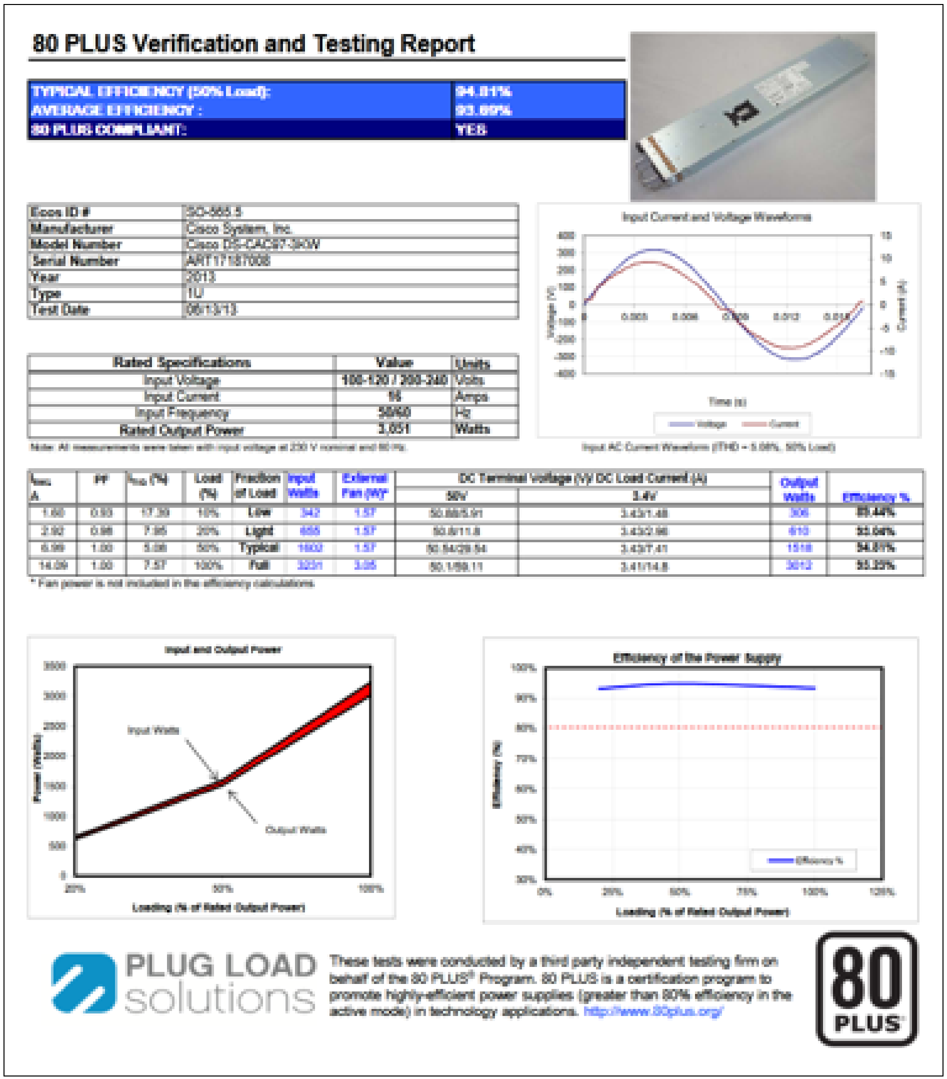

80 PLUS Verification and Testing Report

TYPICAL EFFICIENCY (50% Load):	94.81%
AVERAGE EFFICIENCY :	93.69%
80 PLUS COMPLIANT:	YES

Ecos ID #	SO-565.5
Manufacturer	Cisco System, Inc.
Model Number	Cisco DS-CAC97-3KW
Serial Number	ART1718700B
Year	2013
Type	1U
Test Date	08/13/13

Rated Specifications	Value	Units
Input Voltage	100-120 / 200-240	Volts
Input Current	16	Amps
Input Frequency	50/60	Hz
Rated Output Power	3,051	Watts

Note: All measurements were taken with input voltage at 230 V nominal and 60 Hz.

Input AC Current Waveform (iTHD = 5.08%, 50% Load)

I_{RMS} A	PF	I_{THD} (%)	Load (%)	Fraction of Load	Input Watts	External Fan (W)*	DC Terminal Voltage (V)/ DC Load Current (A)		Output Watts	Efficiency %
							50V	3.4V		
1.60	0.93	17.39	10%	Low	342	1.57	50.88/5.91	3.43/1.48	306	89.44%
2.92	0.98	7.95	20%	Light	655	1.57	50.8/11.8	3.43/2.96	610	93.04%
6.99	1.00	5.08	50%	Typical	1602	1.57	50.54/29.54	3.43/7.41	1518	94.81%
14.09	1.00	7.57	100%	Full	3231	3.05	50.1/59.11	3.41/14.8	3012	93.23%

* Fan power is not included in the efficiency calculations.

These tests were conducted by a third party independent testing firm on behalf of the 80 PLUS® Program. 80 PLUS is a certification program to promote highly-efficient power supplies (greater than 80% efficiency in the active mode) in technology applications. http://www.80plus.org/

Figure 5-61 80Plus testing report for IBM c-type Directors 3 KW AC power supplies

> **Note:** IBM Storage Networking SAN50C-R does not offer 80Plus Platinum certified power supplies.

The maximum rating of the supported power supplies should not be considered as the maximum value for switch power consumption. In many cases, power supplies are oversized to accommodate specific engineering and manufacturing requirements. Moreover, the number of installed power supplies is determined by the power redundancy schema that is assumed. As a result, you should not be surprised if an IBM Storage Networking SAN384C-6 switch comes with six 3-KW AC power supplies but reaches a maximum power consumption of 5 KW and a typical power consumption of only 2.5 KW when fully populated.

All IBM c-type switches can be hosted in standard 2-post or 4-post 19” racks. Some clearance is required in the front and back of the chassis to facilitate serviceability. Director-class switches may accommodate many ports, and appropriate brackets ensure that fiber paths are optimized.

Figure 5-62 shows a professional installation for an IBM Storage Networking SAN384C-6 Director.

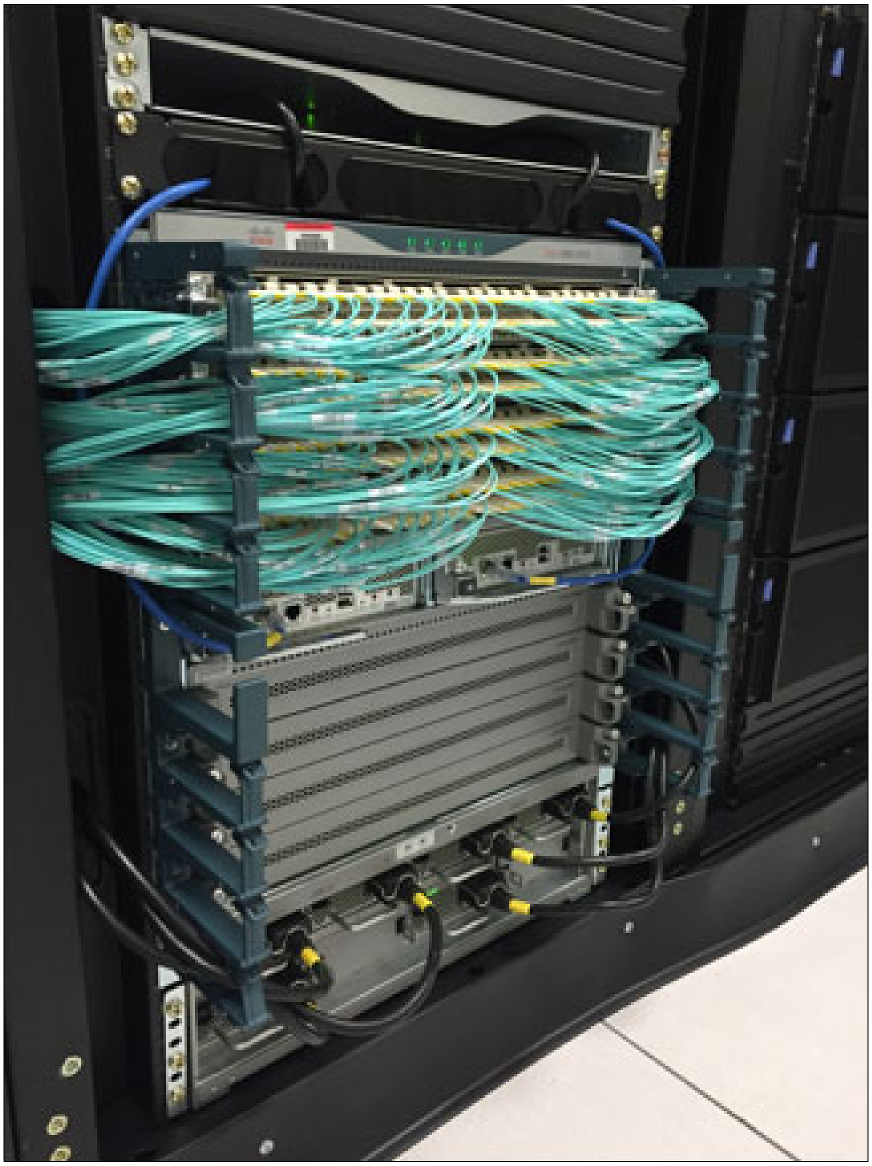

Figure 5-62 A professionally deployed IBM Storage Networking SAN384C-6 Director with plastic brackets in the front

For more information cabling and other best practices, see Cabling Considerations in Storage Area Networks.

5.10 Migration strategy

The IBM Z and IBM LinuxONE supported NX-OS upgrade path to 8.4(1a) is from 8.1(1b). Switches at any prior level must be upgraded to 8.1(1b) before upgrading to 8.4(1a).

FICON migrations can lead to performance improvements and money savings in a mainframe environment. There are several considerations, which are described in the following sections.

5.10.1 FICON cabling

FICON typically uses 1310 nm long wave optical components, which require 9/125 micron, single-mode fiber optic cabling. It is far less common, but it is possible to order FICON channels, switch ports, and CU adapters in 850 nm short wave versions that use 62.5/125 OM1 or 50/125 OM2, OM3, OM4, or OM5 multi-mode fiber optic cabling. OM4 cabling is prevalent these days. Long-wave transceivers that are commonly used for FICON are more expensive than their short-wave counterparts. The cost difference is even more noticeable at higher bit rates.

Both ends of any connection must have the same wavelength. The same fiber can carry different bit rates, so there is no requirement or investment for the cabling when migrating from 8 Gb to 16 Gb, for example. Reusing an existing infrastructure and migrating only one end of an existing connection are strong reasons to continue using the long wave (or short wave) technology.

The SFP optical components are individually replaceable. SFPs for FCP and FICON can auto-negotiate their speed from their maximum to two speed levels below it. For example, a 4 GbE SFP could connect to a 4 GbE, a 2 GbE, or a 1 GbE SFP. A 32 GbE SFP can connect to a 32 GbE, a 16 GbE, or an 8 GbE SFP. This capability allows for migration of only one end of a connection, rather than requiring that both ends are replaced concurrently, which would make migrations more difficult. So, a switch can be upgraded to allow a higher speed, but all the processors and storage can keep their current speed and be upgraded independently later.

Also, the SFPs in a processor, switch, or CU are independent and can be different speeds or wavelengths. A FICON switch could have a mix of 8 GbE, 16 GbE and 32 GbE SFPs of both 850 nm shortwave and 1310 nm longwave. A 16 GbE-capable FICON switch could be upgraded to a 32 GbE capable switch, but continue to use mostly 16 GbE SFPs, and upgrade to 32 GbE only when the channel or device can accept that speed.

Although it is not a best practice, it is possible to move the SFPs from an older device to the new replacement to cut costs. The cost benefit must be balanced by considerations on the maximum achievable speed and degradation of components. In fact, SFPs are constantly transmitting “idle to keep in sync when they are not involved in active I/O operations, so they might fail sooner than the modules into which they are plugged.

5.10.2 FICON Directors

Large mainframe environments generally use FICON Directors for connectivity. Direct attachment of FICON DASD frames to the mainframe makes sense only for a few FICON attachments. FICON Directors come at a cost, and sometime organizations try to minimize their expenses by combining their open systems SAN with their FICON connectivity and operate a FICON/FCP intermix environment.

5.10.3 More considerations

Here are some considerations for DASD and tape environments.

Only the switch, processor, storage, or all at once

FICON has been available for 20+ years and has evolved in both performance and reliability. Born primarily as a solution to overcome device access and performance limitations that are associated with ESCON, FICON technology is now natively supported by CPUs (mainframe), storage (DASD), and networks (director). A FICON migration does not need to be all or nothing. It is possible, and common to refresh only a portion of the infrastructure instead of all of it. The path that you choose depends on the available budget and real needs.

In general, to get the best use from a new feature (such as a speed increase from 8 GbE to 16 GbE), all components must be upgraded. However, rolling upgrades are also a possibility, particularly when there is one infrastructure element with a growth trend higher than others, or when the lease life is different among elements. FICON Directors usually support features such as higher speeds before channels and CU adapters, so that they are not a bottleneck that slows the migration of processors and storage, and they can be purchased to support either processors, storage, or both when they are upgraded, either at the same time or later than the switch.

Tape

As an example (and ignoring security and privacy exposures), imagine you are a bank, and every evening you read in the day's charges from different credit cards, searching for charges from people who have accounts in your bank so that you can deduct the charges from their accounts. The charge history is coming as a flow from a source that you cannot control, so you cannot pause for the deduct process and then resume the search. You have a “bucket of accounts” for the program that is processing Visa, another bucket for the program that is processing Mastercard, and another for AMEX. After all the programs end, you can look through the buckets and process the charges for your customers. If these buckets are on disk storage, how much space should you allocate? Using disks would require constantly monitoring of the size of the allocated data set and managing reallocation and movement to a larger data set if it filled up. Tape provides a large amount of storage, which is useful when the amount of required storage cannot be predetermined.

There are many programs that used the reconciliation processing model. IBM mainframe OSs enable old programs to still run, so many companies still run those processes with a “why fix what isn't broken” attitude. However, with advances in disk capacity and a massive decrease in cost, virtual tape servers (VTSs) or virtual tape libraries (VTLs) were created that appeared to the program as a tape drive because they respond to tape I/O commands, but the device is a large amount of managed disk storage in front of physical tape drives. Disk storage is smaller than tape and much faster in responding to commands (for example, the rewind command does not have to physically move a long tape). Other uses of tapes such as backups or offloading archived files are still common but much faster with VTLs. Cybernetics and StorageTek made many VTLs. IBM had several virtual tape drive systems starting in 1996 with the 3495-B16 MagStar, and the latest one is IBM VS7700, which no longer has any physical tape that is attached, and flash storage replaced rotating disk storage to make it even faster and more capacious.

Thus, tape devices are still vital in mainframe environments and are attached by using FICON adapters, and their migration must be considered, and DASD, processors, and switches.

6

Initial connectivity and setup

This chapter describes the tools and the procedures that are required to make the first connection to an IBM c-type switch, perform the initial setup, and retrieve and install licenses.

By the end of this chapter you will be making a remote connection and continuing the remainder of the switch configuration by using the Data Center Network Manager (DCNM) GUI or NX-OS command-line interface (CLI).

The following topics are covered in this chapter:

- Cable requirements and workstation configuration
- Initial SAN switch setup
- Role-based access control
- Date and time configuration
- Network Time Protocol
- Installing Data Center Network Manager
- Logging in to DCNM
- Device Manager
- Installing switch licenses

6.1 Cable requirements and workstation configuration

For the initial manual setup of a switch, you can use a workstation running any operating system (OS) if it can run a terminal session with VT100 emulation. Some examples are listed below:

- Windows
- Linux (Red Hat, SuSE, and Ubuntu)
- Mac OS
- IBM AIX®

For the example in this book, we use a Windows 10 workstation running PuTTY. PuTTY is open source software that is available to download and use for no charge within the terms of its license.

IBM c-type switches do not come with a default IP address. To make an initial connection from a workstation to a switch requires a serial (RS-232) port. Modern workstations no longer have built-in serial ports, so you must obtain a USB-to-serial port adapter with a USB male interface on one end, and a DB-9 male interface on the other end. Figure 6-1 shows an example of such an adapter.

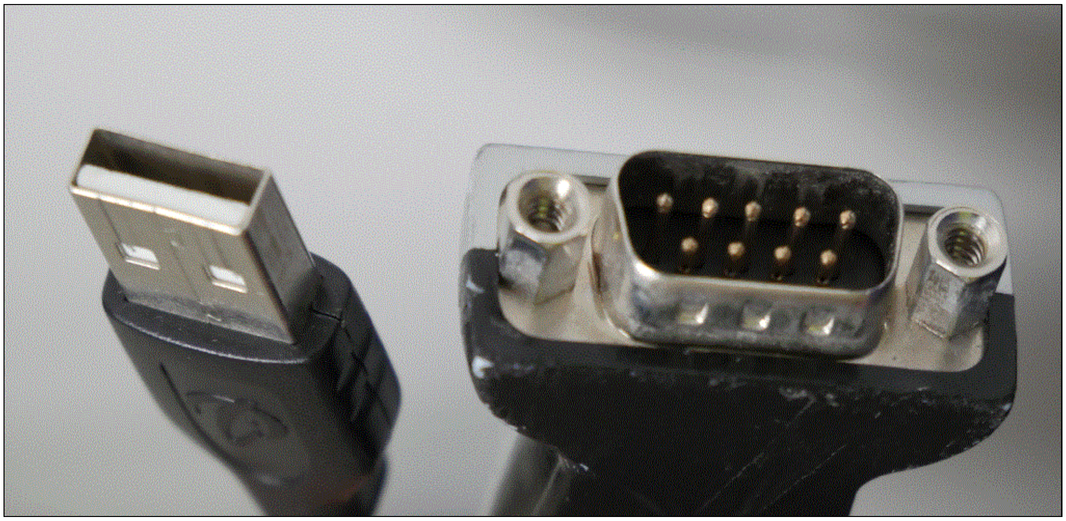

Figure 6-1 USB-to-serial port adapter

To install the adapter, follow the instructions that are provided with it. Then, you must connect a rollover cable between the adapter on your workstation and the switch.

The accessory kit that is provided with the switch contains essential cables, including the rollover cable. One end has a DB-9 female interface, and the other end has an RJ-45 male interface. A rollover cable is not a straight-through or crossover cable. Neither of these cables will work, so you must use a rollover cable with the correct pin-outs.

A rollover cable gets its name from the fact that they have opposite pin assignments on each end of the cable. Typically, they come with a turquoise jacket, although other colors do exist. Figure 6-2 on page 185 shows an example of a rollover cable.

> **Note:** A straight-through cable is often written as *straight-thru*. A crossover cable is sometime referred to as a *null modem cable*.

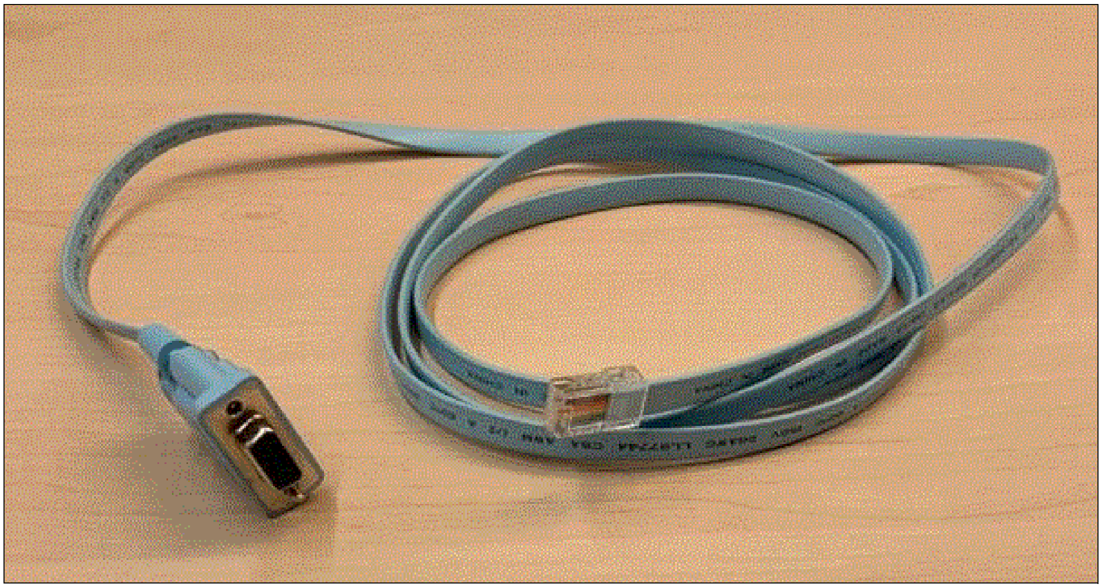

Figure 6-2 Rollover cable

Connect one end of the rollover cable to the serial port on your workstation and the other end to the storage area network (SAN) switch. Figure 6-3 shows the serial port on an IBM Storage Networking SAN50C-R switch. This port is also referred to as the console port (RS-232 port).

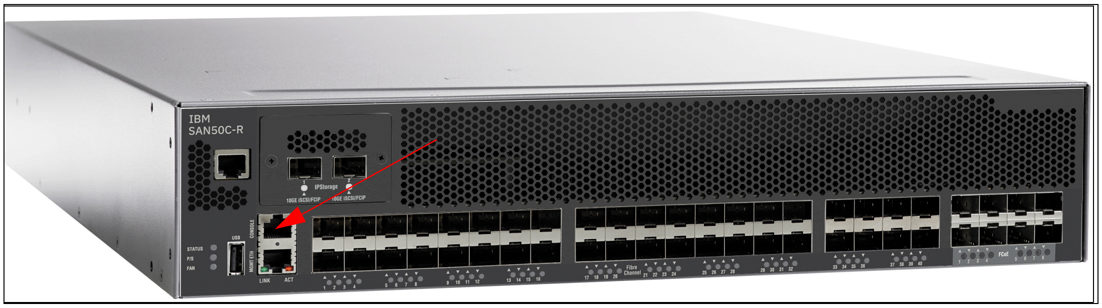

Figure 6-3 IBM Storage Networking SAN50C-R serial port

Note: Ensure that you connect to the correct port on the SAN switch because both the serial (console) and Ethernet (management) ports are the same shape.

The serial port on your workstation receives a COM*n* number, where *n* has a numeric value, for example, 1, 2, 3, and so on. Use the required tools within your OS to determine what COM*n* port number is assigned. In Windows, this task can be done by using Device Manager (DM).

Figure 6-4 shows that the COM port is assigned to COM7.

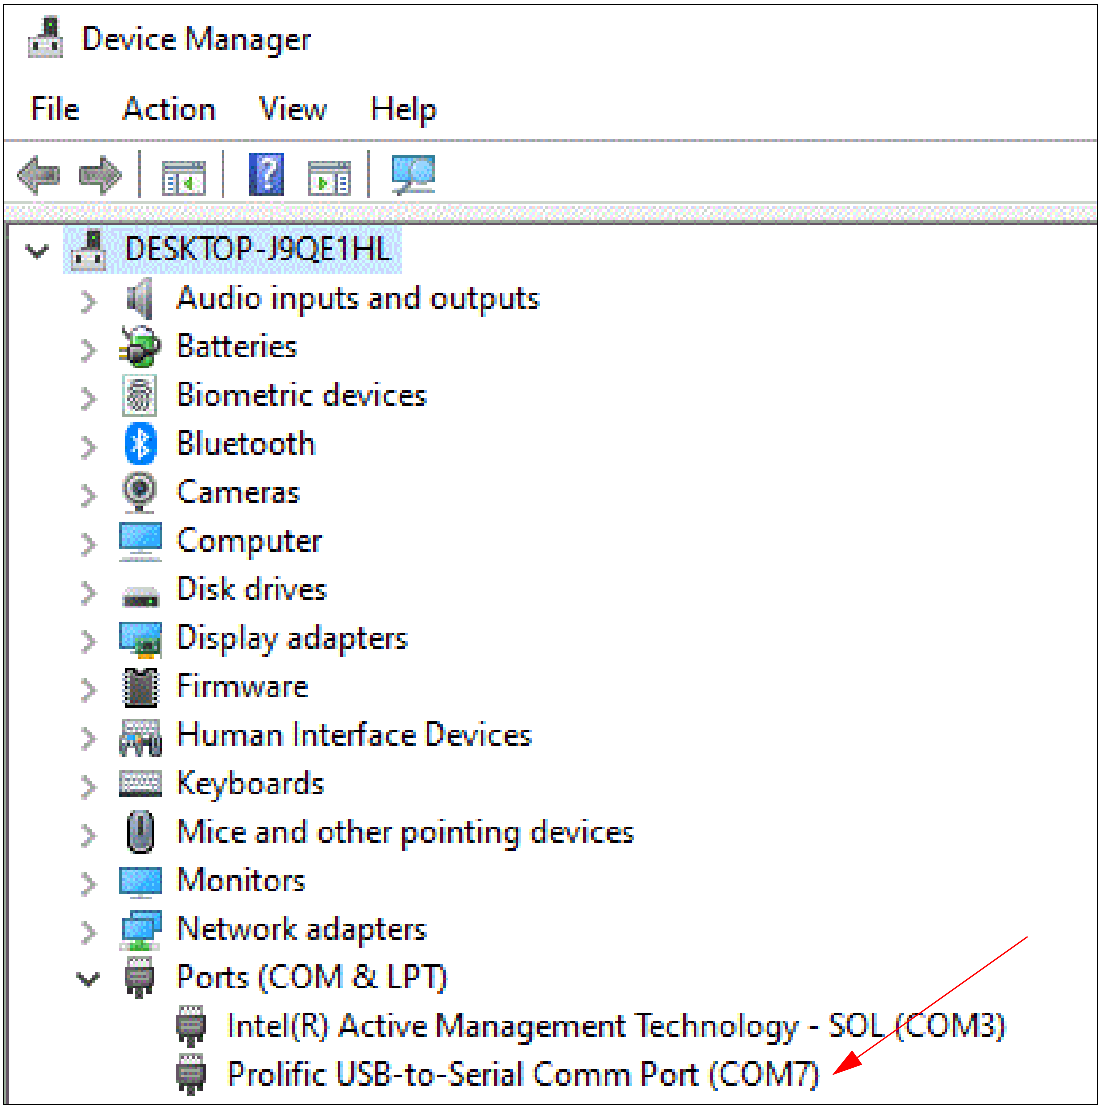

DESKTOP-J9QE1HL
- Audio inputs and outputs
- Batteries
- Biometric devices
- Bluetooth
- Cameras
- Computer
- Disk drives
- Display adapters
- Firmware
- Human Interface Devices
- Keyboards
- Mice and other pointing devices
- Monitors
- Network adapters
- Ports (COM & LPT)
  - Intel(R) Active Management Technology - SOL (COM3)
  - Prolific USB-to-Serial Comm Port (COM7)

Figure 6-4 Com port number

Use your terminal session client to set the serial port with the following parameters:

- Speed = 9600
- Data bits = 8
- Stop bits = 1
- Parity = None
- Flow control = None

Note: The above settings work if the switch is set to its default parameters, which will be the case if it is the first connection. If the settings do not work on an existing switch, then check with your SAN administrator to discover the correct settings.

Figure 6-5 on page 187 shows the configuration in PuTTY. To get to these options, click **Serial**, and then click **Open**.

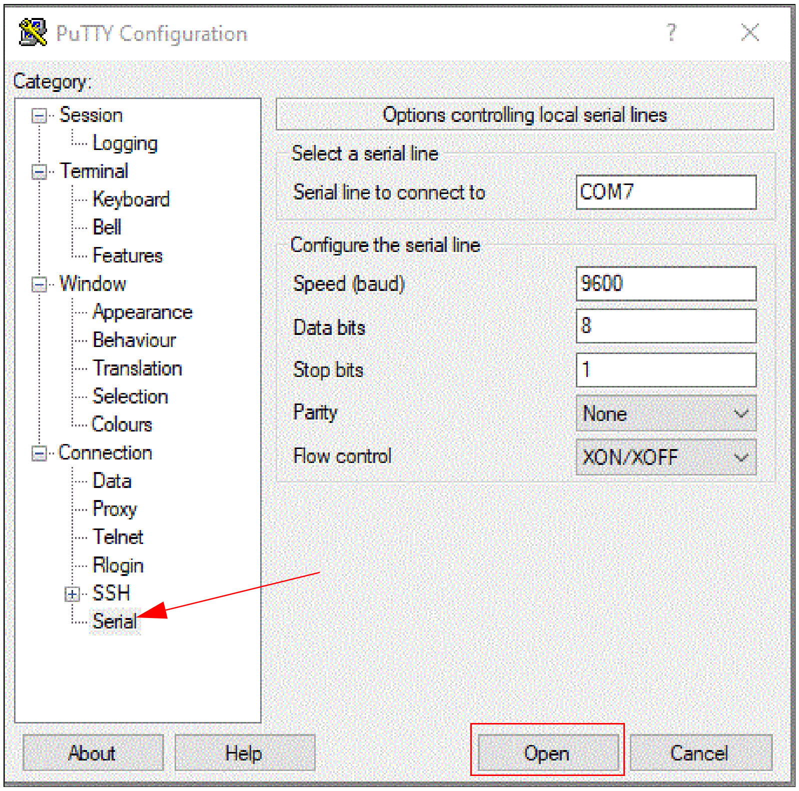

Options controlling local serial lines

Select a serial line

Serial line to connect to: COM7

Configure the serial line

Speed (baud): 9600

Data bits: 8

Stop bits: 1

Parity: None

Flow control: XON/XOFF

Figure 6-5 PuTTY terminal session configuration

6.2 Initial SAN switch setup

After you have the correct settings, click **Open**, which opens an NX-OS CLI session. What appears on the window depends on the state of the switch. The switch might be in one of several states, assuming that it is in working condition with no errors. The typical states are listed below:

- It is turned off, so the window will be blank.
- It is starting.
- It has started and has never been configured or the previous configuration was deleted.
- It has started and has a configuration that was applied. If so but you want to go through the initial setup again, delete the current configuration and restart the switch.

If the switch is not turned on, then turn it on by connecting it to the power line. When the switch initially starts, it goes through numerous checks, which scroll by on the screen until the first prompt appears. The following examples show the initial setup on an unconfigured IBM Storage Networking SAN384-C switch that started at the first prompt after the boot procedure completes.

Other switches have a similar start procedure. Most of the prompts show the default answer in square brackets [y] or [n]. If the Enter key is pressed without you explicitly typing **`yes`** or **`no`**, then the default answer is selected. To explicitly select **`yes`** or **`no`**, type **`yes`**, **`y`**, **`no`**, or **`n`**, and then press Enter.

Further explanations of the answers that are selected in the following examples are not provided unless they differ from the default or if further explanation is required. A basic setup was deliberately selected because the remainder of the configuration is shown elsewhere in this book.

Example 6-1 shows the first three initial setup prompts.

Example 6-1 First three initial setup prompts

```
Do you want to enforce secure password standard (yes/no) [y]: yes

Enter the password for "admin": ********
Confirm the password for "admin": ********

         ---- Basic System Configuration Dialog ----

This setup utility guides you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Register Cisco Multilayer Director Switch (MDS) 9000 Family devices promptly with
your supplier. Failure to register might affect response times for initial
service calls. MDS devices must be registered to receive entitled
support services.

Press Enter at any time to skip a dialog. Use ctrl-c at any time
to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

After you enter the basic configuration dialog, you get a series of prompts, as shown in Example 6-2. Depending on the version of code and options that you select, the prompts might differ slightly from the ones in this example.

Example 6-2 Basic configuration dialog

```
Create another login account (yes/no) [n]: no

Configure read-only SNMP community string (yes/no) [n]: no

Configure read/write SNMP community string (yes/no) [n]: no

Enter the switch name: SAN384C-6

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: yes

    Mgmt0 IPv4 address: 10.122.107.95

    Mgmt0 IPv4 netmask: 255.255.255.0

Configure the default gateway? (yes/no) [y]: yes

    IPv4 address of the default gateway: 10.122.107.1

Configure advanced IP options? (yes/no) [n]: no
```

```
Enable the ssh service? (yes/no) [y]: yes

Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa

Number of rsa key bits <1024-4096> [1024]: 1024

Enable the telnet service? (yes/no) [n]: no

Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: yes

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: con

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
    in range (<200-500>/default), where default is 500.  [d]: 500

Enable the http-server? (yes/no) [y]: no

Configure clock? (yes/no) [n]: no

Configure time zone? (yes/no) [n]: no

Configure summertime? (yes/no) [n]: no

Configure the ntp server? (yes/no) [n]: no

Configure default switchport interface state (shut/noshut) [shut]: shut

Configure default switchport trunk mode (on/off/auto) [on]: on

Configure default switchport port mode F (yes/no) [n]: no

Configure default zone policy (permit/deny) [deny]: deny

Enable full zone set distribution? (yes/no) [n]: no

Configure default zone mode (basic/enhanced) [basic]: basic
```

A summary of the basic configuration to be applied is shown in Example 6-3. If this summary is not correct, select **yes** to go through the options again. Otherwise, select **no** to accept the configuration and continue.

Example 6-3 Basic configuration summary

```
The following configuration will be applied:
  password strength-check
  switchname SAN384C-6
  interface mgmt0
    IP address 10.122.107.95 255.255.255.0
    no shutdown
  ip default-gateway 10.122.107.1
  ssh key rsa 1024 force
  feature ssh
  no feature telnet
  system timeout congestion-drop default logical-type edge
  system timeout congestion-drop default logical-type core
  no feature http-server
  system default switchport shutdown
  system default switchport trunk mode on
  no system default zone default-zone permit
  no system default zone distribute full
  no system default zone mode enhanced

Would you like to edit the configuration? (yes/no) [n]: no
```

Example 6-4 shows the configuration being saved and applied, and then the login prompt appears.

Example 6-4 Basic configuration confirmation

```
Use this configuration and save it? (yes/no) [y]: yes

[########################################] 100%
Copy complete.

User Access Verification
SAN384C-6 login: admnin
Password: ***********
```

The initial setup is now complete. Log in to the switch and run the **show hardware** command to verify that everything is as expected. Example 6-5 shows a truncated version of this command and output.

Example 6-5 The show hardware command

```
Nexus Operating System (NX-OS) Software
IBM support: http://ibm.com/support [ibm.com]
Products: https://www.ibm.com/it-infrastructure/storage/san/c-type [www.ibm.com]
Copyright (c) 2002-2019, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works that are contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

SAN384C-6# show hardware

Software
  BIOS:      version 2.6.0
  kickstart: version 8.4(1a)
  system:    version 8.4(1a)
  BIOS compile time:       05/17/2019
  kickstart image file is: bootflash:///m9700-sf4ek9-kickstart-mz.8.4.1a.bin
  kickstart compile time:  10/31/2019 12:00:00 [11/30/2019 19:14:41]
  system image file is:    bootflash:///m9700-sf4ek9-mz.8.4.1a.bin
  system compile time:     10/31/2019 12:00:00 [11/30/2019 20:38:44]

Hardware
  IBM SAN384C-6 8978-E08 (8 Module) Chassis ("Supervisor Module-4")
  Intel(R) Xeon(R) CPU D-1548  with 14270340 kB of memory.
  Processor Board ID JAE230512MR

  Device name: SAN384C-6
  bootflash:    3932160 kB
  slot0:              0 kB (expansion flash)

Kernel uptime is 0 days, 1 hours, 28 minutes, 2 seconds

Last reset
  Reason: Unknown
  System version: 8.4(1a)
  Service:

plugin
  Core plug-in, Ethernet plug-in
--------------------------------
Switch hardware ID information
--------------------------------

Switch is booted up
  Switch type is SAN384C-6 8978-E08 (8 Module) Chassis
  Model number is 8978-E08
  H/W version is 1.2
  Part Number is 01FT565 E08
  Part Revision is A0
  Manufacture Date is Year 0 Week 13
  Serial number is 000013C506E
  CLEI code is CMM3N00ARA

--------------------------------
Chassis has 10 Module slots and 6 Fabric slots
--------------------------------

Module1  ok
  Module type is 4/8/16/32 Gbps Advanced FC Module
  0 submodules are present
  Model number is 01FT644 48x32 FC
  H/W version is 1.1
  Part Number is 01FT644 48x32 FC
  Part Revision is B0
  Manufacture Date is Year 23 Week 12
  Serial number is JAE23120F1A
SAN384C-6#
```

The remainder of the switch configuration can be continued remotely from your desk. Configuration can be performed by using the GUI or CLI tools.

To use the CLI, you need an SSH client, which you use to connect to the switch by using the IP address and user credentials that were configured in the initial setup.

For a GUI configuration, use the licensed version of DCNM because all the GUI configuration examples that are shown in this book are done so by using the licensed version of DCNM. The no charge version of DCNM starting at release 11.3 has a 60-day trial of all licensed features. Starting at DCNM release 11.5, the trial period is up to 120 days.

Table 6-1 and Table 6-2 on page 193 show lists of the features that are available in DCNM.

Table 6-1 DCNM features

Feature	DCNM Unlicensed Mode	DCNM Licensed Mode
FC/Fibre Channel over Ethernet (FCoE), IBM Fibre Connection (FICON), internet Small Computer Systems Interface (iSCSI) topology view	Yes	Yes
Fabric, device, and summary views	Yes	Yes
Port, Switch, and fabric-level configuration	Yes	Yes
Event and security management	Yes	Yes
Configuration analysis tools	Yes	Yes
Network diagnostic and troubleshooting tools	Yes	Yes
Real-time performance monitoring	Yes	Yes
One command multi-switch CLI access	Yes	Yes
DM	Yes	Yes
Template-based provisioning	Yes	Yes
Generic Online Diagnostics (GOLD)	Yes	Yes
Heterogeneous storage array discovery	-	Yes
Scale-out federation architecture	-	Yes
SAN Host Path Redundancy Analysis	-	Yes
Automatic fabric failover	-	Yes
VMware vCenter plug-in	-	Yes
Multiple fabric management	-	Yes
Centralized management server with discovery	-	Yes
Continuous health and event monitoring	-	Yes
Historical performance monitoring and reporting	-	Yes
Event forwarding	-	Yes
DCNM proxy services	-	Yes
Configuration backup, archive, and compare	-	Yes

Feature	DCNM Unlicensed Mode	DCNM Licensed Mode
Roaming user profiles	-	Yes
VMpath analytics	-	Yes
Domain Dashboards	-	Yes
Capacity Manager	-	Yes
Event Snooze	-	Yes
Reporting	-	Yes

Table 6-2 shows the DCNM advanced features.

Table 6-2 DCNM advanced features

Feature	Works only with licensed fabrics (at least one switch that is licensed on the fabric)	Licensed switch (applies to each switch individually)	Every fabric licensed (at least one switch in every fabric)
VMpath Analytics (VMware discovery)	Yes	-	-
Storage array discovery through Storage Management Initiative Specification (SMI-S)	Yes	-	-
SMI-S APIs -Northbound	Yes	-	-
Performance Monitoring (dashboards, views, and reports)	-	Yes	-
Backup configuration	-	Yes	-
Event forwarding	-	Yes	-
Port capacity manager	-	Yes	
Automatic fabric failover (requires federation)	Yes	-	-
SAN Host Path Redundancy	-	-	Yes
Health Score	-	-	-
Storage Media Encryption	-	Yes	-
Slow-drain analysis	-	-	Yes

6.2.1 Power On Auto Provisioning

Power On Auto Provisioning (POAP) is a way to automatically configure a switch by using one of the following methods:

- A USB drive
- A server that contains the configuration files and software images

POAP is best suited to medium (USB drive) and large (the server method) environments, where the time and effort that are required to set up the environment are realized, and a greater level of consistency is achieved because POAP is less prone to human error. This section provides a basic explanation of POAP. For more information, see Using Power On Auto Provisioning.

There are several prerequisites for POAP to work, and the main ones are listed below. For a full list, see the Cisco documentation.

- A USB drive that is formatted with FAT32 or a combination of DHCP server and TFTP/SCP server. Either of these formats must have the configuration files and software images.
- A switch that supports NX-OS Release 8.1(1b) or later.
- No existing configuration on the switch (otherwise, it boots from this configuration).

The POAP feature is on by default. When a switch starts, if there is no onboard configuration file to boot from, it checks for a USB drive in USB1. If the switch finds one and all the conditions are met, it configures itself and starts from the configuration therein. If there is no USB drive or the correct conditions are not met, then the switch looks for a DHCP server and TFTP/SCP server. If these servers are found and the conditions are met, the switch configures itself and starts from the configuration therein. If a POAP environment was configured but the switch cannot configure itself from either method, then it is necessary to troubleshoot until any problems are resolved.

6.3 Role-based access control

After the initial setup completes on a switch, the next step is to go into the more detailed features and configure it to work within the rules and policies of the environment. One of these features you should configure is role-based access control (RBAC). RBAC authentication allows you to restrict management operations based on role membership.

6.3.1 Users, user roles, and rules

It is important to distinguish between users, user roles, and rules. A *user* or username is assigned to the person who logs in to the switch. A *user role* determines what actions are allowed or denied for a user to whom the user role is assigned. Multiple user roles can be assigned to an individual user. User roles are cumulative, for example, if user1 is assigned role3 and role4, then user1 is allowed to do everything that role3 and role4 allow.

> **Note:** Allow access takes priority over deny access. For example, if a user2 has role4 assigned that denies access to debug commands but also has role5 assigned that allows access to debug commands, then user2 has access to debug commands.

A *rule* is a breakdown of specific actions. Up to 16 rules can be assigned to a user role.

There are two main default user roles on a switch:

- network-admin: This user role has complete read/write access to the entire switch.
- network-operator: This user role has complete read access to the entire switch.

Example 6-6 on page 195 shows the default roles that are configured along with the default rules that are configured for each role.

Example 6-6 The show user role command

```
SAN384C-6# show role
Role: network-admin
  Description: Predefined Network Admin group. This role cannot be modified.
  VSAN policy: permit (default)
  -------------------------------------------------
  Rule    Type    Command-type    Feature
  -------------------------------------------------
  1       permit  clear           *
  2       permit  config          *
  3       permit  debug           *
  4       permit  exec            *
  5       permit  show            *

Role: network-operator
  Description: Predefined Network Operator group. This role cannot be modified.
  VSAN policy: permit (default)
  -------------------------------------------------
  Rule    Type    Command-type    Feature
  -------------------------------------------------
  1       permit  show            *
  2       permit  exec            copy licenses
  3       permit  exec            dir
  4       permit  exec            ssh
  5       permit  exec            terminal
  6       permit  config          username
```

There are more default roles that are not shown here that are system-defined privileged roles.

When a custom user role is created, it does not in itself permit access to functions. Rules must be configured for the user role before it becomes functional.

6.3.2 Creating a user

This section guides you through creating a user. By default, a switch has one default user that is named admin that is defined during the setup. There is no default password that is set. Example 6-7 shows the account information for the admin user.

Example 6-7 Basic user account information

```
SAN384C-6# show user-account
user:admin
        this user account has no expiry date
        roles:network-admin
```

Example 6-8 shows how to create a user, add a password, and set an expiry date.

Example 6-8 Creating a user

```
SAN384C-6# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.

SAN384C-6(config)# username itsouser password Redb00k! role network-operator
expire 2020-12-31.
```

```
SAN384C-6(config)# show user-account
user:admin
        this user account has no expiry date
        roles:network-admin
user:itsouser
        expires on Thu Dec 31 23:59:59 2020
        roles:network-operator
```

By default, a new user is assigned the role of network-operator, but it is always best to explicitly enforce it. For more information about creating users and user advanced security such as password policies, see the online IBM or Cisco documentation.

6.3.3 Creating a user role and adding rules

This section guides you through creating a custom user role and assigning rules to the role. Example 6-9 shows how to create a custom user role that is called *developer*.

Example 6-9 Creating a user role

```
SAN384C-6# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.

SAN384C-6(config)# role name developer

SAN384C-6(config-role)# description development team

SAN384C-6(config-role)# show role
.
.
<output truncated>
.
.
Role: developer
  Description: development team
  VSAN policy: permit (default)
```

Now, the developer role is created, but there are no rules that are assigned. Example 6-10 shows how to add rules to the developer role.

Example 6-10 Adding rules to a user role

```
SAN384C-6# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# role name developer

SAN384C-6(config-role)# rule 1 permit config
SAN384C-6(config-role)# rule 2 permit debug feature zone
SAN384C-6(config-role)# rule 3 deny config feature fspf

SAN384C-6(config-role)# show role
.
.
<output truncated>
.
.
```

```
Role: developer
  Description: development team
  VSAN policy: permit (default)
  ---------------------------------------------
  Rule    Type    Command-type    Feature
  ---------------------------------------------
  1       permit  config          *
  2       permit  debug           zone
  3       deny    config          fspf
```

Changes to role-based configuration should be committed to the configuration database to ensure that they persist across restarts and also are distributed to all switches in the same fabric. Example 6-11 shows how to commit changes to the database and then distribute them to the entire fabric.

Example 6-11 Committing changes to the database and distributing them to the entire fabric

```
SAN384C-6# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.

SAN384C-6(config)# role commit
SAN384C-6(config)# role distribute
```

6.3.4 Assigning a user role to a user

This section guides you through assigning a user role to a user. Example 6-12 shows the procedure.

Example 6-12 Assigning a user role to a user

```
SAN384C-6# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.

SAN384C-6(config)# username itsouser role developer

SAN384C-6(config)# show user-account
user:admin
        this user account has no expiry date
        roles:network-admin
user:itsouser
        expires on Thu Dec 31 23:59:59 2020
        roles:network-operator developer
```

The itsouser user now has both network-operator and developer privileges.

6.3.5 Managing a user by using DCNM

When using DCNM, new roles are not created or defined to access DCNM. Instead, the roles are assigned from existing roles that are supported on devices that are monitored, for example, IBM c-type SAN switches or Cisco Nexus switches. The role that is assigned to a user determines what operations that the user can perform on a particular device.

Table 6-3 shows some of the main roles that are supported. The roles that appear in your DCNM environment might differ depending on the devices that are supported and monitored.

Table 6-3 DCNM user roles

Role	Description
global-admin	Introduced in Cisco Nexus 5000 series switches, this role administers local area network (LAN) and SAN features.
san-network-admin	Administers SAN features.
lan-network-admin	Administers LAN features.
network-admin	Administers LAN features.
san-admin	Administers SAN features, this role was introduced with the Cisco Nexus 5000 series switches.
server-admin	Administers the FC server host feature.
network-operator	General network operator role.

6.3.6 Adding users and user roles by using DCNM

It is assumed that you already installed DCNM and know how to log in to it. For more information about how to install DCNM and log in to it, see 6.6, “Installing Data Center Network Manager” on page 208.

After you are logged in to the DCNM web client user interface, you see a window like Figure 6-6. Your window might look slightly different depending on your version of DCNM and access privileges. Click **Administration**.

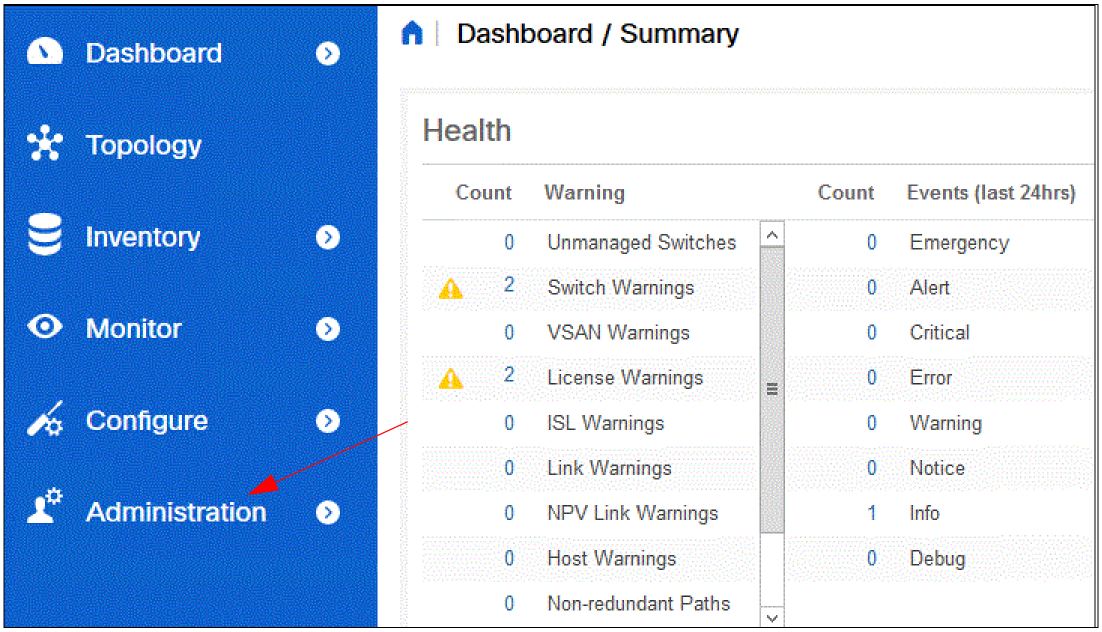

Count	Warning	Count	Events (last 24hrs)
0	Unmanaged Switches	0	Emergency
⚠ 2	Switch Warnings	0	Alert
0	VSAN Warnings	0	Critical
⚠ 2	License Warnings	0	Error
0	ISL Warnings	0	Warning
0	Link Warnings	0	Notice
0	NPV Link Warnings	1	Info
0	Host Warnings	0	Debug
0	Non-redundant Paths		

Figure 6-6 User administration from within the DCNM web client user interface

The window that is shown in Figure 6-7 on page 199 opens. Click **Local**.

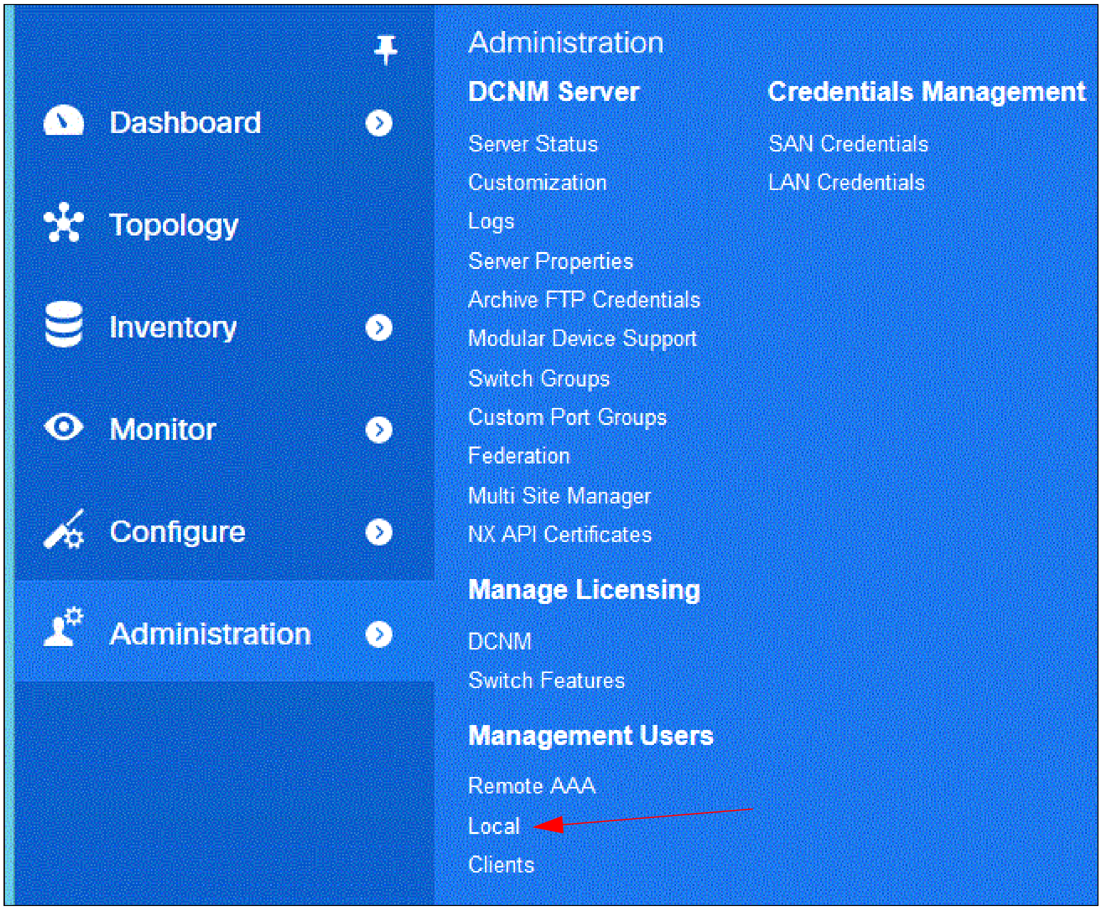

Figure 6-7 Selecting Local under Management Users

The window that is shown in Figure 6-8 opens. Click the **+** icon.

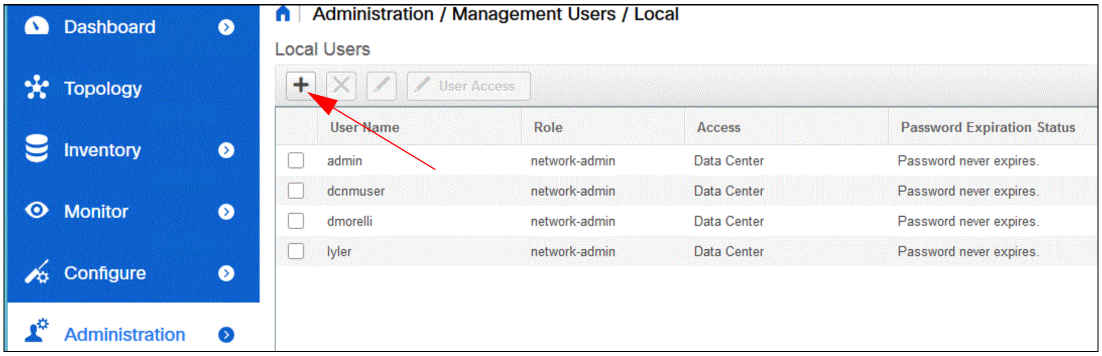

	User Name	Role	Access	Password Expiration Status
□	admin	network-admin	Data Center	Password never expires.
□	dcnmuser	network-admin	Data Center	Password never expires.
□	dmorelli	network-admin	Data Center	Password never expires.
□	lyler	network-admin	Data Center	Password never expires.

Figure 6-8 Adding a user

The window that is shown in Figure 6-9 opens. Add the username, select one of available roles, and provide a password.

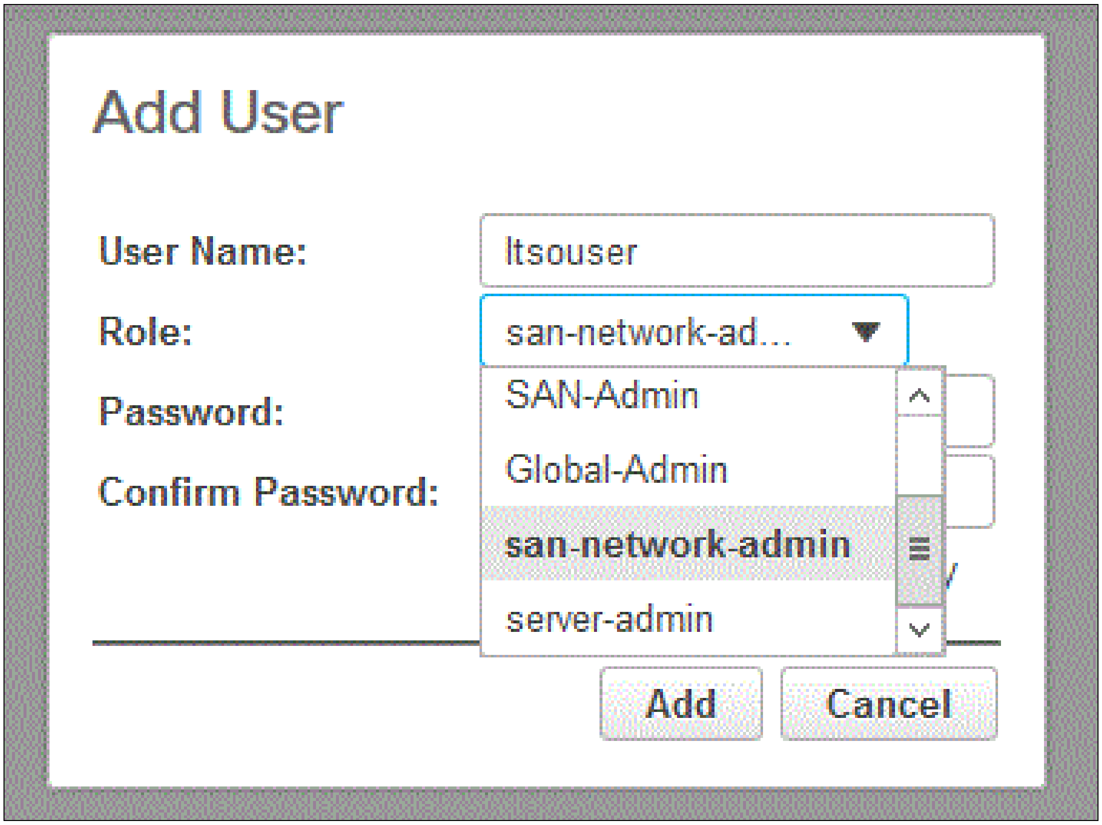

Add User

User Name: ltsouser

Role: san-network-ad...

Password:

Confirm Password:

Figure 6-9 Adding username, selecting a role, and providing a password

The window that is shown in Figure 6-10 on page 201 opens. Confirm that all the options are correct, and then click **Add**.

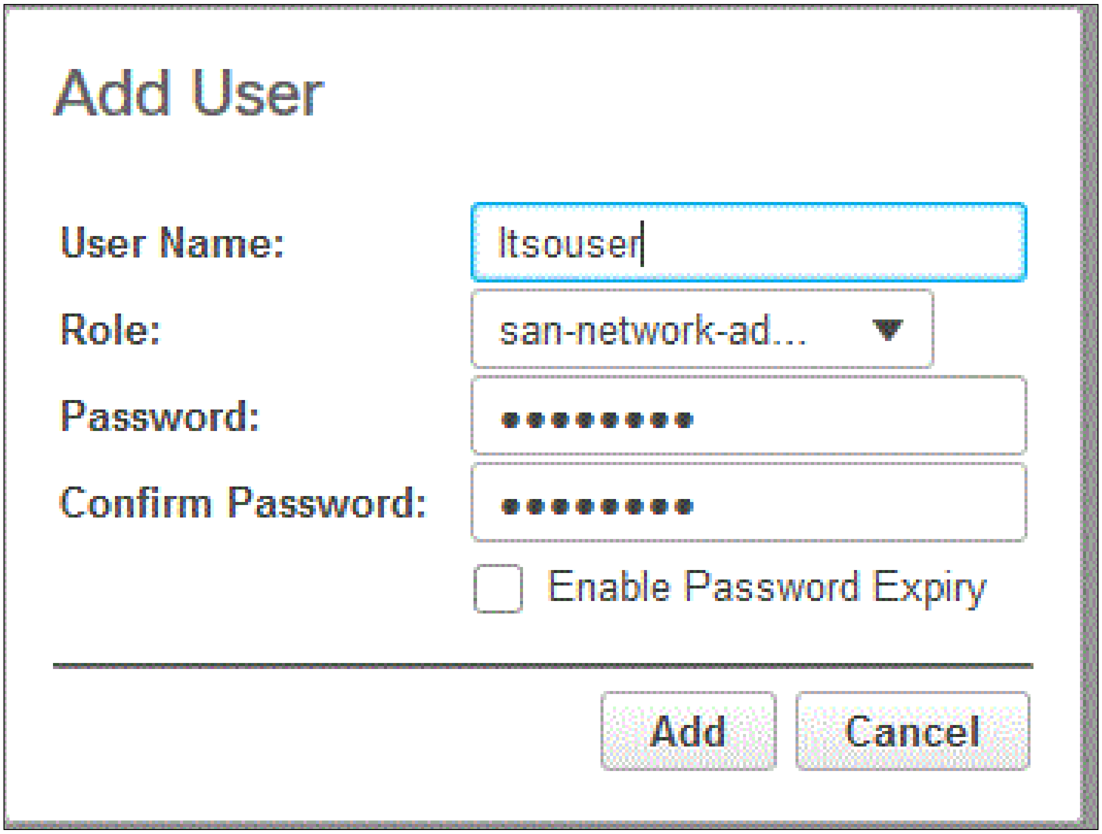

Add User

User Name: ltsouser

Role: san-network-ad...

Password: •••••••••

Confirm Password: •••••••••

□ Enable Password Expiry

Figure 6-10 Confirming user details

The user is added and shows up in the list of users within DCNM, as shown in Figure 6-11.

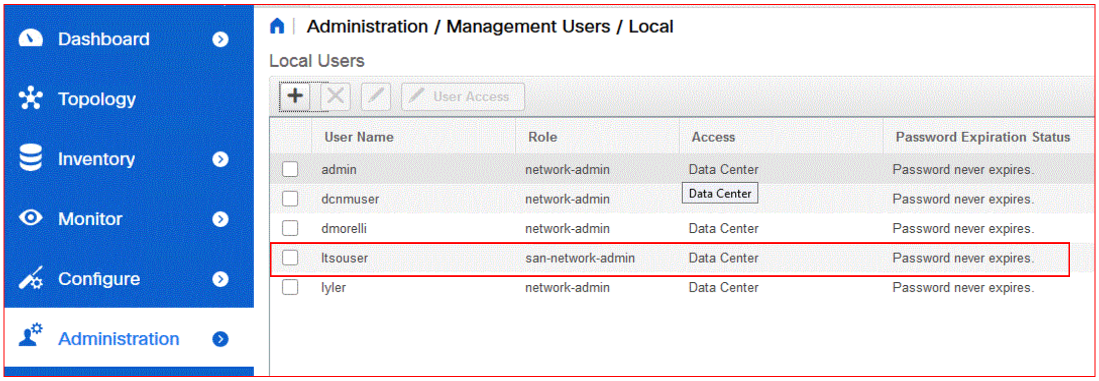

Administration / Management Users / Local

Local Users

	User Name	Role	Access	Password Expiration Status
□	admin	network-admin	Data Center	Password never expires.
□	dcnmuser	network-admin	Data Center	Password never expires.
□	dmorelli	network-admin	Data Center	Password never expires.
□	ltsouser	san-network-admin	Data Center	Password never expires.
□	lyler	network-admin	Data Center	Password never expires.

Figure 6-11 User added to list of users

More local users can be added by using the same procedure.

Local users that are configured on DCNM can access only DCNM. The same local users can also access the c-type switches only if they are explicitly configured on the switches themselves.

6.4 Date and time configuration

Ideally, the date and time should be set by using a Network Time Protocol (NTP) server. If an NTP server is not available or the environment does not require one, such as a test lab, then the date and time should be set manually.

6.4.1 Configuring the date and time by using the CLI

When setting the date and time, it is also necessary to set the time zone at the same time.

Configuring the date and time

Example 6-13 shows how to configure the date and time by using the CLI.

Example 6-13 Configuring the date and time

```
SAN384C-6# clock set 18:12:15 24 January 2020
Fri Jan 24 18:13:03 UTC 2020
```

The time that is specified is relative to the configured time zone.

Showing the data and time

Example 6-14 shows how to display the date and time by using the CLI.

Example 6-14 Showing the data and time

```
SAN384C-6# show clock
Fri Jan 24 18:13:23 UTC 2020
```

Setting the time zone

Example 6-15 shows how to set the time zone so that it is offset from Greenwich Mean Time (GMT) by minus five (-5.00) hours. By default, a switch is set to Universal Coordinated Time (UTC).

Example 6-15 Setting the time zone

```
SAN384C-6# clock time zone GMT -5 0
```

6.5 Network Time Protocol

Accurate time is essential, and for some environments it is critical. Accurate time ensures the efficient running of tasks and events, for example, batch jobs, logging, tape runs, and other such activities. Accurate time aids troubleshooting, and it is mandatory for accurate monitoring.

The date and time can be set manually on individual devices, or it can be gathered from a reliable time source by using the NTP. Your environment can use either method or a mix of both. Setting and maintaining time manually on individual devices can be time-consuming, and it is more prone to error and likely to lead to inaccurate time (at some point) because the devices are not synchronized with each other. It is better to have a reliable time source that all devices synchronize to. To improve reliability, there should be multiple reliable time sources.

A *time source* is the device that provides the accurate time for other devices, and it is referred to as the *time server*. A device that receives the time is called the *time client*. An IBM c-type switch can be either a time server, a time client, or both. It also can act as a time peer, and when acting as a peer it operates primarily as a time client, but if the primary time server is not available, then it can operate as a time server. If there are other time peers in the environment, then they negotiate over which one takes on the role of time server. Up to 64 NTP entities (servers or peers) can be configured on a switch. To participate in NTP, switches must have IP connectivity to the servers, clients, and peers as required. NTP does not require a license.

When using NTP, it is important to work within certain guidelines. The list below (not exhaustive) shows some of these guidelines. The word *server* refers to both servers and peers.

- Always get permission to use upstream time servers.
- Use time servers that are as close as possible to the time clients.
- Spread the workload between different time servers so that no single one is overloaded.
- Try to ensure that different upstream time servers get their time from different sources in case that source (further upstream) goes down.

For a full listing of NTP configuration options, see Configuring NTP.

6.5.1 Configuring NTP by using the GUI

This section shows how to configure NTP on switches by using the time zone option in DCNM.

Configuring NTP

Start the DCNM web client and select **Device Manager** → **Admin** → **NTP**, as shown in Figure 6-12.

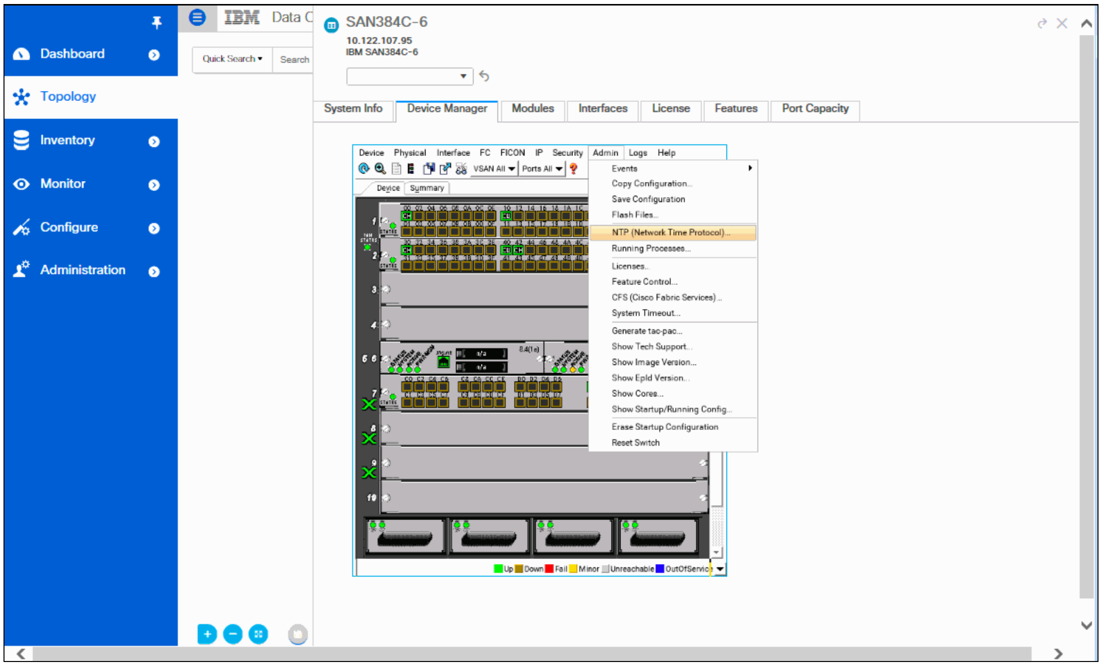

Figure 6-12 Configuring NTP

In Figure 6-13, click **Create**.

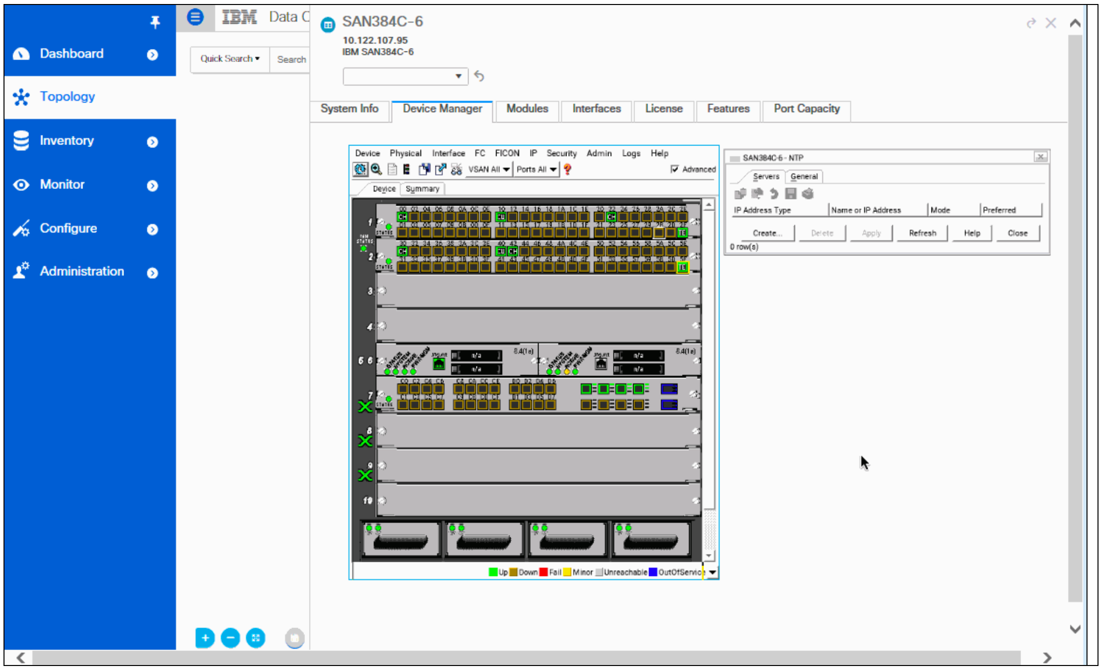

Figure 6-13 Creating an NTP Server or Peer

In Figure 6-14, enter the IP address, mode (Server or Peer), and whether it is preferred. Click **Create**.

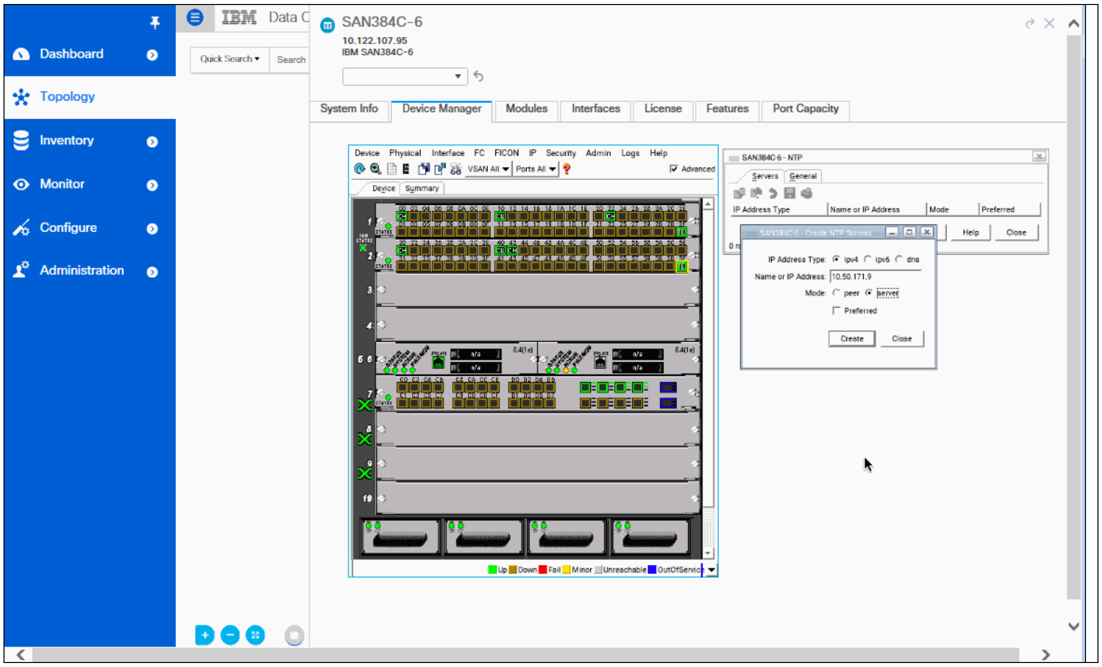

IP Address Type: ● ipv4 ○ ipv6 ○ dns

Name or IP Address: 10.50.171.9

Mode: ○ peer ● server

□ Preferred

Figure 6-14 Configuring the IP address, Mode, and Preferred

Figure 6-15 on page 205 shows the newly created NTP Server or Peer. The dialog box remains open to allow more NTP Servers or Peers to be created. Click **Create** or **Close** as required.

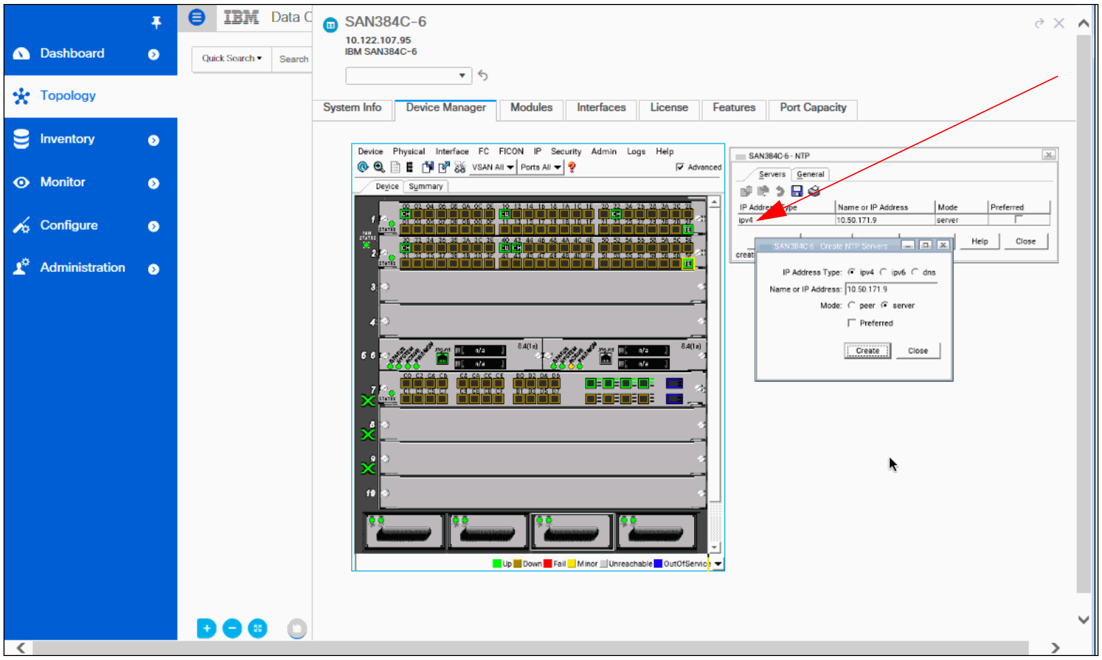

SAN384C-6 - NTP

Servers | General

IP Address Type	Name or IP Address	Mode	Preferred
ipv4	10.50.171.9	server	□

SAN384C-6 - Create NTP Servers

IP Address Type: ● ipv4 ○ ipv6 ○ dns

Name or IP Address: 10.50.171.9

Mode: ○ peer ● server

□ Preferred

Create | Close

Help | Close

Figure 6-15   NTP Server or Peer displayed

Configuring the time zone

Start DCNM and select **Device Manager** → **Physical** → **System**, as shown in Figure 6-16.

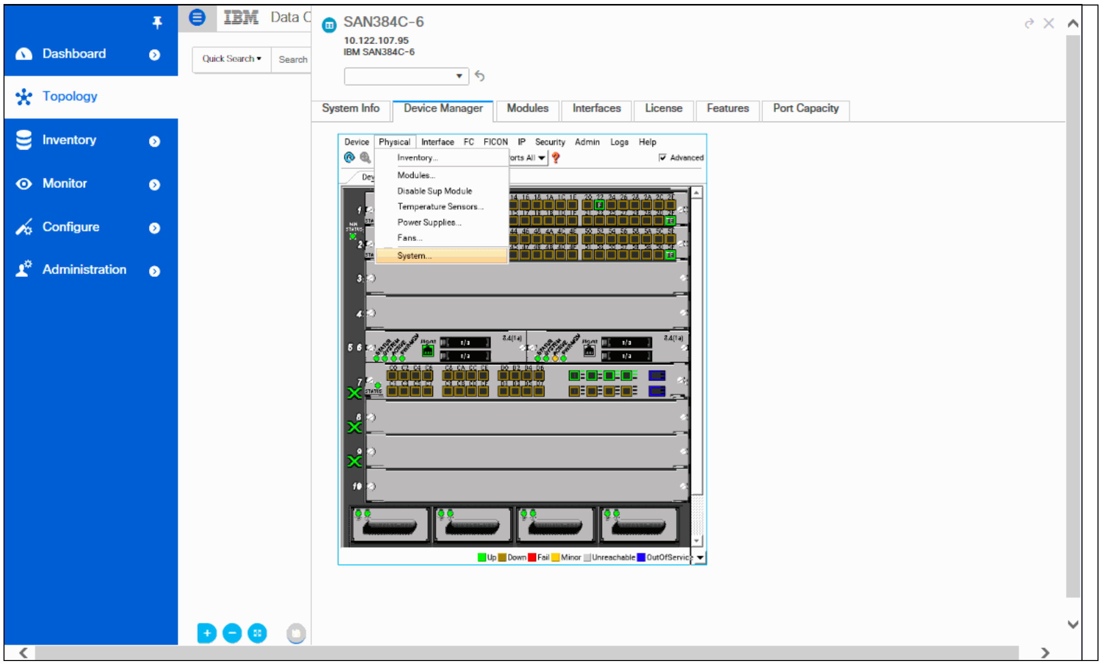

Figure 6-16   Configuring the time zone

Enter the time zone parameters, as shown in Figure 6-17.

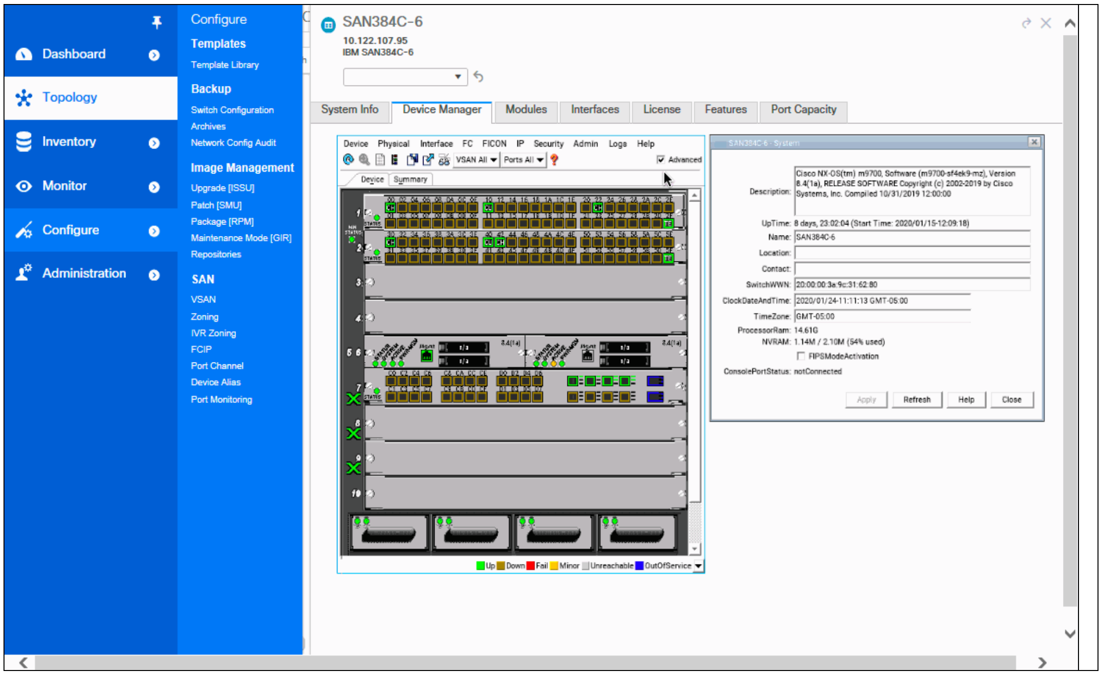

Description: Cisco NX-OS(tm) m9700, Software (m9700-sf4ek9-mz), Version 8.4(1a), RELEASE SOFTWARE Copyright (c) 2002-2019 by Cisco Systems, Inc. Compiled 10/31/2019 12:00:00

UpTime: 8 days, 23:02:04 (Start Time: 2020/01/15-12:09:18)

Name: SAN384C-6

Location:

Contact:

SwitchWWN: 20:00:00:3a:9c:31:62:80

ClockDateAndTime: 2020/01/24-11:11:13 GMT-05:00

TimeZone: GMT-05:00

ProcessorRam: 14.61G

NVRAM: 1.14M / 2.10M (54% used)

☐ FIPSModeActivation

ConsolePortStatus: notConnected

Apply | Refresh | Help | Close

Figure 6-17 Entering the time zone parameters

6.5.2 Configuring NTP by using the CLI

This section shows how to configure NTP by using the CLI. All the examples were performed on a Windows 10 workstation by using PuTTY.

Enabling the NTP feature

NTP is enabled by default, but if you ever need to enable it manually, Example 6-16 shows how to enable NTP.

Example 6-16 Enabling the NTP feature

```
SAN384C-6# configure terminal
SAN384C-6(config)# feature ntp
```

Disabling the NTP feature

Example 6-17 shows how to disable NTP.

Example 6-17 Disabling the NTP feature

```
SAN384C-6# configure terminal
SAN384C-6(config)# no feature ntp
```

Enabling the NTP server

Example 6-18 on page 207 shows how to configure a switch so that it gets its time from an NTP server. The syntax for the command is:

```
ntp server {ip-address | ipv6-address | dns-name} [prefer] [maxpoll interval]
[minpoll interval]
```

The options are as follows:

- **prefer**: Make this server the preferred time server.
- **maxpoll**: The maximum interval in which to poll the time server. The maximum is 16 seconds, and the default is 6 seconds.
- **minpoll**: The minimal interval in which to poll the time server. The minimum is 4 seconds, and the default is 4 seconds.

Example 6-18 Enabling the NTP server

```
SAN384C-6# configure terminal
SAN384C-6(config)# ntp server 10.122.107.100
```

Disabling the NTP server

Example 6-19 shows how to configure a switch so it does not get its time from an NTP server. The syntax for the command is:

```
no ntp server {ip-address | ipv6-address | dns-name}
```

Example 6-19 Disabling the NTP server

```
SAN384C-6# configure terminal
SAN384C-6(config)# ntp server 10.122.107.100
```

Enabling an NTP peer

Example 6-20 shows how to configure a switch so that it receives its time from an NTP peer or sends time to an NTP peer. The syntax for the command is:

```
ntp peer {ip-address | ipv6-address | dns-name} [prefer] [maxpoll interval]
[minpoll interval]
```

The options are as follows:

- **prefer**: Make this server the preferred time server.
- **maxpoll**: The maximum interval in which to poll the time server. The maximum is 16 seconds, and the default is 6 seconds.
- **minpoll**: The minimal interval in which to poll the time server. The minimum is 4 seconds, and the default is 4 seconds.

Example 6-20 Enabling an NTP peer

```
SAN384C-6# configure terminal
SAN384C-6(config)# ntp peer 10.122.107.101
```

Disabling an NTP peer

Example 6-21 shows how to configure a switch so it does not get its time from an NTP peer or send time to an NTP peer. The syntax for the command is:

```
no ntp peer {ip-address | ipv6-address | dns-name}
```

Example 6-21 Disabling an NTP peer

```
SAN384C-6# configure terminal
SAN384C-6(config)# no ntp peer 10.122.107.101
```

Resynchronizing NTP

If a switch has lost its time synchronization with a server or peer, it might need to be resynchronized. Example 6-22 shows how to restart the synchronization process.

Example 6-22 Resynchronizing NTP

```
SAN384C-6# ntp sync-retry
```

Verifying NTP

There are several commands that can be run to determine the status of NTP. Example 6-23 shows one of them.

Example 6-23 Verifying NTP

```
SAN384C-6# show running-config ntp
!Command: show running-config ntp
!Time: Fri Jan 24 6:52:47 2020

version 8.4(1a)
logging level ntp 6
ntp server 10.122.107.100 prefer
ntp server 10.122.107.101
ntp peer 10.122.107.96 prefer
```

6.6 Installing Data Center Network Manager

To use DCNM to manage an IBM c-type switch requires the correct workstation hardware, workstation software, and switch licenses.

6.6.1 License requirements

DCNM allows for two types of licensing schemes; server-based licensing and switch-based licensing. IBM c-type switches support only switch-based licensing. Consider the following items regarding the choice of switch-based licensing:

- All the licenses that are configured for a switch may be shown by using the **`show license`** command. With server-based licensing, the DCNM license is not listed.
- It is possible to have multiple DCNM instances managing the same switch. For example, a customer might want to have one instance for the support team who configures the switch and another instance for the monitoring team who only monitors the switch. This situation would not incur any additional license costs.

Important: IBM c-type switches support only switch-based licenses. Server-based licenses are not supported.

Licenses can be purchased with the switch or later.

6.6.2 Downloading DCNM

To download the DCNM software, you must have the correct website access credentials and profile. If you are unable to download the DCNM software because you do not have the correct access, then apply for it at the Cisco website. There are several access types and levels. You must apply for access according to your status. The DCNM software, release notes, and documentation can be downloaded from Cisco Nexus Dashboard Fabric Controller (Formerly DCNM).

6.6.3 DCNM server requirements

DCNM server requirements change over time as new versions are released. This section covers the latest version at the time of writing, which is Version 11.3(1). The following 64-bit server OSs are supported for SAN deployments:

- Red Hat Enterprise Linux releases 7.3, 7.4, 7.6, and 7.7
- Open Virtual Appliance (OVA) with an integrated CentOS Linux release 7.6
- ISO Virtual Appliance with an integrated CentOS Linux release 7.6
- Microsoft Windows Server 2016
- Microsoft Windows Server 2012 R2

At the time of writing, DCNM Server is distributed with Java Runtime Engine (JRE) 11.0.2. It installs into the following folder:

```
<dcnm_root>/java/jdk11
```

The following databases are supported:

- PostgreSQL 9.4.5 (included with DCNM)
- Oracle 12c Real Application Cluster (RAC) (conventional nonpluggable installation)
- Oracle 12c Enterprise Edition (conventional non-pluggable installation)
- Oracle 11g Express (XE), Standard and Enterprise Editions, and Oracle 11g RACs
- ElasticSearch 5.6 release (included with DCNM, with performance and telemetry data)

Note: DCNM version 11.3(1) does not support the Oracle 12c pluggable database version.

6.6.4 DCNM installation

This section shows the process to install DCNM 11.3(1) on a Windows 2016 server. In the future, it is expected that most deployments will happen by using a prepackaged OVA. The installation shows a first-time installation. If this installation is an upgrade, then follow the upgrade procedure in 6.6.5, “Upgrading to DCNM 11.3x” on page 224.

Note: Before installing DCNM, it is a best practice to create a DCNM Admin user locally on all the switches. This user will be used by DCNM to log in and manage the switch infrastructure.

The DCNM 11.3(1) installer comes packaged with a PostgreSQL database. The installer file for the version that is being installed in this example is `dcnm-installer-x64-windows.11.3.1.exe.zip`. Use PostgreSQL for production enterprise environments unless you have a supported installation of Oracle database that is available, and you consider it to be a better option.

Start the installer. Figure 6-18 shows the initial window loading.

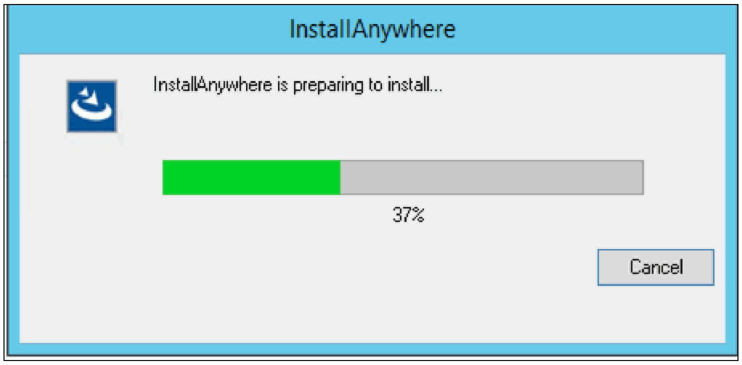

InstallAnywhere is preparing to install...

37%

Figure 6-18 DCNM installation loading dialog

Figure 6-19 shows the OEM vendor options. Select **IBM**, and then click **Next**.

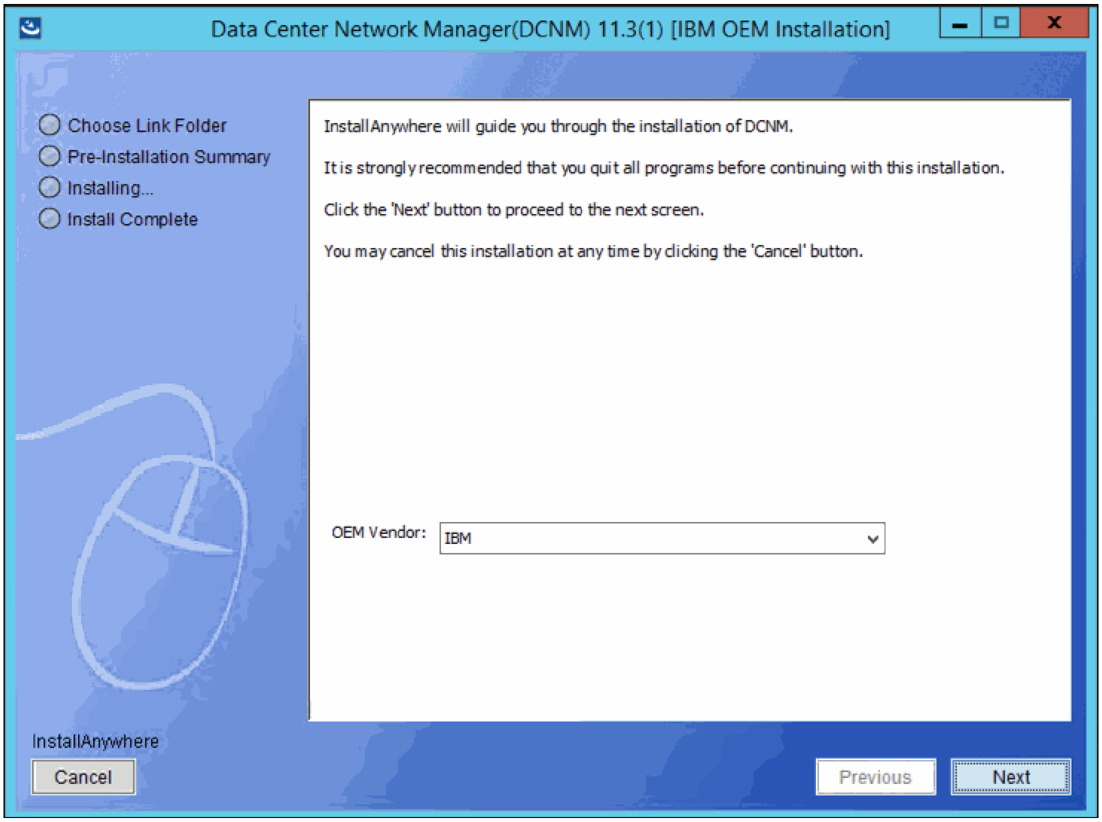

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]
○ Choose Link Folder
○ Pre-Installation Summary
○ Installing...
○ Install Complete

InstallAnywhere will guide you through the installation of DCNM.

It is strongly recommended that you quit all programs before continuing with this installation.

Click the 'Next' button to proceed to the next screen.

You may cancel this installation at any time by clicking the 'Cancel' button.

OEM Vendor: IBM

InstallAnywhere
Cancel
Previous
Next

Figure 6-19 Vendor selection

In Figure 6-20, the default path is shown. If this path needs to be changed, then type in the new path. If the **Secure Ciphers** checkbox is selected, then only switches with strong ciphers will be discovered by DCNM. In our example, this checkbox remains cleared. Click **Next**.

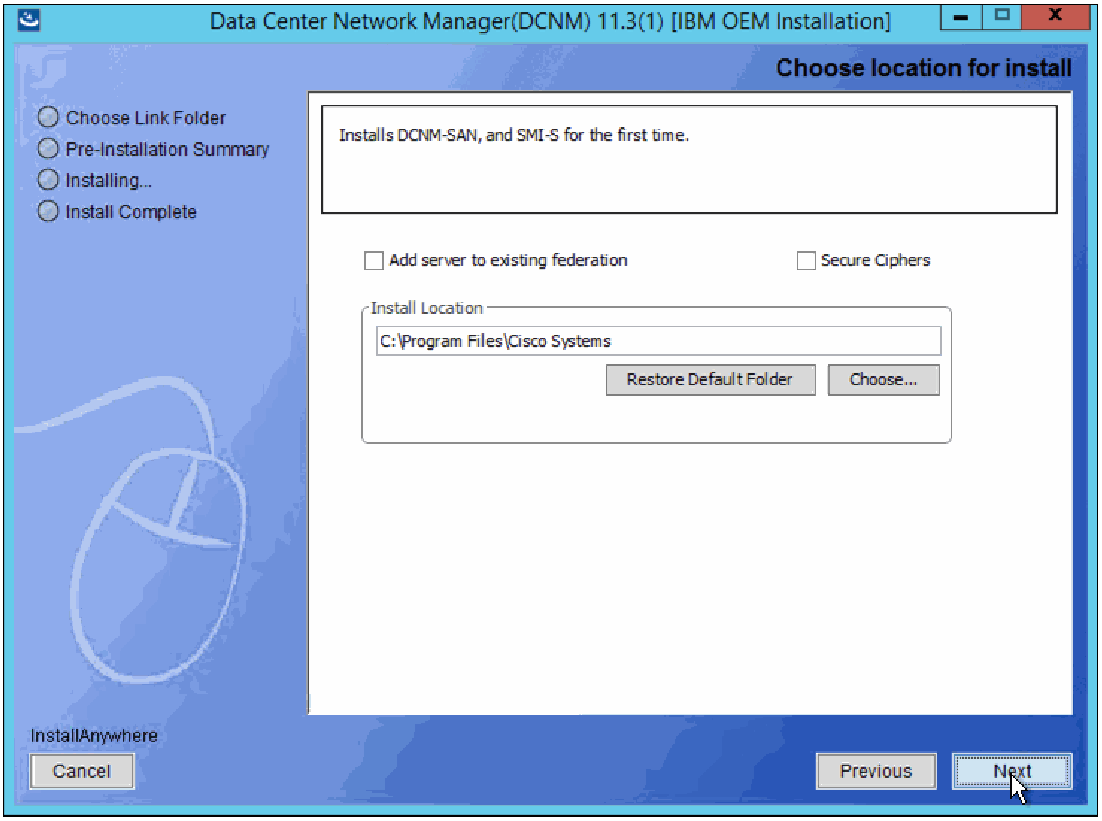
Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]
Choose location for install
○ Choose Link Folder
○ Pre-Installation Summary
○ Installing...
○ Install Complete
InstallAnywhere
Cancel
Previous
Next

Installs DCNM-SAN, and SMI-S for the first time.

□ Add server to existing federation   □ Secure Ciphers

Install Location: C:\Program Files\Cisco Systems

Restore Default Folder
Choose...

Figure 6-20 Installation path

Figure 6-21 shows the options that are available for the following items:

- The Relational Database Management System (RDBMS)
- DCNM Database (DB) User
- Install Location

Enter the details as required, and then click **Next**.

> **Note:** The DCNM DB User password is a standard Windows user, so the password must be compliant with the password policy.

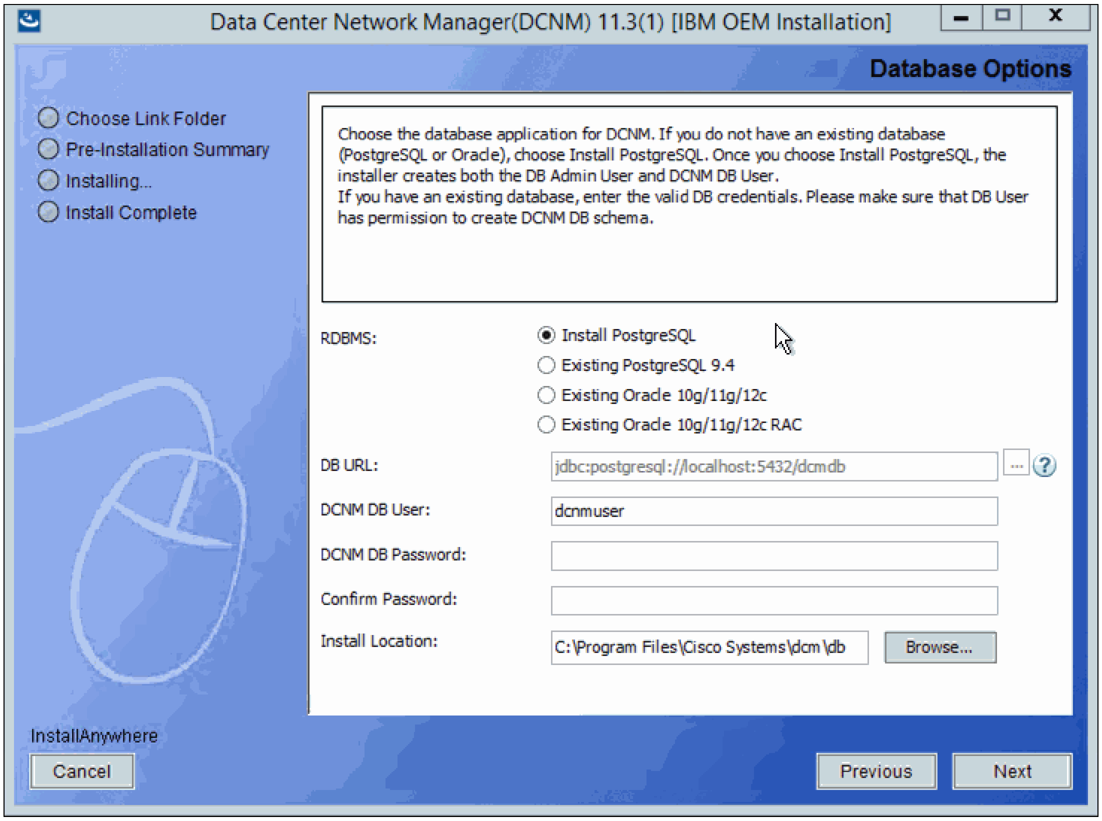

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Database Options

Choose the database application for DCNM. If you do not have an existing database (PostgreSQL or Oracle), choose Install PostgreSQL. Once you choose Install PostgreSQL, the installer creates both the DB Admin User and DCNM DB User.
If you have an existing database, enter the valid DB credentials. Please make sure that DB User has permission to create DCNM DB schema.

RDBMS: ● Install PostgreSQL
○ Existing PostgreSQL 9.4
○ Existing Oracle 10g/11g/12c
○ Existing Oracle 10g/11g/12c RAC

DB URL: jdbc:postgresql://localhost:5432/dcmdb

DCNM DB User: dcnmuser

DCNM DB Password:

Confirm Password:

Install Location: C:\Program Files\Cisco Systems\dcm\db

Figure 6-21 Selecting the database

In Figure 6-22, click **Yes** for the DCNM Installer dialog box. The Secondary Logon Service is required for the installer to initialize the database.

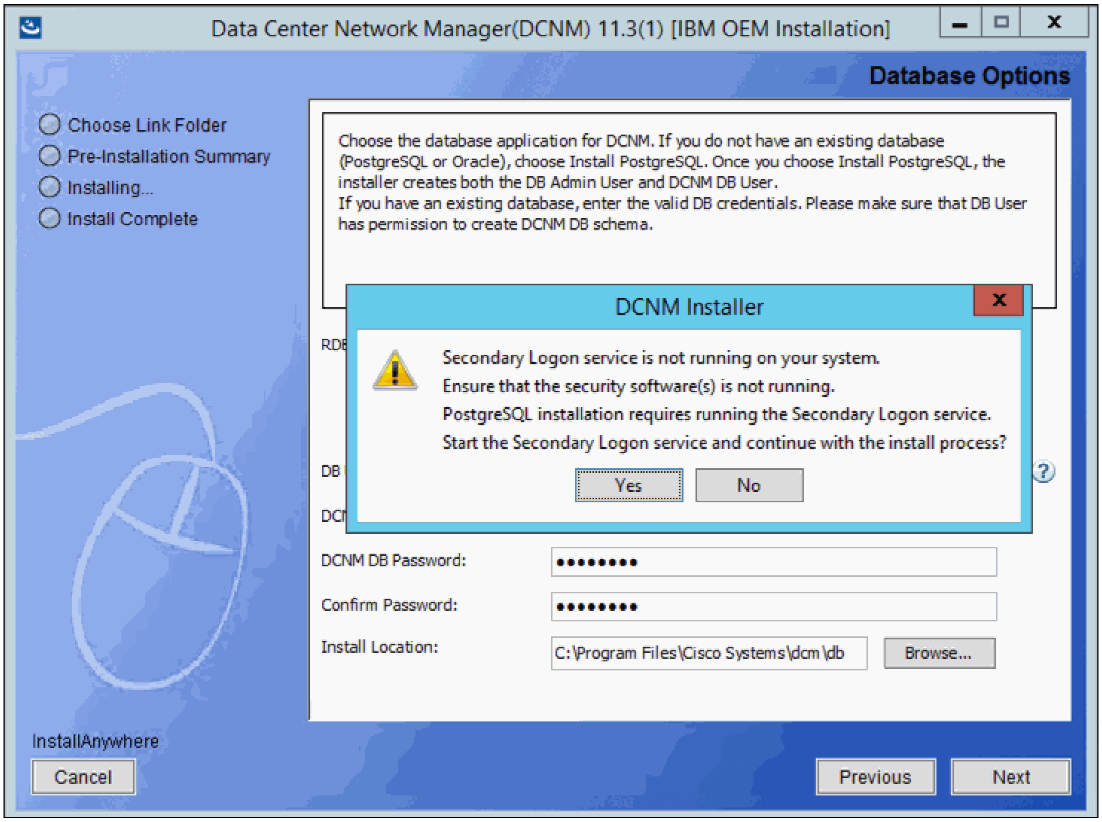

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Database Options

Choose the database application for DCNM. If you do not have an existing database (PostgreSQL or Oracle), choose Install PostgreSQL. Once you choose Install PostgreSQL, the installer creates both the DB Admin User and DCNM DB User.
If you have an existing database, enter the valid DB credentials. Please make sure that DB User has permission to create DCNM DB schema.

DCNM Installer

Secondary Logon service is not running on your system.
Ensure that the security software(s) is not running.
PostgreSQL installation requires running the Secondary Logon service.
Start the Secondary Logon service and continue with the install process?

DCNM DB Password: ••••••••

Confirm Password: ••••••••

Install Location: C:\Program Files\Cisco Systems\dcm\db

Figure 6-22 Secondary Logon Service for PostgreSQL installation

In Figure 6-23, select the correct **Server IP address** and **SAN Webserver Port**. For our example, we used the default SAN Webserver Port, which is 443. Click **Next**. The IP address and port are used by your web browser to log in to DCNM. If the default 443 port is used, then in your web browser you need to put only the IP address into the address field. If the port is changed, then both the IP address and port number are required.

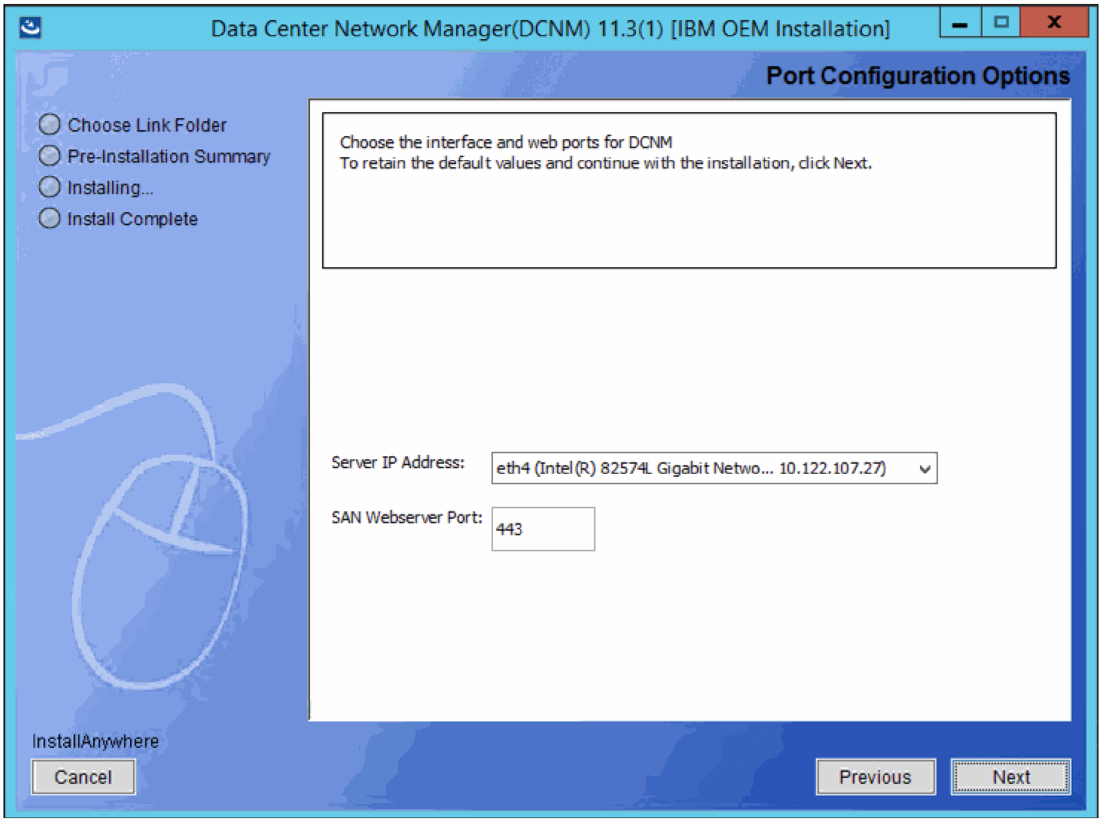

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Port Configuration Options

Choose the interface and web ports for DCNM
To retain the default values and continue with the installation, click Next.

Server IP Address: eth4 (Intel(R) 82574L Gigabit Netwo... 10.122.107.27)

SAN Webserver Port: 443

Figure 6-23 Selecting the network adapter

In Figure 6-24, select the **Archive Folder Location**, which is the file path to a location where you want to store and back up the switch configuration file and user preferences. This folder location is typically on the local server. However, when a federated DCNM deployment is wanted, the folder location should be on a shared remote directory. Click **Next**.

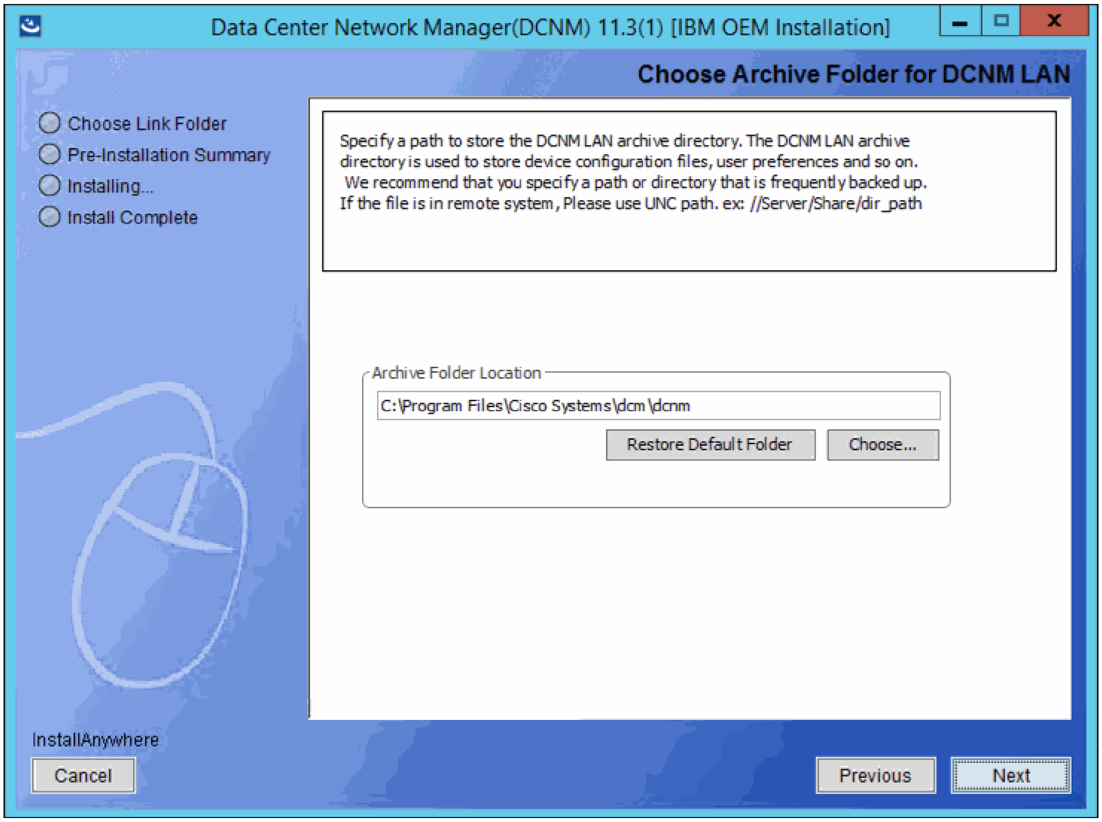

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Choose Archive Folder for DCNM LAN

Specify a path to store the DCNM LAN archive directory. The DCNM LAN archive directory is used to store device configuration files, user preferences and so on.
We recommend that you specify a path or directory that is frequently backed up.
If the file is in remote system, Please use UNC path. ex: //Server/Share/dir_path

Archive Folder Location: C:\Program Files\Cisco Systems\dcm\dcnm

Figure 6-24   Archive Folder Location

In Figure 6-25, set a local **Admin Username** and **Password** for the DCNM-SAN. This combination of username and password is locally stored, and is required when accessing the DCNM GUI. This password does not need to be compliant with the Windows password policy, but it is a best practice if it is. Click **Next**.

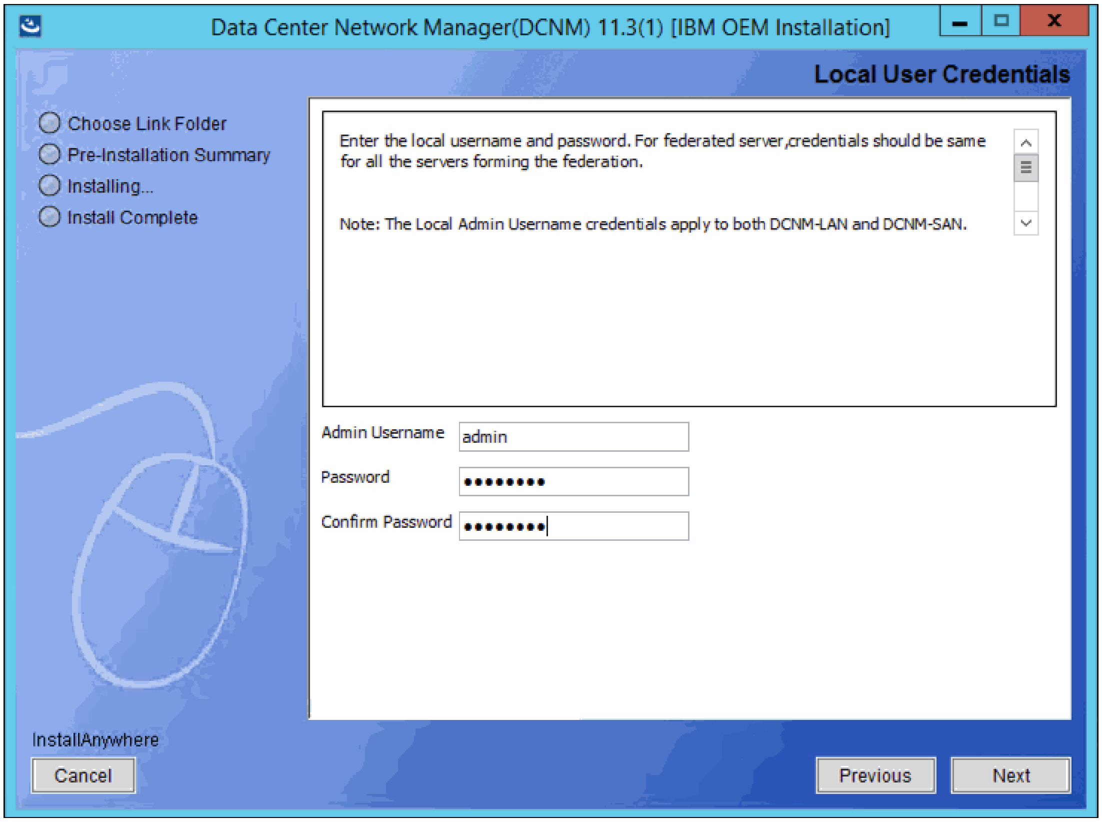

Enter the local username and password. For federated server,credentials should be same for all the servers forming the federation.

Note: The Local Admin Username credentials apply to both DCNM-LAN and DCNM-SAN.

Admin Username: admin

Password: ••••••••

Confirm Password: ••••••••

Figure 6-25 Database username

In Figure 6-26, if your environment supports a centralized authentication solution, you can configure the necessary parameters here. If remote authentication is not configured at this stage, it can be configured later. Remote authentication methods include the following ones:

- Remote Authentication Dial-In User Service (RADIUS)
- Terminal Access Controller Access Control System Plus (TACACS+)

A third remote authentication method is the Lightweight Directory Access Protocol (LDAP). This method can be configured only after DCNM is running, but not during deployment.

After remote authentication is configured, the Local DCNM credentials do not work. In our example, we selected the local database for user authentication. Click **Next**.

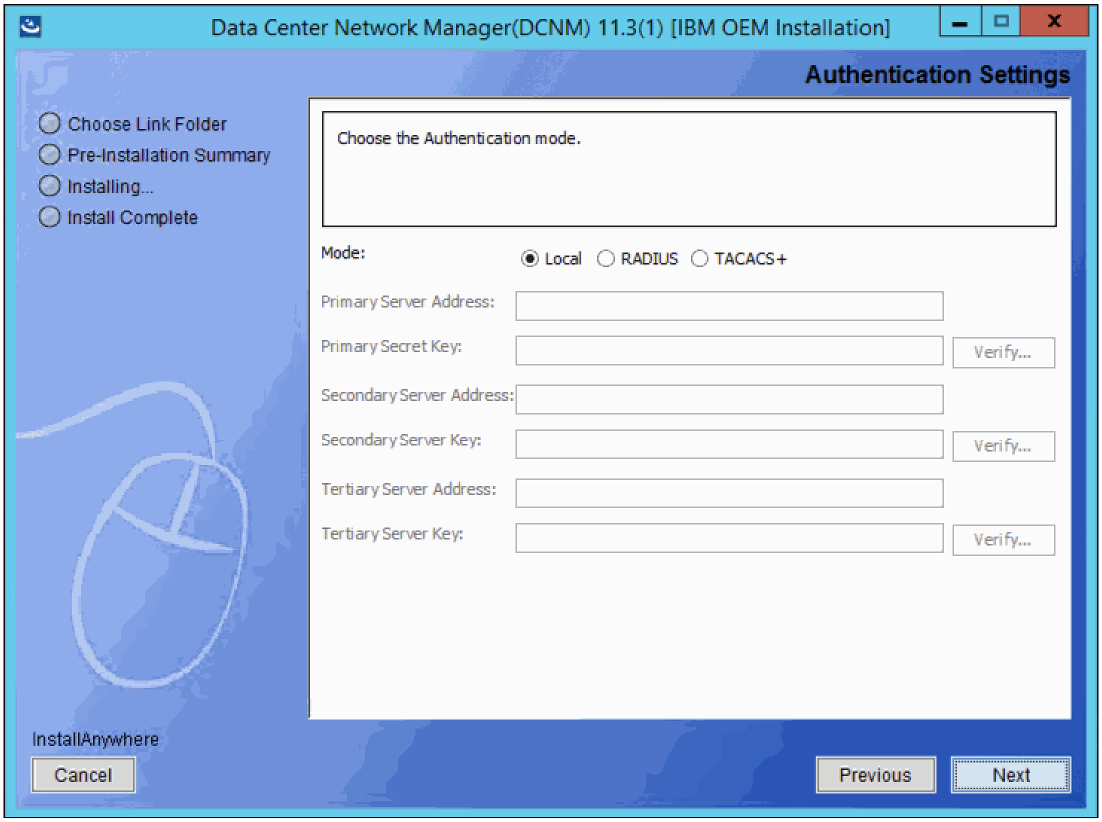

Choose the Authentication mode.

Mode: ● Local ○ RADIUS ○ TACACS+

Primary Server Address:

Primary Secret Key:

Secondary Server Address:

Secondary Server Key:

Tertiary Server Address:

Tertiary Server Key:

Figure 6-26 Authentication Mode

In Figure 6-27, select the product icon preferences. Click **Next**.

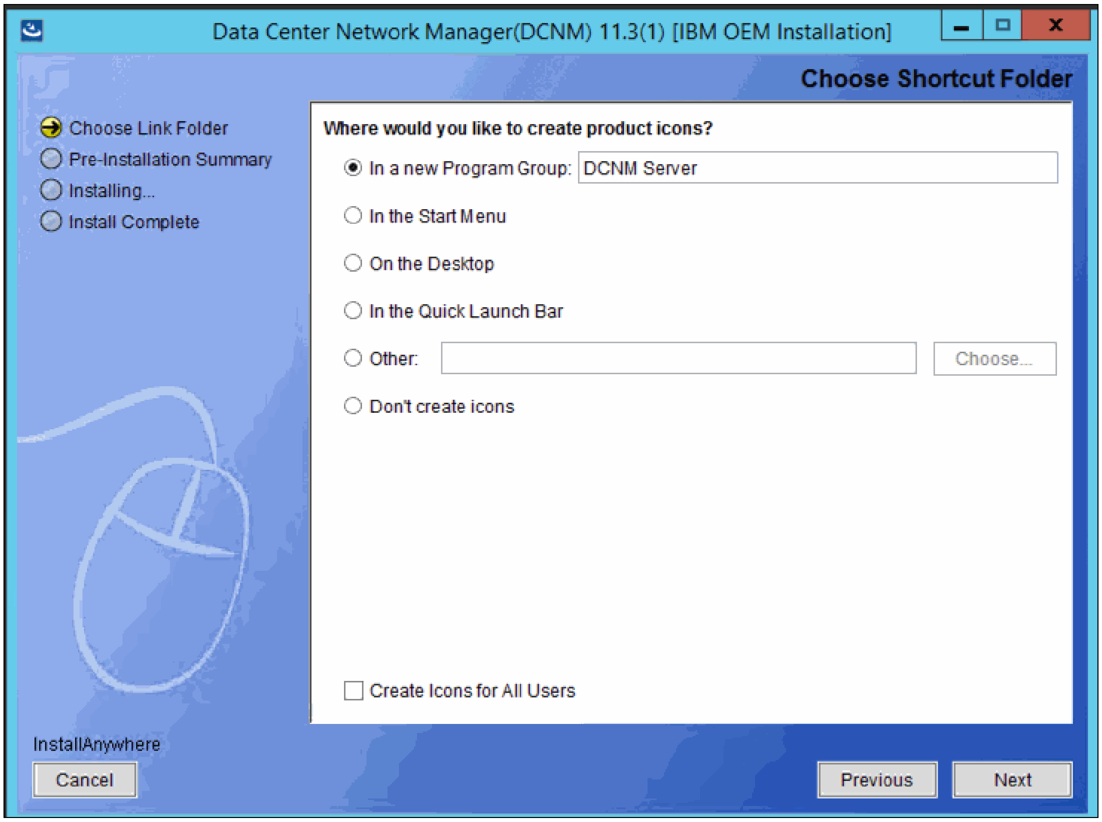

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Choose Shortcut Folder

Where would you like to create product icons?

● In a new Program Group: DCNM Server

○ In the Start Menu

○ On the Desktop

○ In the Quick Launch Bar

○ Other:

○ Don't create icons

□ Create Icons for All Users

Figure 6-27 Product icons

In Figure 6-28 and Figure 6-29 on page 220, review the summary of the installation. The option **Create Icons for All users** allows all users who log in to the workstation access to the DCNM icons, but it still requires valid user credentials to log in to DCNM. Click **Next**.

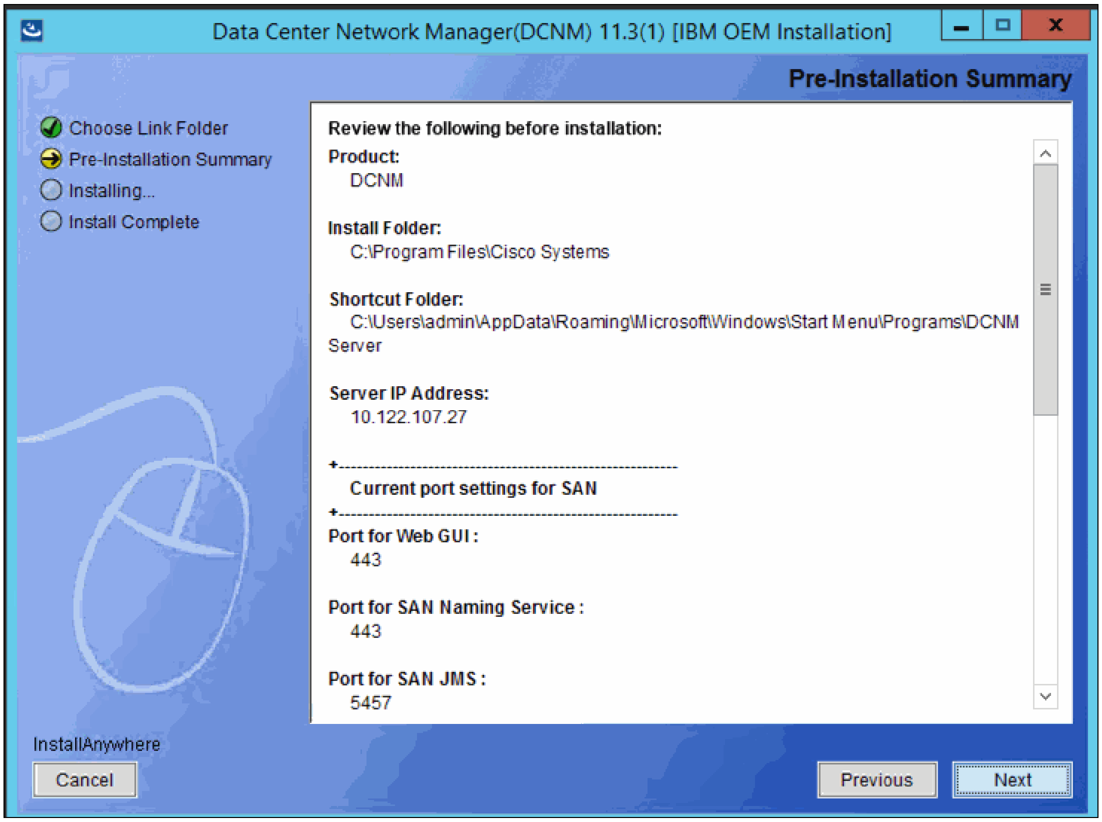

Review the following before installation:
Product:
DCNM

Install Folder:
C:\Program Files\Cisco Systems

Shortcut Folder:
C:\Users\admin\AppData\Roaming\Microsoft\Windows\Start Menu\Programs\DCNM Server

Server IP Address:
10.122.107.27

```
+-------------------------------------------------------
   Current port settings for SAN
+-------------------------------------------------------
```

Port for Web GUI :
443

Port for SAN Naming Service :
443

Port for SAN JMS :
5457

Figure 6-28   Configuration summary window 1

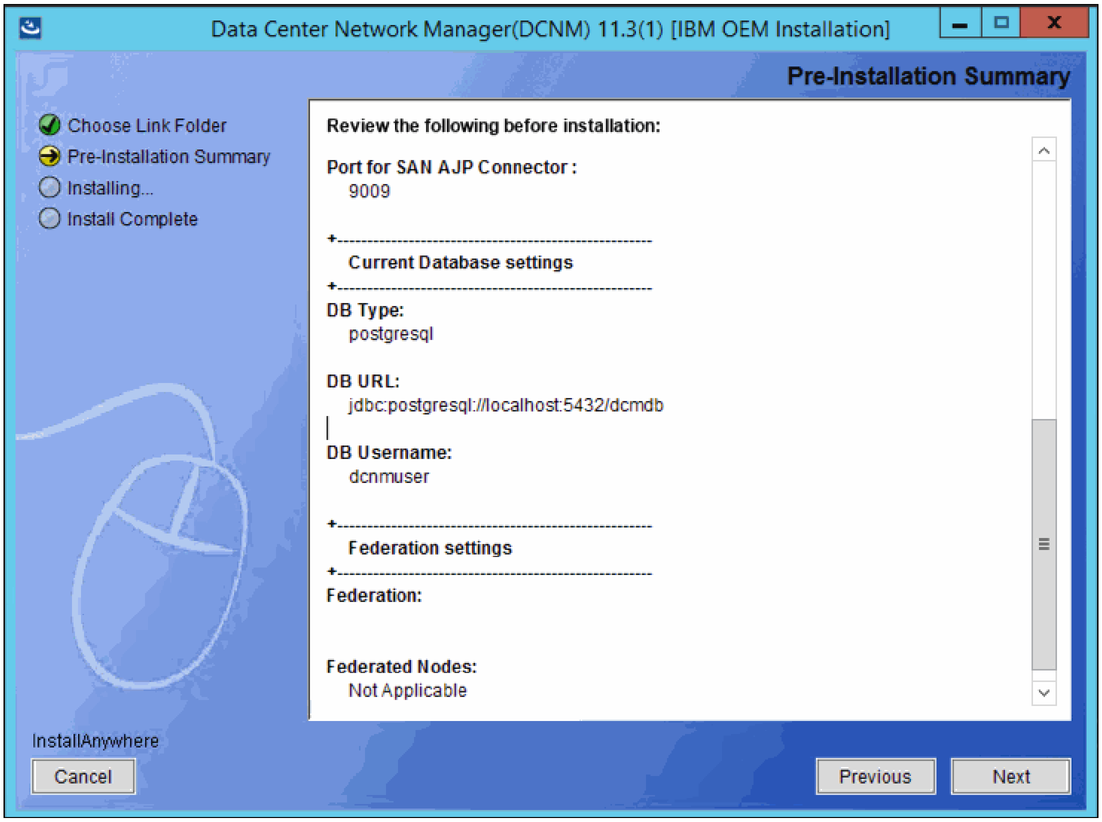

```
Review the following before installation:

Port for SAN AJP Connector :
    9009

+----------------------------------------------------
    Current Database settings
+----------------------------------------------------
DB Type:
    postgresql

DB URL:
    jdbc:postgresql://localhost:5432/dcmdb

DB Username:
    dcnmuser

+----------------------------------------------------
    Federation settings
+----------------------------------------------------
Federation:


Federated Nodes:
    Not Applicable
```

Figure 6-29 Configuration summary window 2

Figure 6-30 shows the confirmation window. Click **Yes**.

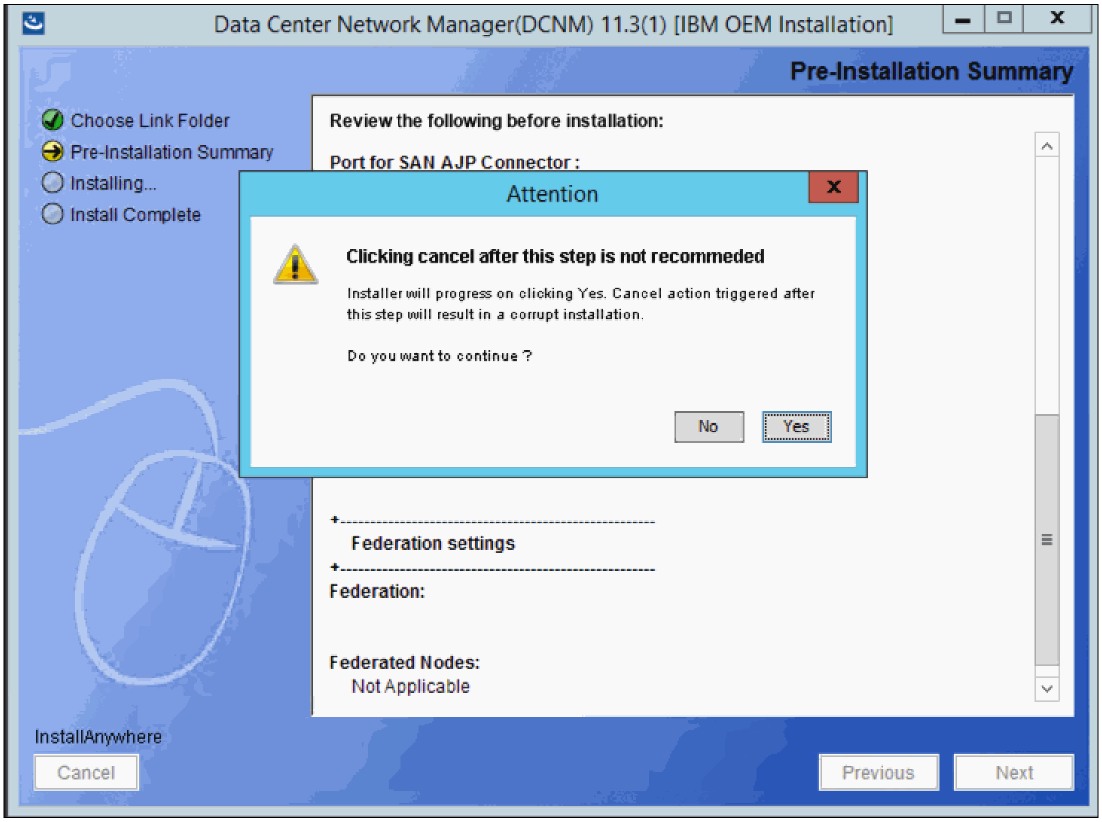

Review the following before installation:

Port for SAN AJP Connector :

Clicking cancel after this step is not recommeded

Installer will progress on clicking Yes. Cancel action triggered after this step will result in a corrupt installation.

Do you want to continue ?

```
+-----------------------------------------------------
   Federation settings
+-----------------------------------------------------
Federation:


Federated Nodes:
   Not Applicable
```

Figure 6-30 Confirmation window

Figure 6-31 and Figure 6-32 on page 223 show the software and configuration options being installed.

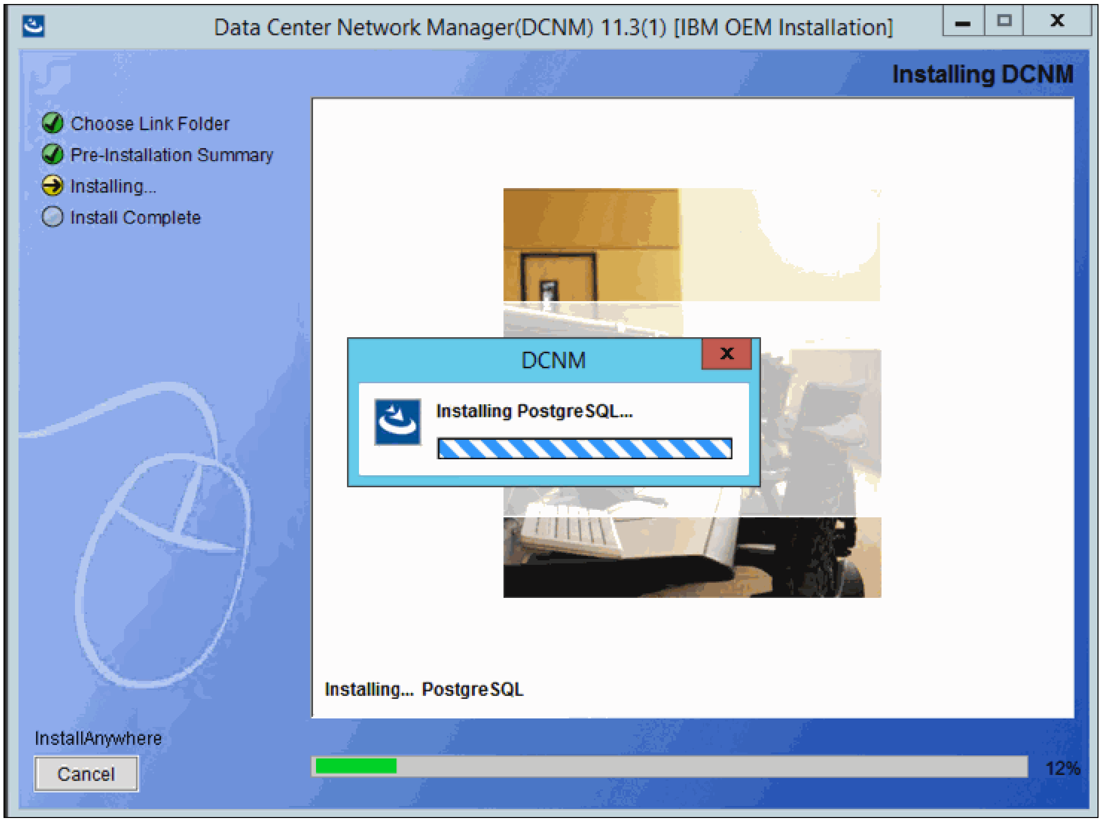

Installing PostgreSQL...

Installing... PostgreSQL

Figure 6-31 Installing PostgreSQL

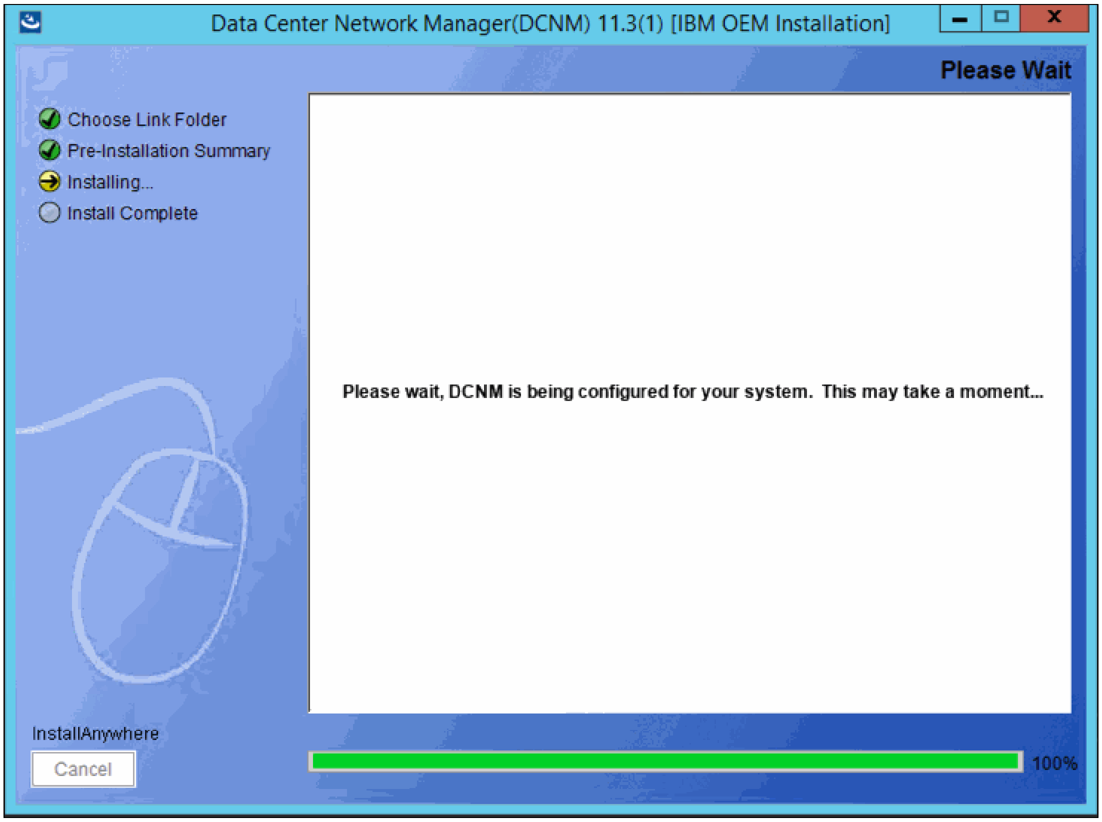

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Please Wait

Please wait, DCNM is being configured for your system. This may take a moment...

Figure 6-32   Configuring the system

Figure 6-33 shows the summary of the installation. Click **Done**.

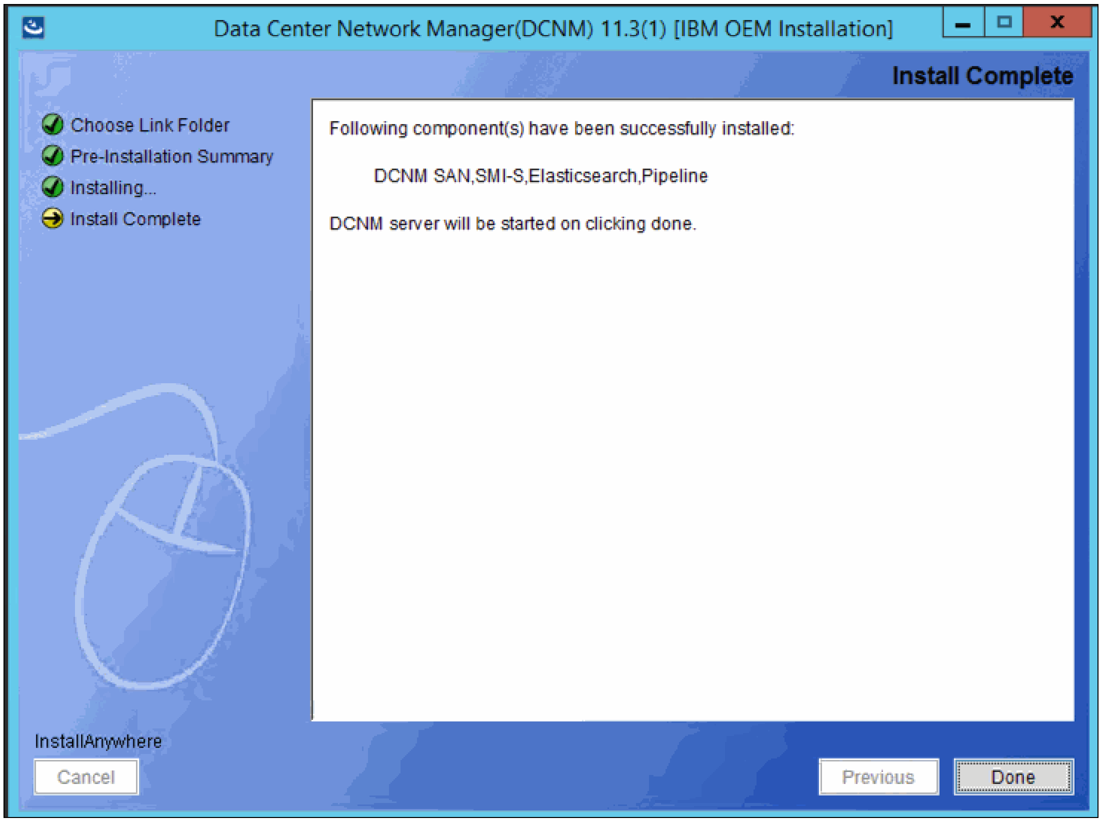

Data Center Network Manager(DCNM) 11.3(1) [IBM OEM Installation]

Install Complete

- ✓ Choose Link Folder
- ✓ Pre-Installation Summary
- ✓ Installing...
- ➔ Install Complete

Following component(s) have been successfully installed:

DCNM SAN,SMI-S,Elasticsearch,Pipeline

DCNM server will be started on clicking done.

Figure 6-33 Installation complete

The installation is now complete.

6.6.5 Upgrading to DCNM 11.3x

To upgrade DCNM 11.3x, Version 11.2x already must be installed. The method to upgrade DCNM depends on the version that is installed and the new version to be installed, and the experience can be different between versions. For a comprehensive guide to upgrading DCNM on all supported OSs, see Cisco Nexus Dashboard Fabric Controller (Formerly DCNM).

6.7 Logging in to DCNM

This section shows the initial launch window and login procedure only. Further examples showing specific tasks are included in relevant sections of this book.

Start your web browser, and in the address field, type the DCNM IP address that was used during installation. The login window opens. If you changed the default (443) port, then you also need to input the port into the address field. If DCNM does not start, read any error messages and check that all the required services are running.

Note: Only specific browsers and versions are fully tested for DCNM, and they are described in the DCNM release notes. Other versions might not work properly. Untested browsers and versions might lead to artifacts on the display.

The login window opens, as shown in Figure 6-34. Enter the username and password that was defined at installation time, and then click **Login**. Our example shows local authentication.

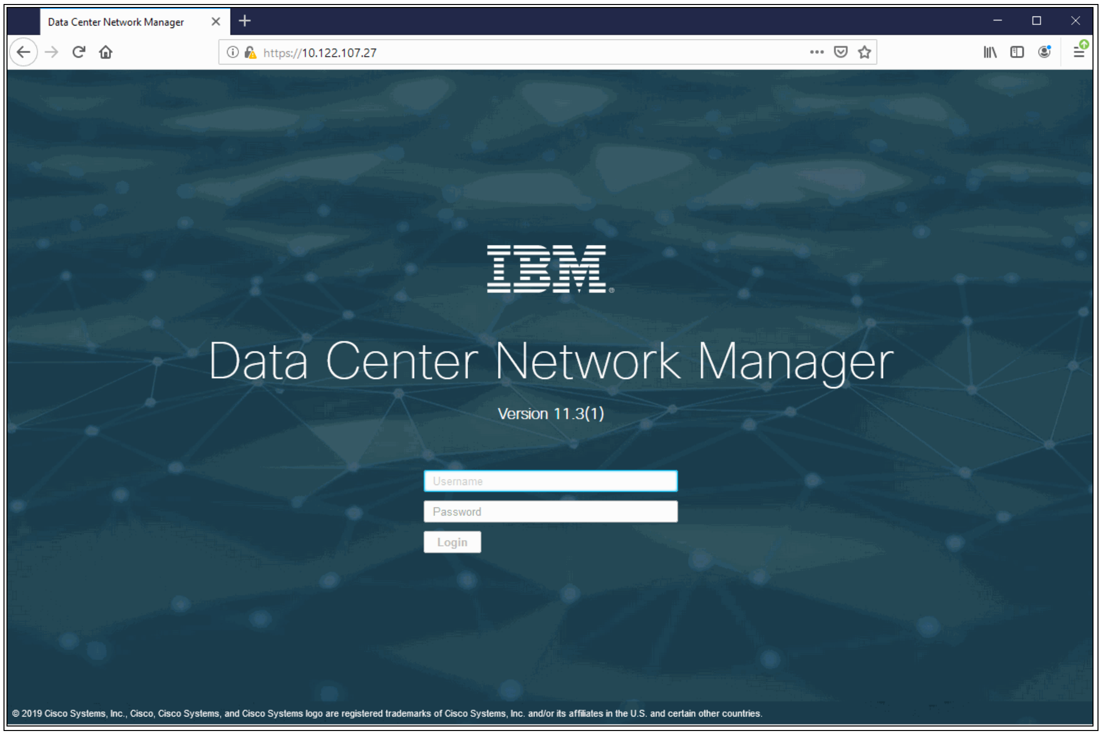
Data Center Network Manager
https://10.122.107.27
IBM
Data Center Network Manager
Version 11.3(1)
Username
Password
Login
© 2019 Cisco Systems, Inc., Cisco, Cisco Systems, and Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Figure 6-34 DCNM login

The opening window opens, as shown in Figure 6-35. The dialog box can be disabled by selecting the **Do not show this message again** checkbox. The dialog box shows text about logging in to LAN switches. If you have only SAN switches, click **No**.

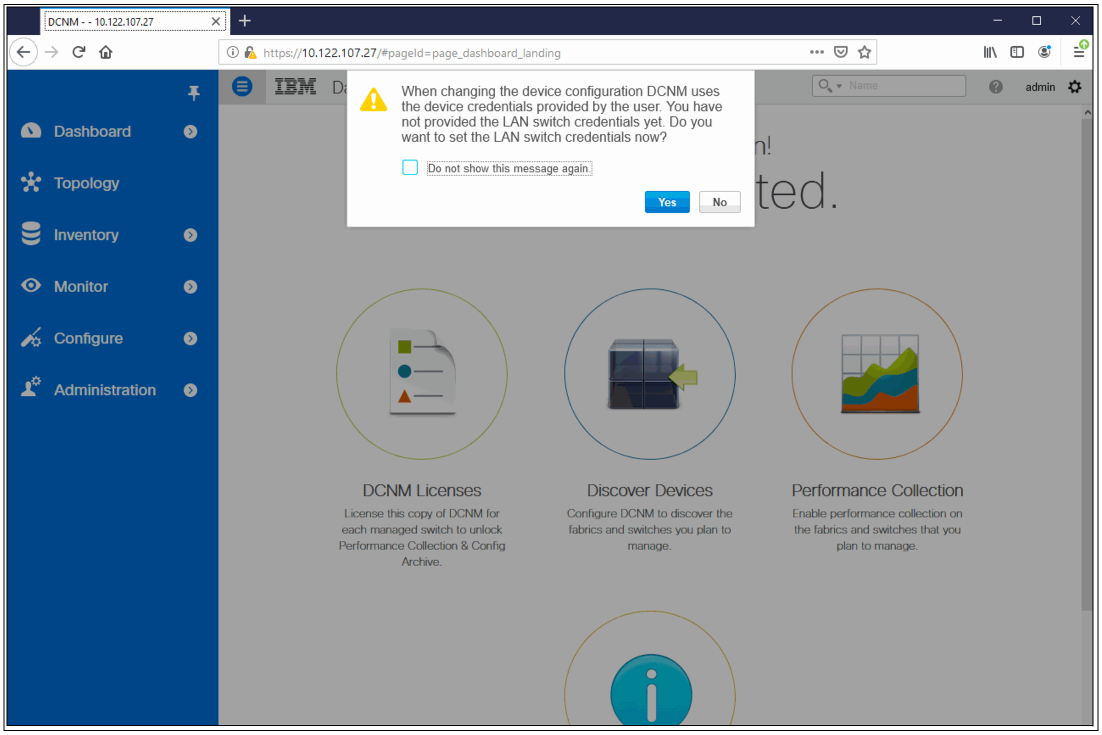

When changing the device configuration DCNM uses the device credentials provided by the user. You have not provided the LAN switch credentials yet. Do you want to set the LAN switch credentials now?

□ Do not show this message again

Yes No

DCNM Licenses

License this copy of DCNM for each managed switch to unlock Performance Collection & Config Archive.

Discover Devices

Configure DCNM to discover the fabrics and switches you plan to manage.

Performance Collection

Enable performance collection on the fabrics and switches that you plan to manage.

Figure 6-35 DCNM opening window

From this point onward, DCNM is ready to use. For a comprehensive guide to using DCNM, see Cisco Nexus Dashboard Fabric Controller (Formerly DCNM).

6.8 Device Manager

DM now comes embedded inside the DCNM web client user interface. To use DM, first install DCNM as described in 6.6.4, “DCNM installation” on page 209, and then start DM from within DCNM by completing the following steps:

1. Log in to DCNM, as described in 6.7, “Logging in to DCNM” on page 224.
2. Select **Topology**, double-click the required device, and select **Show more details** → **Device Manager**.
3. Log in to Device Manager with your user credentials. The logon window is shown in Figure 6-36 on page 227.

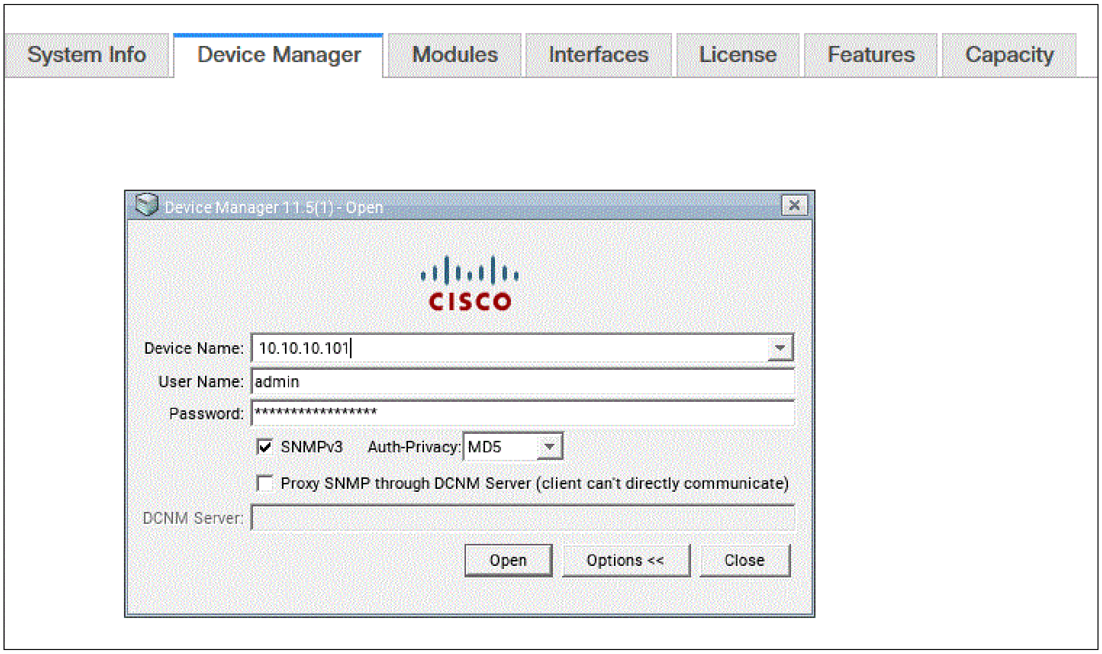

Device Name: 10.10.10.101
User Name: admin
Password: *****************
☑ SNMPv3 Auth-Privacy: MD5
□ Proxy SNMP through DCNM Server (client can't directly communicate)
DCNM Server:

Figure 6-36 Device Manager

6.9 Installing switch licenses

On the IBM c-type switches, there are some features and functions that are part of the base support that is provided in the NX-OS software, and there are other features and functions that are provided only when extra licenses are purchased and installed. For example, there is a mainframe license that is required for mainframe environments. This license enables mainframe Channel Path IDs (CHPIDs) to log in to the FICON virtual storage area networks (VSANs) on the switches. The mainframe license also enables the IBM Control Unit Port (CUP) FICON features. In our lab environment, we are using three licenses:

- Mainframe license: Allows the creation of FICON VSANs and enables certain security flows that are required for mainframes to log in to a SAN fabric. The CUP inband management feature is also enabled.
- Enterprise license: Has several advanced features that it enables, but in our lab environment, it is required so that we can show the use of encryption for Fibre Channel over IP (FCIP) tunnels. This license also enables RBAC on a per VSAN basis.
- DCNM SAN Advanced license: Allows for the historical collection of SAN statistics and the management of multiple SAN fabrics at the same time.

The installation of the licenses for an IBM c-type switch is a two-step process. The first step is to use the Product Activation Key (PAK) that is provided with the system to create the license key files. After this task is done, these license key files must be moved to the switches and installed by using either DCNM or the CLI.

6.9.1 Using the PAK letter to create license keys

To start the process, we need a PAK letter, and then we need to log in to your account:

Login to your account

We start from the initial window for this website, as shown in Figure 6-37. Select **Add New PAK / Token**.

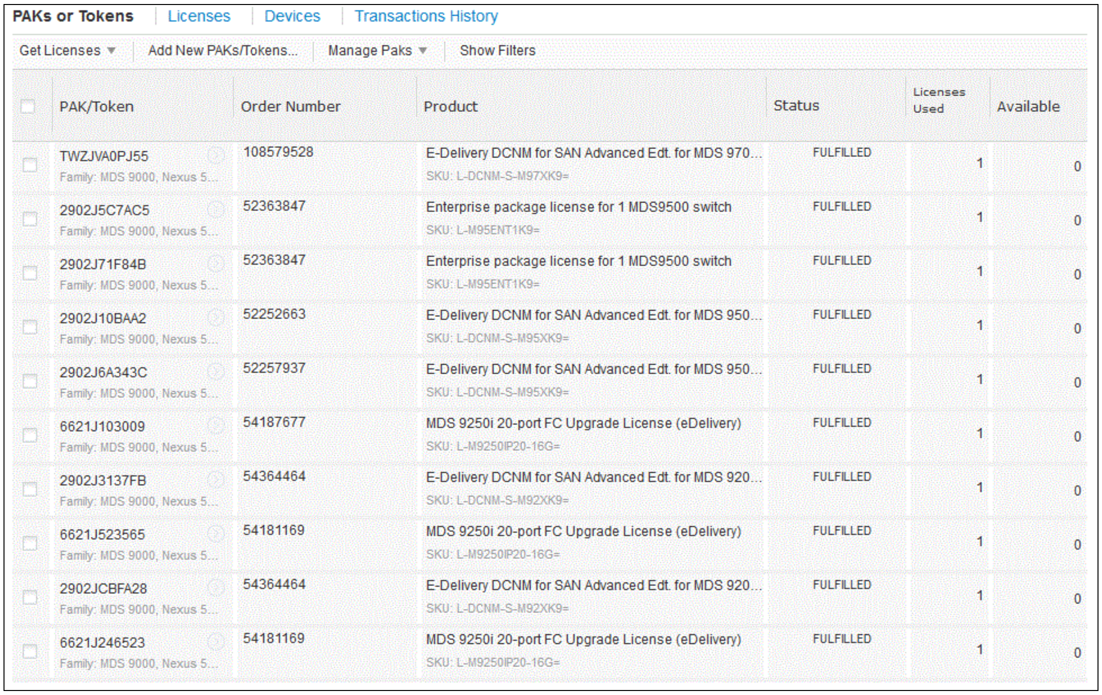

□	PAK/Token	Order Number	Product	Status	Licenses Used	Available
□	TWZJVA0PJ55 Family: MDS 9000, Nexus 5...	108579528	E-Delivery DCNM for SAN Advanced Edt. for MDS 970... SKU: L-DCNM-S-M97XK9=	FULFILLED	1	0
□	2902J5C7AC5 Family: MDS 9000, Nexus 5...	52363847	Enterprise package license for 1 MDS9500 switch SKU: L-M95ENT1K9=	FULFILLED	1	0
□	2902J71F84B Family: MDS 9000, Nexus 5...	52363847	Enterprise package license for 1 MDS9500 switch SKU: L-M95ENT1K9=	FULFILLED	1	0
□	2902J10BAA2 Family: MDS 9000, Nexus 5...	52252663	E-Delivery DCNM for SAN Advanced Edt. for MDS 950... SKU: L-DCNM-S-M95XK9=	FULFILLED	1	0
□	2902J6A343C Family: MDS 9000, Nexus 5...	52257937	E-Delivery DCNM for SAN Advanced Edt. for MDS 950... SKU: L-DCNM-S-M95XK9=	FULFILLED	1	0
□	6621J103009 Family: MDS 9000, Nexus 5...	54187677	MDS 9250i 20-port FC Upgrade License (eDelivery) SKU: L-M9250IP20-16G=	FULFILLED	1	0
□	2902J3137FB Family: MDS 9000, Nexus 5...	54364464	E-Delivery DCNM for SAN Advanced Edt. for MDS 920... SKU: L-DCNM-S-M92XK9=	FULFILLED	1	0
□	6621J523565 Family: MDS 9000, Nexus 5...	54181169	MDS 9250i 20-port FC Upgrade License (eDelivery) SKU: L-M9250IP20-16G=	FULFILLED	1	0
□	2902JCBFA28 Family: MDS 9000, Nexus 5...	54364464	E-Delivery DCNM for SAN Advanced Edt. for MDS 920... SKU: L-DCNM-S-M92XK9=	FULFILLED	1	0
□	6621J246523 Family: MDS 9000, Nexus 5...	54181169	MDS 9250i 20-port FC Upgrade License (eDelivery) SKU: L-M9250IP20-16G=	FULFILLED	1	0

Figure 6-37 The cisco.com/go/license page

Figure 6-38 shows the window where we select the **Enter PAK or Token ID** radio button and enter the PAK number that is printed on the PAK letter. Click **OK** to continue.

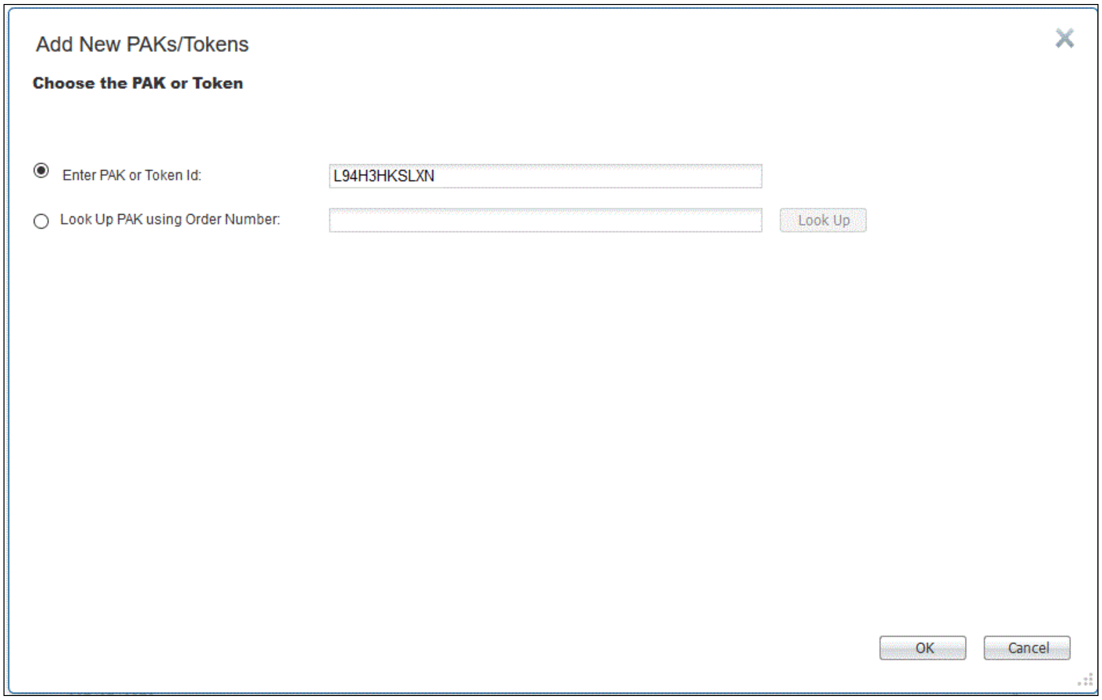

Add New PAKs/Tokens

Choose the PAK or Token

● Enter PAK or Token Id: L94H3HKSLXN

○ Look Up PAK using Order Number:

Figure 6-38 Entering the PAK number

The system requests the PIN that is provided on the PAK letter to validate the license ownership for this PAK, as shown in Figure 6-39. Enter the information that is provided on the letter, and click **OK**.

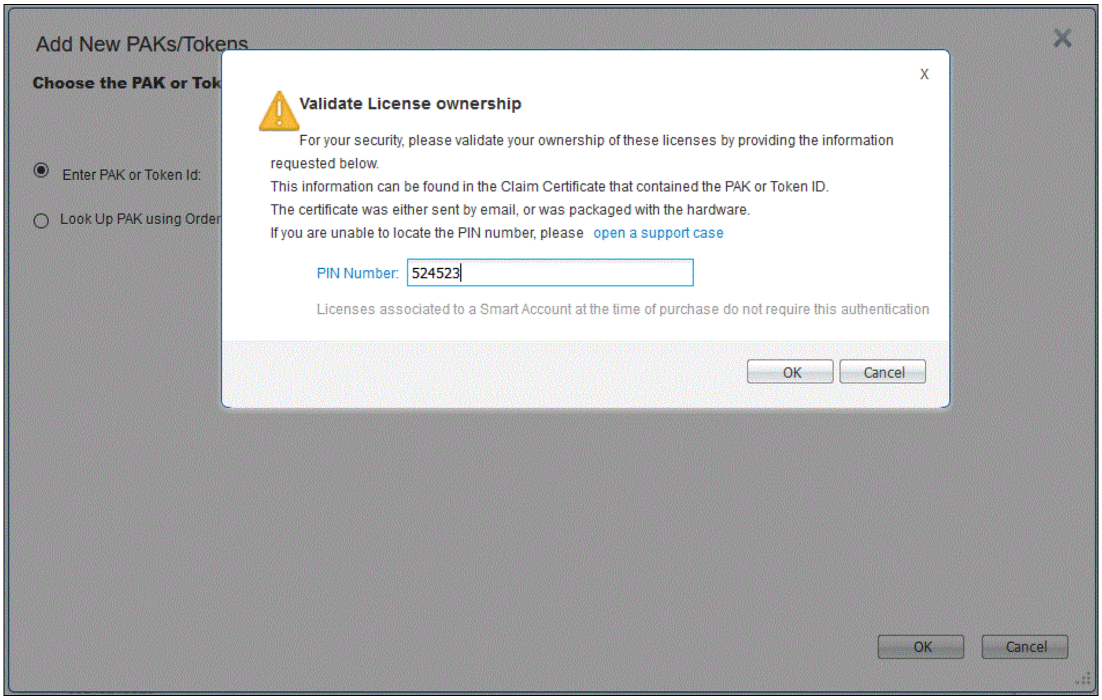

Validate License ownership

For your security, please validate your ownership of these licenses by providing the information requested below.
This information can be found in the Claim Certificate that contained the PAK or Token ID.
The certificate was either sent by email, or was packaged with the hardware.
If you are unable to locate the PIN number, please open a support case

PIN Number: 524523

Licenses associated to a Smart Account at the time of purchase do not require this authentication

Figure 6-39 Enter PIN

Now, you see the newly added PAK in the window, as shown in Figure 6-40. Select this PAK by checking the box in the first column and click **Get Licenses**.

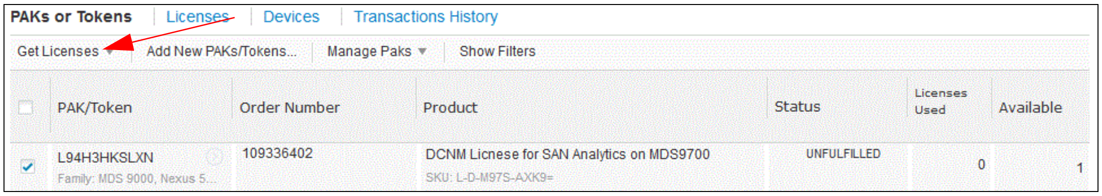

☐	PAK/Token	Order Number	Product	Status	Licenses Used	Available
☑	L94H3HKSLXN Family: MDS 9000, Nexus 5...	109336402	DCNM Licnese for SAN Analytics on MDS9700 SKU: L-D-M97S-AXK9=	UNFULFILLED	0	1

Figure 6-40 Added PAK

On the Get Licenses selection dialog box, we click **From Selected PAKs...**, as shown in Figure 6-41.

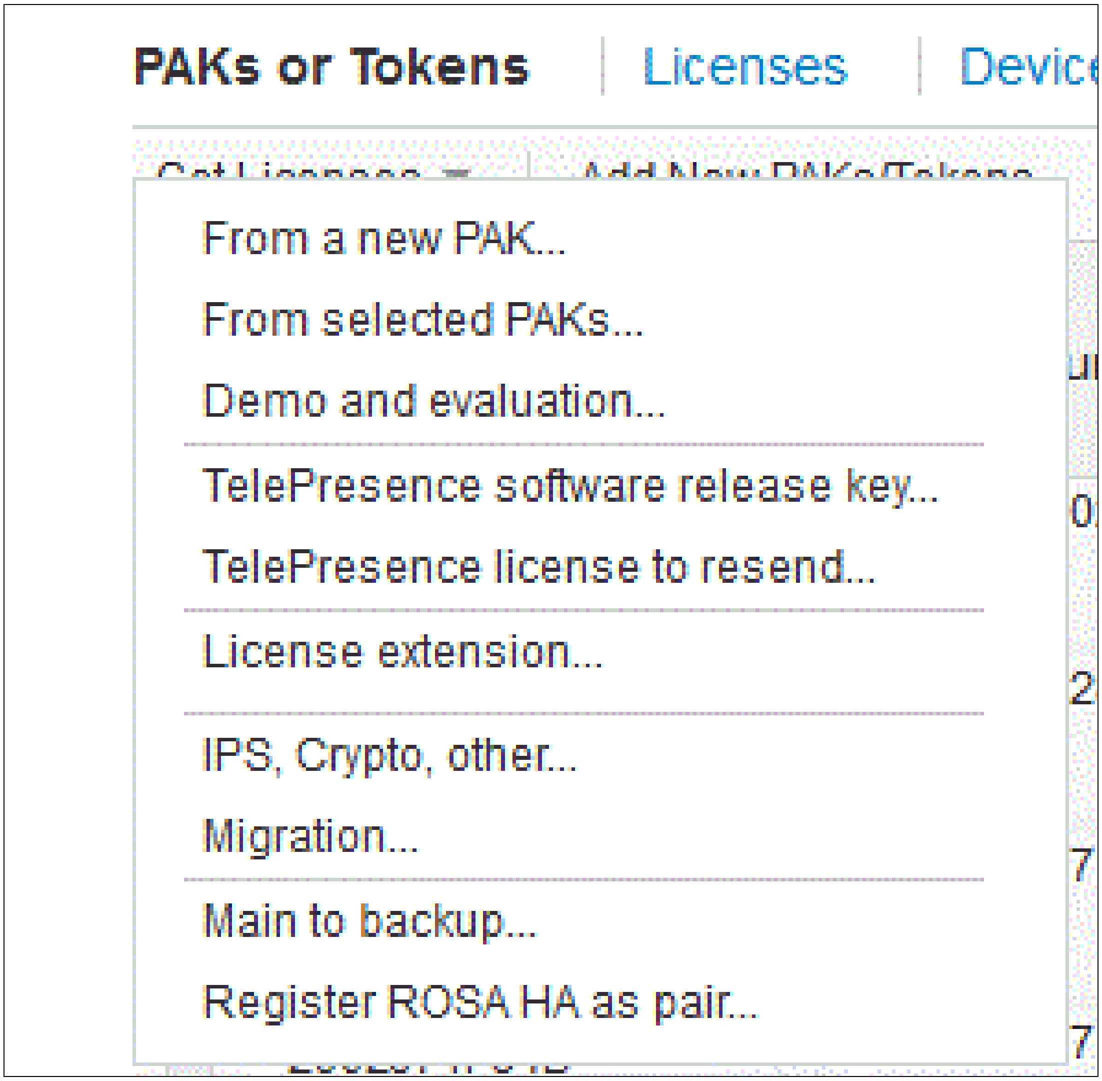

Figure 6-41 Get License selection

Now, you see the Get New Licenses from Single PAK/Token window, as shown in Figure 6-42. Click **Next**.

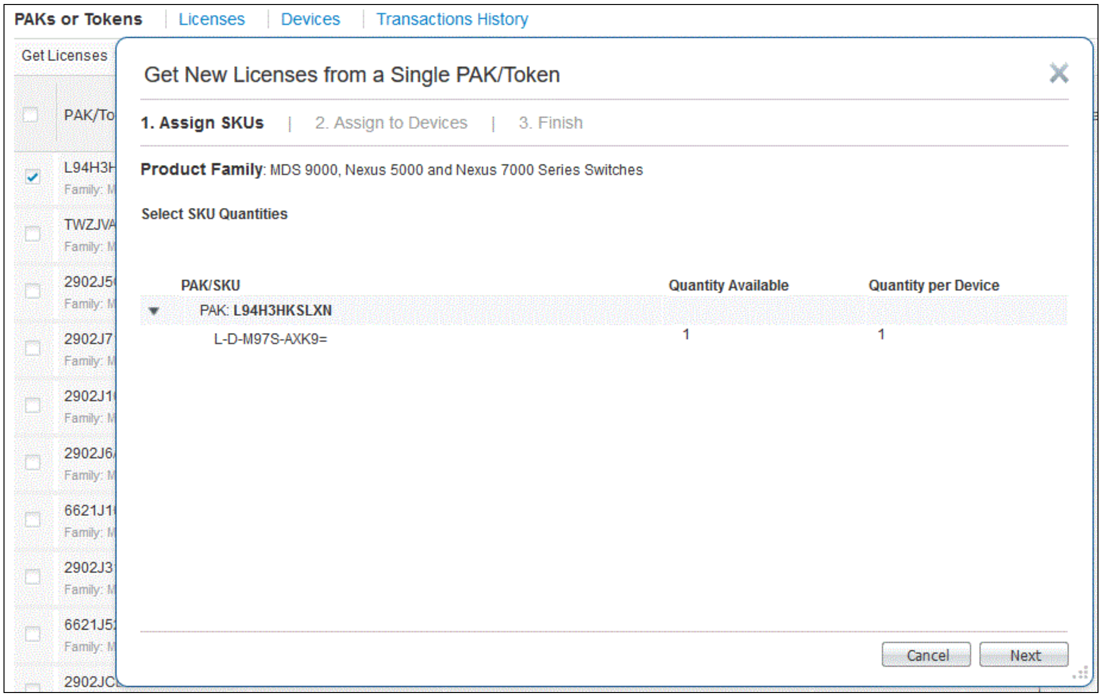

Get New Licenses from a Single PAK/Token

1. Assign SKUs | 2. Assign to Devices | 3. Finish

Product Family: MDS 9000, Nexus 5000 and Nexus 7000 Series Switches

Select SKU Quantities

PAK/SKU	Quantity Available	Quantity per Device
PAK: **L94H3HKSLXN**		
L-D-M97S-AXK9=	1	1

Figure 6-42 Getting new licenses

You must enter the serial number of the IBM c-type switch in to the window that is shown in Figure 6-43. The serial number can be found in one of two ways:

- When looking at the system, the serial number is printed on a small white label with “IBM S/N: 0000xxxxxxx” on either the top or front of the chassis or casing.
- From the CLI, running the command **`show license host-id`** also provides the serial number. Click **Next**.

Important: Enter the full serial number, including any leading zeros. It must be 11 digits. Failure to do this task prevents the license from being installed.

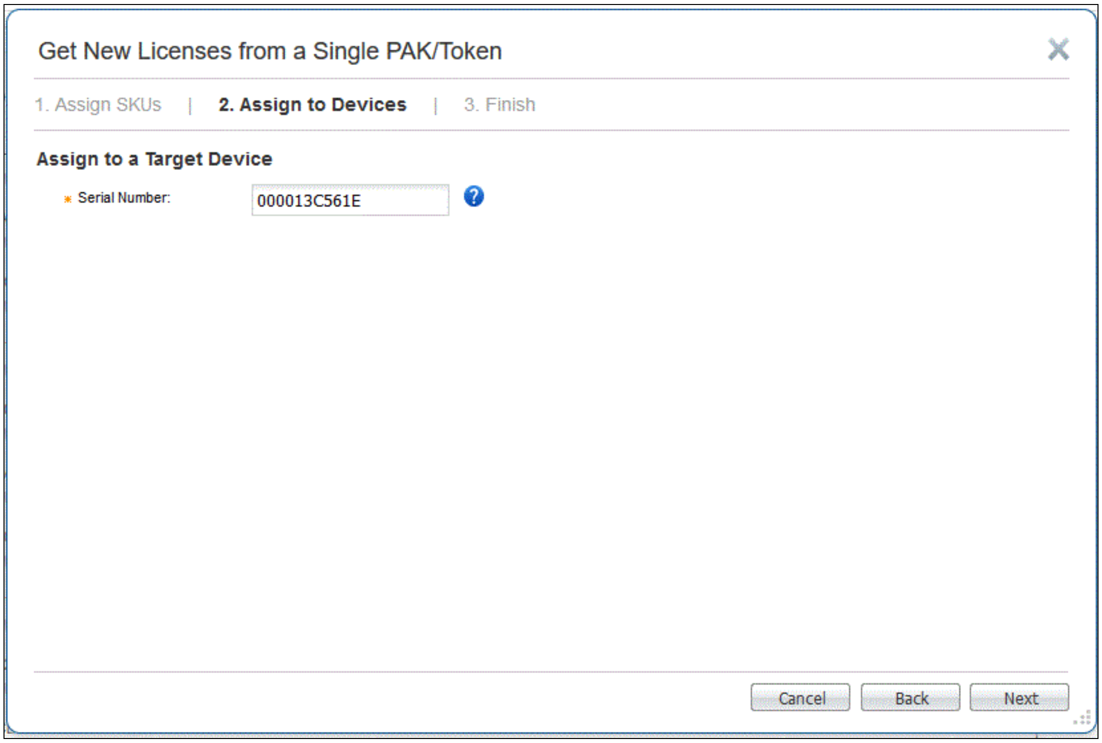

Get New Licenses from a Single PAK/Token

1. Assign SKUs | **2. Assign to Devices** | 3. Finish

Assign to a Target Device

* Serial Number: 000013C561E

Figure 6-43 Entering the serial number

In Figure 6-44, you must enter the user information and email address for the recipient who should get the license file for the switch. Click **Submit**.

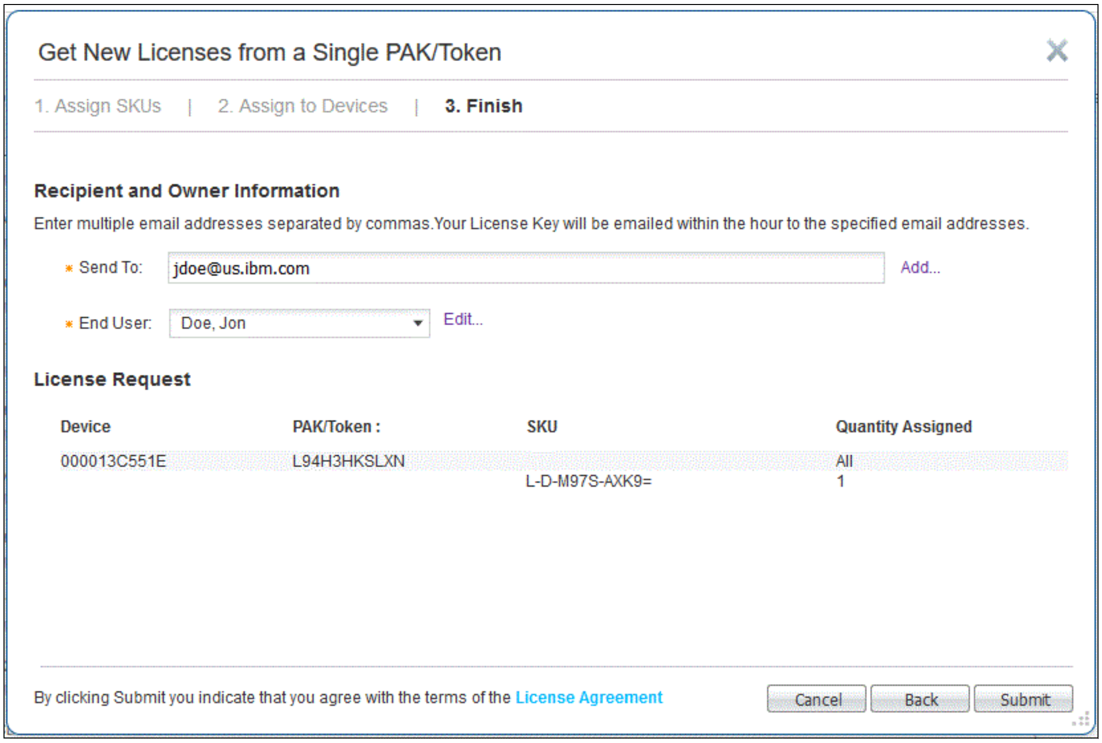

Get New Licenses from a Single PAK/Token

Recipient and Owner Information

Enter multiple email addresses separated by commas.Your License Key will be emailed within the hour to the specified email addresses.

* Send To: jdoe@us.ibm.com
* End User: Doe, Jon

License Request

Device	PAK/Token :	SKU	Quantity Assigned
000013C551E	L94H3HKSLXN		All
		L-D-M97S-AXK9=	1

By clicking Submit you indicate that you agree with the terms of the License Agreement

Figure 6-44 Entering the license recipient

In Figure 6-45, you see that an email was sent to the provided email address. Now, you may directly download the license file from the web page. Click **Download** and provide the location for the license file to be saved locally.

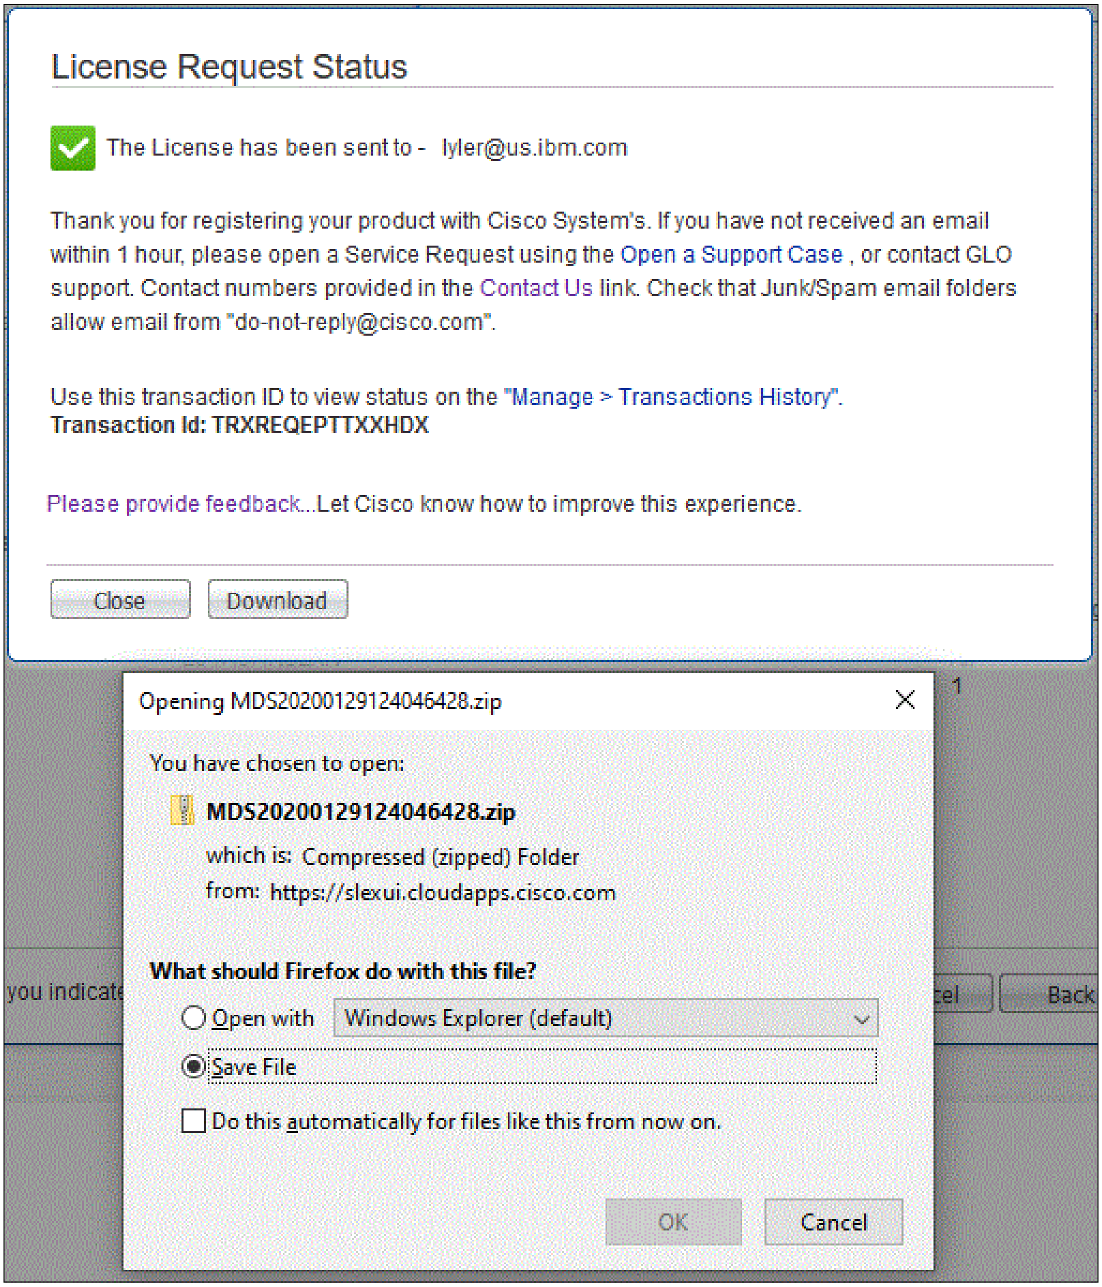

License Request Status

✓ The License has been sent to - lyler@us.ibm.com

Thank you for registering your product with Cisco System's. If you have not received an email within 1 hour, please open a Service Request using the Open a Support Case , or contact GLO support. Contact numbers provided in the Contact Us link. Check that Junk/Spam email folders allow email from "do-not-reply@cisco.com".

Use this transaction ID to view status on the "Manage > Transactions History".
Transaction Id: TRXREQEPTTXXHDX

Please provide feedback...Let Cisco know how to improve this experience.

Close Download

Opening MDS20200129124046428.zip

You have chosen to open:

MDS20200129124046428.zip

which is: Compressed (zipped) Folder

from: https://slexui.cloudapps.cisco.com

What should Firefox do with this file?

○ Open with Windows Explorer (default)

● Save File

□ Do this automatically for files like this from now on.

OK Cancel

Figure 6-45 Downloading the license file

As you see in Figure 6-46 on page 235, the PAK now shows as fulfilled because the license file was delivered. The license file is packaged in compressed format, so before the file is moved to the switch, it must be decompressed.

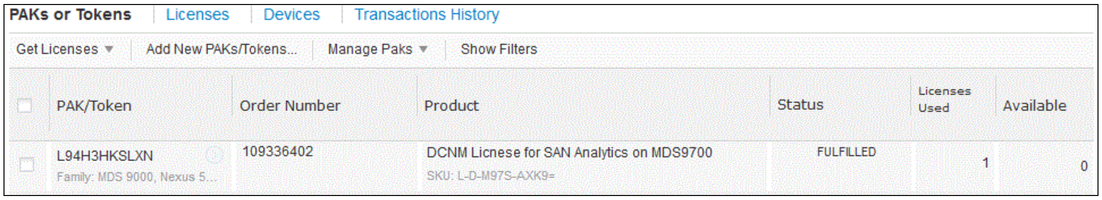

□	PAK/Token	Order Number	Product	Status	Licenses Used	Available
□	L94H3HKSLXN Family: MDS 9000, Nexus 5...	109336402	DCNM Licnese for SAN Analytics on MDS9700 SKU: L-D-M97S-AXK9=	FULFILLED	1	0

Figure 6-46 PAK fulfilled

6.9.2 Transferring license files to the switch

Now that you have the license file for the IBM c-type switch, you must install it. The license file has a name like `MDS20210127200818220.lic` and is keyed to be installed on the specific serial number for which it was generated. In the following sections regarding the installation of licenses, the license file names might change. It is important to recognize the license file itself rather than its specific name because licenses were generated on different systems in the lab.

You must move the license file to the bootflash on the IBM c-type switch. There are several applications that can be used to transfer files, but in this example we use WinSCP. Whichever file transfer tool that you use, you should be familiar with it and know how to transfer files. Launch WinSCP, as shown in Figure 6-47, and log in to your switch. There are several protocols that you can use, and the one that you choose might depend on the policies in place. In many cases, unsecure protocols such as TFTP are banned. In this example, we use Secure File Transfer Protocol (SFTP).

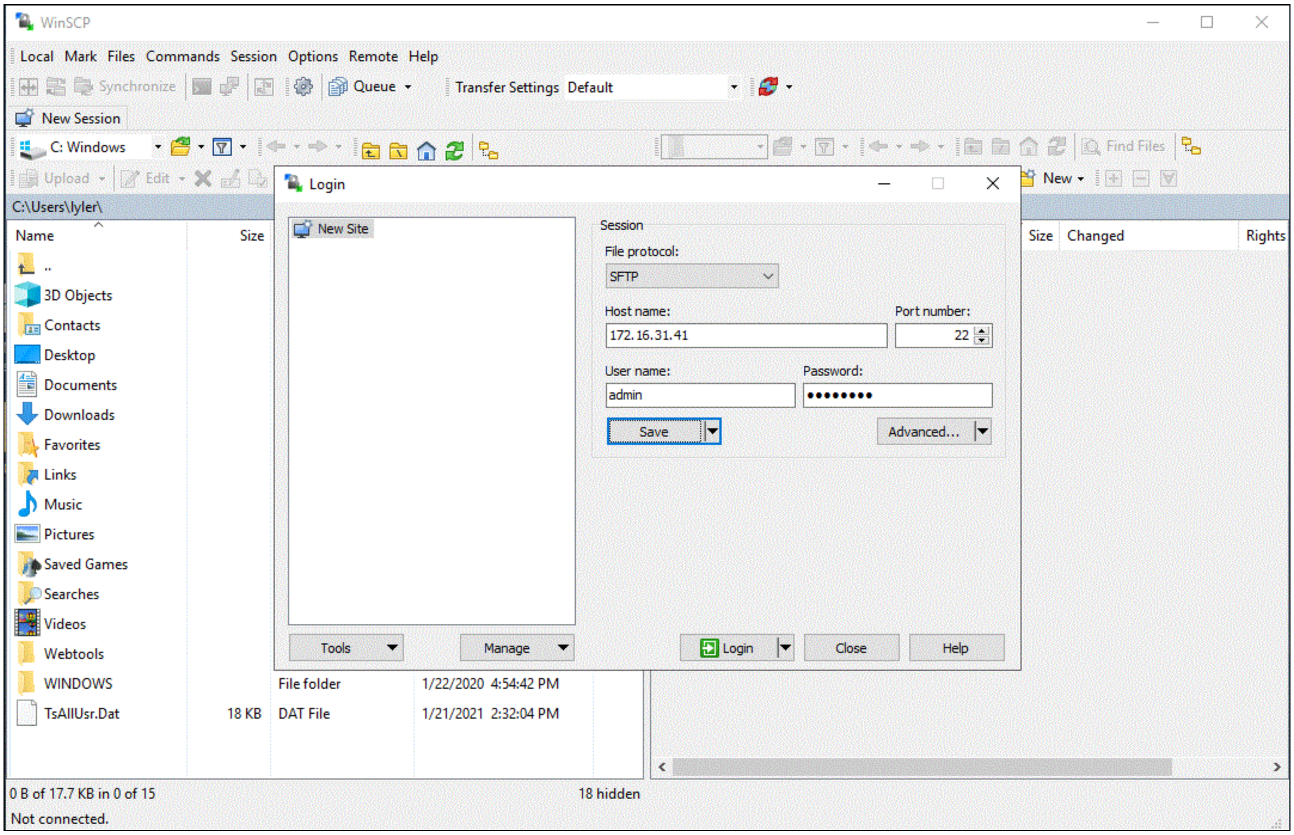

Session

File protocol: SFTP

Host name: 172.16.31.41

Port number: 22

User name: admin

Password: ••••••••

..
3D Objects
Contacts
Desktop
Documents
Downloads
Favorites
Links
Music
Pictures
Saved Games
Searches
Videos
Webtools
WINDOWS — File folder — 1/22/2020 4:54:42 PM
TsAllUsr.Dat — 18 KB — DAT File — 1/21/2021 2:32:04 PM

0 B of 17.7 KB in 0 of 15 — 18 hidden

Not connected.

Figure 6-47 Logging in to WinSCP

After you enter the correct details, click **Login**. The window that is shown in Figure 6-48 opens.

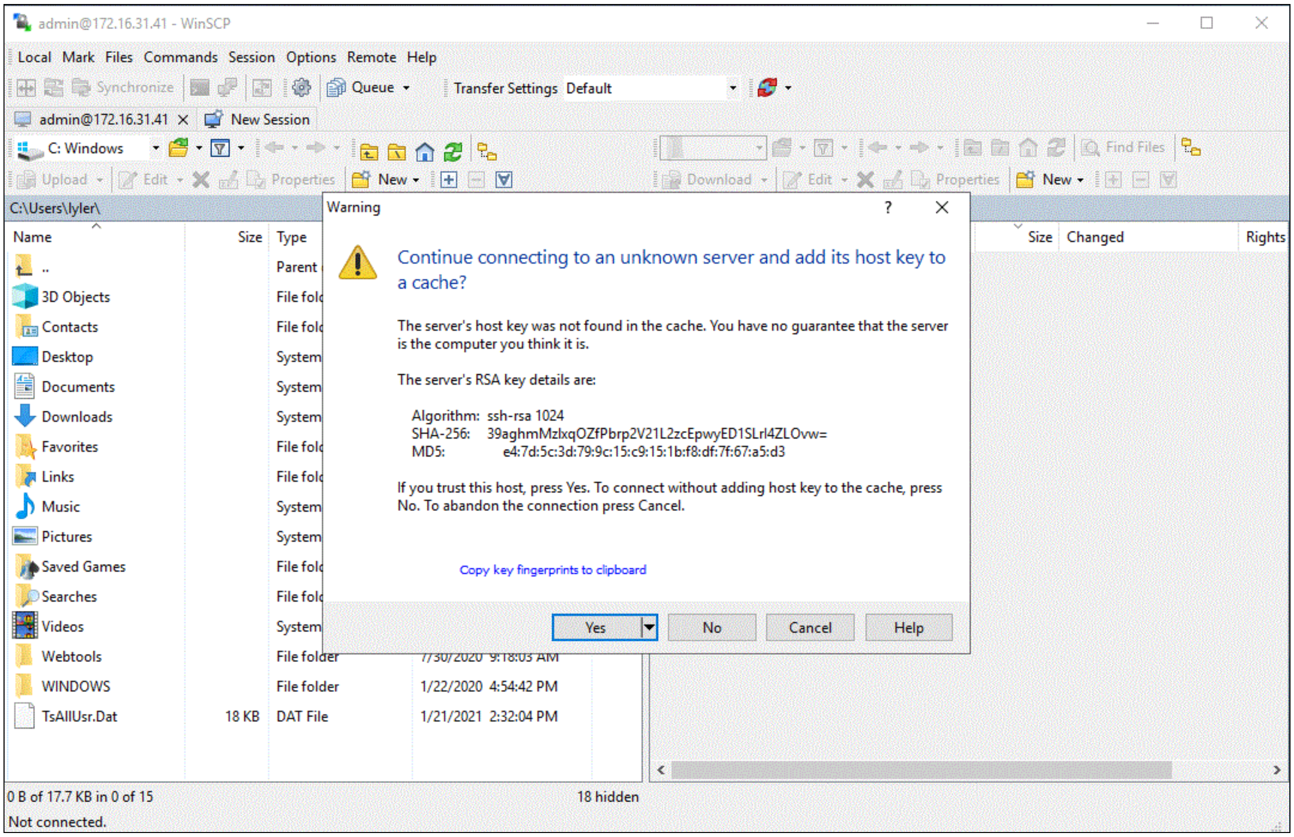

Warning

Continue connecting to an unknown server and add its host key to a cache?

The server's host key was not found in the cache. You have no guarantee that the server is the computer you think it is.

The server's RSA key details are:

```
Algorithm:  ssh-rsa 1024
SHA-256:    39aghmMzlxqOZfPbrp2V21L2zcEpwyED1SLrl4ZLOvw=
MD5:          e4:7d:5c:3d:79:9c:15:c9:15:1b:f8:df:7f:67:a5:d3
```

If you trust this host, press Yes. To connect without adding host key to the cache, press No. To abandon the connection press Cancel.

Copy key fingerprints to clipboard

Name	Size	Type	
..		Parent	
3D Objects		File fold	
Contacts		File fold	
Desktop		System	
Documents		System	
Downloads		System	
Favorites		File fold	
Links		File fold	
Music		System	
Pictures		System	
Saved Games		File fold	
Searches		File fold	
Videos		System	
Webtools		File folder	7/30/2020 9:18:03 AM
WINDOWS		File folder	1/22/2020 4:54:42 PM
TsAllUsr.Dat	18 KB	DAT File	1/21/2021 2:32:04 PM

0 B of 17.7 KB in 0 of 15

18 hidden

Not connected.

Figure 6-48 Continue to log in

This window or one like it enables you to double-check that you are connecting to the correct device. Check that the device is correct, and then click **Yes**. The window that is shown in Figure 6-49 opens.

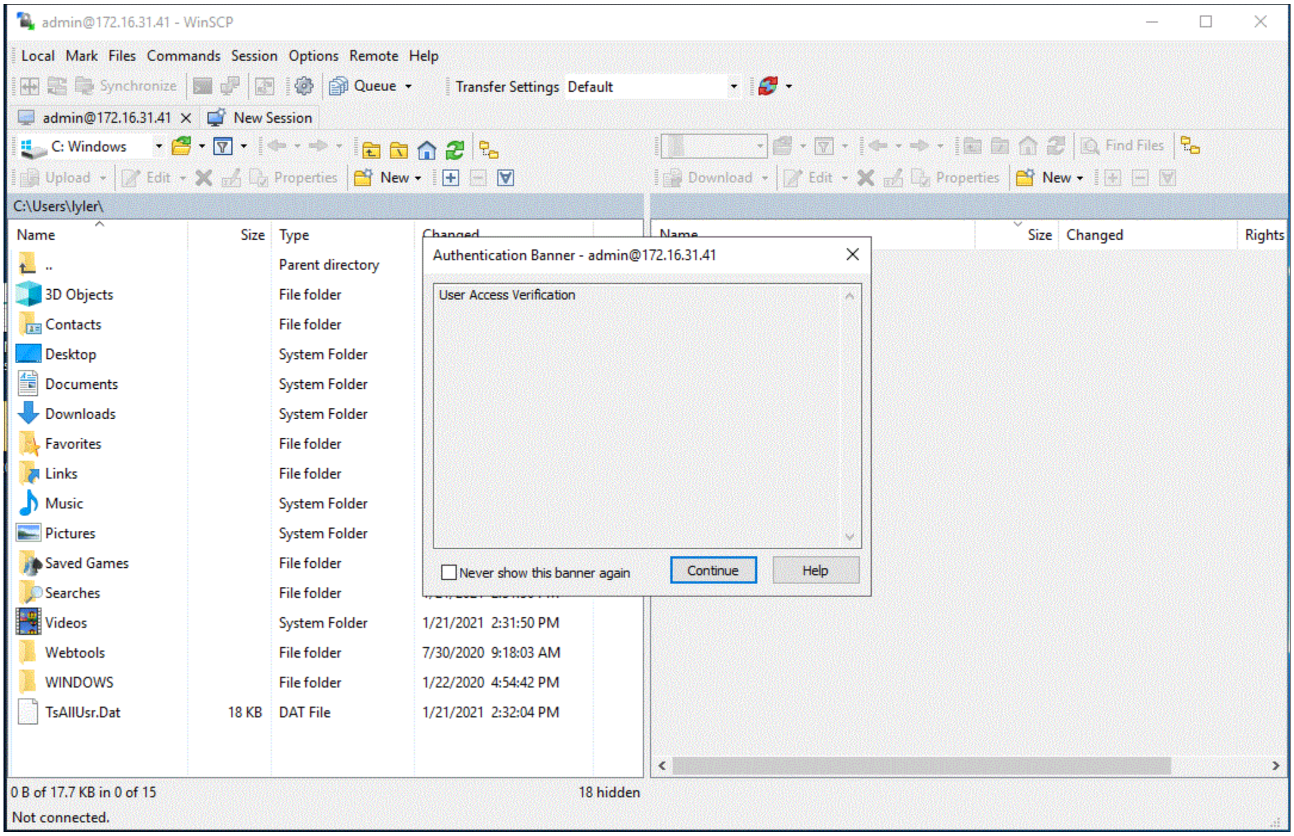
admin@172.16.31.41 - WinSCP
Local Mark Files Commands Session Options Remote Help
Synchronize
Queue
Transfer Settings Default
admin@172.16.31.41
New Session
C: Windows
Find Files
Upload
Edit
Properties
New
Download
Edit
Properties
New
Authentication Banner - admin@172.16.31.41
Never show this banner again
Continue
Help

C:\Users\lyler\

Name	Size	Type	Changed
..		Parent directory	
3D Objects		File folder	
Contacts		File folder	
Desktop		System Folder	
Documents		System Folder	
Downloads		System Folder	
Favorites		File folder	
Links		File folder	
Music		System Folder	
Pictures		System Folder	
Saved Games		File folder	
Searches		File folder	
Videos		System Folder	1/21/2021 2:31:50 PM
Webtools		File folder	7/30/2020 9:18:03 AM
WINDOWS		File folder	1/22/2020 4:54:42 PM
TsAllUsr.Dat	18 KB	DAT File	1/21/2021 2:32:04 PM

Name	Size	Changed	Rights

User Access Verification

0 B of 17.7 KB in 0 of 15 18 hidden

Not connected.

Figure 6-49 Login verification

A login verification opens (unless the option to not show it was selected). Click **Continue**. The window that is shown in Figure 6-50 opens.

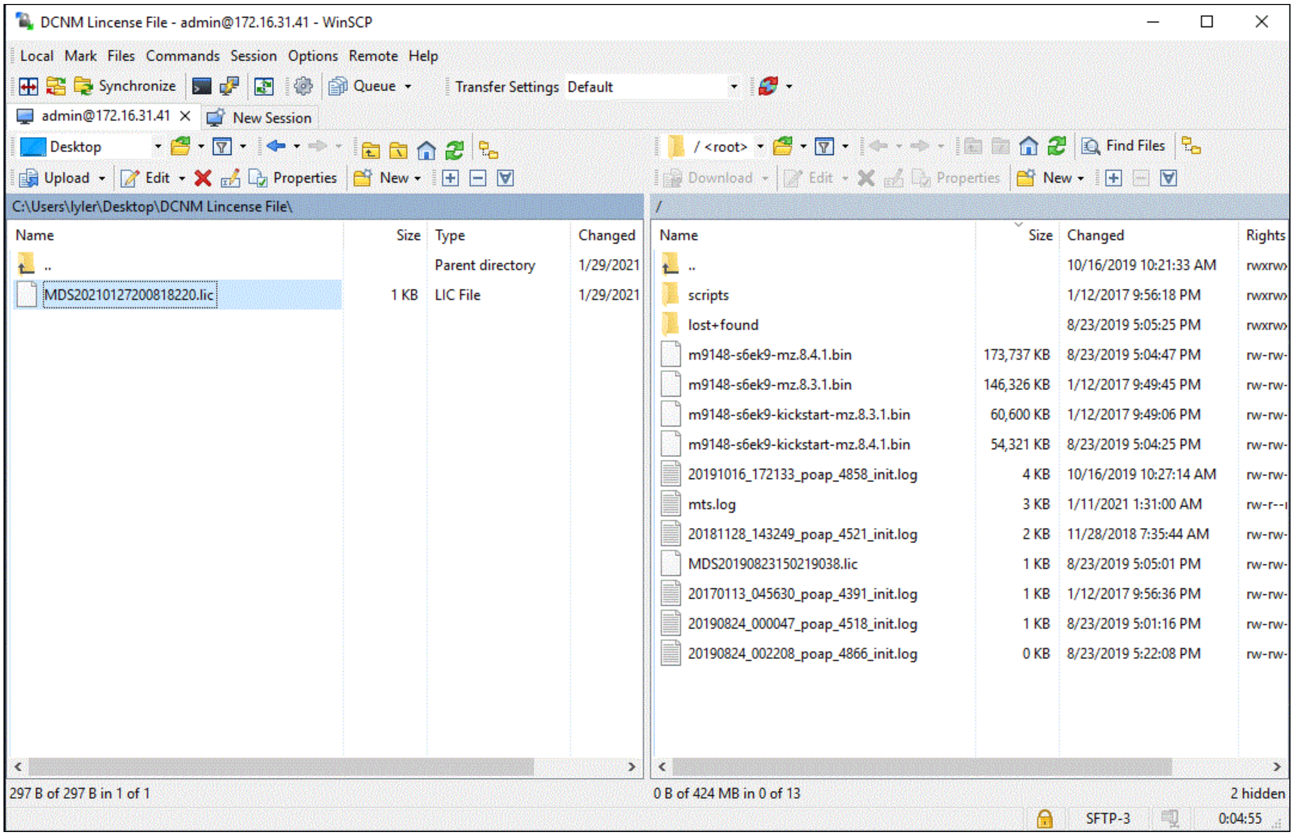

DCNM Lincense File - admin@172.16.31.41 - WinSCP
Local Mark Files Commands Session Options Remote Help
Synchronize
Queue
Transfer Settings Default
admin@172.16.31.41
New Session
Desktop
Upload
Edit
Properties
New
/ <root>
Find Files
Download

C:\Users\lyler\Desktop\DCNM Lincense File\

Name	Size	Type	Changed
..		Parent directory	1/29/2021
MDS20210127200818220.lic	1 KB	LIC File	1/29/2021

/

Name	Size	Changed	Rights
..		10/16/2019 10:21:33 AM	rwxrw
scripts		1/12/2017 9:56:18 PM	rwxrw
lost+found		8/23/2019 5:05:25 PM	rwxrw
m9148-s6ek9-mz.8.4.1.bin	173,737 KB	8/23/2019 5:04:47 PM	rw-rw-
m9148-s6ek9-mz.8.3.1.bin	146,326 KB	1/12/2017 9:49:45 PM	rw-rw-
m9148-s6ek9-kickstart-mz.8.3.1.bin	60,600 KB	1/12/2017 9:49:06 PM	rw-rw-
m9148-s6ek9-kickstart-mz.8.4.1.bin	54,321 KB	8/23/2019 5:04:25 PM	rw-rw-
20191016_172133_poap_4858_init.log	4 KB	10/16/2019 10:27:14 AM	rw-rw-
mts.log	3 KB	1/11/2021 1:31:00 AM	rw-r--
20181128_143249_poap_4521_init.log	2 KB	11/28/2018 7:35:44 AM	rw-rw-
MDS20190823150219038.lic	1 KB	8/23/2019 5:05:01 PM	rw-rw-
20170113_045630_poap_4391_init.log	1 KB	1/12/2017 9:56:36 PM	rw-rw-
20190824_000047_poap_4518_init.log	1 KB	8/23/2019 5:01:16 PM	rw-rw-
20190824_002208_poap_4866_init.log	0 KB	8/23/2019 5:22:08 PM	rw-rw-

297 B of 297 B in 1 of 1

0 B of 424 MB in 0 of 13

2 hidden

SFTP-3

0:04:55

Figure 6-50 WinSCP navigation window

The WinSCP navigation window opens. Go to the correct locations on both the local and remote devices. Highlight the file to be transferred and drag it from the local to the remote location. In Figure 6-50, this action will be from the pane on the left to the pane on the right. As a best practice, copy the license file to the bootflash, and after it is copied, the licenses automatically are duplicated to both supervisors if this device is a director (a switch has only one supervisor).

The window that is show in Figure 6-51 opens.

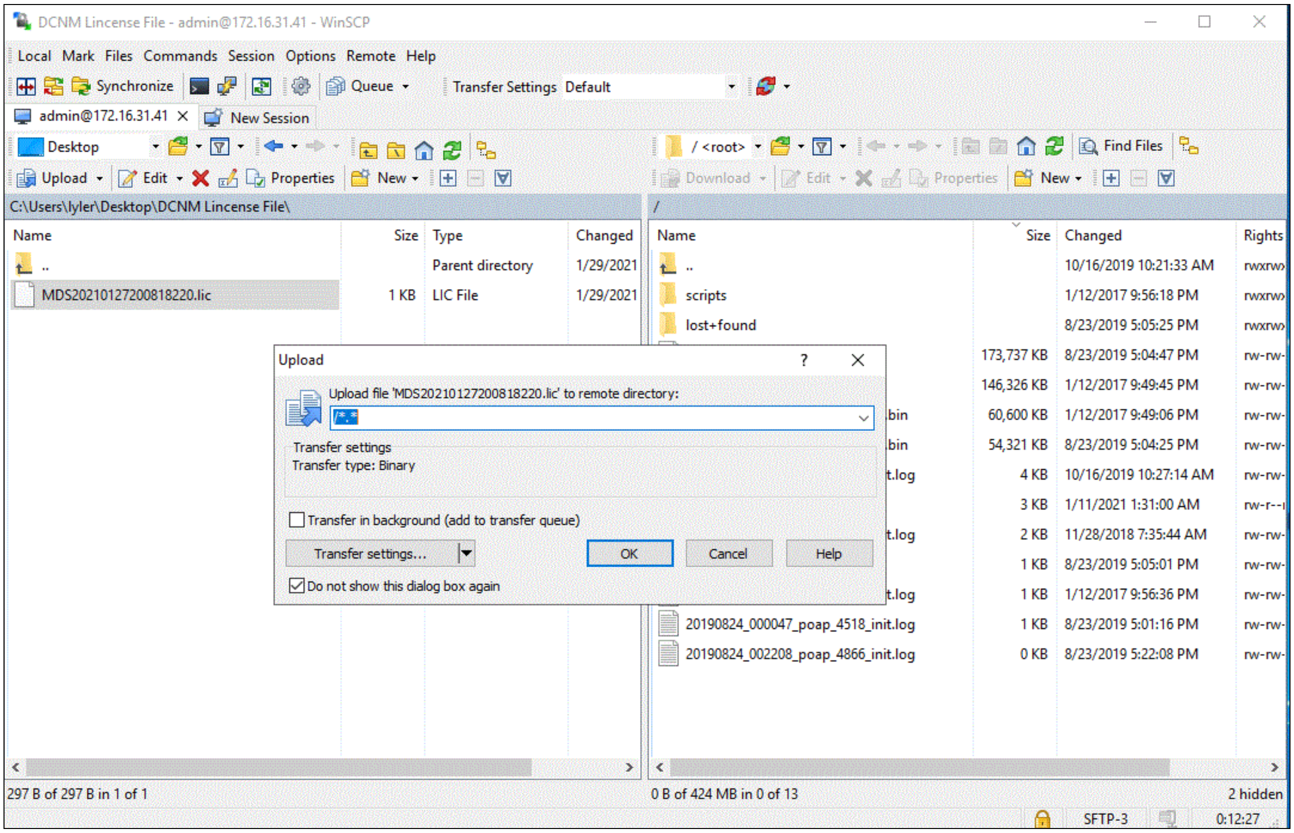
DCNM Lincense File - admin@172.16.31.41 - WinSCP
Local Mark Files Commands Session Options Remote Help
Synchronize
Queue
Transfer Settings Default
admin@172.16.31.41
New Session
Desktop
Upload
Edit
Properties
New
/ <root>
Find Files
Download
Edit
Properties
New

C:\Users\lyler\Desktop\DCNM Lincense File\

Name	Size	Type	Changed
..		Parent directory	1/29/2021
MDS20210127200818220.lic	1 KB	LIC File	1/29/2021

/

Name	Size	Changed	Rights
..		10/16/2019 10:21:33 AM	rwxrw
scripts		1/12/2017 9:56:18 PM	rwxrw
lost+found		8/23/2019 5:05:25 PM	rwxrw
	173,737 KB	8/23/2019 5:04:47 PM	rw-rw-
	146,326 KB	1/12/2017 9:49:45 PM	rw-rw-
bin	60,600 KB	1/12/2017 9:49:06 PM	rw-rw-
bin	54,321 KB	8/23/2019 5:04:25 PM	rw-rw-
t.log	4 KB	10/16/2019 10:27:14 AM	rw-rw-
	3 KB	1/11/2021 1:31:00 AM	rw-r--I
t.log	2 KB	11/28/2018 7:35:44 AM	rw-rw-
	1 KB	8/23/2019 5:05:01 PM	rw-rw-
t.log	1 KB	1/12/2017 9:56:36 PM	rw-rw-
20190824_000047_poap_4518_init.log	1 KB	8/23/2019 5:01:16 PM	rw-rw-
20190824_002208_poap_4866_init.log	0 KB	8/23/2019 5:22:08 PM	rw-rw-

Upload

Upload file 'MDS20210127200818220.lic' to remote directory:

/*.*

Transfer settings
Transfer type: Binary

☐ Transfer in background (add to transfer queue)

☑ Do not show this dialog box again

297 B of 297 B in 1 of 1

0 B of 424 MB in 0 of 13

2 hidden

SFTP-3

0:12:27

Figure 6-51 WinSCP file upload

Check that the file is correct, and then click **Yes**. The window that is shown in Figure 6-52 opens.

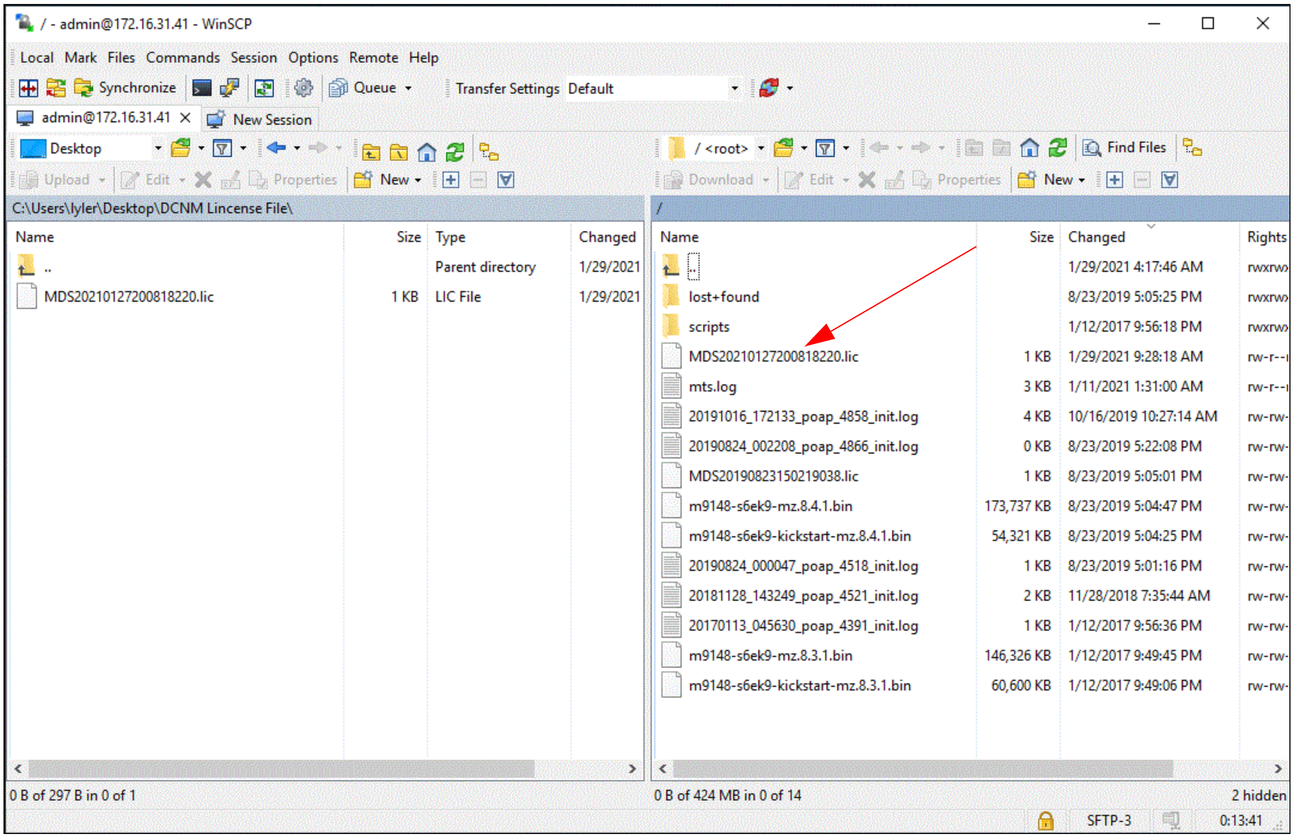

C:\Users\lyler\Desktop\DCNM Lincense File\

Name	Size	Type	Changed
..		Parent directory	1/29/2021
MDS20210127200818220.lic	1 KB	LIC File	1/29/2021

/

Name	Size	Changed	Rights
..		1/29/2021 4:17:46 AM	rwxrwx
lost+found		8/23/2019 5:05:25 PM	rwxrwx
scripts		1/12/2017 9:56:18 PM	rwxrwx
MDS20210127200818220.lic	1 KB	1/29/2021 9:28:18 AM	rw-r--r
mts.log	3 KB	1/11/2021 1:31:00 AM	rw-r--r
20191016_172133_poap_4858_init.log	4 KB	10/16/2019 10:27:14 AM	rw-rw-
20190824_002208_poap_4866_init.log	0 KB	8/23/2019 5:22:08 PM	rw-rw-
MDS20190823150219038.lic	1 KB	8/23/2019 5:05:01 PM	rw-rw-
m9148-s6ek9-mz.8.4.1.bin	173,737 KB	8/23/2019 5:04:47 PM	rw-rw-
m9148-s6ek9-kickstart-mz.8.4.1.bin	54,321 KB	8/23/2019 5:04:25 PM	rw-rw-
20190824_000047_poap_4518_init.log	1 KB	8/23/2019 5:01:16 PM	rw-rw-
20181128_143249_poap_4521_init.log	2 KB	11/28/2018 7:35:44 AM	rw-rw-
20170113_045630_poap_4391_init.log	1 KB	1/12/2017 9:56:36 PM	rw-rw-
m9148-s6ek9-mz.8.3.1.bin	146,326 KB	1/12/2017 9:49:45 PM	rw-rw-
m9148-s6ek9-kickstart-mz.8.3.1.bin	60,600 KB	1/12/2017 9:49:06 PM	rw-rw-

0 B of 297 B in 0 of 1

0 B of 424 MB in 0 of 14 2 hidden

SFTP-3 0:13:41

Figure 6-52 WinSCP transferred file

The file is now transferred.

As a best practice, back up the license key file to a remote server in case it will be needed again.

6.9.3 Installing license files from the DCNM

Before a license can be installed, ensure that it has been transferred to the switch, as described in 6.9.2, “Transferring license files to the switch” on page 235.

To install the licenses from DM, start from the DCNM window. Add the first switch by selecting **Inventory** → **Discovery** → **SAN Switches**, as shown in Figure 6-53.

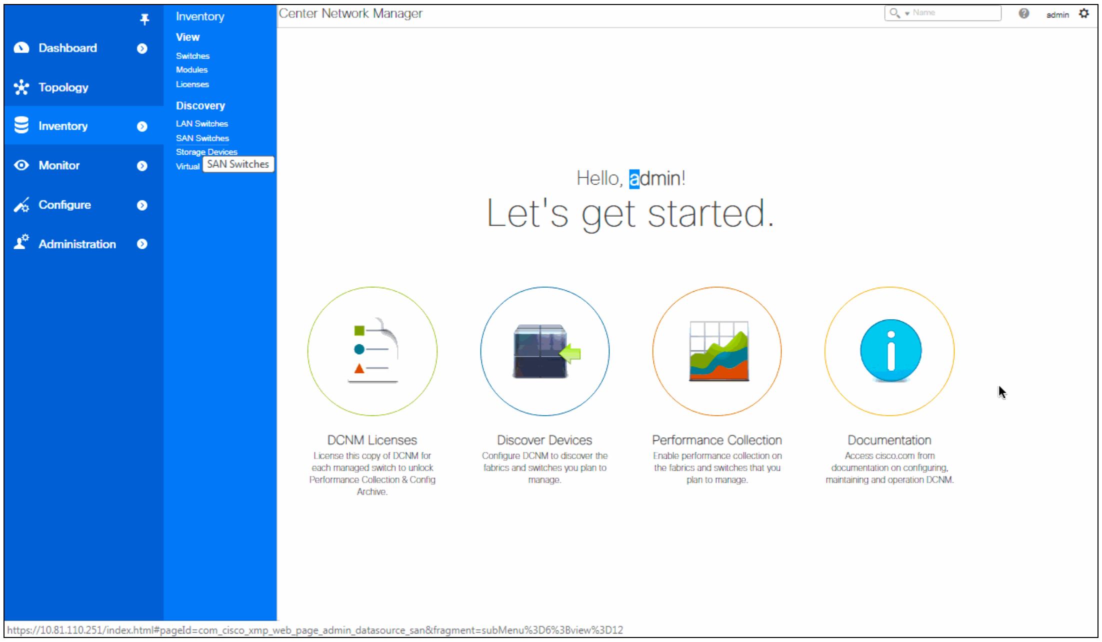

Hello, admin!

Let's get started.

License this copy of DCNM for each managed switch to unlock Performance Collection & Config Archive.

Configure DCNM to discover the fabrics and switches you plan to manage.

Enable performance collection on the fabrics and switches that you plan to manage.

Access cisco.com from documentation on configuring, maintaining and operation DCNM.

https://10.81.110.251/index.html#pageId=com_cisco_xmp_web_page_admin_datasource_san&fragment=subMenu%3D6%3Bview%3D12

Figure 6-53 SAN Switch Inventory

The Inventory / Discovery / SAN Switches window opens, as shown in Figure 6-54. Click **+** to add the first switch, which is also known as the *seed switch*.

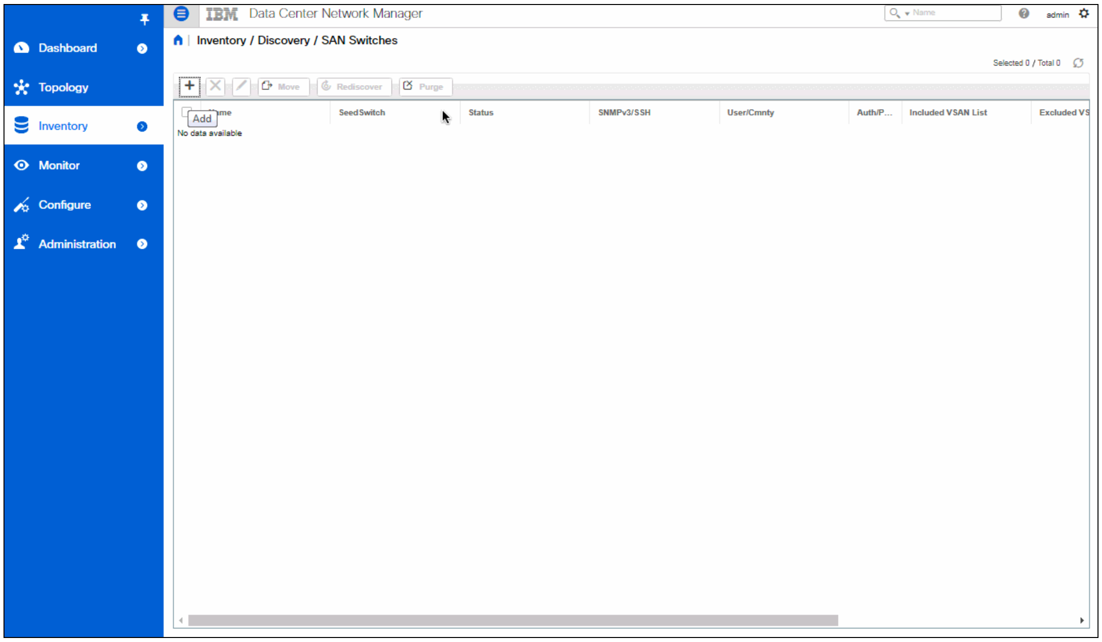

Figure 6-54 Adding a SAN switch

The Add Fabric dialog box opens. Enter the information for the first of the lab switches into the dialog box, as shown in Figure 6-55. Click **Add**.

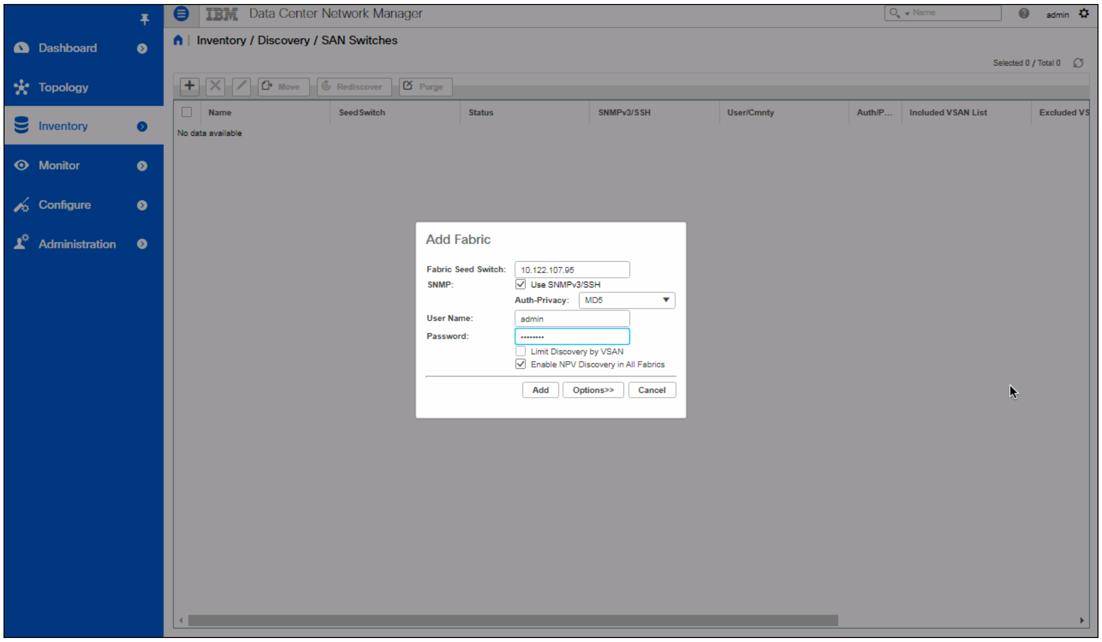

☐	Name	SeedSwitch	Status	SNMPv3/SSH	User/Cmnty	Auth/P...	Included VSAN List	Excluded VS

No data available

Add Fabric

Fabric Seed Switch: 10.122.107.95
SNMP: ☑ Use SNMPv3/SSH
Auth-Privacy: MD5
User Name: admin
Password: ••••••••
☐ Limit Discovery by VSAN
☑ Enable NPV Discovery in All Fabrics

Figure 6-55 Add Fabric

After a short delay for discovery, the new switch appears in the inventory, as shown in Figure 6-56. Click **Topology**.

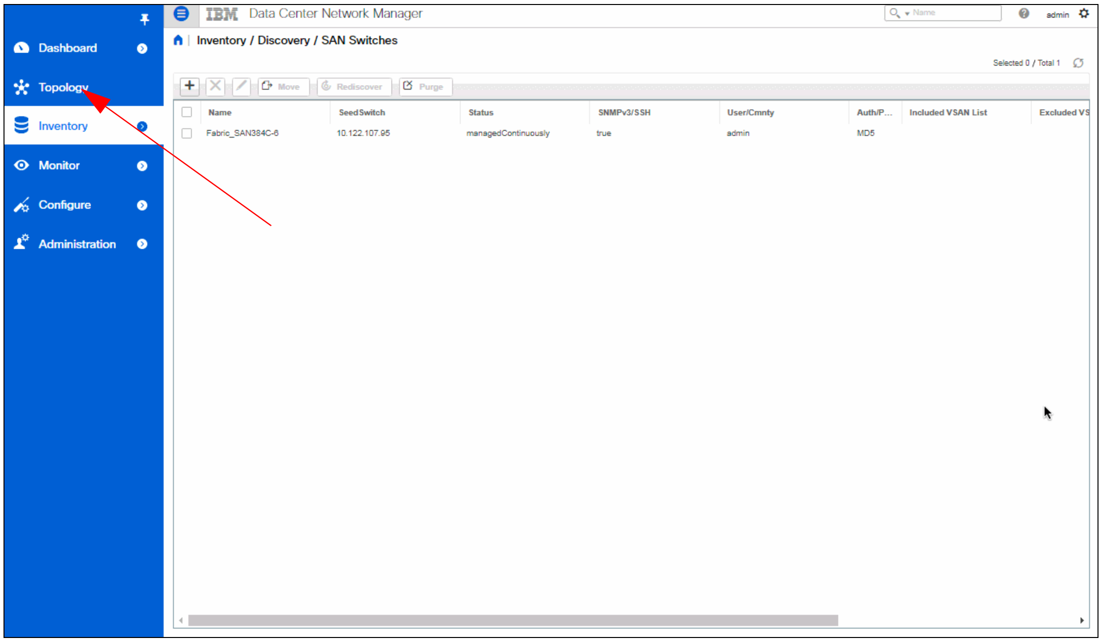

☐	Name	SeedSwitch	Status	SNMPv3/SSH	User/Cmnty	Auth/P...	Included VSAN List	Excluded VS
☐	Fabric_SAN384C-6	10.122.107.95	managedContinuously	true	admin	MD5		

Figure 6-56 SAN switch inventory

On the SAN topology window that is shown in Figure 6-57 on page 243, you see the first switch for the lab environment. Double-click the icon for the switch.

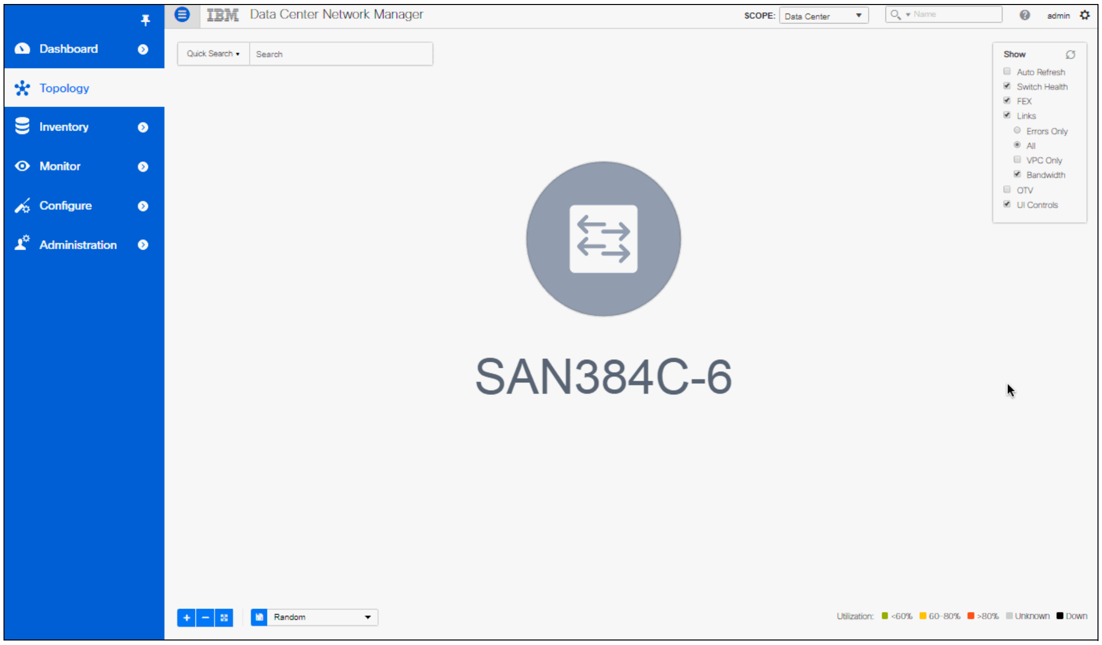

Show

- ☐ Auto Refresh
- ☑ Switch Health
- ☑ FEX
- ☑ Links
  - ○ Errors Only
  - ● All
  - ☐ VPC Only
  - ☑ Bandwidth
- ☐ OTV
- ☑ UI Controls

Figure 6-57 SAN Topology

Now, you see the summary information for this switch in Figure 6-58. Click **Show More Details**.

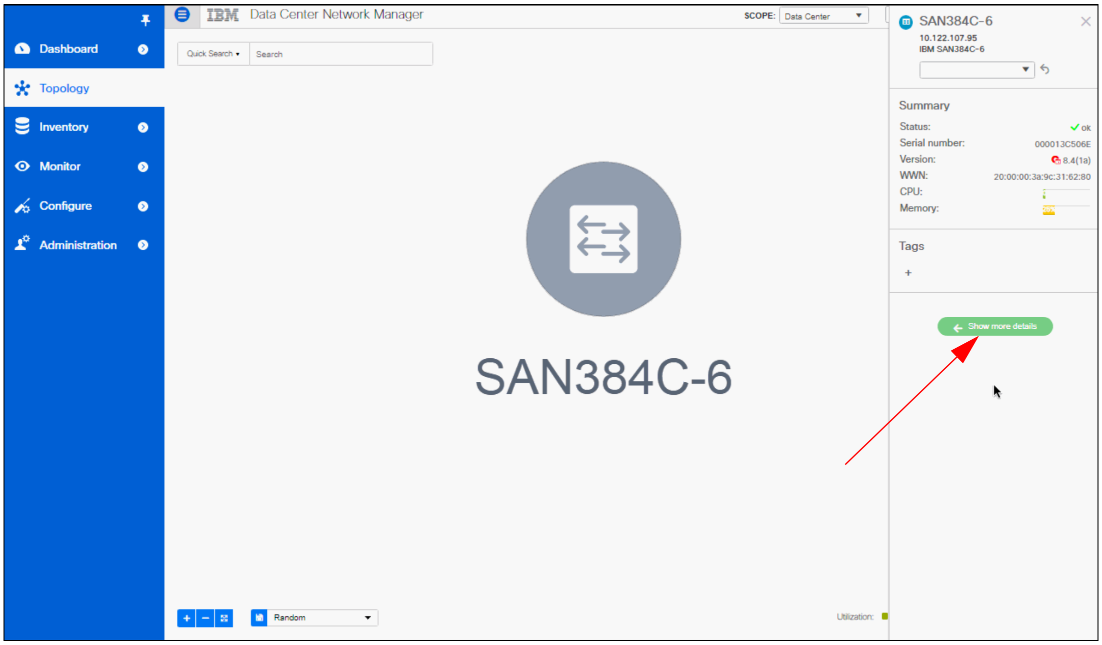

SAN384C-6
10.122.107.95
IBM SAN384C-6

Summary

Status: ✓ ok
Serial number: 000013C506E
Version: 8.4(1a)
WWN: 20:00:00:3a:9c:31:62:80
CPU:
Memory: 25%

Figure 6-58 Switch summary

In Figure 6-59, you see further information about the switch. Click **Device Manager**.

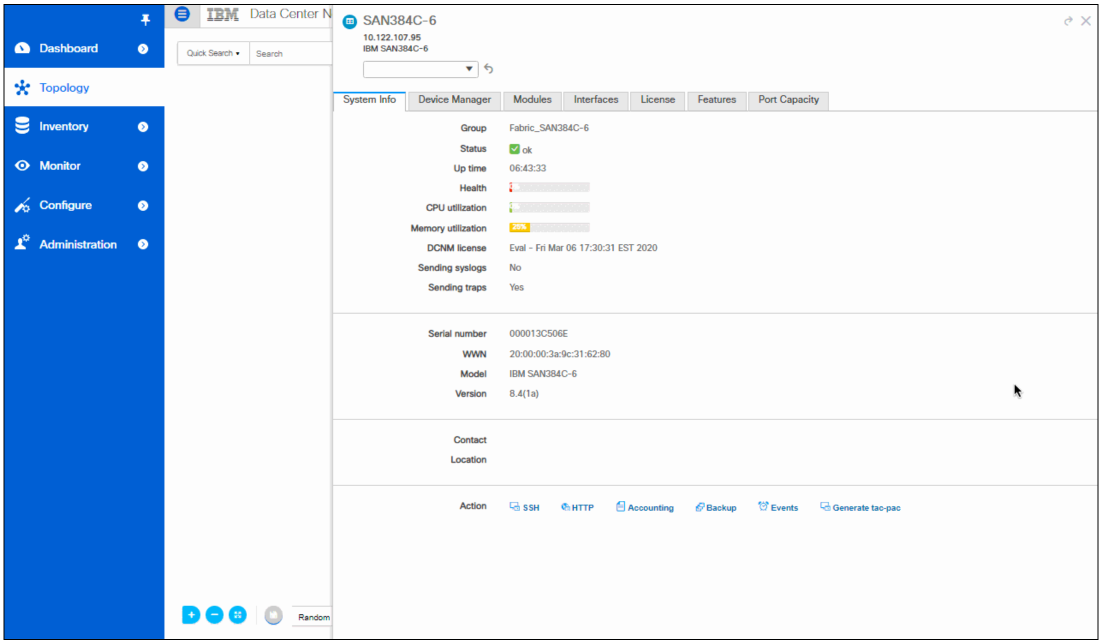

Group Fabric_SAN384C-6
Status ✓ ok
Up time 06:43:33
DCNM license Eval - Fri Mar 06 17:30:31 EST 2020
Sending syslogs No
Sending traps Yes

Serial number 000013C506E
WWN 20:00:00:3a:9c:31:62:80
Model IBM SAN384C-6
Version 8.4(1a)

Contact
Location

Figure 6-59 Switch information

You see the Device view in DM for the switch. Select **Admin** → **Licenses**, as shown in Figure 6-60.

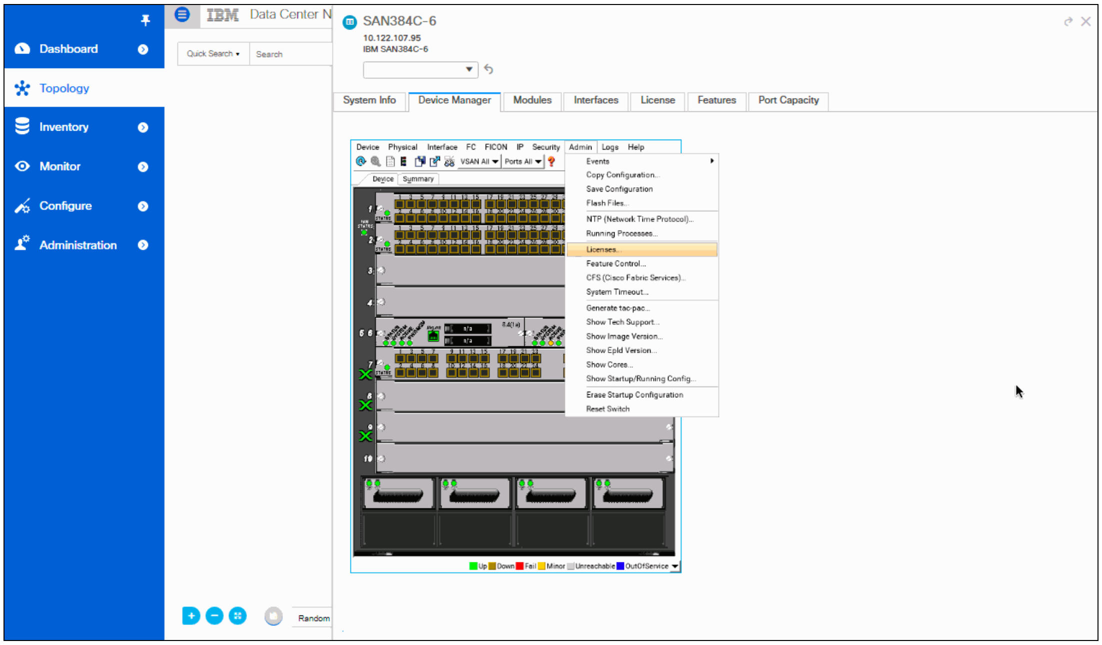

Figure 6-60 Opening the Licenses window from the DM

in Figure 6-61 on page 245, you see the Licenses window. You can see all the licenses that are applicable to this model of IBM c-type switch, but there are no licenses that are installed. Click the **Install** tab.

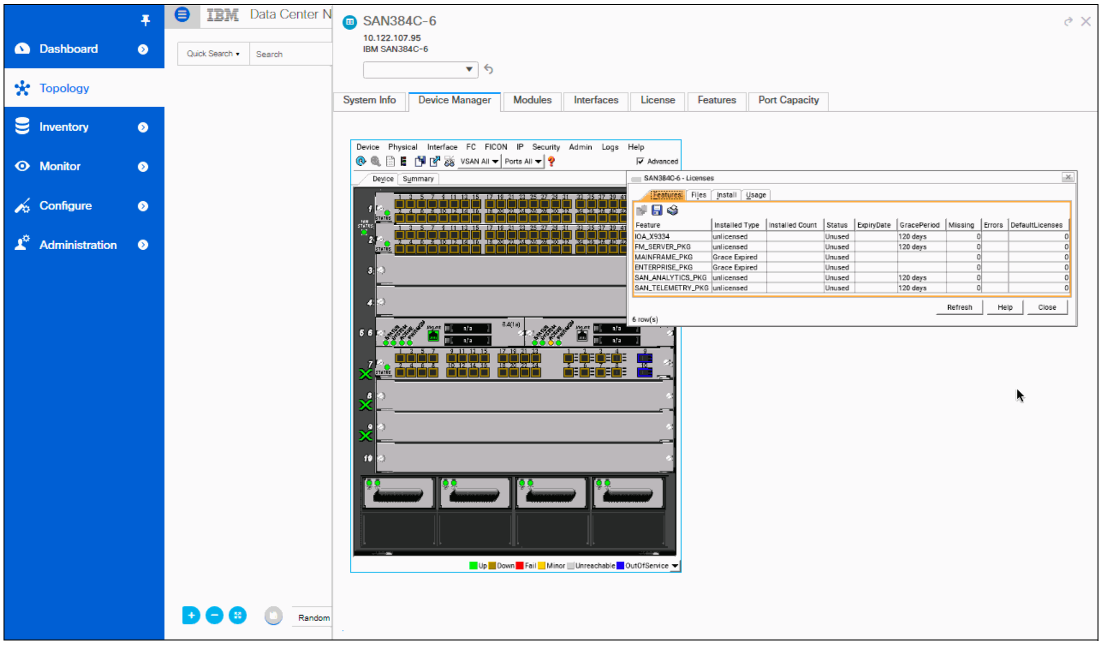

Feature	Installed Type	Installed Count	Status	ExpiryDate	GracePeriod	Missing	Errors	DefaultLicenses
IOA_X9334	unlicensed		Unused		120 days	0		0
FM_SERVER_PKG	unlicensed		Unused		120 days	0		0
MAINFRAME_PKG	Grace Expired		Unused			0		0
ENTERPRISE_PKG	Grace Expired		Unused			0		0
SAN_ANALYTICS_PKG	unlicensed		Unused		120 days	0		0
SAN_TELEMETRY_PKG	unlicensed		Unused		120 days	0		0

6 row(s)

Figure 6-61 Licenses window

Figure 6-62 shows the Install tab for the Licenses window. Select the drop-down menu for **URI** to see the licenses that are on the switch and available for installation. Select one of them as required.

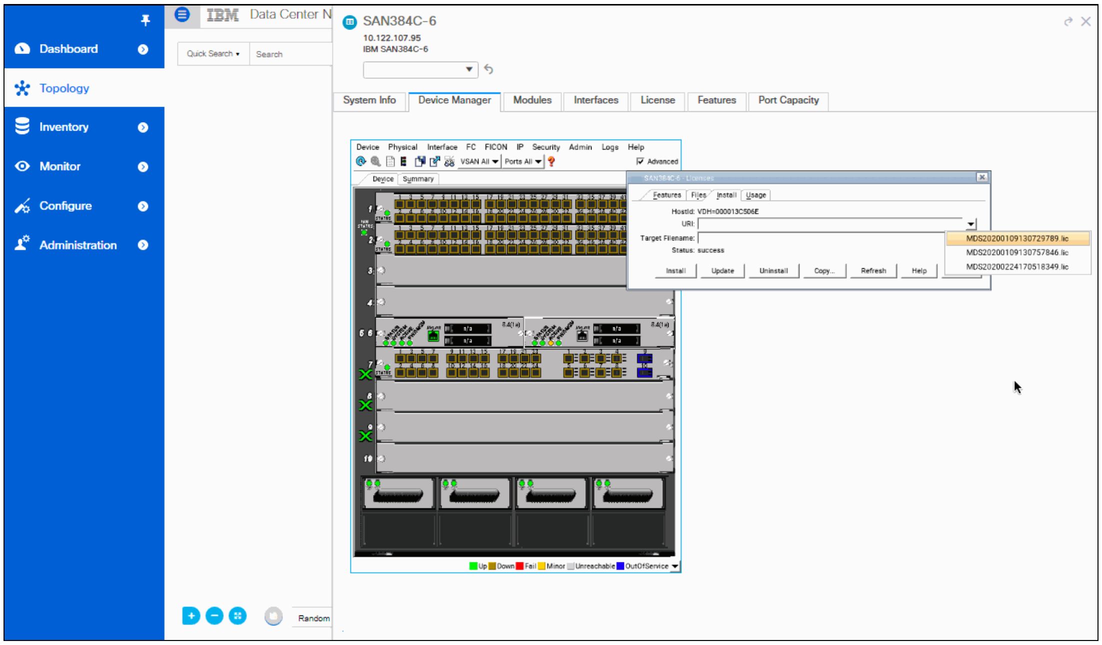

HostId: VDH=000013C506E
URI:
Target Filename:
Status: success

MDS20200109130729789.lic
MDS20200109130757846.lic
MDS20200224170518349.lic

Figure 6-62 Selecting a license

After the license is selected, the remaining fields on the window are automatically completed from the license file, as shown in Figure 6-63. Click **Install**.

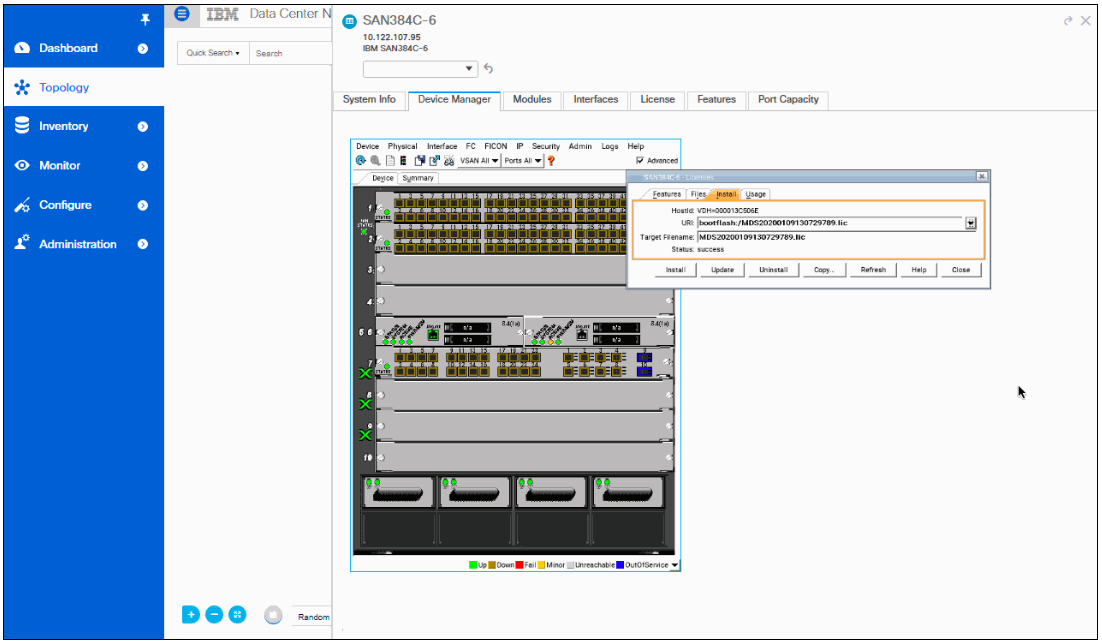

HostId: VDH=000013C506E
URI: bootflash:/MDS20200109130729789.lic
Target Filename: MDS20200109130729789.lic
Status: success

Figure 6-63 Installing the license

After a few seconds, the license installation is successful, as shown in Figure 6-64. Repeat the same steps to install all the licenses for this switch.

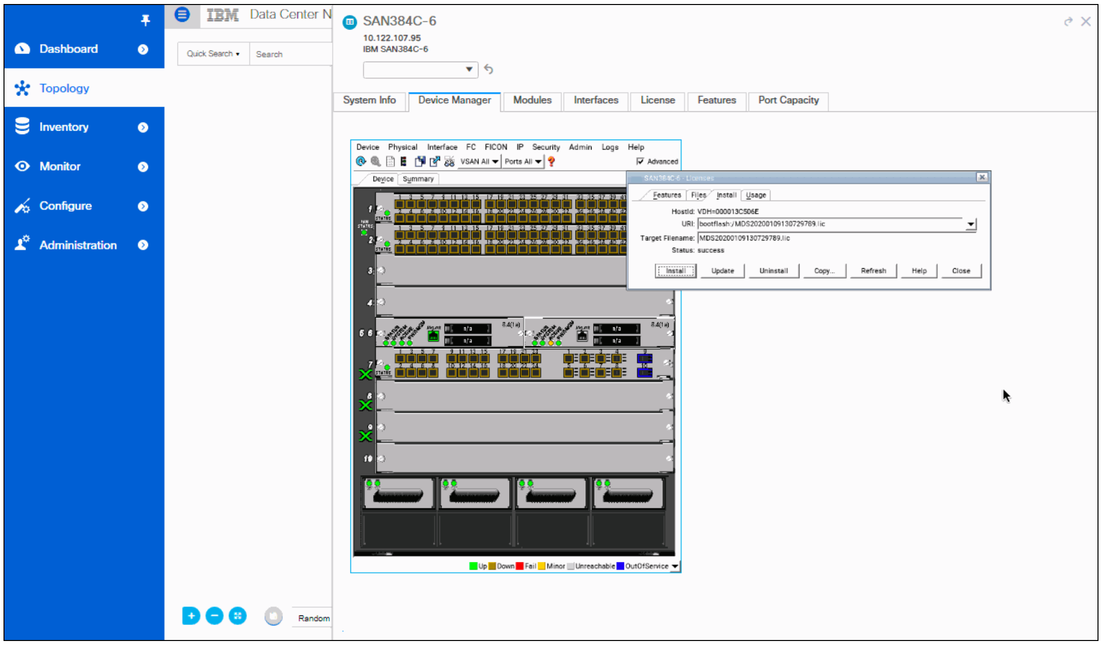

HostId: VDH=000013C506E
URI: bootflash:/MDS20200109130729789.lic
Target Filename: MDS20200109130729789.lic
Status: success

Figure 6-64 License installed

Now, click the **Features** tab. In Figure 6-65 on page 247, you now see the three installed licenses for the IBM Storage Networking SAN384C-6 switch for the lab environment.

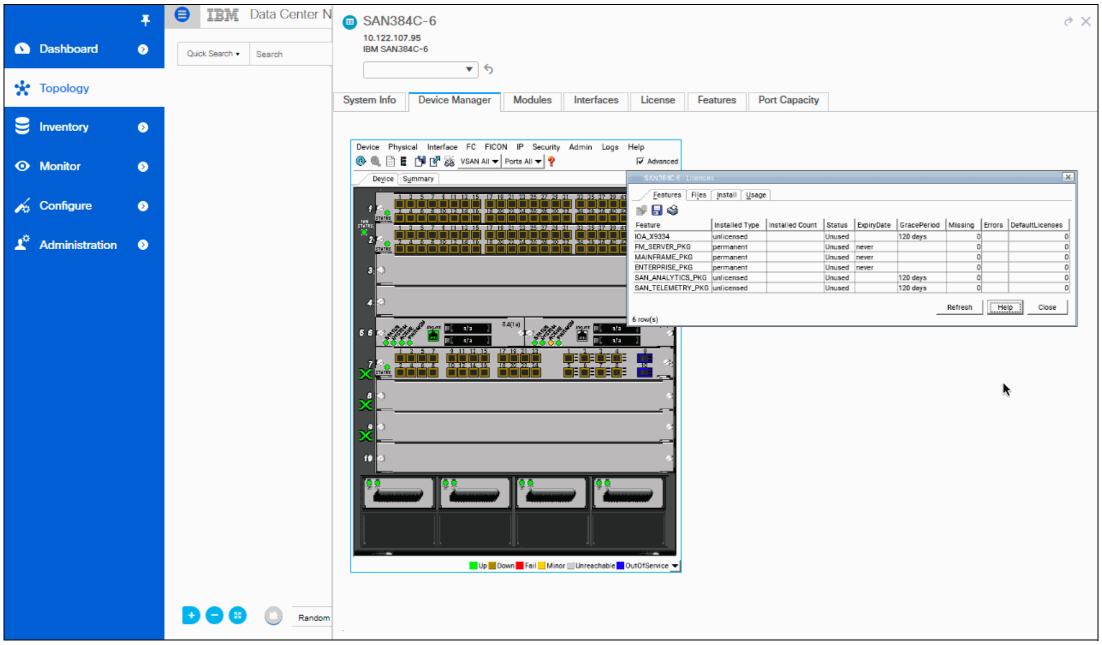

Feature	Installed Type	Installed Count	Status	ExpiryDate	GracePeriod	Missing	Errors	DefaultLicenses
IOA_X9334	unlicensed		Unused		120 days	0		0
FM_SERVER_PKG	permanent		Unused	never		0		0
MAINFRAME_PKG	permanent		Unused	never		0		0
ENTERPRISE_PKG	permanent		Unused	never		0		0
SAN_ANALYTICS_PKG	unlicensed		Unused		120 days	0		0
SAN_TELEMETRY_PKG	unlicensed		Unused		120 days	0		0

6 row(s)

Figure 6-65 Installed licenses

To build the lab environment, add all the licenses for the other switches.

6.9.4 Installing bulk licenses by using the DCNM

Since release 11.3(1), DCNM includes a feature that is called *bulk license installation*, where administrators can upload multiple licenses at a single instance and manage file-based licenses that are installed on the switches. DCNM parses the license files and extracts the switch serial numbers. It maps the serial numbers in the license files with the discovered fabric to install the licenses on each switch. License files are moved to bootflash and installed.

To do bulk incense installations, log in to DCNM and select **Administration** → **Manage Licensing** → **Switch Features**, as shown in Figure 6-66.

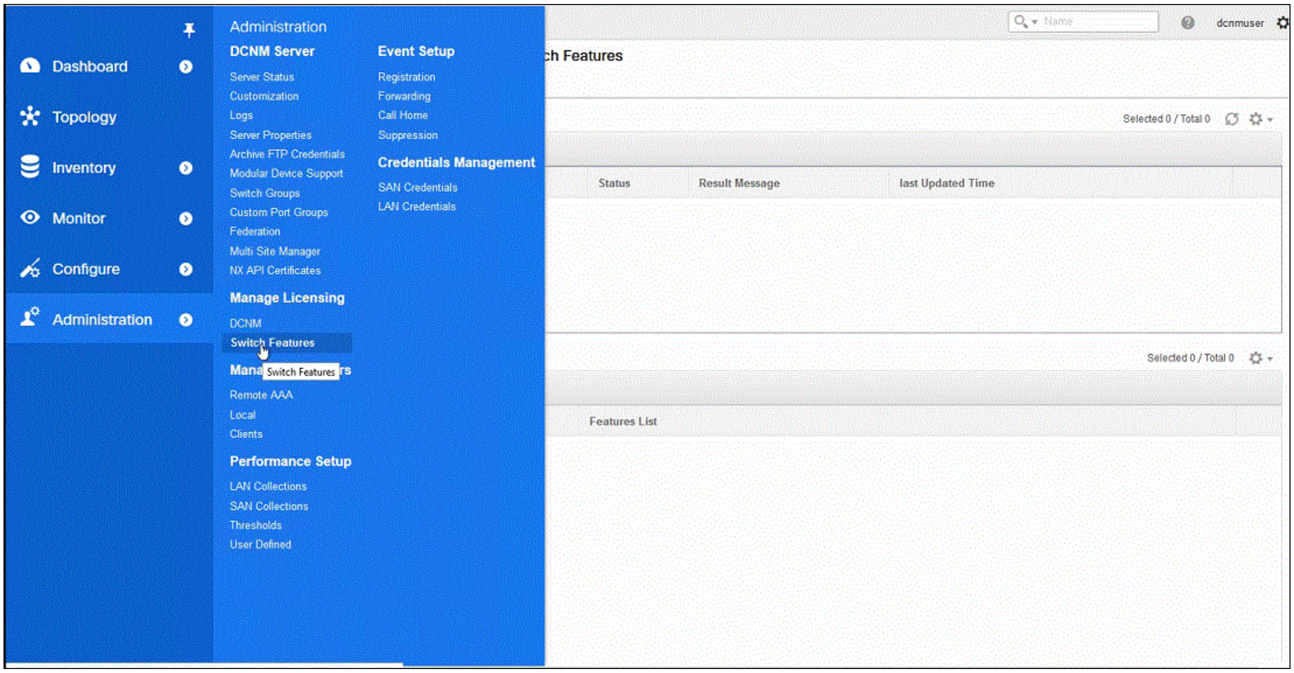

Figure 6-66   DCNM switch features

The window that is shown in Figure 6-67 opens.

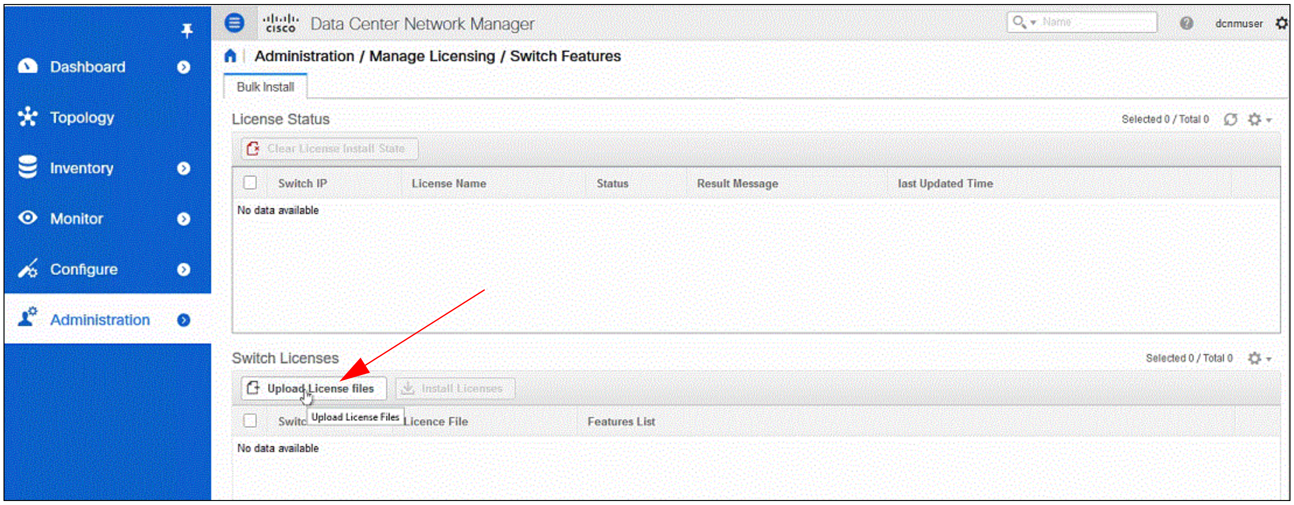

Figure 6-67   Bulk installation

In the Switch Licenses window, click **Upload License files**. The window that is shown in Figure 6-68 on page 249 opens.

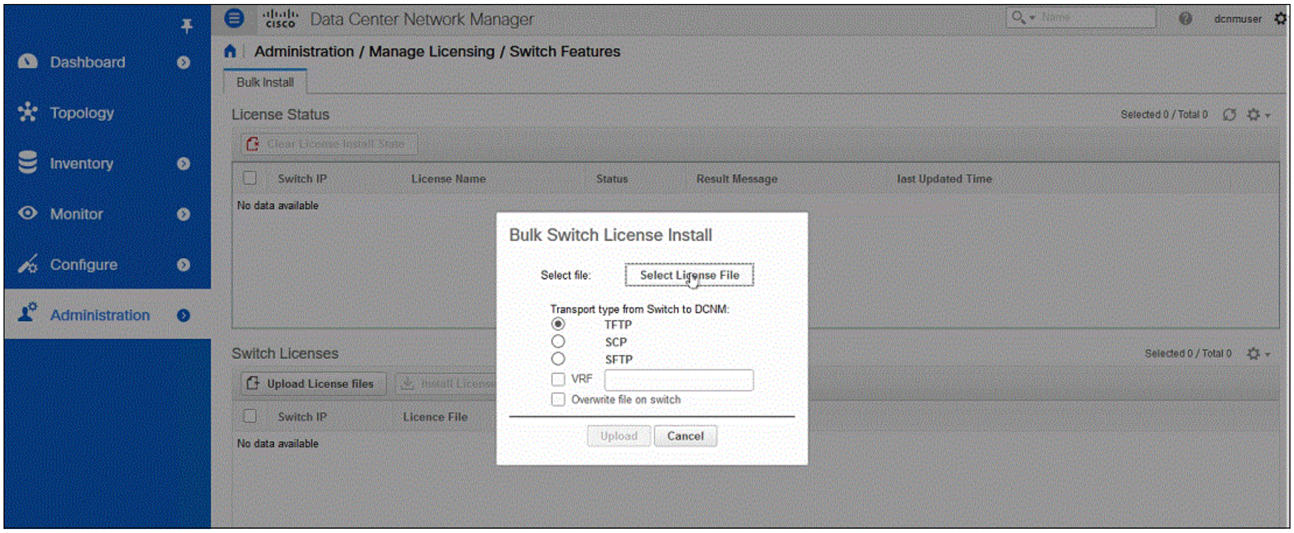

Figure 6-68   Uploading a license file

In the Bulk Switch License Install window, ensure that the correct file transfer protocol is selected. Select either **TFTP**, **SCP**, or **SFTP** to upload the license file. Not all protocols are supported for all platforms. TFTP is supported for Windows or RHEL DCNM SAN installation, but only SFTP and SCP are supported for all installation types. Click **Select License File**. The window that is shown in Figure 6-69 opens.

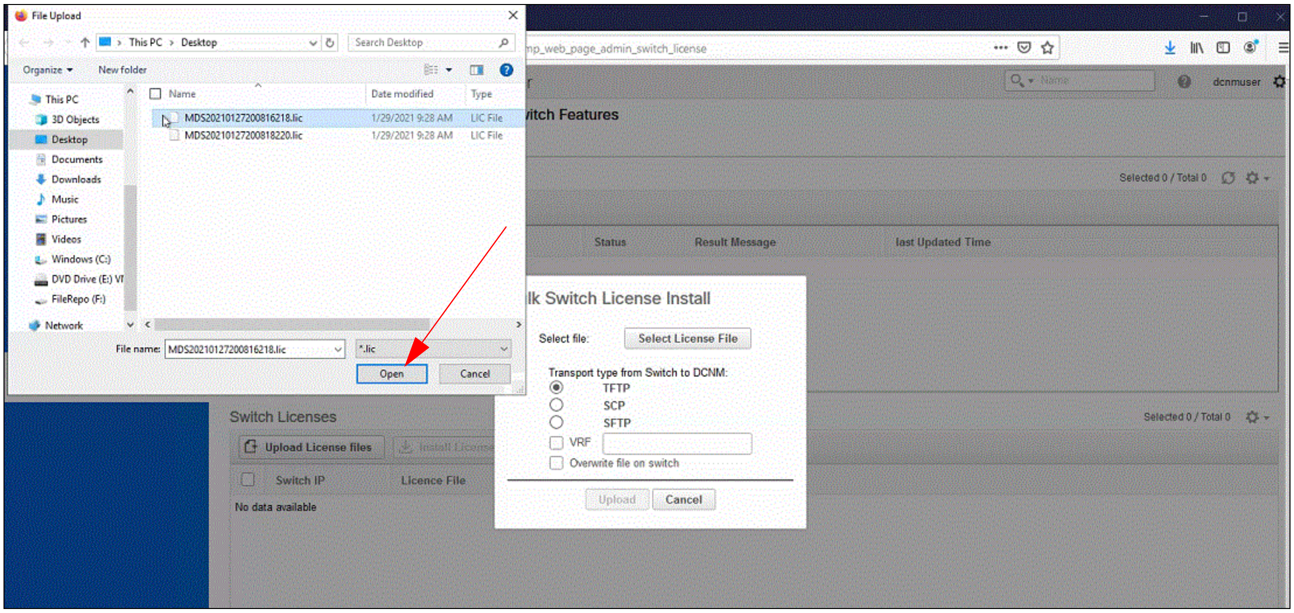

Figure 6-69   Selecting a license file

Select one or more license files. In our example, we select only one file because the other one was previously installed. After the file or files are selected, click **Open**.

The window that is shown in Figure 6-70 opens.

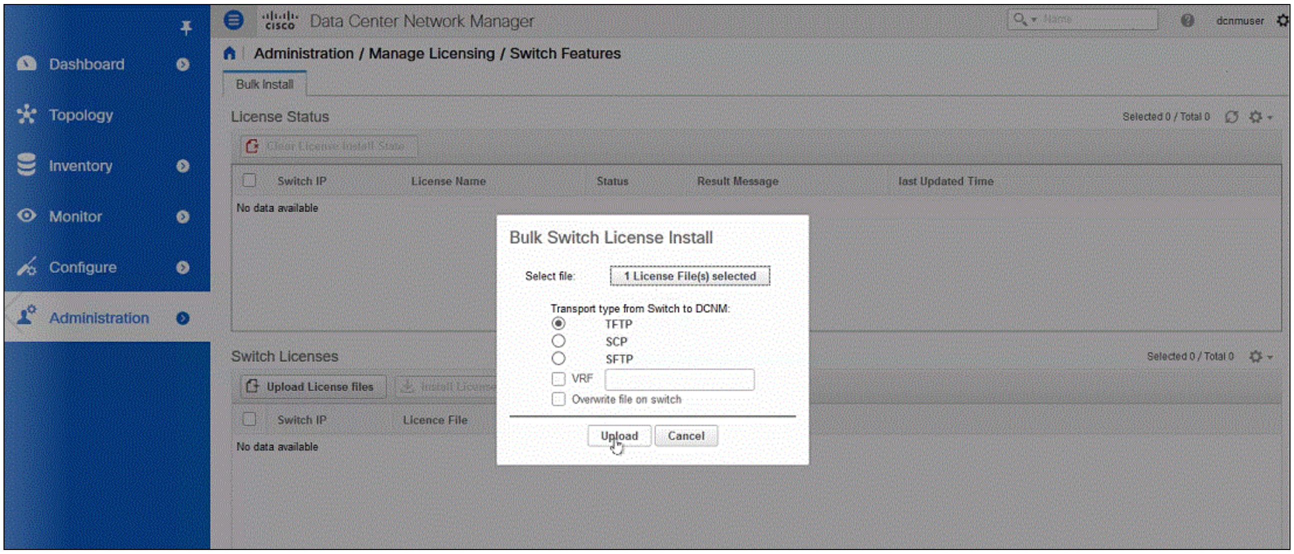

Switch IP	License Name	Status	Result Message	last Updated Time
No data available				

Bulk Switch License Install

Select file: 1 License File(s) selected

Transport type from Switch to DCNM:
● TFTP
○ SCP
○ SFTP
□ VRF
□ Overwrite file on switch

Switch IP	Licence File
No data available	

Figure 6-70 Uploading a license file

Click **Upload** to upload the selected file or files. The license file is uploaded, and the switch IP address to which the license is assigned is extracted along with the file name and feature list. The window that is shown in Figure 6-71 shows this information.

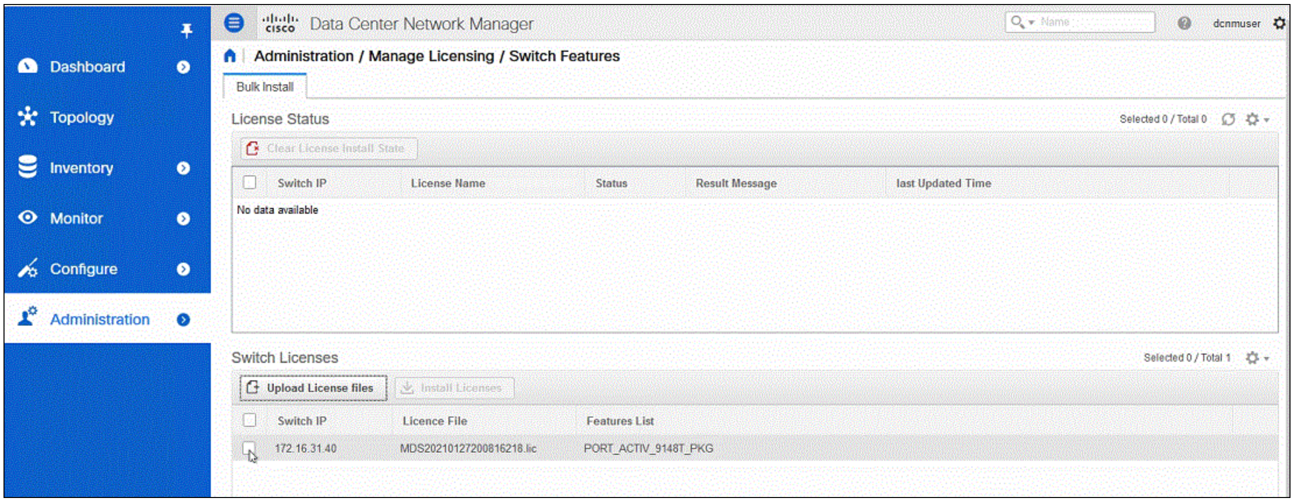

Switch IP	License Name	Status	Result Message	last Updated Time
No data available				

Switch IP	Licence File	Features List
172.16.31.40	MDS20210127200816218.lic	PORT_ACTIV_9148T_PKG

Figure 6-71 Uploaded license file

Select the licenses to be installed. In our example, we have only one file, but we could have selected multiple files. After the license is selected, the **Install** button becomes active, as shown in Figure 6-72 on page 251.

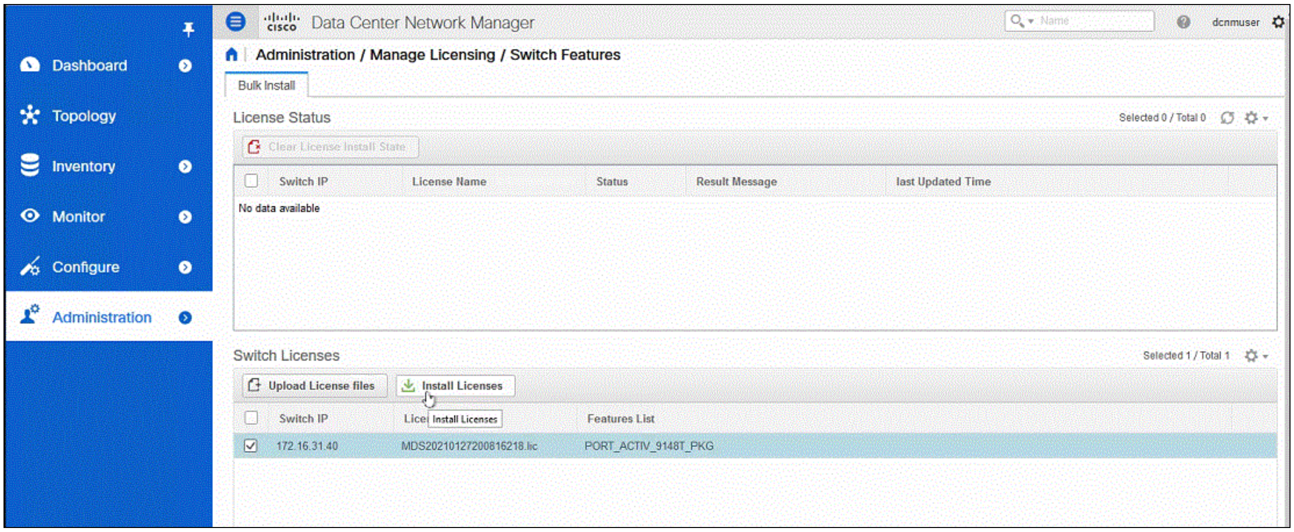

Administration / Manage Licensing / Switch Features

License Status

□	Switch IP	License Name	Status	Result Message	last Updated Time
No data available					

Switch Licenses

□	Switch IP	Lice	Features List
☑	172.16.31.40	MDS20210127200816218.lic	PORT_ACTIV_9148T_PKG

Figure 6-72 License selected and ready for installation

Click **Install Licenses**. The window that is shown in Figure 6-73 opens. The initial status is INSTALLING.

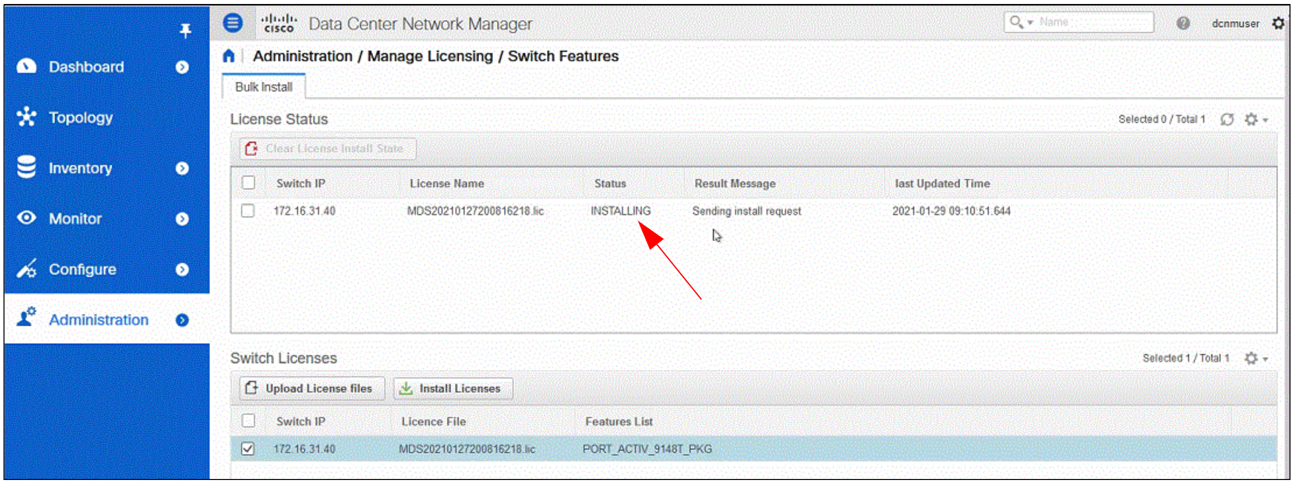

Administration / Manage Licensing / Switch Features

License Status

□	Switch IP	License Name	Status	Result Message	last Updated Time
□	172.16.31.40	MDS20210127200816218.lic	INSTALLING	Sending install request	2021-01-29 09:10:51.644

Switch Licenses

□	Switch IP	Licence File	Features List
☑	172.16.31.40	MDS20210127200816218.lic	PORT_ACTIV_9148T_PKG

Figure 6-73 License installing

After the installation is complete, the status changes to **DONE**, as shown in Figure 6-74.

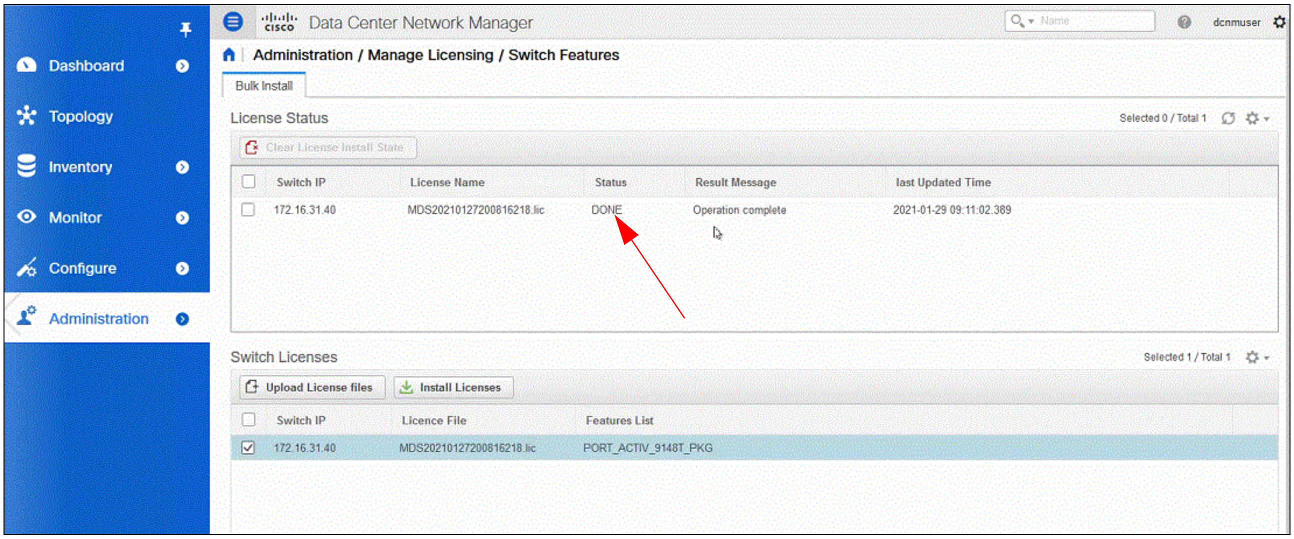

□	Switch IP	License Name	Status	Result Message	last Updated Time
□	172.16.31.40	MDS20210127200816218.lic	DONE	Operation complete	2021-01-29 09:11:02.389

□	Switch IP	Licence File	Features List
☑	172.16.31.40	MDS20210127200816218.lic	PORT_ACTIV_9148T_PKG

Figure 6-74 License installation complete

The installation of licenses on the switch or switches is now complete.

6.9.5 Installing license files by using the CLI

Before a license can be installed, ensure that it has been transferred to the switch, as described in 6.9.2, “Transferring license files to the switch” on page 235.

Log in to the switch by using an SSH client. In our examples, we used PuTTY. Example 6-24 lists the contents of the bootflash folder. We ensure that the correct file, which in our example is `MDS20210127200818220.lic`, is there.

Example 6-24 Listing the contents of the bootflash

```
CiscoFabB# dir bootflash:

       4096    Jan 13 04:55:32 2017  .partner/
       4096    Jan 13 04:55:43 2017  .patch/
        299    Jan 13 04:56:36 2017  20170113_045630_poap_4391_init.log
       1695    Nov 28 14:35:44 2018  20181128_143249_poap_4521_init.log
        168    Aug 24 00:01:16 2019  20190824_000047_poap_4518_init.log
          0    Aug 24 00:22:08 2019  20190824_002208_poap_4866_init.log
       3671    Oct 16 17:27:14 2019  20191016_172133_poap_4858_init.log
        840    Aug 24 00:05:01 2019  MDS20190823150219038.lic
        297    Jan 29 16:28:18 2021  MDS20210127200818220.lic
       4096    Aug 24 00:05:25 2019  lost+found/
   62054400    Jan 13 04:49:06 2017  m9148-s6ek9-kickstart-mz.8.3.1.bin
   55624192    Aug 24 00:04:25 2019  m9148-s6ek9-kickstart-mz.8.4.1.bin
  149837604    Jan 13 04:49:45 2017  m9148-s6ek9-mz.8.3.1.bin
  177906426    Aug 24 00:04:47 2019  m9148-s6ek9-mz.8.4.1.bin
       2893    Jan 11 08:31:00 2021  mts.log
       4096    Jan 13 04:56:18 2017  scripts/

Usage for bootflash://sup-local
  626331648 bytes used
```

```
 2889486336 bytes free
 3515817984 bytes total
```

Example 6-25 shows how to run the license installation and view the license information.

Example 6-25 License installation

```
CiscoFabB# install license bootflash:MDS20210127200818220.lic

Installing license .....done

CiscoFabB# show license

MDS20190823150219038.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT SAN_ANALYTICS_PKG cisco 1.0 22-aug-2022 uncounted \
        VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>L-D-M91S-AXK9</SKU> \
        HOSTID=VDH=JPG224600GU \
        NOTICE="<LicFileID>20190823150219038</LicFileID><LicLineID>1</LicLineID> \
        <PAK></PAK>" SIGN=38DF102A41F0
INCREMENT FM_SERVER_PKG cisco 1.0 permanent uncounted \
        VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M91ENTDCNMX-K9</SKU>
\
        HOSTID=VDH=JPG224600GU \
        NOTICE="<LicFileID>20190823150219038</LicFileID><LicLineID>2</LicLineID> \
        <PAK></PAK>" SIGN=9BC81CC4EDE0
INCREMENT ENTERPRISE_PKG cisco 1.0 permanent uncounted \
        VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M91ENTDCNMX-K9</SKU>
\
        HOSTID=VDH=JPG224600GU \
        NOTICE="<LicFileID>20190823150219038</LicFileID><LicLineID>3</LicLineID> \
        <PAK></PAK>" SIGN=0067A8188CE6

MDS20210127200818220.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT PORT_ACTIV_9148T_PKG cisco 1.0 permanent 24 \
        VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M9148T-PL12</SKU> \
        HOSTID=VDH=JPG224600GU \
        NOTICE="<LicFileID>20210127200818220</LicFileID><LicLineID>1</LicLineID> \
        <PAK></PAK>" SIGN=1EB58EDC2C26



CiscoFabB# show license brief

MDS20210127200818220.lic

CiscoFabB# show license usage

Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 Yes   -   Unused never       -
ENTERPRISE_PKG                Yes   -   Unused never       -
```

```
SAN_ANALYTICS_PKG             Yes    -   Unused 22 Aug 2022 -
PORT_ACTIV_9148T_PKG          Yes  48    In use never       -
--------------------------------------------------------------------------------
```

7

IBM Storage Networking c-type configuration

In this chapter, we describe the IBM Fibre Connection (FICON) configuration on IBM c-type switches. We describe local switched, local administrator, IP Security (IPsec), and Internet Key Exchange (IKE) protocol prerequisites when configuring Fibre Channel over IP (FCIP). We use a combination of Data Center Network Manager (DCNM), Device Manager (DM), and command-line interface (CLI) to perform the configurations in this chapter.

> **Note:** DCNM and DM installation is described in Chapter 6, “Initial connectivity and setup” on page 183.

The FICON mainframe interface capabilities benefit the IBM c-type family by supporting FICON storage area network (SAN) environments.

The following topics are covered in this chapter:

- Hardware Configuration Definition
- IBM Z Hardware Configuration Definition

7.1 Hardware Configuration Definition

With a Hardware Configuration Definition (HCD), you can define I/O configurations for software and hardware on a single interface. When defining configurations, use HCD to create an input/output definition file (IODF).

The following major tasks can be performed when using HCD:

- Define a new IODF or parts of an IDOF configuration.
- Activate configuration data.
- View and modify configuration data.
- Maintain IODF copies, imports, and exports.
- Use query and print configuration data.
- Migrate configuration data.
- Leverage automatic I/O configuration to compare a defined or active I/O configuration to an available I/O configuration.

7.2 IBM Z Hardware Configuration Definition

In this section, we describe what is required to define IBM c-type FICON switches to the IBM Z platform. We assume that the reader has a working knowledge of HCD, so we focus on the details to ensure that the necessary values match between the mainframe-based configuration and the configuration that is done within the FICON switches.

For all the examples in the following sections and the HCD definitions, we use the topology that is shown in Figure 4-1 on page 92.

Figure 7-1 shows the Hardware Configuration screen. We select option 1, Design, Modify, or View Configuration Data.

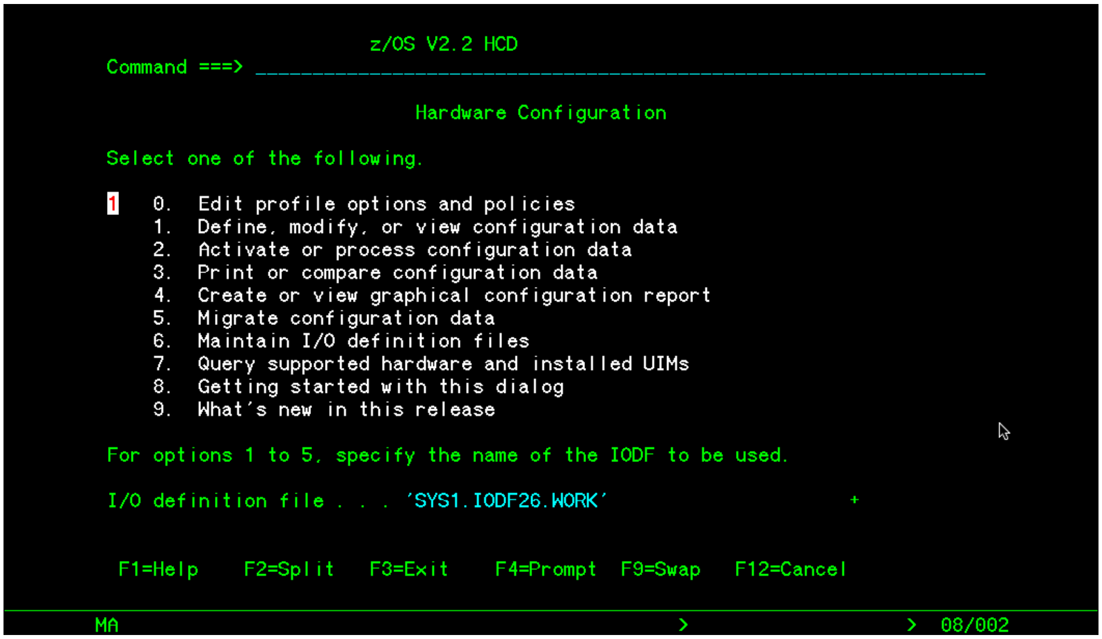

```
                       z/OS V2.2 HCD
Command ===> ______________________________________________________________

                           Hardware Configuration

Select one of the following.

1   0.  Edit profile options and policies
    1.  Define, modify, or view configuration data
    2.  Activate or process configuration data
    3.  Print or compare configuration data
    4.  Create or view graphical configuration report
    5.  Migrate configuration data
    6.  Maintain I/O definition files
    7.  Query supported hardware and installed UIMs
    8.  Getting started with this dialog
    9.  What's new in this release

For options 1 to 5, specify the name of the IODF to be used.

I/O definition file . . . 'SYS1.IODF26.WORK'                +


 F1=Help    F2=Split   F3=Exit    F4=Prompt  F9=Swap   F12=Cancel

MA                                               >                   >  08/002
```

Figure 7-1 Hardware Configuration screen

7.2.1 Defining the FICON switches to IBM Z

Define each FICON virtual storage area network (VSAN) as a switch within HCD. The IBM Z hardware configuration does not have a concept of a VSAN, so the mainframe thinks that each FICON VSAN is a unique switch.

Starting at the Design, Modify, or View Configuration Data screen, select option 2 - Switches, as shown in Figure 7-2.

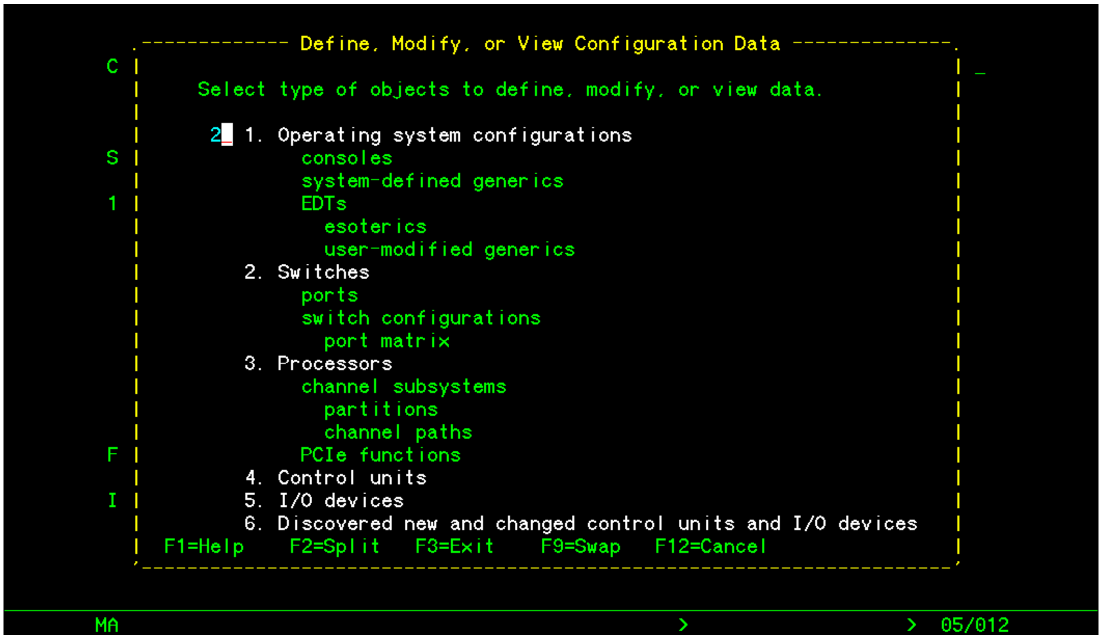

```
   .------------- Define, Modify, or View Configuration Data --------------.
 C |                                                                       | _
   |     Select type of objects to define, modify, or view data.           |
   |                                                                       |
   |      2_ 1. Operating system configurations                            |
 S |              consoles                                                 |
   |              system-defined generics                                  |
 1 |              EDTs                                                     |
   |                esoterics                                              |
   |                user-modified generics                                 |
   |         2. Switches                                                   |
   |              ports                                                    |
   |              switch configurations                                    |
   |                port matrix                                            |
   |         3. Processors                                                 |
   |              channel subsystems                                       |
   |                partitions                                             |
   |                channel paths                                          |
 F |              PCIe functions                                           |
   |         4. Control units                                              |
 I |         5. I/O devices                                                |
   |         6. Discovered new and changed control units and I/O devices   |
   |  F1=Help     F2=Split    F3=Exit     F9=Swap    F12=Cancel            |
   '-----------------------------------------------------------------------'

  MA                                                 >                   >  05/012
```

Figure 7-2   Define, Modify, or View Configuration Data

Go to the Add Switch screen by selecting F11, as shown in Figure 7-3.

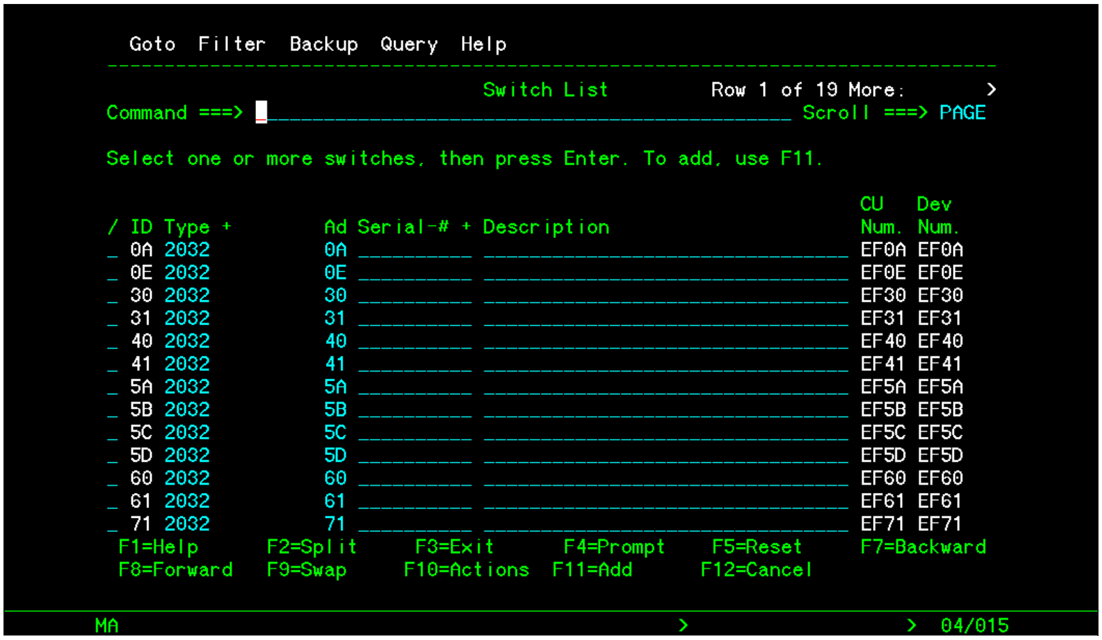

```
   Goto  Filter  Backup  Query  Help
 ------------------------------------------------------------------------------
                                  Switch List         Row 1 of 19 More:       >
 Command ===> _____________________________________________ Scroll ===> PAGE

 Select one or more switches, then press Enter. To add, use F11.

                                                                    CU   Dev
 / ID Type +          Ad Serial-# + Description                     Num. Num.
 _ 0A 2032            0A __________ _______________________________ EF0A EF0A
 _ 0E 2032            0E __________ _______________________________ EF0E EF0E
 _ 30 2032            30 __________ _______________________________ EF30 EF30
 _ 31 2032            31 __________ _______________________________ EF31 EF31
 _ 40 2032            40 __________ _______________________________ EF40 EF40
 _ 41 2032            41 __________ _______________________________ EF41 EF41
 _ 5A 2032            5A __________ _______________________________ EF5A EF5A
 _ 5B 2032            5B __________ _______________________________ EF5B EF5B
 _ 5C 2032            5C __________ _______________________________ EF5C EF5C
 _ 5D 2032            5D __________ _______________________________ EF5D EF5D
 _ 60 2032            60 __________ _______________________________ EF60 EF60
 _ 61 2032            61 __________ _______________________________ EF61 EF61
 _ 71 2032            71 __________ _______________________________ EF71 EF71
  F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
  F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

  MA                                                 >                   >  04/015
```

Figure 7-3   Switch List

The first switch that we define is for VSAN 40 on switch IBM Storage Networking SAN384C-6. We specify both the Switch ID field and the Switch Address field first. Both of these fields must match the VSAN domain ID. In HCD, this domain ID is entered in hexadecimal format, but on the switch, it can be entered either in hexadecimal or decimal format, with the default being decimal format. We describe how to ensure that this field matches in 7.2.9, “Creating virtual storage area networks” on page 293, when the VSAN is defined on the switch.

We specify the switch type as 2032, which means that we are defining a FICON Director.

For the installed port range, it is a best practice to enter the complete architectural port range that is possible (0x00 - 0xFD). This range may be smaller, but there is no particular advantage to making it smaller, and because FICON Port Addresses on IBM c-type switches are all virtual, you have flexibility for the future.

We enter the switch control unit (CU) number and the switch device number. These items are the definitions of the IBM Control Unit Port (CUP) device. In our example network, we set the FICON CUP CU number and the CUP device number as the same, but this approach is not required.

After completing all the fields as shown in Figure 7-4, we press Enter.

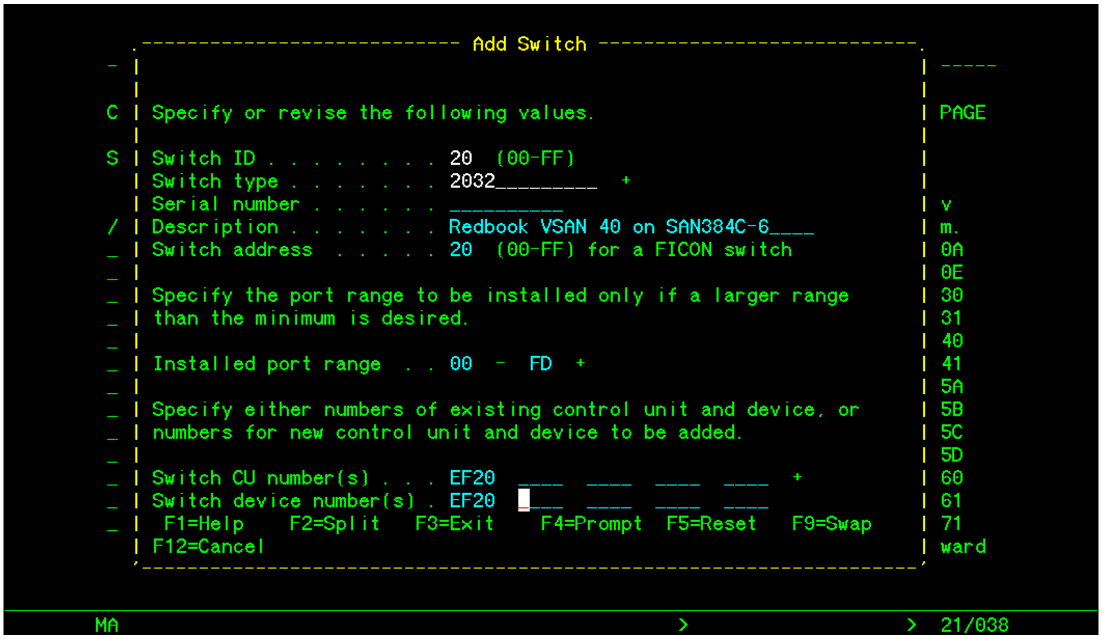

```
      .--------------------------- Add Switch ---------------------------.
    - |                                                                  | -----
      |                                                                  |
    C | Specify or revise the following values.                          | PAGE
      |                                                                  |
    S | Switch ID . . . . . . . . 20  (00-FF)                            |
      | Switch type . . . . . . . 2032_________  +                       |
      | Serial number . . . . . . __________                             | v
    / | Description . . . . . . . Redbook VSAN 40 on SAN384C-6____       | m.
    _ | Switch address  . . . . . 20  (00-FF) for a FICON switch         | 0A
    _ |                                                                  | 0E
    _ | Specify the port range to be installed only if a larger range    | 30
    _ | than the minimum is desired.                                     | 31
    _ |                                                                  | 40
    _ | Installed port range  . . 00  -  FD  +                           | 41
    _ |                                                                  | 5A
    _ | Specify either numbers of existing control unit and device, or   | 5B
    _ | numbers for new control unit and device to be added.             | 5C
    _ |                                                                  | 5D
    _ | Switch CU number(s) . . . EF20  ____  ____  ____  ____  +        | 60
    _ | Switch device number(s) . EF20  ____  ____  ____  ____           | 61
    _ |  F1=Help    F2=Split    F3=Exit    F4=Prompt  F5=Reset   F9=Swap | 71
      | F12=Cancel                                                       | ward
      '------------------------------------------------------------------'

       MA                                              >                   >  21/038
```

Figure 7-4 Add Switch

Figure 7-5 on page 259 shows the results of defining the first switch. It is expected that after this switch is defined, HCD sends a message that the CUP CU and device definitions are created, but we still have work to do on them. We describe this work after we define the FICON channels and CUs.

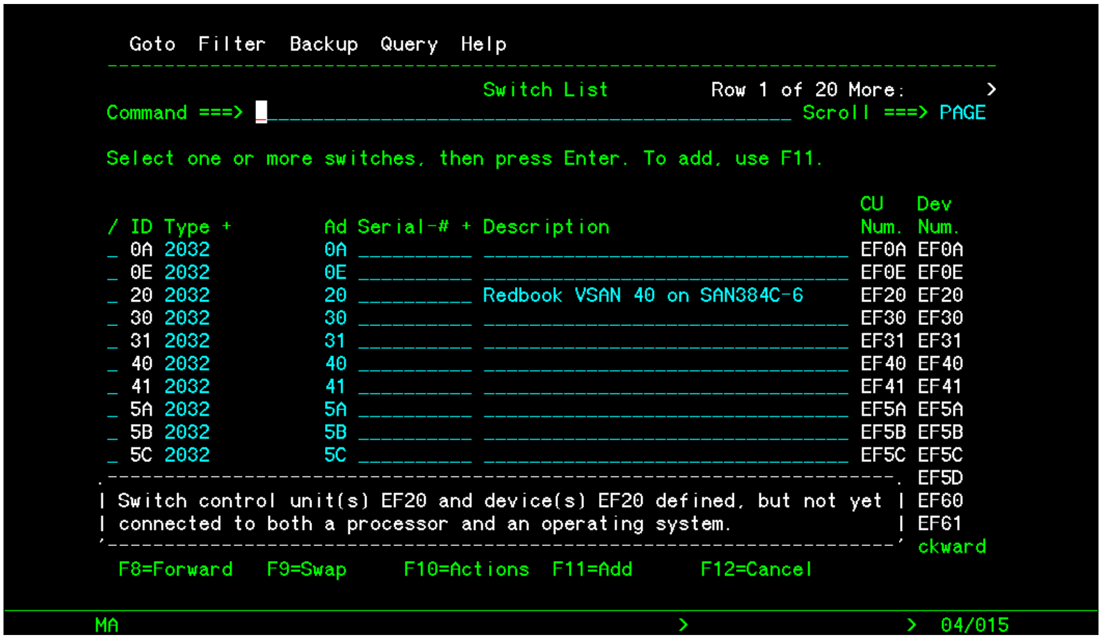

```
  Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                                Switch List          Row 1 of 20 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more switches, then press Enter. To add, use F11.

                                                                   CU   Dev
/ ID Type +          Ad Serial-# + Description                     Num. Num.
_ 0A 2032            0A __________ ________________________________ EF0A EF0A
_ 0E 2032            0E __________ ________________________________ EF0E EF0E
_ 20 2032            20 __________ Redbook VSAN 40 on SAN384C-6     EF20 EF20
_ 30 2032            30 __________ ________________________________ EF30 EF30
_ 31 2032            31 __________ ________________________________ EF31 EF31
_ 40 2032            40 __________ ________________________________ EF40 EF40
_ 41 2032            41 __________ ________________________________ EF41 EF41
_ 5A 2032            5A __________ ________________________________ EF5A EF5A
_ 5B 2032            5B __________ ________________________________ EF5B EF5B
_ 5C 2032            5C __________ ________________________________ EF5C EF5C
.---------------------------------------------------------------------. EF5D
| Switch control unit(s) EF20 and device(s) EF20 defined, but not yet | EF60
| connected to both a processor and an operating system.              | EF61
'---------------------------------------------------------------------' ckward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                                   >                   >  04/015
```

Figure 7-5   Defined first switch

Repeat this process for each of the FICON VSANs on the two switches in the test environment. Add the following IDs:

- Switch ID 0x21 for FICON VSAN 40 on switch IBM Storage Networking SAN192C-6
- Switch ID 0x10 for FICON VSAN 50 on switch IBM Storage Networking SAN384C-6
- Switch ID 0x11 for FICON VSAN 50 on switch IBM Storage Networking SAN192C-6

We do not need to create a switch definition within the mainframe hardware configuration for VSAN 100 on either the IBM Storage Networking SAN384C-6 or IBM Storage Networking SAN50C-R switches that are being used for disk replication because the data connection that is used for this replication is Fibre Channel Protocol (FCP), not FICON.

When all the switches are defined, the screen that is shown in Figure 7-6 opens.

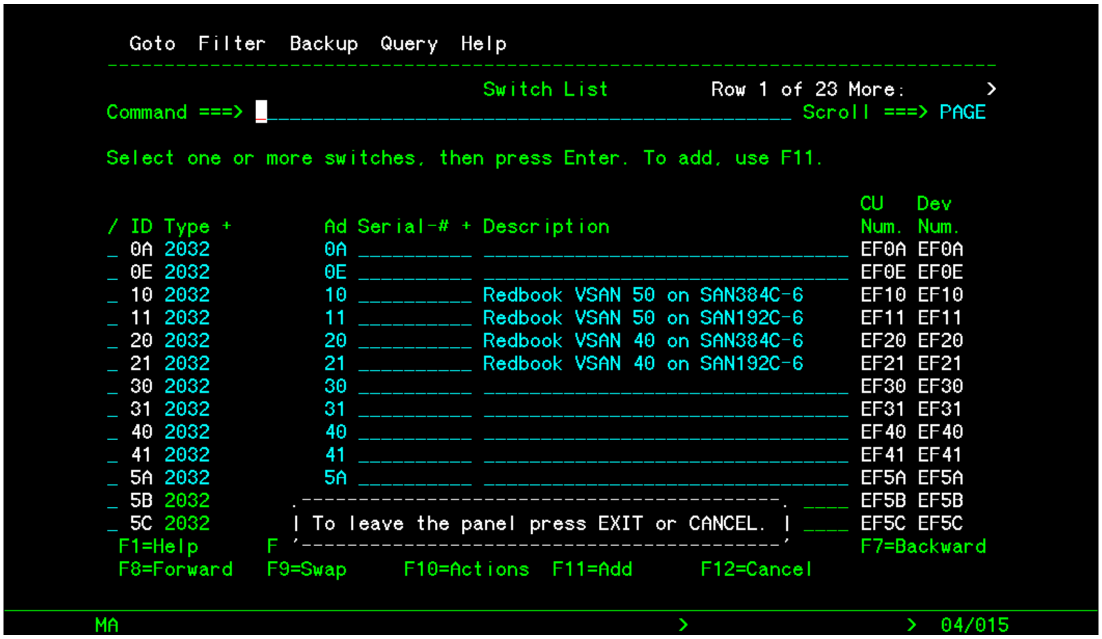

```
  Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                                 Switch List        Row 1 of 23 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more switches, then press Enter. To add, use F11.

                                                                 CU   Dev
/ ID Type +        Ad Serial-# + Description                     Num. Num.
_ 0A 2032          0A __________ ________________________________ EF0A EF0A
_ 0E 2032          0E __________ ________________________________ EF0E EF0E
_ 10 2032          10 __________ Redbook VSAN 50 on SAN384C-6     EF10 EF10
_ 11 2032          11 __________ Redbook VSAN 50 on SAN192C-6     EF11 EF11
_ 20 2032          20 __________ Redbook VSAN 40 on SAN384C-6     EF20 EF20
_ 21 2032          21 __________ Redbook VSAN 40 on SAN192C-6     EF21 EF21
_ 30 2032          30 __________ ________________________________ EF30 EF30
_ 31 2032          31 __________ ________________________________ EF31 EF31
_ 40 2032          40 __________ ________________________________ EF40 EF40
_ 41 2032          41 __________ ________________________________ EF41 EF41
_ 5A 2032          5A __________ ________________________________ EF5A EF5A
_ 5B 2032        .------------------------------------------. ____ EF5B EF5B
_ 5C 2032        | To leave the panel press EXIT or CANCEL. | ____ EF5C EF5C
 F1=Help      F '------------------------------------------'       F7=Backward
 F8=Forward   F9=Swap      F10=Actions  F11=Add      F12=Cancel

MA                                                  >                   >  04/015
```

Figure 7-6   All switches defined

7.2.2 Defining Channel Path IDs that are connected to the FICON switches

Now that all the FICON switches are defined to the hardware configuration, we must either add or modify the existing Channel Path IDs (CHPIDs) to represent how they are connected in our example network. Starting from the Design, Modify, or View Configuration Data screen, we select option 3 Processors, and the resulting screen is shown in Figure 7-7.

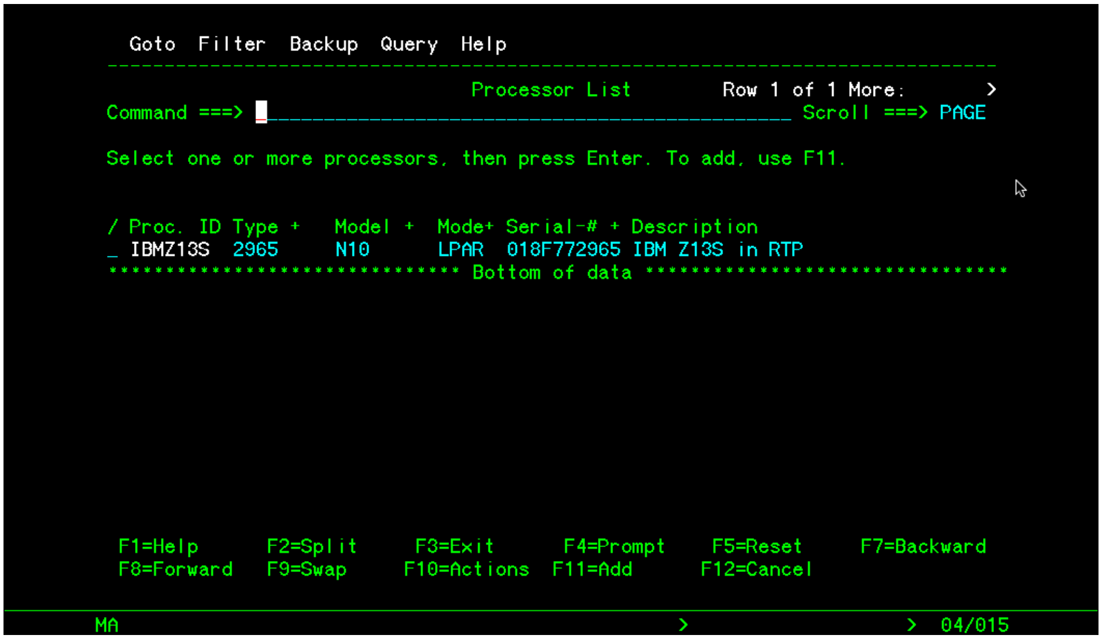

```
  Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                                Processor List       Row 1 of 1 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more processors, then press Enter. To add, use F11.


/ Proc. ID Type +   Model +   Mode+ Serial-# + Description
_ IBMZ13S  2965     N10       LPAR  018F772965 IBM Z13S in RTP
****************************** Bottom of data *******************************









 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap      F10=Actions  F11=Add      F12=Cancel

MA                                                  >                   >  04/015
```

Figure 7-7   Processor List

We select the appropriate processor that the channels are on. In our environment, we have only one processor, so we select it by placing an s to the left of the row, and press Enter. We see a list of the applicable channel subsystems (CSSs), as shown in Figure 7-8.

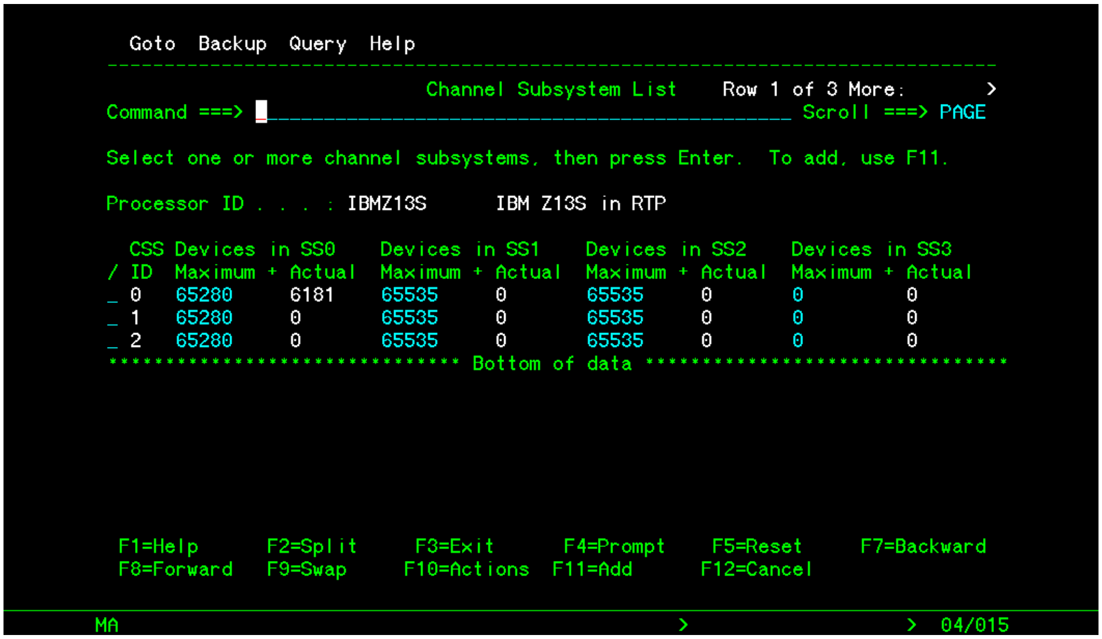

```
  Goto  Backup  Query  Help
 ------------------------------------------------------------------------------
                            Channel Subsystem List     Row 1 of 3 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more channel subsystems, then press Enter.  To add, use F11.

Processor ID . . . : IBMZ13S      IBM Z13S in RTP

  CSS Devices in SS0    Devices in SS1    Devices in SS2    Devices in SS3
/ ID  Maximum + Actual  Maximum + Actual  Maximum + Actual  Maximum + Actual
_ 0   65280     6181    65535     0       65535     0       0         0
_ 1   65280     0       65535     0       65535     0       0         0
_ 2   65280     0       65535     0       65535     0       0         0
******************************* Bottom of data ********************************





 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                                    >                   >  04/015
```

Figure 7-8   Channel Subsystem List

We select the appropriate CSS number by putting an s next to it (in our environment, it is 0), and press Enter. The resulting screen is shown in Figure 7-9.

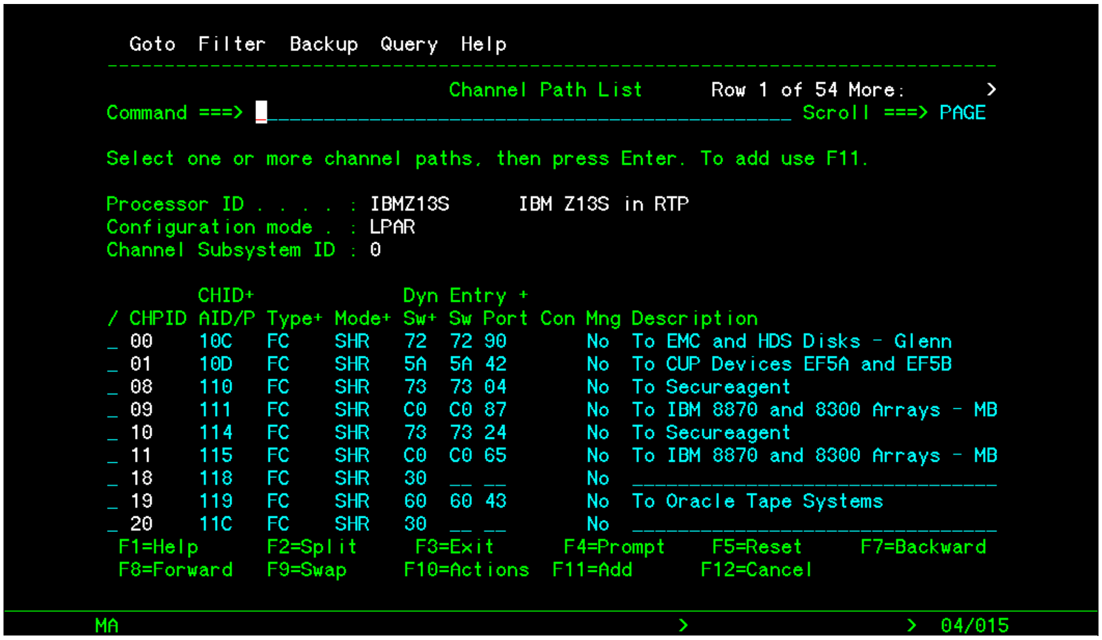

```
  Goto  Filter  Backup  Query  Help
 ------------------------------------------------------------------------------
                              Channel Path List      Row 1 of 54 More:        >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : IBMZ13S      IBM Z13S in RTP
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+                 Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 00    10C   FC    SHR   72  72 90       No  To EMC and HDS Disks - Glenn
_ 01    10D   FC    SHR   5A  5A 42       No  To CUP Devices EF5A and EF5B
_ 08    110   FC    SHR   73  73 04       No  To Secureagent
_ 09    111   FC    SHR   C0  C0 87       No  To IBM 8870 and 8300 Arrays - MB
_ 10    114   FC    SHR   73  73 24       No  To Secureagent
_ 11    115   FC    SHR   C0  C0 65       No  To IBM 8870 and 8300 Arrays - MB
_ 18    118   FC    SHR   30  __ __       No  ________________________________
_ 19    119   FC    SHR   60  60 43       No  To Oracle Tape Systems
_ 20    11C   FC    SHR   30  __ __       No  ________________________________
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                                    >                   >  04/015
```

Figure 7-9   Channel Path List

Looking at the example network, the first channels that we must define are CHPIDs 18 and 20, which are used to access the IBM DS8870 system. In our example, these CHPIDs exist and are being reused for this network. If new channels must be added, use the same dialog.

For more information about adding new channels, see *z/OS V2R2 Hardware Configuration Definition (HCD) Users Guide*, SC34-2669-02.

We type a / next to CHPID 18 and press Enter. Now, we see the action selection dialog for this CHPID, as shown in Figure 7-10.

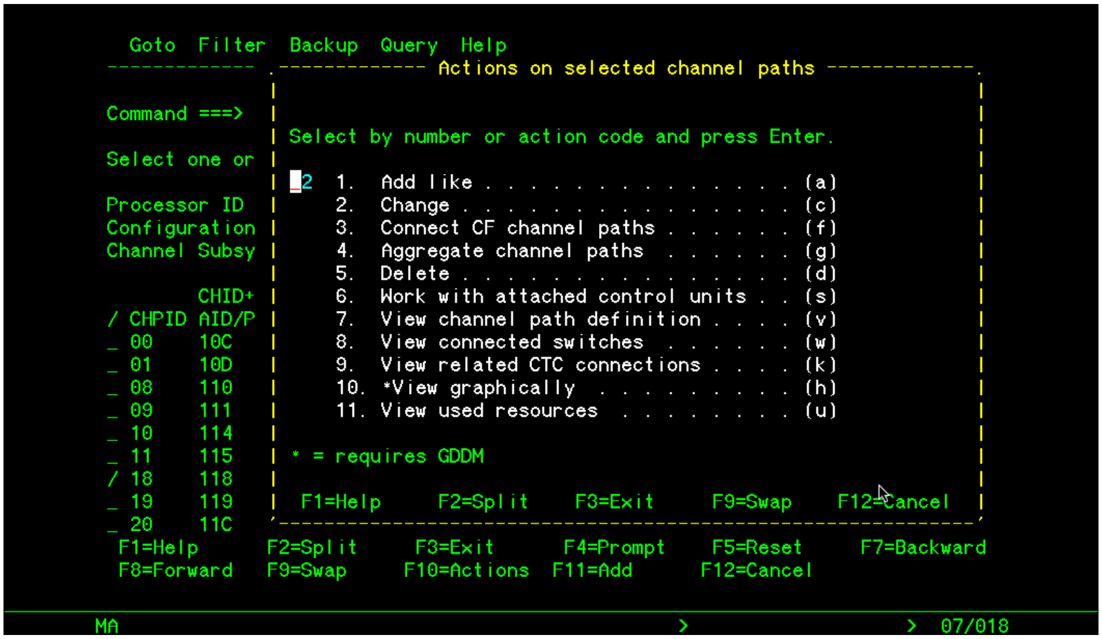

```
    Goto  Filter  Backup  Query  Help
  ------------- .------------- Actions on selected channel paths -------------.
                |                                                              |
  Command ===>  |                                                              |
                | Select by number or action code and press Enter.             |
  Select one or |                                                              |
                | _2  1.  Add like . . . . . . . . . . . . . . (a)             |
  Processor ID  |     2.  Change . . . . . . . . . . . . . . . (c)             |
  Configuration |     3.  Connect CF channel paths . . . . . . (f)             |
  Channel Subsy |     4.  Aggregate channel paths  . . . . . . (g)             |
                |     5.  Delete . . . . . . . . . . . . . . . (d)             |
          CHID+ |     6.  Work with attached control units . . (s)             |
  / CHPID AID/P |     7.  View channel path definition . . . . (v)             |
  _ 00    10C   |     8.  View connected switches  . . . . . . (w)             |
  _ 01    10D   |     9.  View related CTC connections . . . . (k)             |
  _ 08    110   |     10. *View graphically  . . . . . . . . . (h)             |
  _ 09    111   |     11. View used resources  . . . . . . . . (u)             |
  _ 10    114   |                                                              |
  _ 11    115   | * = requires GDDM                                            |
  / 18    118   |                                                              |
  _ 19    119   |  F1=Help     F2=Split     F3=Exit     F9=Swap    F12=Cancel  |
  _ 20    11C   '--------------------------------------------------------------'
   F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
   F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

 MA                                                  >                  >  07/018
```

Figure 7-10   Action selection screen

We type 2 for the Change option to get to the screen that is shown in Figure 7-11.

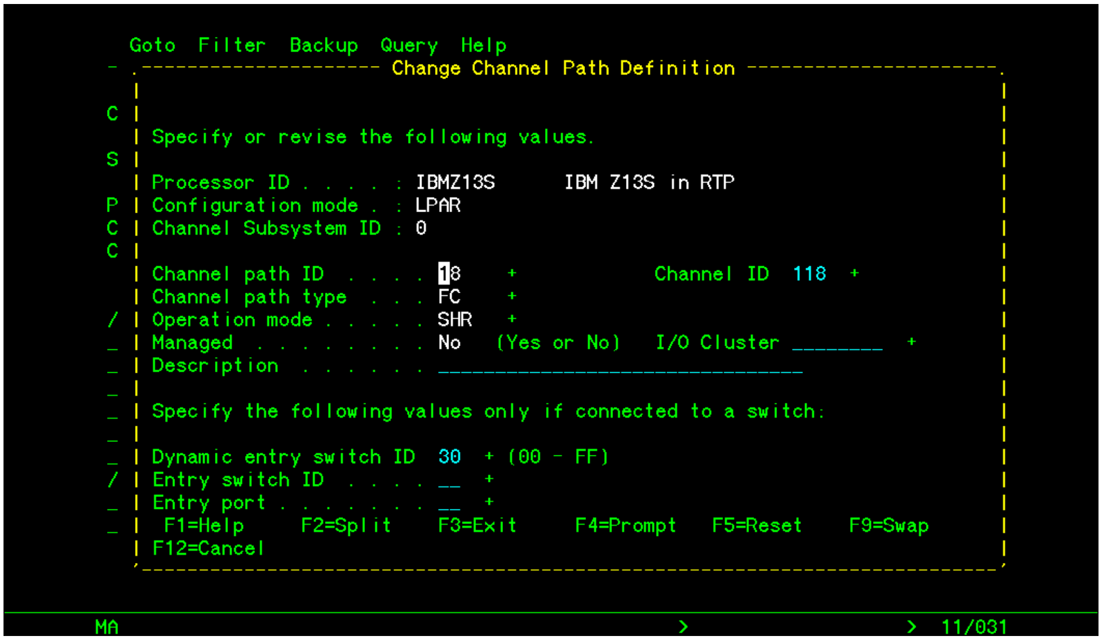

```
    Goto  Filter  Backup  Query  Help
  - .-------------------- Change Channel Path Definition ----------------------.
    |                                                                           |
  C |                                                                           |
    | Specify or revise the following values.                                   |
  S |                                                                           |
    | Processor ID . . . . : IBMZ13S      IBM Z13S in RTP                       |
  P | Configuration mode . : LPAR                                               |
  C | Channel Subsystem ID : 0                                                  |
  C |                                                                           |
    | Channel path ID  . . . . 18    +            Channel ID  118  +            |
    | Channel path type  . . . FC    +                                          |
  / | Operation mode . . . . . SHR   +                                          |
  _ | Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
  _ | Description  . . . . . . ________________________________                 |
  _ |                                                                           |
  _ | Specify the following values only if connected to a switch:               |
  _ |                                                                           |
  _ | Dynamic entry switch ID  30  + (00 - FF)                                  |
  / | Entry switch ID  . . . . __  +                                            |
  _ | Entry port . . . . . . . __  +                                            |
  _ |  F1=Help     F2=Split     F3=Exit     F4=Prompt    F5=Reset     F9=Swap    |
    | F12=Cancel                                                                |
    '---------------------------------------------------------------------------'

 MA                                                  >                  >  11/031
```

Figure 7-11   Change Channel Path Definition

For FICON CHPIDs, we must confirm that the Channel path type is defined as FC. Because in our environment we are sharing the channels between multiple logical partitions (LPARs), we need the Operation mode to be SHR. We can optionally add a description.

The Dynamic entry switch ID and Entry switch ID must have the same value, and they must match the domain ID for the VSAN to which the CHPID will be connected. This value is also the same as the switch ID for this VSAN that was defined in 7.2.1, “Defining the FICON switches to IBM Z” on page 257. As before, in HCD this value is entered in hexadecimal format.

Finally, we specify the Entry port to as the FICON Port Address that is assigned to the interface on the switch to which it will be connected. The Entry port is also specified in hexadecimal format. In our example network, the switch ID that is used for CHPID 18 is 0x20, and the Entry port value is 0x00. We show later how to verify that the FICON Port Address on the switch definition matches.

Figure 7-12 shows the CHPID change dialog after completion. Press Enter.

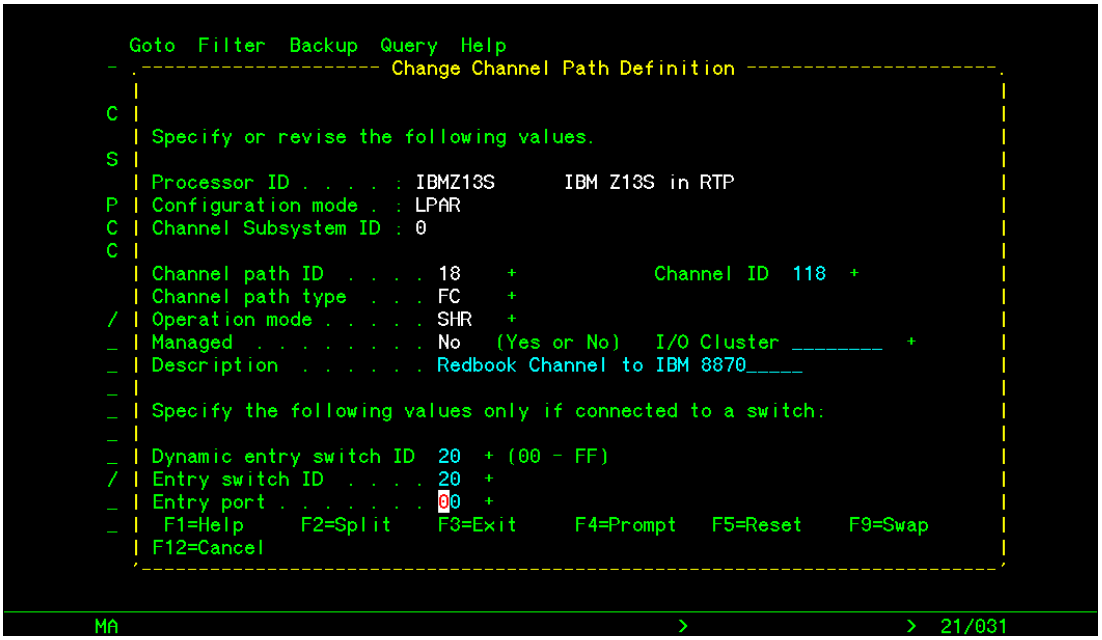

```
    Goto  Filter  Backup  Query  Help
  - .--------------------- Change Channel Path Definition ----------------------.
    |                                                                           |
  C |                                                                           |
    | Specify or revise the following values.                                   |
  S |                                                                           |
    | Processor ID . . . . : IBMZ13S      IBM Z13S in RTP                       |
  P | Configuration mode . : LPAR                                               |
  C | Channel Subsystem ID : 0                                                  |
  C |                                                                           |
    | Channel path ID  . . . . 18    +            Channel ID  118  +            |
    | Channel path type  . . . FC    +                                          |
  / | Operation mode . . . . . SHR   +                                          |
  _ | Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
  _ | Description  . . . . . . Redbook Channel to IBM 8870_____                 |
  _ |                                                                           |
  _ | Specify the following values only if connected to a switch:               |
  _ |                                                                           |
  _ | Dynamic entry switch ID  20  + (00 - FF)                                  |
  / | Entry switch ID  . . . . 20  +                                            |
  _ | Entry port . . . . . . . 00  +                                            |
  _ |  F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F9=Swap  |
    | F12=Cancel                                                                |
    '---------------------------------------------------------------------------'

      MA                                              >                   >  21/031
```

Figure 7-12   CHPID change dialog completed

We now see the screen to define the LPAR access list for this CHPID, which defines which LPARs are available to access this CHPID. We select the appropriate partitions and press Enter. Figure 7-13 shows the completed screen.

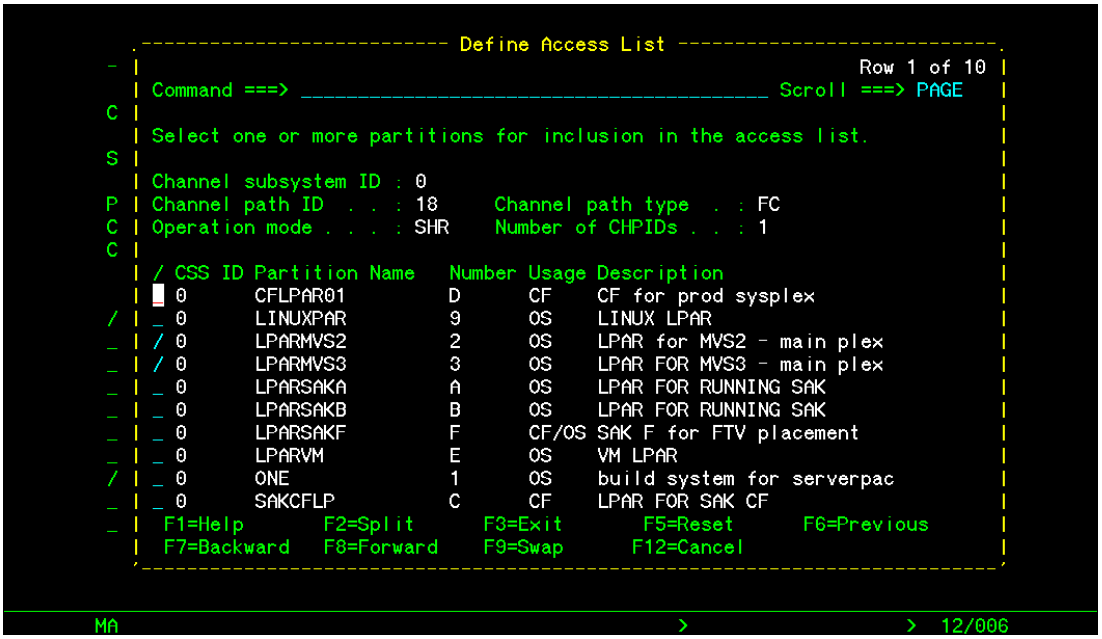

```
         .-------------------------- Define Access List ---------------------------.
       - |                                                          Row 1 of 10 |
         | Command ===> ________________________________________ Scroll ===> PAGE |
       C |                                                                      |
         | Select one or more partitions for inclusion in the access list.      |
       S |                                                                      |
         | Channel subsystem ID : 0                                             |
       P | Channel path ID . . : 18     Channel path type  . : FC               |
       C | Operation mode . . . : SHR    Number of CHPIDs . . : 1               |
       C |                                                                      |
         | / CSS ID Partition Name   Number Usage Description                   |
         | _ 0      CFLPAR01         D      CF    CF for prod sysplex           |
       / | _ 0      LINUXPAR         9      OS    LINUX LPAR                    |
       _ | / 0      LPARMVS2         2      OS    LPAR for MVS2 - main plex     |
       _ | / 0      LPARMVS3         3      OS    LPAR FOR MVS3 - main plex     |
       _ | _ 0      LPARSAKA         A      OS    LPAR FOR RUNNING SAK          |
       _ | _ 0      LPARSAKB         B      OS    LPAR FOR RUNNING SAK          |
       _ | _ 0      LPARSAKF         F      CF/OS SAK F for FTV placement       |
       _ | _ 0      LPARVM           E      OS    VM LPAR                       |
       / | _ 0      ONE              1      OS    build system for serverpac    |
       _ | _ 0      SAKCFLP          C      CF    LPAR FOR SAK CF               |
       _ |  F1=Help      F2=Split     F3=Exit      F5=Reset     F6=Previous     |
         |  F7=Backward  F8=Forward   F9=Swap     F12=Cancel                    |
         '----------------------------------------------------------------------'

      MA                                                   >                   >  12/006
```

Figure 7-13 Define Access List

The next screen that we see is for the selection of the candidate LPARs, which is another selection mechanism that is related to which LPARs are allowed to use a particular CHPID. For our environment, we press Enter. Figure 7-14 shows this screen.

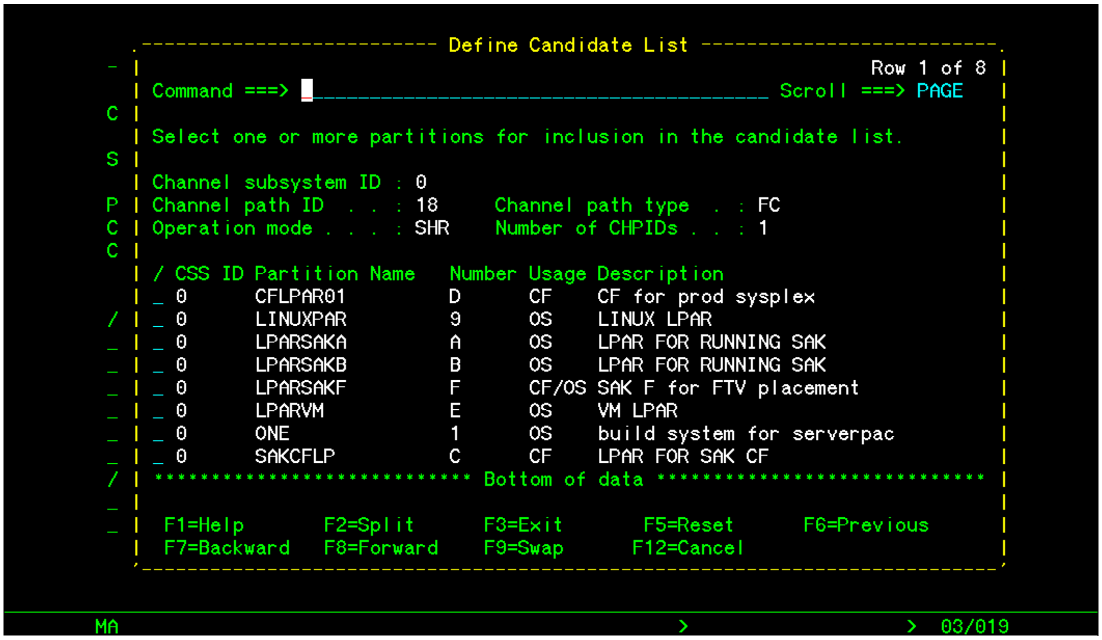

```
         .------------------------- Define Candidate List -------------------------.
       - |                                                           Row 1 of 8 |
         | Command ===> _________________________________________ Scroll ===> PAGE |
       C |                                                                      |
         | Select one or more partitions for inclusion in the candidate list.   |
       S |                                                                      |
         | Channel subsystem ID : 0                                             |
       P | Channel path ID . . : 18     Channel path type  . : FC               |
       C | Operation mode . . . : SHR    Number of CHPIDs . . : 1               |
       C |                                                                      |
         | / CSS ID Partition Name   Number Usage Description                   |
         | _ 0      CFLPAR01         D      CF    CF for prod sysplex           |
       / | _ 0      LINUXPAR         9      OS    LINUX LPAR                    |
       _ | _ 0      LPARSAKA         A      OS    LPAR FOR RUNNING SAK          |
       _ | _ 0      LPARSAKB         B      OS    LPAR FOR RUNNING SAK          |
       _ | _ 0      LPARSAKF         F      CF/OS SAK F for FTV placement       |
       _ | _ 0      LPARVM           E      OS    VM LPAR                       |
       _ | _ 0      ONE              1      OS    build system for serverpac    |
       _ | _ 0      SAKCFLP          C      CF    LPAR FOR SAK CF               |
       / | *************************** Bottom of data **************************** |
       _ |                                                                      |
       _ |  F1=Help      F2=Split     F3=Exit      F5=Reset     F6=Previous     |
         |  F7=Backward  F8=Forward   F9=Swap     F12=Cancel                    |
         '----------------------------------------------------------------------'

      MA                                                   >                   >  03/019
```

Figure 7-14 Define Candidate List

We have completed the modifications for CHPID 18 so that it is correctly defined to match our example network. On the screen that is shown in Figure 7-15, we can verify that the Dynamic and Entry Switch IDs and the Entry Port values are correct. All these values are in hexadecimal format.

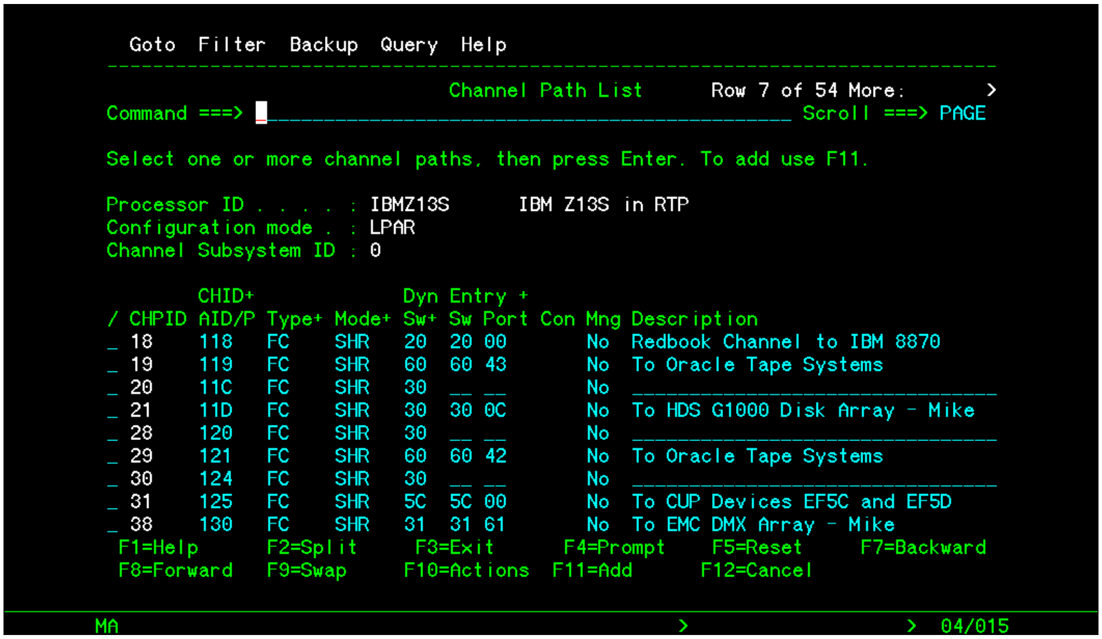

```
  Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                              Channel Path List      Row 7 of 54 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : IBMZ13S      IBM Z13S in RTP
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 18    118   FC    SHR   20  20 00       No  Redbook Channel to IBM 8870
_ 19    119   FC    SHR   60  60 43       No  To Oracle Tape Systems
_ 20    11C   FC    SHR   30  __ __       No  ________________________________
_ 21    11D   FC    SHR   30  30 0C       No  To HDS G1000 Disk Array - Mike
_ 28    120   FC    SHR   30  __ __       No  ________________________________
_ 29    121   FC    SHR   60  60 42       No  To Oracle Tape Systems
_ 30    124   FC    SHR   30  __ __       No  ________________________________
_ 31    125   FC    SHR   5C  5C 00       No  To CUP Devices EF5C and EF5D
_ 38    130   FC    SHR   31  31 61       No  To EMC DMX Array - Mike
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap      F10=Actions  F11=Add      F12=Cancel

MA                                                   >                   >  04/015
```

Figure 7-15   CHPID 18 modifications complete

Now, we repeat this process for each FICON CHPID that we will be using in the example network:

- CHPID 20 is attached to FICON Port Address 0x30 on FICON VSAN 40, which is on an IBM Storage Networking SAN384C-6 switch that uses switch ID 0x20.
- CHPID 70 is attached to FICON Port Address 0x02 on FICON VSAN 50, which is on an IBM Storage Networking SAN384C-6 switch that uses switch ID 0x10.
- CHPID 78 is attached to FICON Port Address 0x07 on FICON VSAN 50, which is on an IBM Storage Networking SAN384C-6 switch that uses switch ID 0x10.
- CHPID 28 is attached to FICON Port Address 0x22 on FICON VSAN 40, which is on an IBM Storage Networking SAN384C-6 switch that uses switch ID 0x20.
- CHPID 30 is attached to FICON Port Address 0x42 on FICON VSAN 40, which is on an IBM Storage Networking SAN384C-6 switch that uses switch ID 0x30.

After all these CHIPIDs are defined, you can verify that they match correctly, as shown in Figure 7-16.

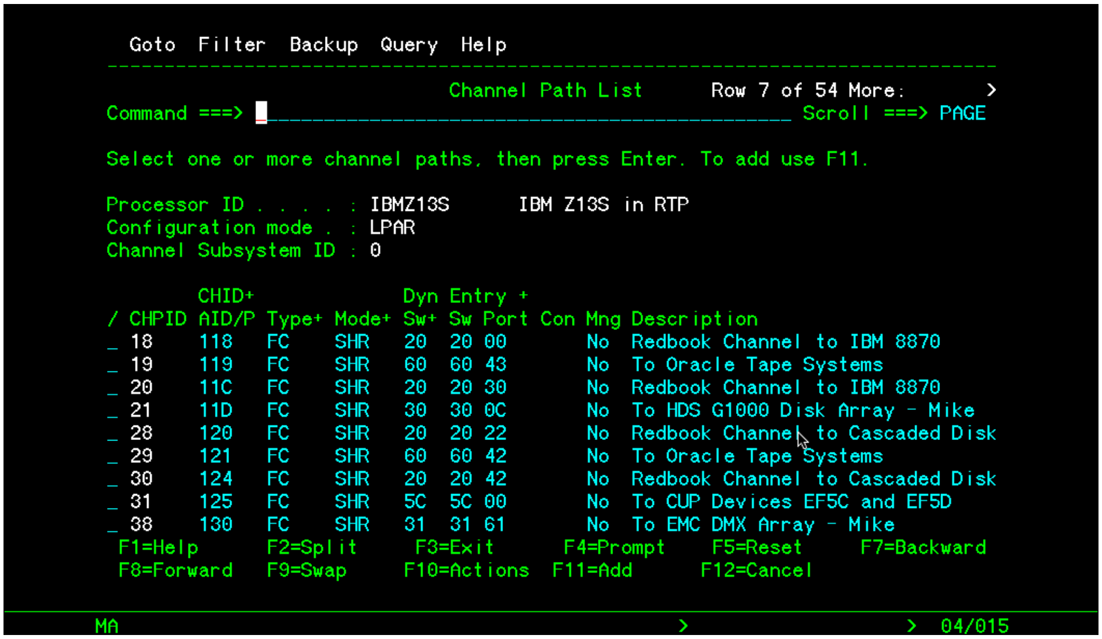

```
  Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                             Channel Path List      Row 7 of 54 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : IBMZ13S      IBM Z13S in RTP
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 18    118   FC    SHR   20  20 00       No  Redbook Channel to IBM 8870
_ 19    119   FC    SHR   60  60 43       No  To Oracle Tape Systems
_ 20    11C   FC    SHR   20  20 30       No  Redbook Channel to IBM 8870
_ 21    11D   FC    SHR   30  30 0C       No  To HDS G1000 Disk Array - Mike
_ 28    120   FC    SHR   20  20 22       No  Redbook Channel to Cascaded Disk
_ 29    121   FC    SHR   60  60 42       No  To Oracle Tape Systems
_ 30    124   FC    SHR   20  20 42       No  Redbook Channel to Cascaded Disk
_ 31    125   FC    SHR   5C  5C 00       No  To CUP Devices EF5C and EF5D
_ 38    130   FC    SHR   31  31 61       No  To EMC DMX Array - Mike
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                              >                   >  04/015
```

Figure 7-16   All CHPIDs modifications complete

CHPIDs 70 and 78 are shown in Figure 7-17. Although these CHPIDs are connected to the same physical chassis (IBM Storage Networking SAN384C-6), the IBM Z hardware thinks that they are connected to two different switches they are in different VSANs.

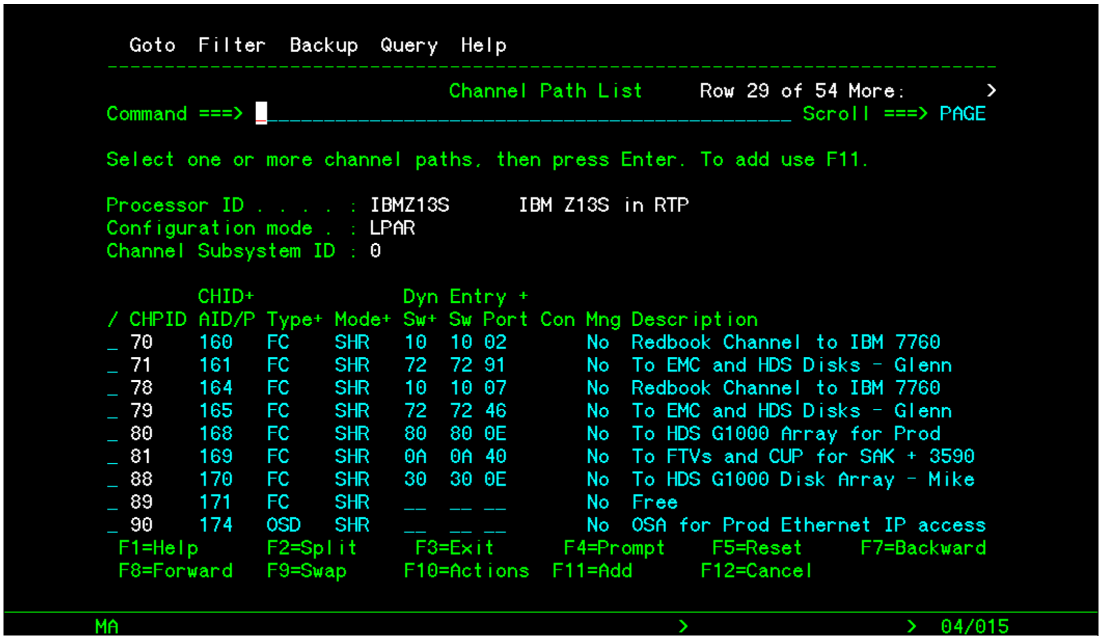

```
  Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                             Channel Path List     Row 29 of 54 More:       >
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : IBMZ13S      IBM Z13S in RTP
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 70    160   FC    SHR   10  10 02       No  Redbook Channel to IBM 7760
_ 71    161   FC    SHR   72  72 91       No  To EMC and HDS Disks - Glenn
_ 78    164   FC    SHR   10  10 07       No  Redbook Channel to IBM 7760
_ 79    165   FC    SHR   72  72 46       No  To EMC and HDS Disks - Glenn
_ 80    168   FC    SHR   80  80 0E       No  To HDS G1000 Array for Prod
_ 81    169   FC    SHR   0A  0A 40       No  To FTVs and CUP for SAK + 3590
_ 88    170   FC    SHR   30  30 0E       No  To HDS G1000 Disk Array - Mike
_ 89    171   FC    SHR   __  __ __       No  Free
_ 90    174   OSD   SHR   __  __ __       No  OSA for Prod Ethernet IP access
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                              >                   >  04/015
```

Figure 7-17   Channel Path List

7.2.3 Defining a local switched control unit that is connected to FICON switches

In mainframe environments, the disks and tapes that are accessed by the host CHPIDs are not dynamically discovered by using the Fibre Channel (FC) name server as they are in open systems. For FICON devices, you must manually define the connections in the IBM Z hardware configuration.

In this section, we define a disk array that is attached to the same switch as the host channels that will be accessing it. This configuration is known as a *locally switched device*. In the example network, there is an IBM DS8870 disk array that is attached to an IBM Storage Networking SAN386C-6 switch with two ports in VSAN 40 (one at FICON Port Address 0x10 and the other at port address 0x40). The host channels that access this disk array are CHPIDs 18 and 20, and they are connected to the same switch and VSAN, and FICON Port Addresses 0x00 and 0x30.

We start at the Design, Modify, or View Configuration Data screen and select option 4 for CUs. The resulting screen is shown in Figure 7-18.

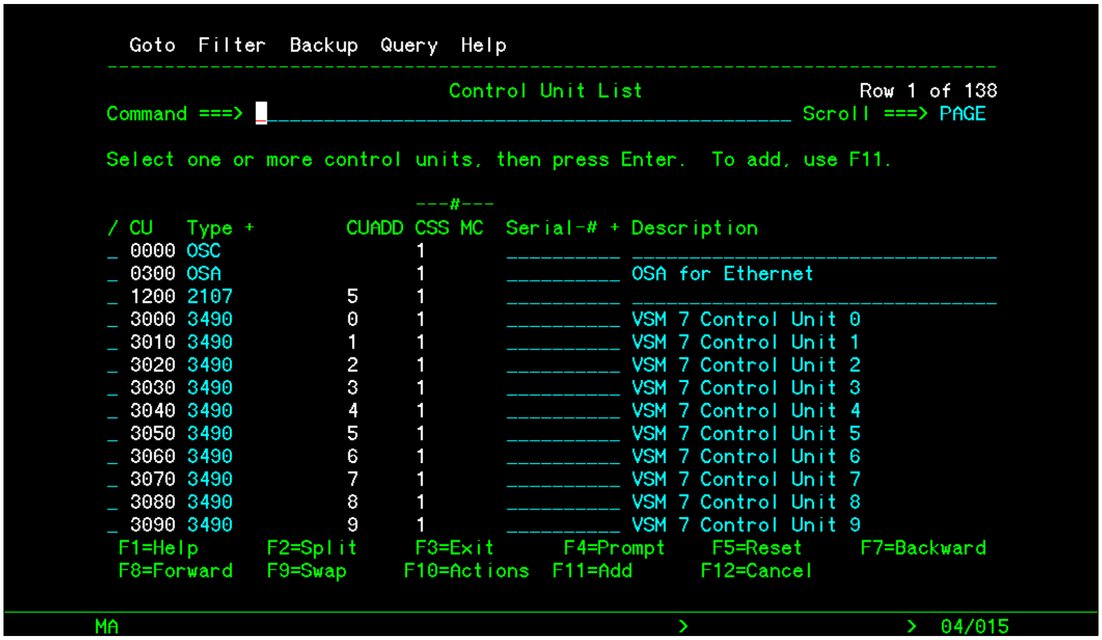

```
  Goto  Filter  Backup  Query  Help
--------------------------------------------------------------------------
                              Control Unit List                Row 1 of 138
Command ===> _____________________________________________ Scroll ===> PAGE

Select one or more control units, then press Enter.  To add, use F11.

                           ---#---
/ CU   Type +        CUADD CSS MC  Serial-# + Description
_ 0000 OSC                  1      __________ ________________________________
_ 0300 OSA                  1      __________ OSA for Ethernet
_ 1200 2107            5    1      __________ ________________________________
_ 3000 3490            0    1      __________ VSM 7 Control Unit 0
_ 3010 3490            1    1      __________ VSM 7 Control Unit 1
_ 3020 3490            2    1      __________ VSM 7 Control Unit 2
_ 3030 3490            3    1      __________ VSM 7 Control Unit 3
_ 3040 3490            4    1      __________ VSM 7 Control Unit 4
_ 3050 3490            5    1      __________ VSM 7 Control Unit 5
_ 3060 3490            6    1      __________ VSM 7 Control Unit 6
_ 3070 3490            7    1      __________ VSM 7 Control Unit 7
_ 3080 3490            8    1      __________ VSM 7 Control Unit 8
_ 3090 3490            9    1      __________ VSM 7 Control Unit 9
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                                >                   >  04/015
```

Figure 7-18 Control Unit List

For our network, we are adding a disk array (as opposed to modifying an existing one), so we press F11 to add it.

For this new CU, we must define the type and characteristics of the device and the connectivity. In our example, we enter the CU number as 9A00 and specify the device type as 2107. We may enter a description and serial number if we want.

Now, we define the connectivity from the host to this disk CU. We have two ports from the host going to this CU, and they are both connected to FICON VSAN 40 on an IBM Storage Networking SAN384C-6 switch, which is defined as domain 0x20. The FICON VSAN number is not specified anywhere in this screen because HCD and the hardware configuration do not have visibility to VSANs. For each place where we enter the switch number, we must input the associated FICON Port Address that is connected to by the disk array.

Note: In our example, we use a single FICON fabric. In practice, there are always at least two, and often four or even eight parallel fabrics that are used to connect between the multiple host channels that access disk and tape CUs. So in the HCD dialog, all the different switches that are used for access would be listed in the “Connected to switches” list along with their associated ports on each of these switches.

The fields in the Add Control Unit dialog are shown in Figure 7-19. Now, we press Enter to move to the next screen.

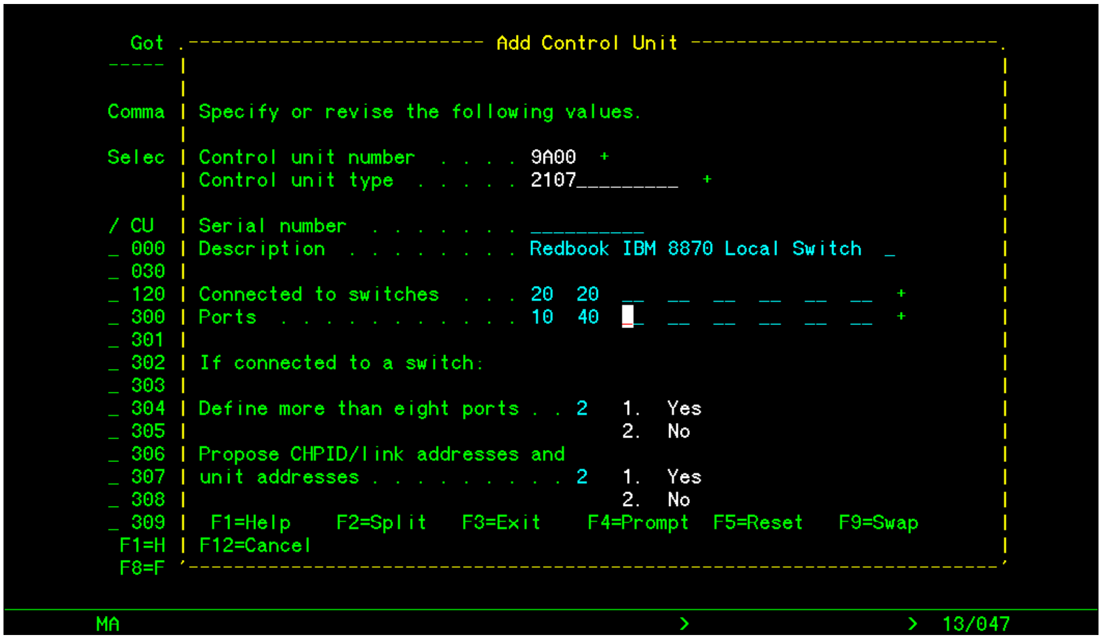

```
   Got .------------------------- Add Control Unit --------------------------.
  ----- |                                                                    |
        |                                                                    |
  Comma | Specify or revise the following values.                            |
        |                                                                    |
  Selec | Control unit number  . . . . 9A00  +                               |
        | Control unit type  . . . . . 2107_________  +                      |
        |                                                                    |
  / CU  | Serial number  . . . . . . . __________                            |
  _ 000 | Description  . . . . . . . . Redbook IBM 8870 Local Switch  _       |
  _ 030 |                                                                    |
  _ 120 | Connected to switches  . . . 20  20  __  __  __  __  __  __  +     |
  _ 300 | Ports  . . . . . . . . . . . 10  40  █_  __  __  __  __  __  +     |
  _ 301 |                                                                    |
  _ 302 | If connected to a switch:                                          |
  _ 303 |                                                                    |
  _ 304 | Define more than eight ports . . 2   1.  Yes                       |
  _ 305 |                                      2.  No                        |
  _ 306 | Propose CHPID/link addresses and                                   |
  _ 307 | unit addresses . . . . . . . . . 2   1.  Yes                       |
  _ 308 |                                      2.  No                        |
  _ 309 |  F1=Help    F2=Split    F3=Exit    F4=Prompt  F5=Reset    F9=Swap  |
   F1=H | F12=Cancel                                                         |
   F8=F '--------------------------------------------------------------------'

  MA                                              >                   >  13/047
```

Figure 7-19 Add Control Unit

Now, we see the Select Processor / CU screen. Here, we tell the host configuration which host CHPIDs will be performing I/O to the disk ports that we defined on the previous screen. For the IBM Z I/O subsystem (IOS), there may be eight parallel paths to each CU, and you can specify them in this screen. These CHPID through switch-to-disk connections are configured as `<CC.SSPP>`, where CC is the 1-byte CHPID number, SS is the CU switch ID, and PP is the FICON Port Address for the switch interface where the disk port is attached.

For our example network, we input 18.2010 and 20.2040, which means that this mainframe can initiate I/O from CHPID 18 and send it to the switch that it is connected to (in this case, switch ID 0x20) and that the frames that are associated with this I/O are directed to the device that is connected to FICON Port Address 0x10 on switch ID 0x20.

Figure 7-20 on page 269 shows the completed dialog. Press Enter.

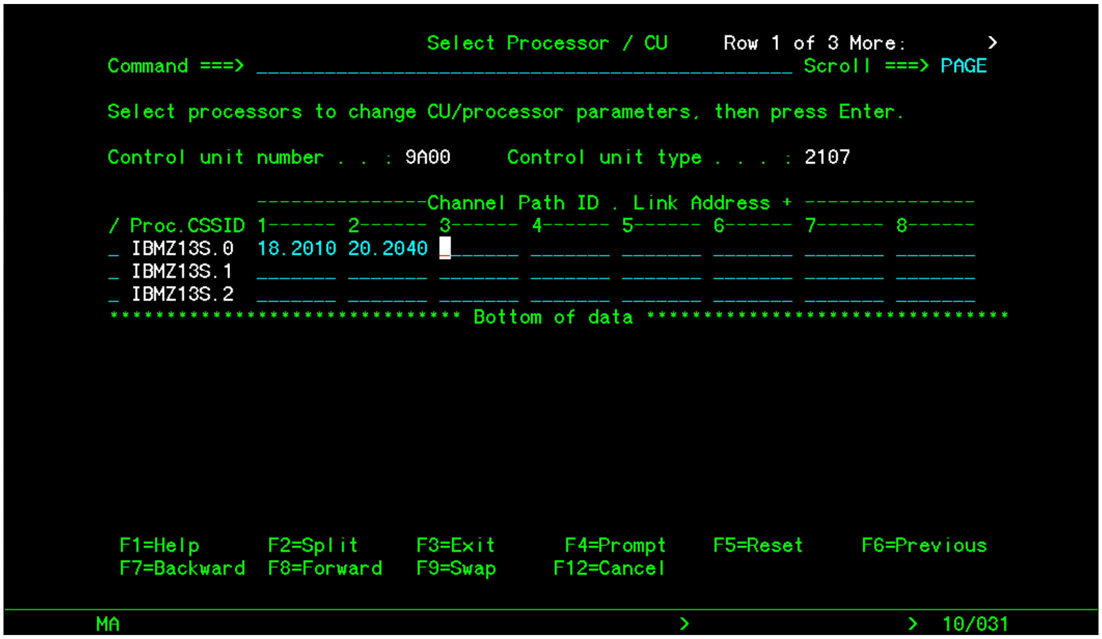

```
                     Select Processor / CU     Row 1 of 3 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : 9A00     Control unit type . . . : 2107

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ IBMZ13S.0  18.2010 20.2040 _______ _______ _______ _______ _______ _______
_ IBMZ13S.1  _______ _______ _______ _______ _______ _______ _______ _______
_ IBMZ13S.2  _______ _______ _______ _______ _______ _______ _______ _______
****************************** Bottom of data *******************************



 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F6=Previous
 F7=Backward  F8=Forward   F9=Swap     F12=Cancel

        MA                                           >                   >  10/031
```

Figure 7-20 Input Channel Path ID and link address

Now, we must define the CU address (CUADDR), base device address, and address range for this CU. These values are provided by your disk administrator. The values for our example network are shown in Figure 7-21. We press Enter twice.

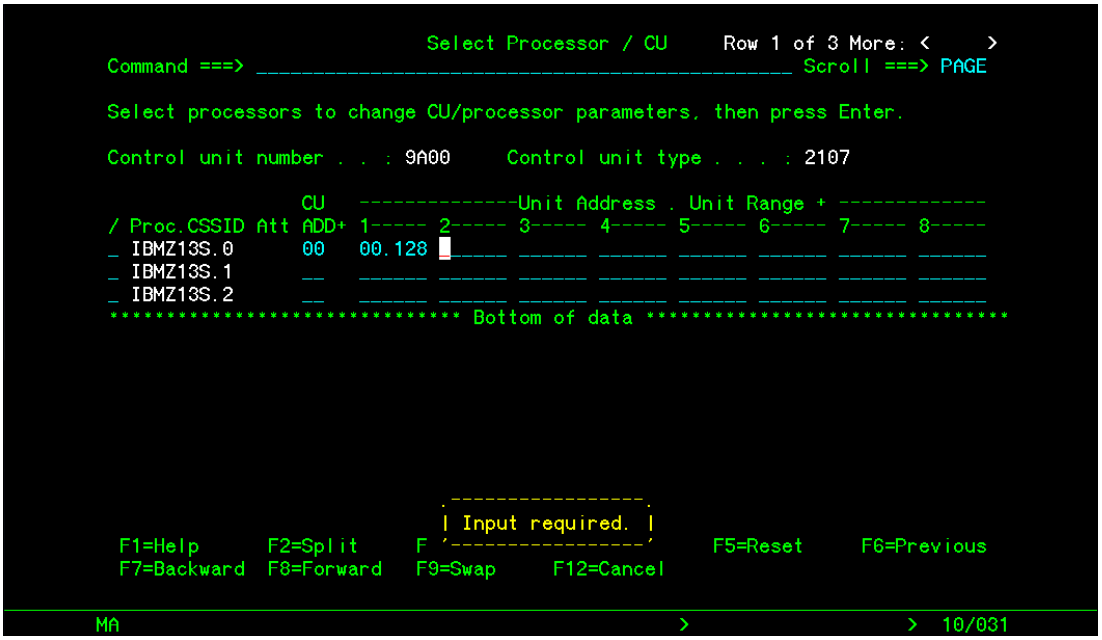

```
                     Select Processor / CU     Row 1 of 3 More: <     >
Command ===> _______________________________________________ Scroll ===> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : 9A00     Control unit type . . . : 2107

                 CU   -------------Unit Address . Unit Range + ------------
/ Proc.CSSID Att ADD+ 1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
_ IBMZ13S.0      00   00.128 ______ ______ ______ ______ ______ ______ ______
_ IBMZ13S.1      __   ______ ______ ______ ______ ______ ______ ______ ______
_ IBMZ13S.2      __   ______ ______ ______ ______ ______ ______ ______ ______
****************************** Bottom of data *******************************



                                   .-----------------.
                                   | Input required. |
 F1=Help      F2=Split     F       '-----------------'   F5=Reset     F6=Previous
 F7=Backward  F8=Forward   F9=Swap     F12=Cancel

        MA                                           >                   >  10/031
```

Figure 7-21 Defining the unit address

On the resulting screen, after we scroll up a line, we can see the CU that we defined, as shown in Figure 7-22.

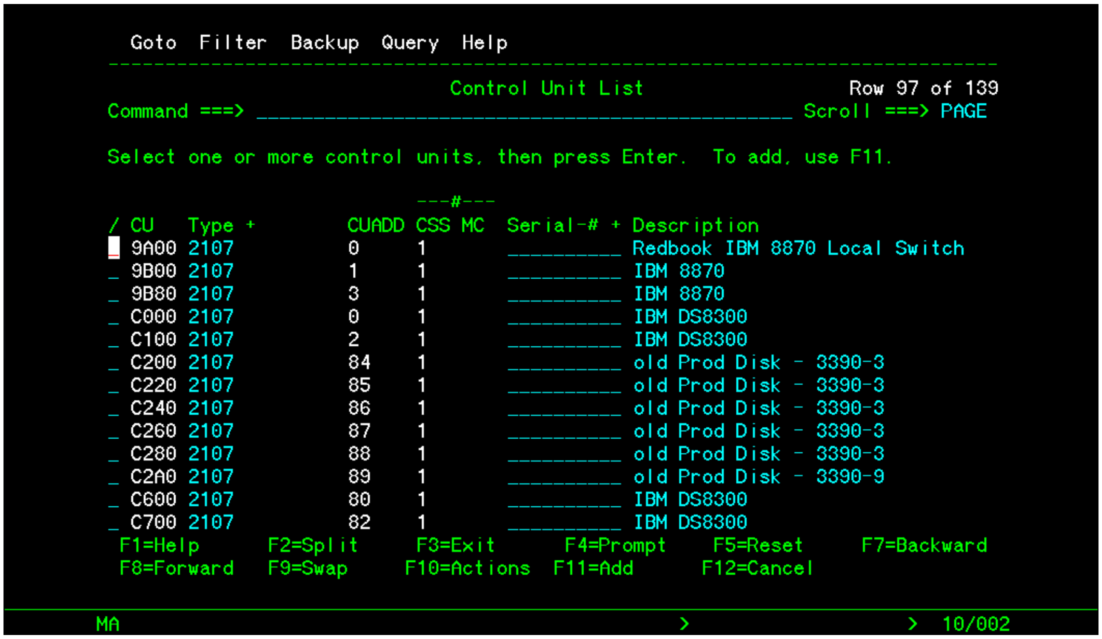

```
  Goto  Filter  Backup  Query  Help
 ------------------------------------------------------------------------------
                             Control Unit List                 Row 97 of 139
 Command ===> ____________________________________________ Scroll ===> PAGE

 Select one or more control units, then press Enter.  To add, use F11.

                             ---#---
 / CU   Type +         CUADD CSS MC  Serial-# + Description
 _ 9A00 2107           0     1       __________ Redbook IBM 8870 Local Switch
 _ 9B00 2107           1     1       __________ IBM 8870
 _ 9B80 2107           3     1       __________ IBM 8870
 _ C000 2107           0     1       __________ IBM DS8300
 _ C100 2107           2     1       __________ IBM DS8300
 _ C200 2107           84    1       __________ old Prod Disk - 3390-3
 _ C220 2107           85    1       __________ old Prod Disk - 3390-3
 _ C240 2107           86    1       __________ old Prod Disk - 3390-3
 _ C260 2107           87    1       __________ old Prod Disk - 3390-3
 _ C280 2107           88    1       __________ old Prod Disk - 3390-3
 _ C2A0 2107           89    1       __________ old Prod Disk - 3390-9
 _ C600 2107           80    1       __________ IBM DS8300
 _ C700 2107           82    1       __________ IBM DS8300
  F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
  F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                                 >                   >  10/002
```

Figure 7-22   Control Unit List

Now that the disk CU is defined, the associated disk devices also must be defined. That topic is beyond the scope of this book, but it is a well-known process, and there are not associations that must be made between the device definition and the definitions for the FICON SAN.

7.2.4 Defining a cascaded control unit that is connected to the FICON switches

A CU is cascaded when there are two or more switches between the host channels and the CU ports. In the example network, there are two cascaded CUs: One is a non-IBM disk array, and the other is for the IBM Virtual Tape Server. For this example, we define the non-IBM disk array so that the differences between the local and cascaded disks are easily discernible.

We use HCD to define the connection between CHPID 28, which is attached to FICON Port Address 0x22 on the IBM Storage Networking SAN384C-6 switch, and the disk port, which is attached to FICON Port Address 0x0A on the IBM Storage Networking SAN192C-6 switch. The fact that this connection crosses one or more Inter-Switch Links (ISLs) will not be defined as part of this process, but it is implied.

We start at the Design, Modify, or View Configuration Data screen and select option 4 for CUs. The resulting screen look like Figure 7-18 on page 267. We press F11 to add a CU.

Like the CU that was configured earlier, we input the CU number as 8A00 and the device type as 2107. We input a description, and for now leave the serial number blank.

Now, we define the connectivity from the host to this cascaded disk CU. On this first screen, we define the two disk ports that are both connected to FICON VSAN 40 on the IBM Storage Networking SAN192C-6 switch, which is defined with domain 0x21. For each of these places, we enter the switch number, we must input the associated FICON Port Address that is being connected to by the disk array.

The Add Control Unit dialog for 8A00 is shown in Figure 7-23. We press Enter.

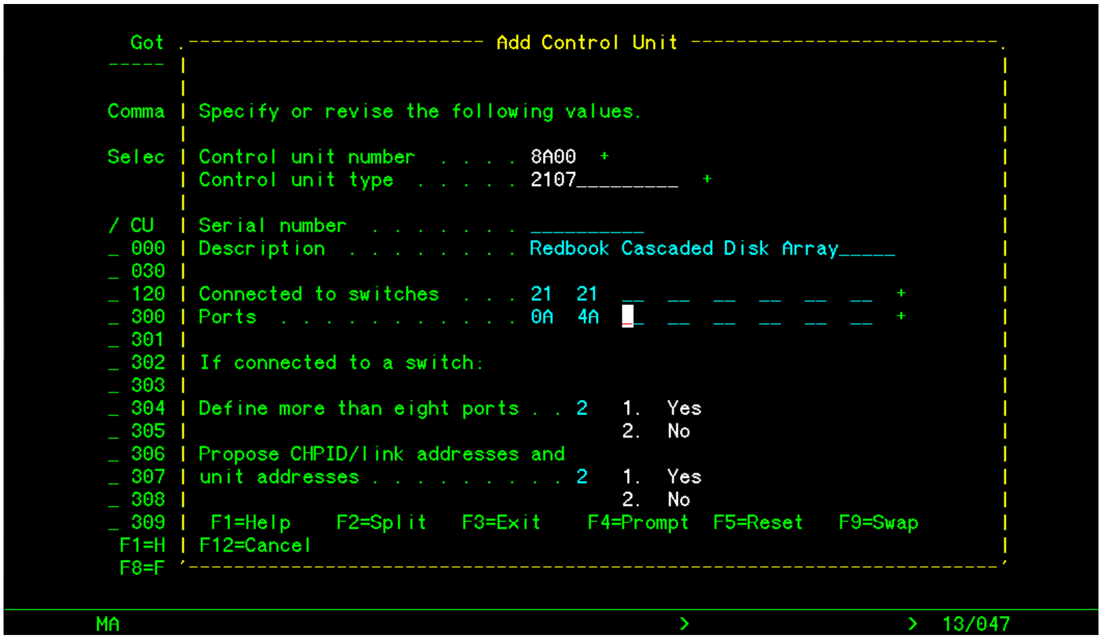

```
  Got .-------------------------- Add Control Unit --------------------------.
 ----- |                                                                    |
       |                                                                    |
 Comma | Specify or revise the following values.                            |
       |                                                                    |
 Selec | Control unit number . . . . 8A00  +                                |
       | Control unit type . . . . . 2107_________   +                      |
       |                                                                    |
 / CU  | Serial number . . . . . . . __________                             |
 _ 000 | Description . . . . . . . . Redbook Cascaded Disk Array_____       |
 _ 030 |                                                                    |
 _ 120 | Connected to switches . . . 21  21  __  __  __  __  __  __  +      |
 _ 300 | Ports . . . . . . . . . . . 0A  4A  __  __  __  __  __  __  +      |
 _ 301 |                                                                    |
 _ 302 | If connected to a switch:                                          |
 _ 303 |                                                                    |
 _ 304 | Define more than eight ports . . 2   1.  Yes                       |
 _ 305 |                                      2.  No                        |
 _ 306 | Propose CHPID/link addresses and                                   |
 _ 307 | unit addresses . . . . . . . . . 2   1.  Yes                       |
 _ 308 |                                      2.  No                        |
 _ 309 |  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap    |
  F1=H | F12=Cancel                                                         |
  F8=F '--------------------------------------------------------------------'

       MA                                          >                   >  13/047
```

Figure 7-23 Add Control Unit

Now, we see the Select Processor/CU screen, which is shown in Figure 7-24.

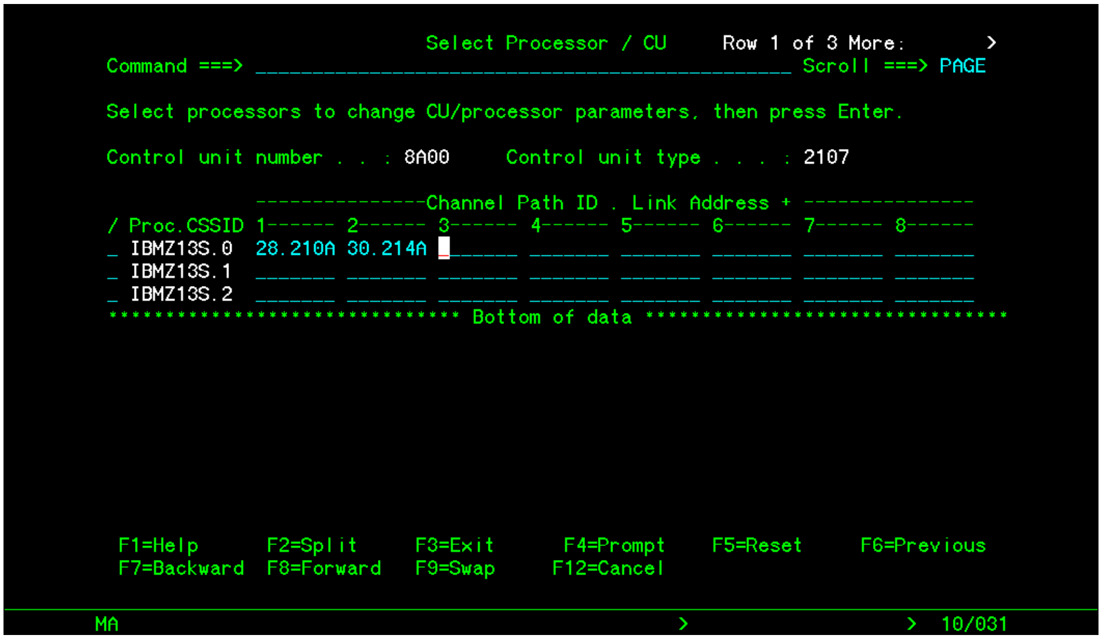

```
                           Select Processor / CU      Row 1 of 3 More:       >
 Command ===> _____________________________________________ Scroll ===> PAGE

 Select processors to change CU/processor parameters, then press Enter.

 Control unit number . . : 8A00      Control unit type . . . : 2107

              --------------Channel Path ID . Link Address + ---------------
 / Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
 _ IBMZ13S.0  28.210A 30.214A _______ _______ _______ _______ _______ _______
 _ IBMZ13S.1  _______ _______ _______ _______ _______ _______ _______ _______
 _ IBMZ13S.2  _______ _______ _______ _______ _______ _______ _______ _______
 ****************************** Bottom of data *******************************








  F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F6=Previous
  F7=Backward  F8=Forward   F9=Swap     F12=Cancel

       MA                                          >                   >  10/031
```

Figure 7-24 Select Processor / CU

We must tell the host configuration that host CHPIDs 28 and 30 will be connecting to these two disk ports. The way that the host configuration knows that this CU is a cascaded CU is that CHPIDs 28 and 30 are defined as being attached to switch ID 0x20, and the new disk ports that we are defining are attached to switch ID 0x21.

Again, we define the CHPID through switches to disk connections by using the format `<CC.SSPP>`, where CC is the CHPID number, SS is the switch ID for the switch where the CU ports are attached, and PP is the FICON Port Address for the interface where the disk ports are connected to the switch. For our example network, we input 28.210A and 30.214A. We do not reference where the CHPIDs are entering the FICON fabric because this information can be learned only from the CHPID definitions.

This mainframe can initiate I/O from CHPID 28 and send it to the switch that it is connected to (in this case, switch ID 0x20) and that the frames that are associated with this I/O will be directed to the device that is connected to FICON Port Address 0x0A on switch ID 0x21. The routing of the frames after they enter the entry switch (where the CHPID is attached) until they exit from the destination switch is handled by the switches, and the mainframe has no visibility of this route.

Press Enter to continue.

Again, we see that we must define the CUADDR, the base device address, and address range for this CU. The values for our example network are shown in Figure 7-25. Press Enter twice to continue.

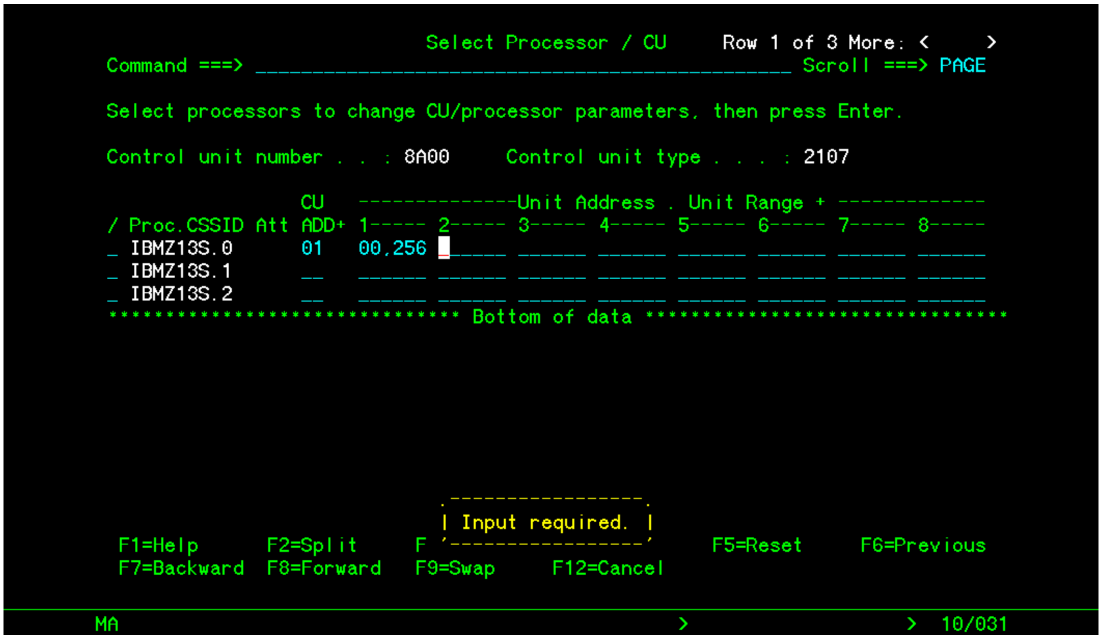

```
                              Select Processor / CU      Row 1 of 3 More: <     >
 Command ===> ______________________________________________ Scroll ===> PAGE

 Select processors to change CU/processor parameters, then press Enter.

 Control unit number . . : 8A00      Control unit type . . . : 2107

                 CU     --------------Unit Address . Unit Range + -------------
 / Proc.CSSID Att ADD+  1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
 _ IBMZ13S.0      01    00,256 ______ ______ ______ ______ ______ ______ ______
 _ IBMZ13S.1      __    ______ ______ ______ ______ ______ ______ ______ ______
 _ IBMZ13S.2      __    ______ ______ ______ ______ ______ ______ ______ ______
 ***************************** Bottom of data ******************************





                                      .------------------.
                                      | Input required.  |
  F1=Help       F2=Split      F       '------------------'   F5=Reset      F6=Previous
  F7=Backward   F8=Forward    F9=Swap      F12=Cancel

        MA                                            >                   >  10/031
```

Figure 7-25 Defining the unit address

The CU is defined. After scrolling through the resulting screen, we can see the CU information. As before with the locally switched CU, the device configuration must be created and linked to the CU.

The CUs for the IBM TS7760 must be created in a like way. Because the device type for the IBM TS7760 is different, there will be minor differences in the definition, but the concepts of how the CHPIDs, switches, and CU FICON Port Address appear in the screens are the same.

> **Note:** There is no difference in defining cascaded CUs that are connected over FCIP links versus the ones that are connected over more conventional FC ISLs because the mainframe configuration does not have any visibility into the switch-to-switch connections.

7.2.5 Complete CUP control unit and device definitions

Next, we complete the definition of the CUP CU and device definitions that were started when each of the switches was defined earlier. When we defined the example switch, we input the switch CU number as 0xEF20 and the switch device number also as 0xEF20.

> **Note:** In most cases, the CU number would not be the same as the device numbers that are defined under it. Because there is a 1:1 relationship between the CUP CU number and the CUP device number, this configuration is acceptable, and it is the convention that is used within the lab that we are using.

We define which FICON channels will be used to communicate with the CUP device on each of the switches. We also define which LPARs will have access to these devices. We start at the Design, Modify, or View Configuration Data screen and select option 4 for CUs. Then, we scroll down until CU EF20 is seen, as shown in Figure 7-26.

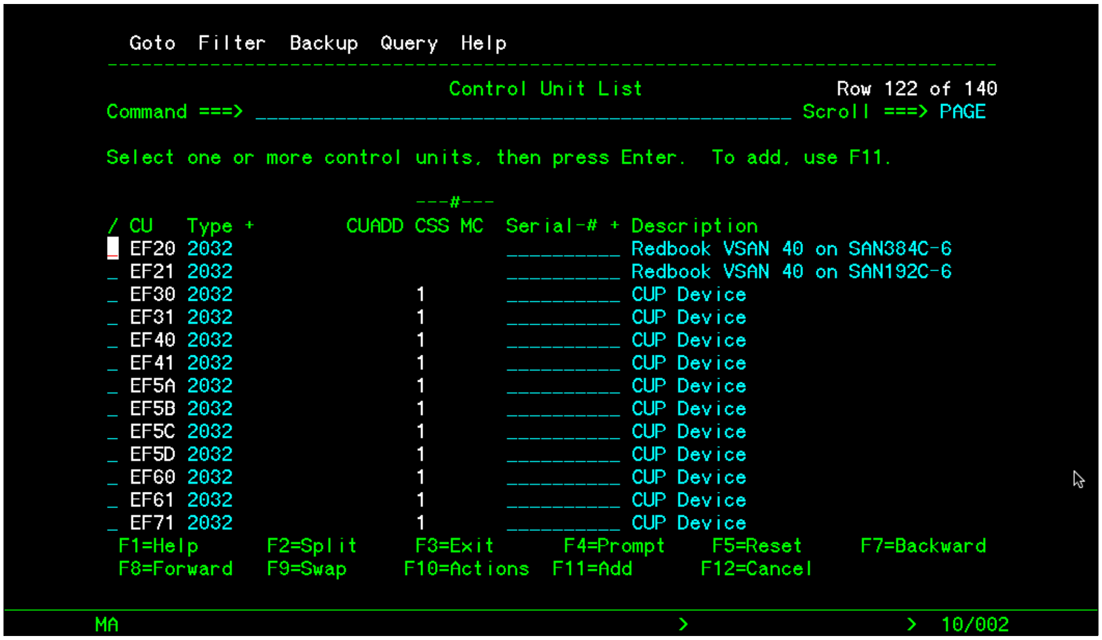

```
  Goto  Filter  Backup  Query  Help
--------------------------------------------------------------------------
                            Control Unit List                Row 122 of 140
Command ===> ______________________________________________ Scroll ===> PAGE

Select one or more control units, then press Enter.  To add, use F11.

                          ---#---
/ CU   Type +        CUADD CSS MC  Serial-# + Description
_ EF20 2032                        __________ Redbook VSAN 40 on SAN384C-6
_ EF21 2032                        __________ Redbook VSAN 40 on SAN192C-6
_ EF30 2032                 1      __________ CUP Device
_ EF31 2032                 1      __________ CUP Device
_ EF40 2032                 1      __________ CUP Device
_ EF41 2032                 1      __________ CUP Device
_ EF5A 2032                 1      __________ CUP Device
_ EF5B 2032                 1      __________ CUP Device
_ EF5C 2032                 1      __________ CUP Device
_ EF5D 2032                 1      __________ CUP Device
_ EF60 2032                 1      __________ CUP Device
_ EF61 2032                 1      __________ CUP Device
_ EF71 2032                 1      __________ CUP Device
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

MA                                             >                   >  10/002
```

Figure 7-26 Control Unit List

We put a / next to EF20 and press Enter. In the resulting dialog, we type 2, as shown in Figure 7-27.

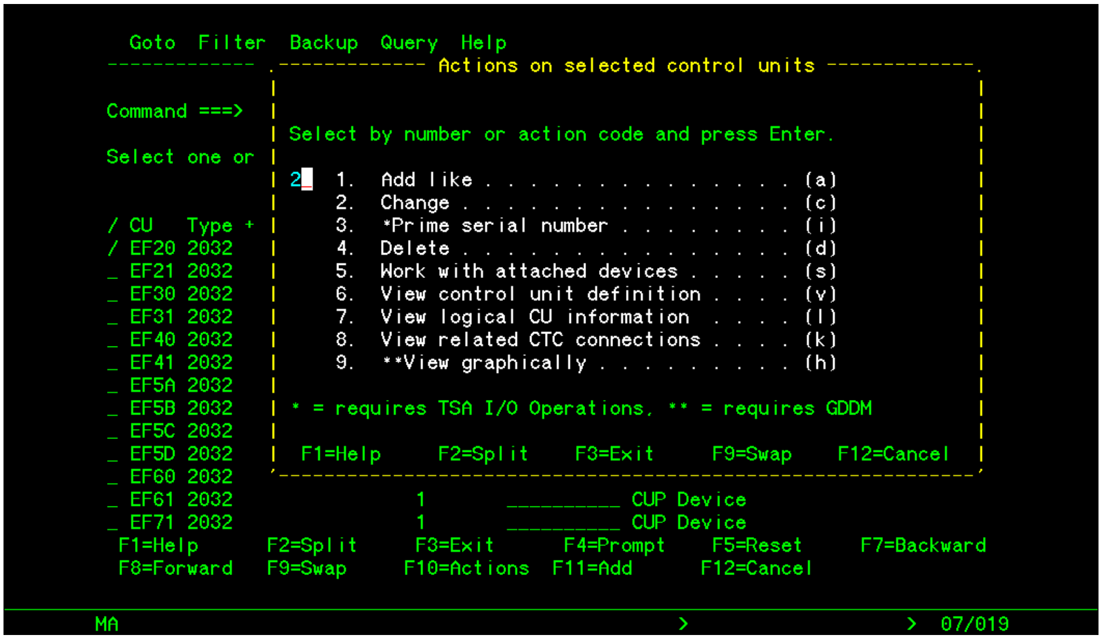

```
  Goto  Filter  Backup  Query  Help
------------- .------------- Actions on selected control units -------------.
              |                                                              |
Command ===>  |                                                              |
              | Select by number or action code and press Enter.             |
Select one or |                                                              |
              | 2_  1.  Add like . . . . . . . . . . . . . . . . (a)         |
              |     2.  Change . . . . . . . . . . . . . . . . . (c)         |
/ CU   Type + |     3.  *Prime serial number . . . . . . . . . (i)         |
/ EF20 2032   |     4.  Delete . . . . . . . . . . . . . . . . . (d)         |
_ EF21 2032   |     5.  Work with attached devices . . . . . (s)         |
_ EF30 2032   |     6.  View control unit definition . . . . (v)         |
_ EF31 2032   |     7.  View logical CU information . . . . (l)         |
_ EF40 2032   |     8.  View related CTC connections . . . . (k)         |
_ EF41 2032   |     9.  **View graphically . . . . . . . . . (h)         |
_ EF5A 2032   |                                                              |
_ EF5B 2032   | * = requires TSA I/O Operations, ** = requires GDDM          |
_ EF5C 2032   |                                                              |
_ EF5D 2032   |  F1=Help     F2=Split     F3=Exit     F9=Swap     F12=Cancel |
_ EF60 2032   '--------------------------------------------------------------'
_ EF61 2032                 1        __________ CUP Device
_ EF71 2032                 1        __________ CUP Device
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

       MA                                              >                 >  07/019
```

Figure 7-27   Action selection dialog

We press Enter and see the information for the CUP CU, as shown in Figure 7-28.

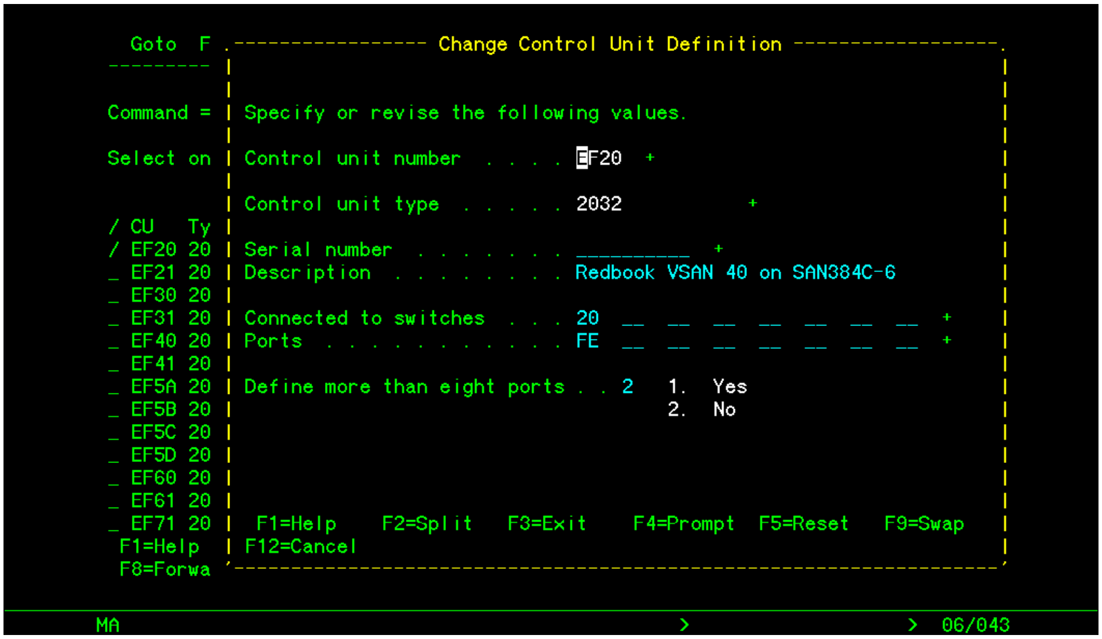

```
  Goto  F .---------------- Change Control Unit Definition -----------------.
 ---------|                                                                  |
          |                                                                  |
Command = | Specify or revise the following values.                         |
          |                                                                  |
Select on | Control unit number . . . . EF20  +                              |
          |                                                                  |
          | Control unit type . . . . . 2032           +                     |
/ CU   Ty |                                                                  |
/ EF20 20 | Serial number . . . . . . . __________  +                        |
_ EF21 20 | Description . . . . . . . . Redbook VSAN 40 on SAN384C-6        |
_ EF30 20 |                                                                  |
_ EF31 20 | Connected to switches . . . 20 __ __ __ __ __ __ __  +          |
_ EF40 20 | Ports . . . . . . . . . . . FE __ __ __ __ __ __ __  +          |
_ EF41 20 |                                                                  |
_ EF5A 20 | Define more than eight ports . . 2   1.  Yes                     |
_ EF5B 20 |                                      2.  No                      |
_ EF5C 20 |                                                                  |
_ EF5D 20 |                                                                  |
_ EF60 20 |                                                                  |
_ EF61 20 |                                                                  |
_ EF71 20 |  F1=Help    F2=Split    F3=Exit    F4=Prompt  F5=Reset   F9=Swap   |
 F1=Help  | F12=Cancel                                                       |
 F8=Forwa '------------------------------------------------------------------'

       MA                                              >                 >  06/043
```

Figure 7-28   Change Control Unit Definition

All the information that was populated as a result of the earlier definition is correct. If you want, you can add the serial number for the FICON VSAN. In our example, we leave this serial number field blank for now and press Enter again, as shown in Figure 7-29 on page 275.

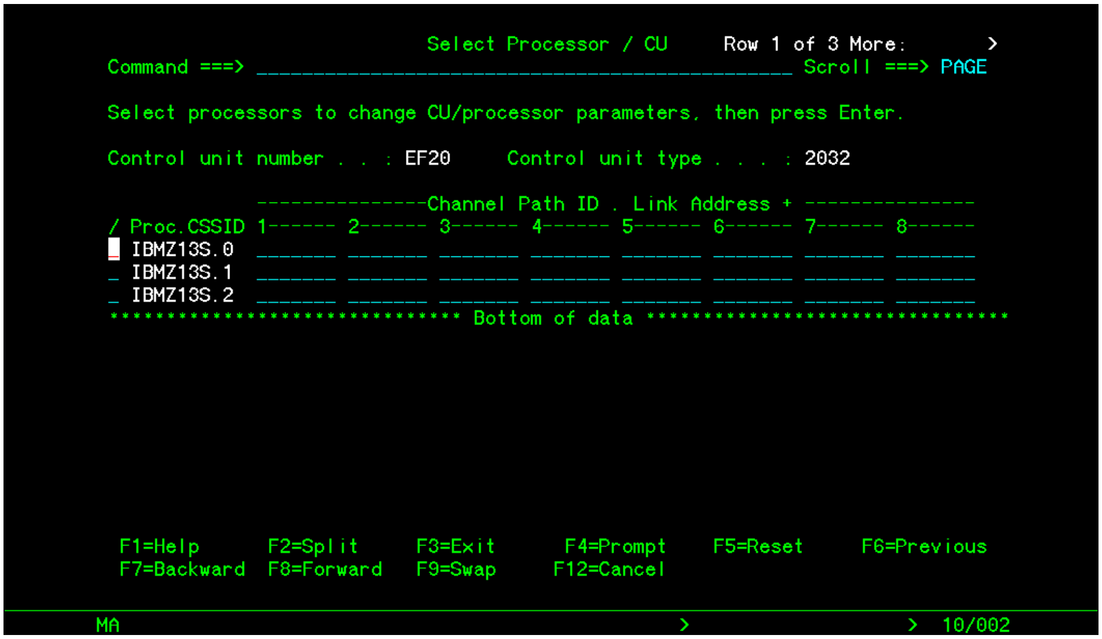

```
                          Select Processor / CU     Row 1 of 3 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : EF20     Control unit type . . . : 2032

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ IBMZ13S.0  _______ _______ _______ _______ _______ _______ _______ _______
_ IBMZ13S.1  _______ _______ _______ _______ _______ _______ _______ _______
_ IBMZ13S.2  _______ _______ _______ _______ _______ _______ _______ _______
****************************** Bottom of data *******************************

 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F6=Previous
 F7=Backward  F8=Forward   F9=Swap     F12=Cancel

MA                                                  >                   >  10/002
```

Figure 7-29 Select Processor / CU

We enter the CHPID and link address paths for this CU. In our example, we define the CUP devices to be accessible from both CHPID 18 and CHPID 20. The same pattern of <CC.SSPP> is used, where CC is the 1-byte CHPID number, SS is the CU switch ID, and PP is the FICON Port Address for the switch internal port. For the CUP CU, we always use the reserved value of 0xFE for the port (PP), and SS is the switch ID that we are defining (0x20). Because both CHPIDs have the same destination, we enter 18.20FE and 20.20FE, as shown in Figure 7-30, and then press Enter.

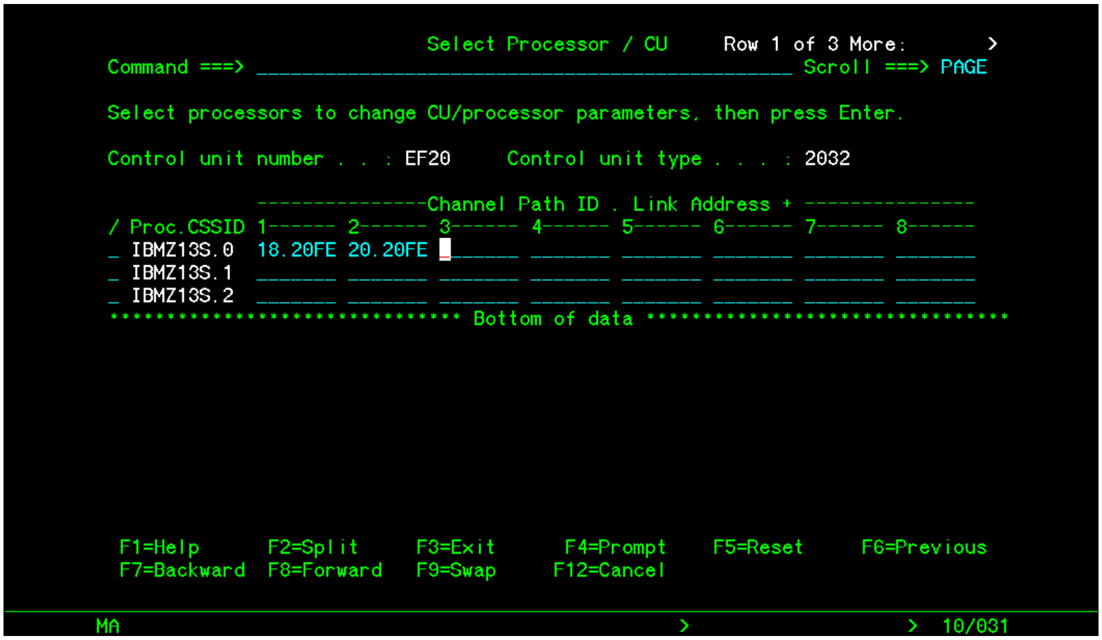

```
                          Select Processor / CU     Row 1 of 3 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : EF20     Control unit type . . . : 2032

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ IBMZ13S.0  18.20FE 20.20FE _______ _______ _______ _______ _______ _______
_ IBMZ13S.1  _______ _______ _______ _______ _______ _______ _______ _______
_ IBMZ13S.2  _______ _______ _______ _______ _______ _______ _______ _______
****************************** Bottom of data *******************************

 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F6=Previous
 F7=Backward  F8=Forward   F9=Swap     F12=Cancel

MA                                                  >                   >  10/031
```

Figure 7-30 Select Processor / CU

We enter the unit address for the CUP device as 00, and no range because there is only a single CUP device. CUP devices must always use address 00 by definition. When the screen looks like what is shown in Figure 7-31, press Enter.

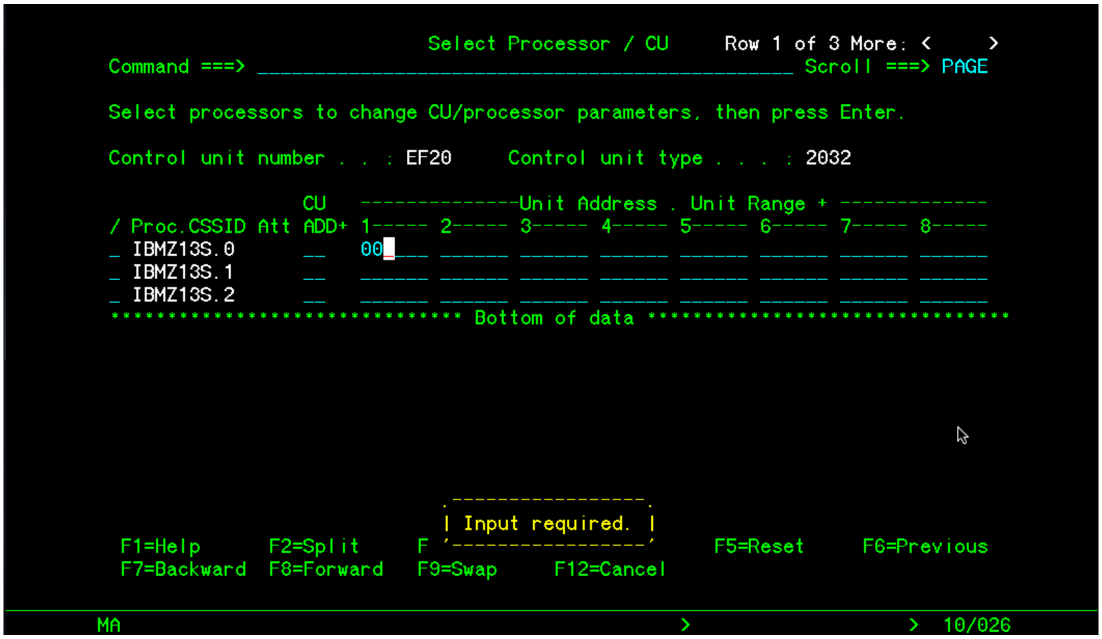

```
                            Select Processor / CU     Row 1 of 3 More: <     >
 Command ===> ______________________________________________ Scroll ===> PAGE

 Select processors to change CU/processor parameters, then press Enter.

 Control unit number . . : EF20     Control unit type . . . : 2032

                  CU   --------------Unit Address . Unit Range + ------------
 / Proc.CSSID Att ADD+ 1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
 _ IBMZ13S.0      __   00____ ______ ______ ______ ______ ______ ______ ______
 _ IBMZ13S.1      __   ______ ______ ______ ______ ______ ______ ______ ______
 _ IBMZ13S.2      __   ______ ______ ______ ______ ______ ______ ______ ______
 ****************************** Bottom of data *******************************









                             .-----------------.
                             | Input required. |
  F1=Help      F2=Split     F '-----------------'     F5=Reset     F6=Previous
  F7=Backward  F8=Forward   F9=Swap      F12=Cancel

     MA                                         >                    >  10/026
```

Figure 7-31   Select Processor / CU

In Figure 7-32, we confirm the device parameter from the earlier screens, so we press Enter.

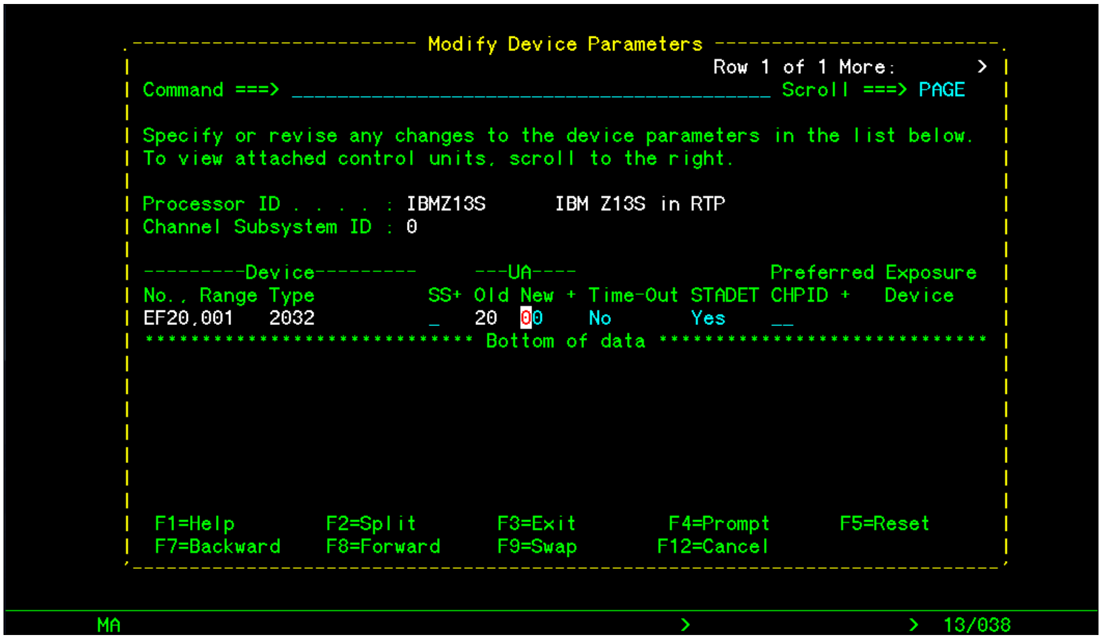

```
  .------------------------- Modify Device Parameters -------------------------.
  |                                                    Row 1 of 1 More:      > |
  | Command ===> __________________________________________ Scroll ===> PAGE   |
  |                                                                            |
  | Specify or revise any changes to the device parameters in the list below.  |
  | To view attached control units, scroll to the right.                       |
  |                                                                            |
  | Processor ID . . . . : IBMZ13S      IBM Z13S in RTP                        |
  | Channel Subsystem ID : 0                                                   |
  |                                                                            |
  | ---------Device---------       ---UA----                Preferred Exposure |
  | No., Range Type              SS+ Old New + Time-Out STADET CHPID +   Device|
  | EF20,001   2032              _   20  00    No       Yes    __              |
  | ***************************** Bottom of data ***************************** |
  |                                                                            |
  |                                                                            |
  |                                                                            |
  |                                                                            |
  |                                                                            |
  |                                                                            |
  |  F1=Help       F2=Split       F3=Exit        F4=Prompt      F5=Reset       |
  |  F7=Backward   F8=Forward     F9=Swap       F12=Cancel                     |
  '----------------------------------------------------------------------------'

     MA                                         >                    >  13/038
```

Figure 7-32   Modify Device Parameters

In Figure 7-33 on page 277, we verify our settings and press Enter.

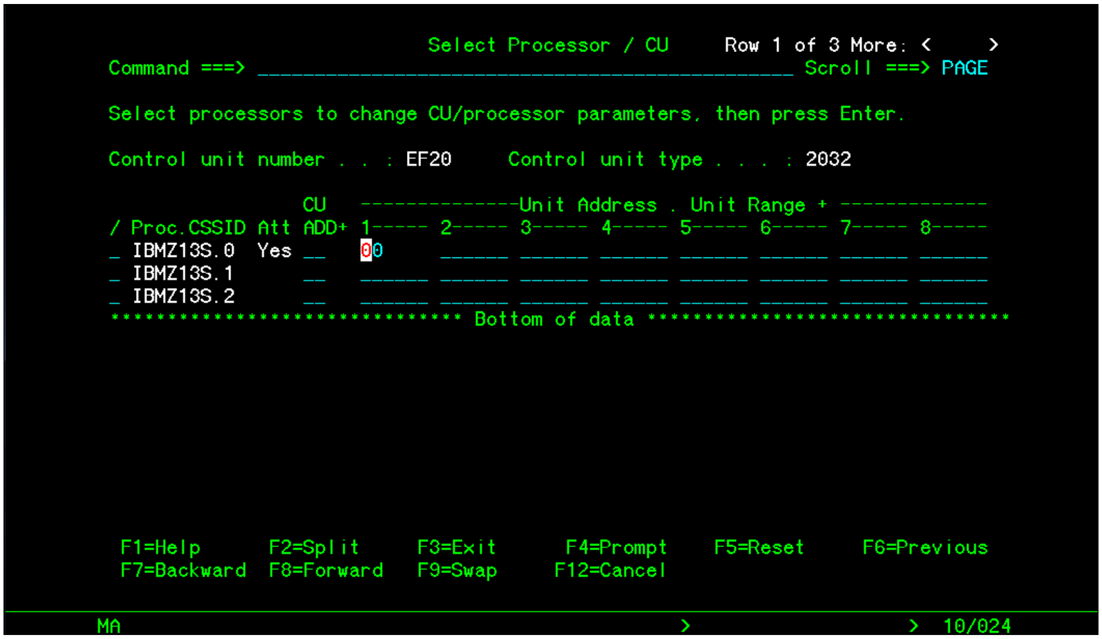

```
                            Select Processor / CU     Row 1 of 3 More: <     >
Command ===> ______________________________________________ Scroll ===> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : EF20     Control unit type . . . : 2032

                 CU   -------------Unit Address . Unit Range + ------------
/ Proc.CSSID Att ADD+ 1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
_ IBMZ13S.0  Yes __   00     ______ ______ ______ ______ ______ ______ ______
_ IBMZ13S.1      __   ______ ______ ______ ______ ______ ______ ______ ______
_ IBMZ13S.2      __   ______ ______ ______ ______ ______ ______ ______ ______
****************************** Bottom of data *******************************









 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F6=Previous
 F7=Backward  F8=Forward   F9=Swap     F12=Cancel

        MA                                                  >                   >  10/024
```

Figure 7-33   Select Processor / CU

We are back on the Control Unit List screen, but now the CUP control CU shows as being attached to a CSS. We must work on the device. We return to the Design, Modify, or View Configuration Data screen and select option 5 for I/O Devices. When we scroll down to device EF20, we see it as shown in Figure 7-34.

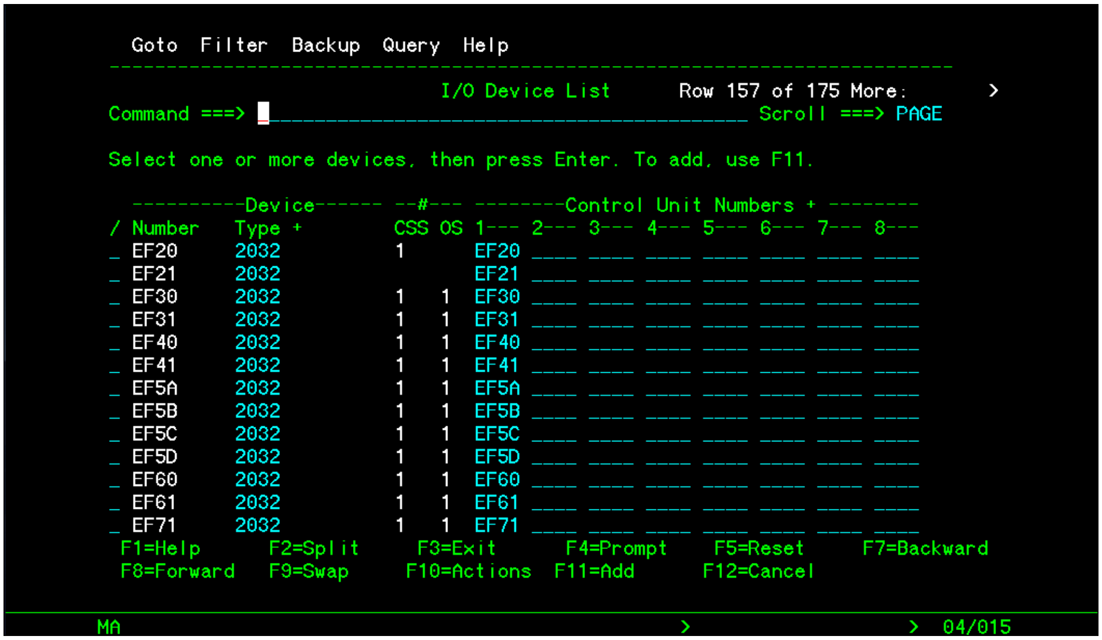

```
  Goto  Filter  Backup  Query  Help
----------------------------------------------------------------------------
                          I/O Device List      Row 157 of 175 More:      >
Command ===> _________________________________________ Scroll ===> PAGE

Select one or more devices, then press Enter. To add, use F11.

  ----------Device------ --#--- --------Control Unit Numbers + --------
/ Number   Type +        CSS OS 1--- 2--- 3--- 4--- 5--- 6--- 7--- 8---
_ EF20     2032          1      EF20 ____ ____ ____ ____ ____ ____ ____
_ EF21     2032                 EF21 ____ ____ ____ ____ ____ ____ ____
_ EF30     2032          1   1  EF30 ____ ____ ____ ____ ____ ____ ____
_ EF31     2032          1   1  EF31 ____ ____ ____ ____ ____ ____ ____
_ EF40     2032          1   1  EF40 ____ ____ ____ ____ ____ ____ ____
_ EF41     2032          1   1  EF41 ____ ____ ____ ____ ____ ____ ____
_ EF5A     2032          1   1  EF5A ____ ____ ____ ____ ____ ____ ____
_ EF5B     2032          1   1  EF5B ____ ____ ____ ____ ____ ____ ____
_ EF5C     2032          1   1  EF5C ____ ____ ____ ____ ____ ____ ____
_ EF5D     2032          1   1  EF5D ____ ____ ____ ____ ____ ____ ____
_ EF60     2032          1   1  EF60 ____ ____ ____ ____ ____ ____ ____
_ EF61     2032          1   1  EF61 ____ ____ ____ ____ ____ ____ ____
_ EF71     2032          1   1  EF71 ____ ____ ____ ____ ____ ____ ____
 F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
 F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

        MA                                                  >                   >  04/015
```

Figure 7-34   I/O Device List

To complete our work on the CUP device, we connect it to an operating system (OS) configuration that will use it. Enter a / to the left of EF20 and press Enter. Select option 2 Change and press Enter. The screen that is shown in Figure 7-35 opens.

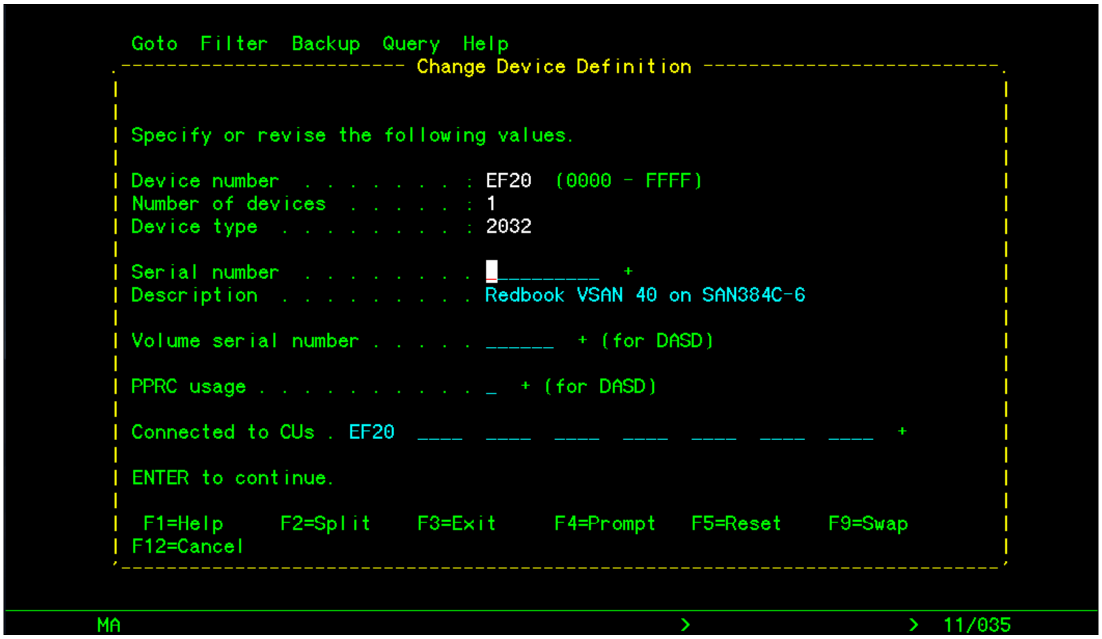

```
  Goto  Filter  Backup  Query  Help
.------------------------ Change Device Definition -------------------------.
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Device number . . . . . . . : EF20  (0000 - FFFF)                         |
| Number of devices . . . . . : 1                                           |
| Device type . . . . . . . . : 2032                                        |
|                                                                           |
| Serial number . . . . . . . . __________  +                               |
| Description . . . . . . . . . Redbook VSAN 40 on SAN384C-6                |
|                                                                           |
| Volume serial number . . . . . ______  + (for DASD)                       |
|                                                                           |
| PPRC usage . . . . . . . . . . _  + (for DASD)                            |
|                                                                           |
| Connected to CUs . EF20 ____  ____  ____  ____  ____  ____  ____  +       |
|                                                                           |
| ENTER to continue.                                                        |
|                                                                           |
|  F1=Help     F2=Split     F3=Exit     F4=Prompt    F5=Reset    F9=Swap    |
| F12=Cancel                                                                |
'---------------------------------------------------------------------------'

       MA                                            >                   >  11/035
```

Figure 7-35   Change Device Definition

On this screen, we verify that the device number is connected to the correct CU, which is EF20 for both values in this case. Press Enter to see the screen that is shown in Figure 7-36.

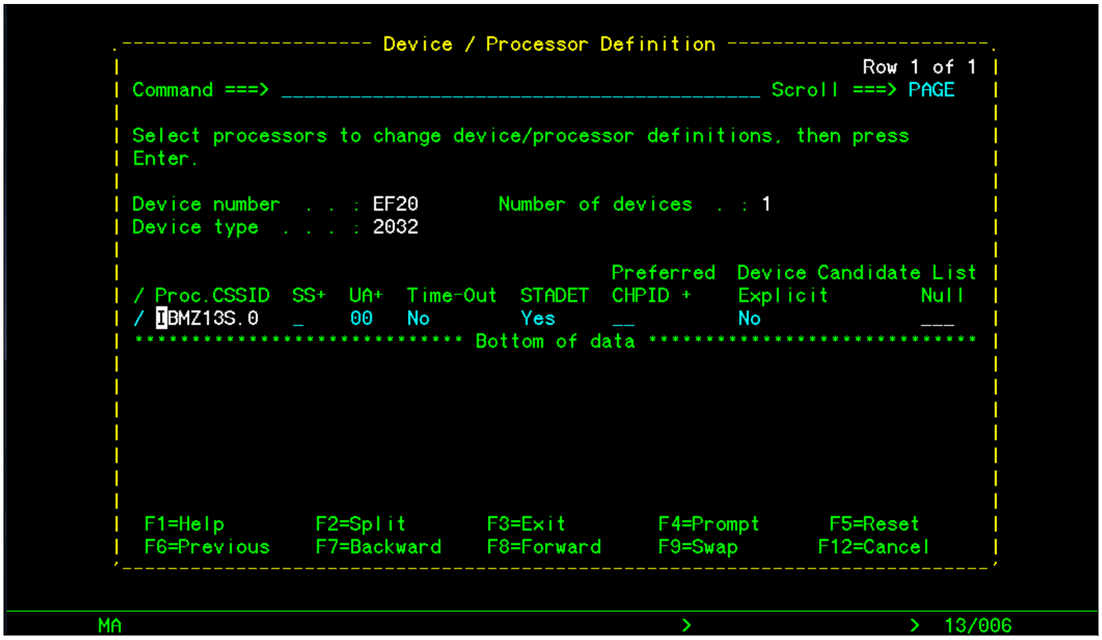

```
.--------------------- Device / Processor Definition -----------------------.
|                                                                Row 1 of 1 |
| Command ===> ________________________________________ Scroll ===> PAGE    |
|                                                                           |
| Select processors to change device/processor definitions, then press      |
| Enter.                                                                    |
|                                                                           |
| Device number . . : EF20      Number of devices . : 1                     |
| Device type . . . : 2032                                                  |
|                                                                           |
|                                         Preferred  Device Candidate List  |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +   Explicit        Null  |
| / IBMZ13S.0   _    00   No        Yes     __        No              ___   |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|  F1=Help        F2=Split       F3=Exit        F4=Prompt      F5=Reset     |
|  F6=Previous    F7=Backward    F8=Forward     F9=Swap       F12=Cancel    |
'---------------------------------------------------------------------------'

       MA                                            >                   >  13/006
```

Figure 7-36   Device / Processor Definition

On this screen, we select the processor and CSS combination with a / next to the line and press Enter to move to the next screen (Figure 7-37 on page 279).

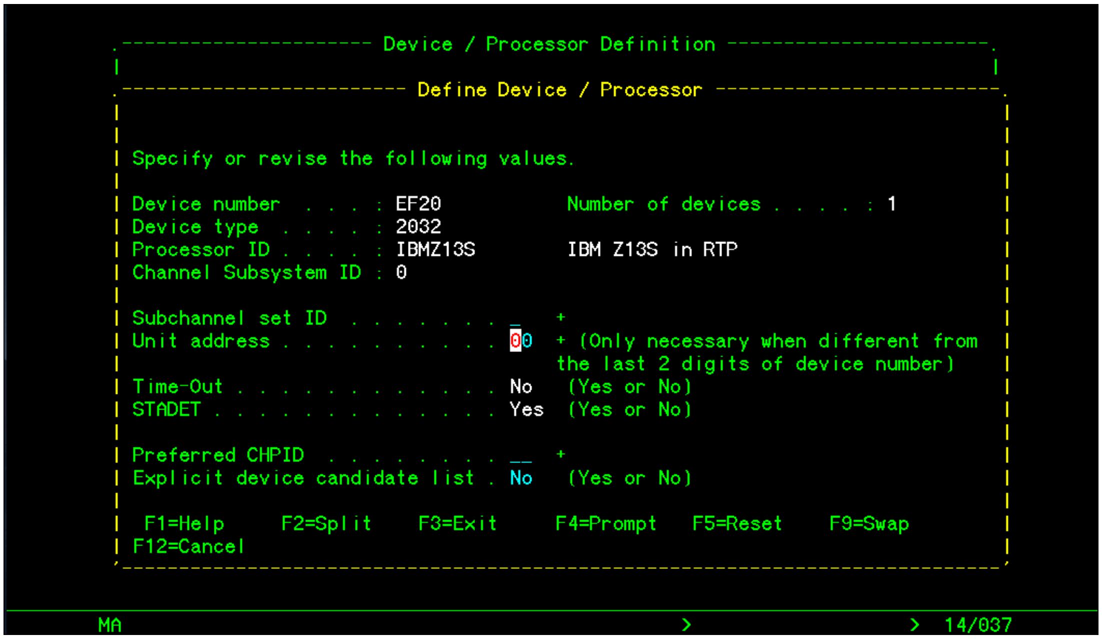

```
 .--------------------- Device / Processor Definition ----------------------.
 |                                                                          |
 .------------------------ Define Device / Processor ------------------------.
 |                                                                           |
 |                                                                           |
 | Specify or revise the following values.                                   |
 |                                                                           |
 | Device number  . . . : EF20          Number of devices . . . . : 1        |
 | Device type  . . . . : 2032                                               |
 | Processor ID . . . . : IBMZ13S       IBM Z13S in RTP                      |
 | Channel Subsystem ID : 0                                                  |
 |                                                                           |
 | Subchannel set ID  . . . . . . . . _   +                                  |
 | Unit address . . . . . . . . . . . 00  + (Only necessary when different from |
 |                                        the last 2 digits of device number) |
 | Time-Out . . . . . . . . . . . . . No   (Yes or No)                       |
 | STADET . . . . . . . . . . . . . . Yes  (Yes or No)                       |
 |                                                                           |
 | Preferred CHPID  . . . . . . . . . __  +                                  |
 | Explicit device candidate list . No   (Yes or No)                         |
 |                                                                           |
 |  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap      |
 | F12=Cancel                                                                |
 '---------------------------------------------------------------------------'

        MA                                              >         >  14/037
```

Figure 7-37   Device / Processor Definition

Because the information is correct for our environment, we press Enter to go to the screen that is shown in Figure 7-38.

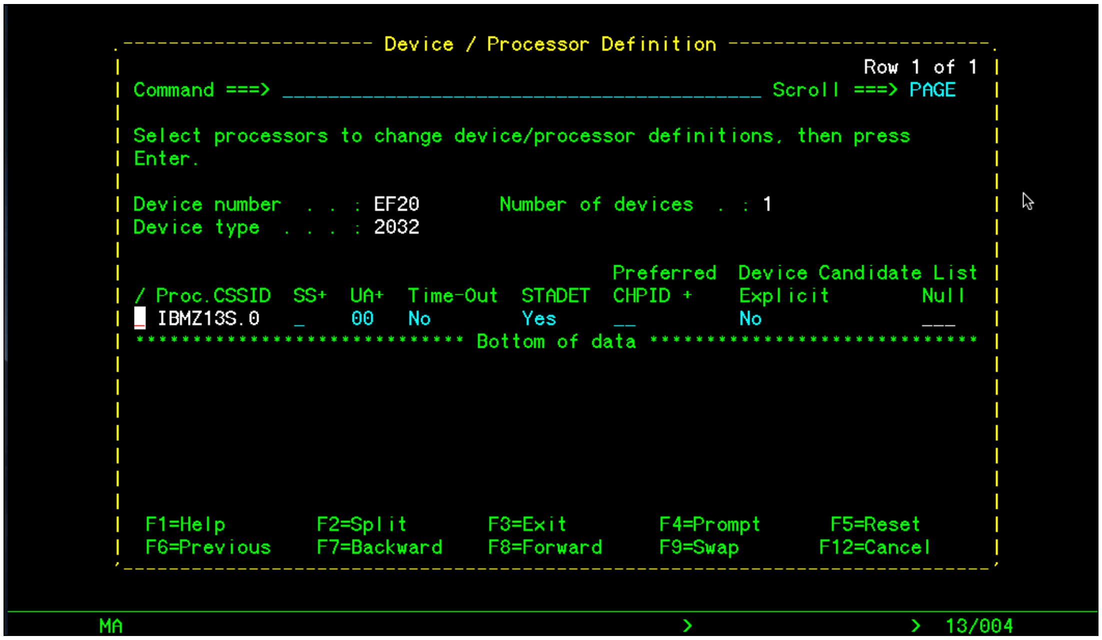

```
 .--------------------- Device / Processor Definition ----------------------.
 |                                                            Row 1 of 1 |
 | Command ===> __________________________________________ Scroll ===> PAGE |
 |                                                                         |
 | Select processors to change device/processor definitions, then press    |
 | Enter.                                                                  |
 |                                                                         |
 | Device number  . . : EF20       Number of devices  . : 1                |
 | Device type  . . . : 2032                                               |
 |                                                                         |
 |                                              Preferred  Device Candidate List |
 | / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null |
 | _ IBMZ13S.0   _    00   No        Yes     __         No             ___  |
 | ***************************** Bottom of data ***************************** |
 |                                                                         |
 |                                                                         |
 |                                                                         |
 |                                                                         |
 |                                                                         |
 |                                                                         |
 |  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset       |
 |  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel      |
 '-------------------------------------------------------------------------'

        MA                                              >         >  13/004
```

Figure 7-38   Device / Processor Definition

We position the cursor next to the same processor and CSS combination and press Enter with no command character. This action opens screen where we connect the CUP device to the required OS configuration, as shown in Figure 7-39.

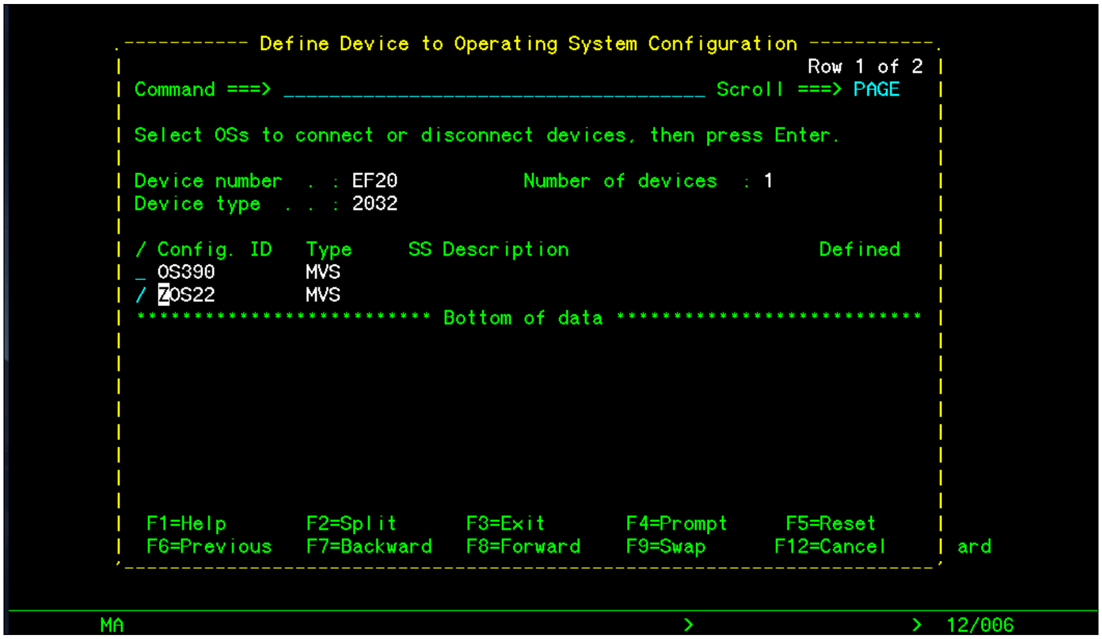

```
 .----------- Define Device to Operating System Configuration -----------.
 |                                                            Row 1 of 2 |
 | Command ===> _____________________________________ Scroll ===> PAGE   |
 |                                                                       |
 | Select OSs to connect or disconnect devices, then press Enter.        |
 |                                                                       |
 | Device number  . : EF20           Number of devices  : 1              |
 | Device type  . . : 2032                                               |
 |                                                                       |
 | / Config. ID   Type     SS Description                        Defined |
 | _ OS390        MVS                                                    |
 | / ZOS22        MVS                                                    |
 | ************************* Bottom of data **************************** |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset     |
 |  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel    | ard
 '-----------------------------------------------------------------------'

        MA                                                  >                   >  12/006
```

Figure 7-39 Define Device to Operating System Configuration

Mark the needed OS configuration with a /, press Enter, and then choose option 1 Select to connect, as shown in Figure 7-40.

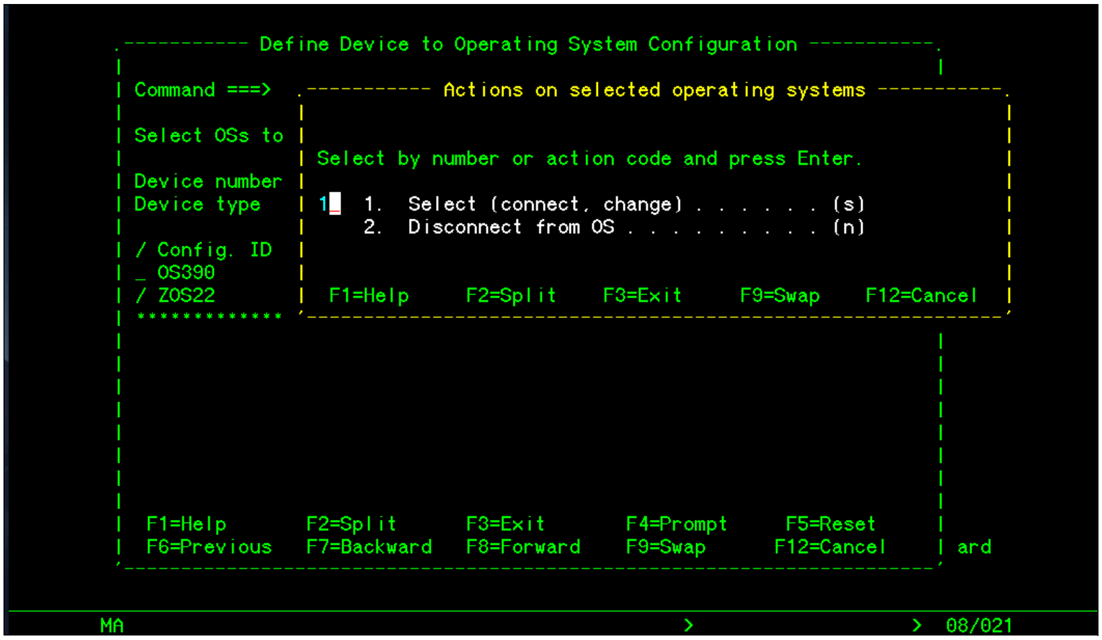

```
 .----------- Define Device to Operating System Configuration -----------.
 |                                                                       |
 | Command ===>  .----------- Actions on selected operating systems -----------.
 |               |                                                             |
 | Select OSs to |                                                             |
 |               | Select by number or action code and press Enter.            |
 | Device number |                                                             |
 | Device type   | 1_  1.  Select (connect, change) . . . . . . (s)            |
 |               |     2.  Disconnect from OS . . . . . . . . . (n)            |
 | / Config. ID  |                                                             |
 | _ OS390       |                                                             |
 | / ZOS22       |  F1=Help     F2=Split     F3=Exit     F9=Swap    F12=Cancel |
 | ************* '-------------------------------------------------------------'
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset     |
 |  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel    | ard
 '-----------------------------------------------------------------------'

        MA                                                  >                   >  08/021
```

Figure 7-40 Actions on selected operating systems

We choose the device parameters and features to configure, which are modified to match our lab standards that are shown in Figure 7-41 on page 281. Press Enter.

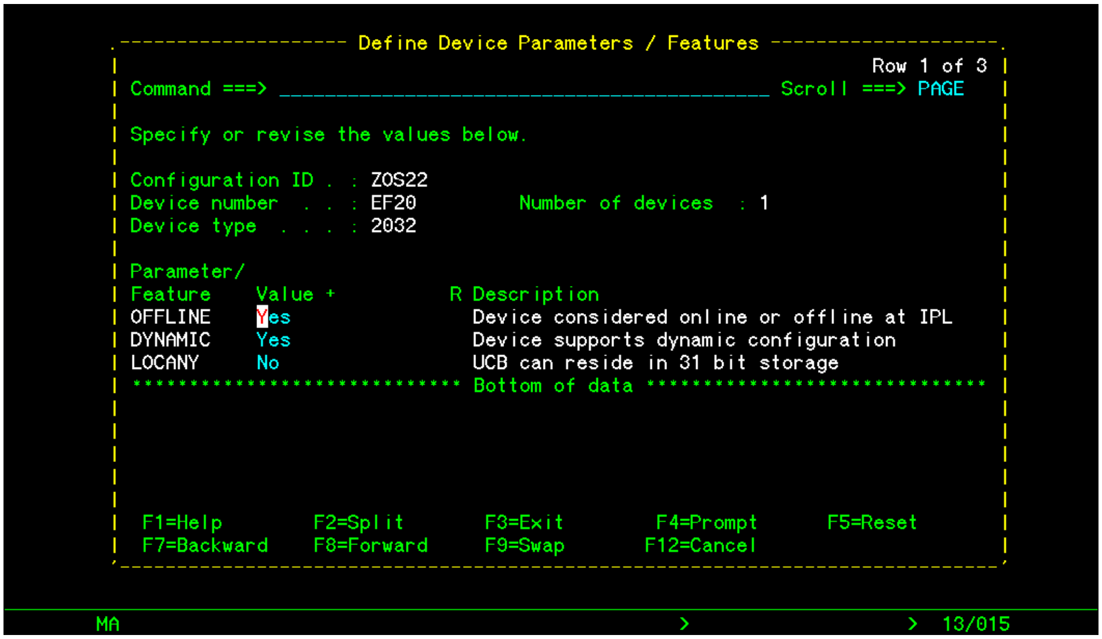

```
 .------------------- Define Device Parameters / Features --------------------.
 |                                                                 Row 1 of 3 |
 | Command ===> ___________________________________________ Scroll ===> PAGE  |
 |                                                                            |
 | Specify or revise the values below.                                        |
 |                                                                            |
 | Configuration ID . : ZOS22                                                 |
 | Device number  . . : EF20          Number of devices  : 1                  |
 | Device type  . . . : 2032                                                  |
 |                                                                            |
 | Parameter/                                                                 |
 | Feature    Value +          R Description                                  |
 | OFFLINE    Yes                Device considered online or offline at IPL   |
 | DYNAMIC    Yes                Device supports dynamic configuration        |
 | LOCANY     No                 UCB can reside in 31 bit storage             |
 | ***************************** Bottom of data ***************************** |
 |                                                                            |
 |                                                                            |
 |                                                                            |
 |                                                                            |
 |                                                                            |
 |  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset          |
 |  F7=Backward   F8=Forward    F9=Swap      F12=Cancel                       |
 '----------------------------------------------------------------------------'

        MA                                                >                   >  13/015
```

Figure 7-41   Define Device Parameters / Features

We see the dialog that is shown in Figure 7-42 for attaching system-defined esoterics to the CUP device. CUP devices are not part of any normal esoteric group, so nothing is selected. Press Enter.

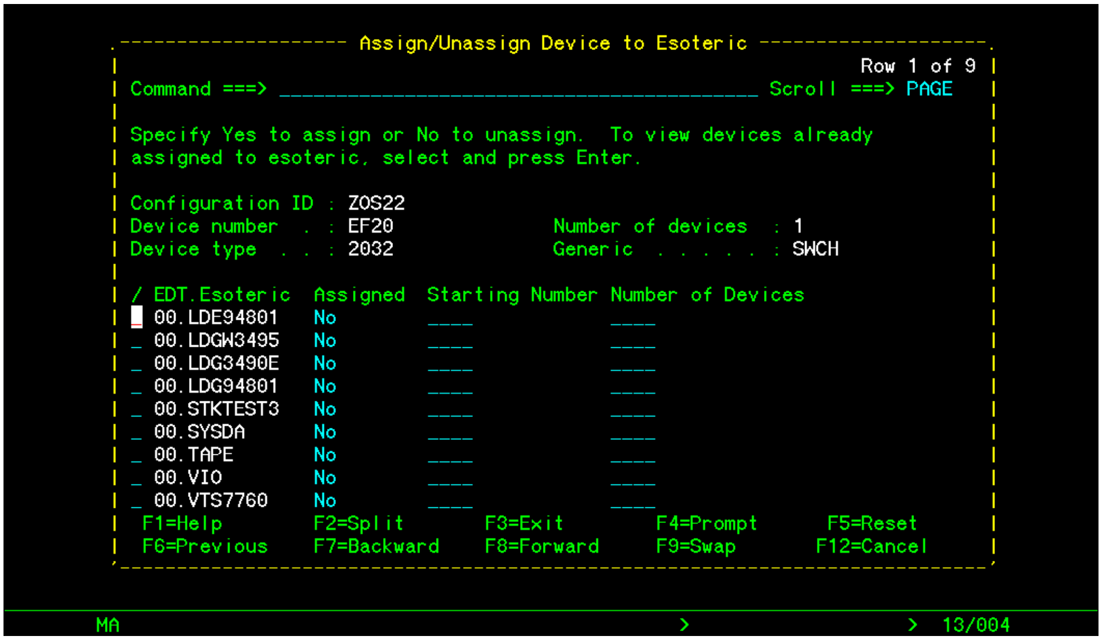

```
 .------------------- Assign/Unassign Device to Esoteric --------------------.
 |                                                                Row 1 of 9 |
 | Command ===> ___________________________________________ Scroll ===> PAGE |
 |                                                                           |
 | Specify Yes to assign or No to unassign.  To view devices already         |
 | assigned to esoteric, select and press Enter.                             |
 |                                                                           |
 | Configuration ID : ZOS22                                                  |
 | Device number  . : EF20                Number of devices  : 1             |
 | Device type  . . : 2032                Generic  . . . . . : SWCH          |
 |                                                                           |
 | / EDT.Esoteric  Assigned  Starting Number  Number of Devices              |
 | _ 00.LDE94801   No        ____             ____                           |
 | _ 00.LDGW3495   No        ____             ____                           |
 | _ 00.LDG3490E   No        ____             ____                           |
 | _ 00.LDG94801   No        ____             ____                           |
 | _ 00.STKTEST3   No        ____             ____                           |
 | _ 00.SYSDA      No        ____             ____                           |
 | _ 00.TAPE       No        ____             ____                           |
 | _ 00.VIO        No        ____             ____                           |
 | _ 00.VTS7760    No        ____             ____                           |
 |  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset         |
 |  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel        |
 '---------------------------------------------------------------------------'

        MA                                                >                   >  13/004
```

Figure 7-42   Assign/Unassign Device

We return to the screen that is shown in Figure 7-39 on page 280, except now we have an OS configuration that is defined that can access the CUP device. From this screen (see Figure 7-43), we press Enter to complete the CUP device definition.

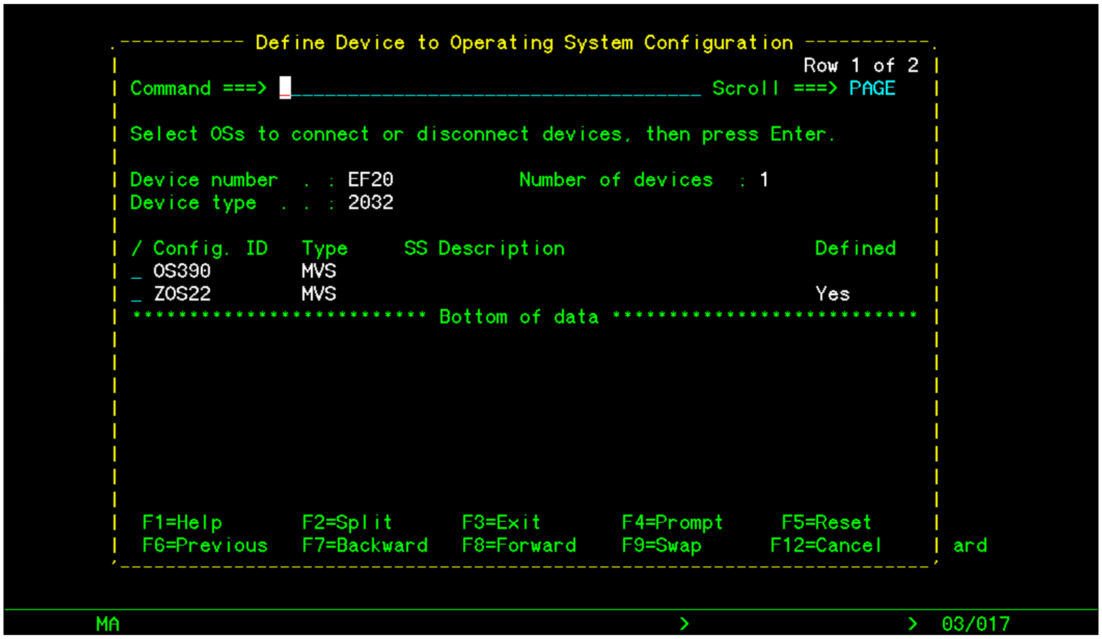

```
 .----------- Define Device to Operating System Configuration -----------.
 |                                                            Row 1 of 2 |
 | Command ===> _____________________________________ Scroll ===> PAGE   |
 |                                                                       |
 | Select OSs to connect or disconnect devices, then press Enter.        |
 |                                                                       |
 | Device number  . : EF20          Number of devices  : 1               |
 | Device type  . . : 2032                                               |
 |                                                                       |
 | / Config. ID   Type     SS Description                    Defined     |
 | _ OS390        MVS                                                    |
 | _ ZOS22        MVS                                         Yes        |
 | ************************* Bottom of data ************************** |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |                                                                       |
 |  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset     |
 |  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel    | ard
 '-----------------------------------------------------------------------'

        MA                                           >                   >  03/017
```

Figure 7-43   Selecting OSs

We are back at the I/O Device list screen, as shown in Figure 7-44. We see that the CUP device is complete because it has both CSS and OS connections. This system is ready to talk to the CUP device for switch ID 0x20.

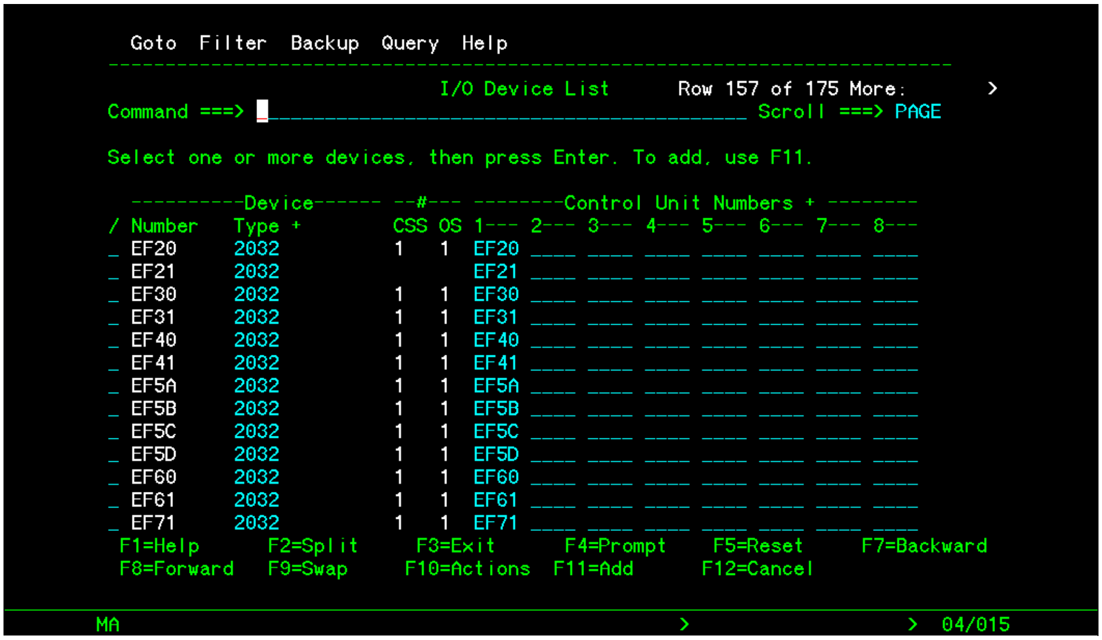

```
   Goto  Filter  Backup  Query  Help
 --------------------------------------------------------------------------
                             I/O Device List      Row 157 of 175 More:       >
 Command ===> __________________________________________ Scroll ===> PAGE

 Select one or more devices, then press Enter. To add, use F11.

   ----------Device------ --#--- --------Control Unit Numbers + --------
 / Number    Type +        CSS OS 1--- 2--- 3--- 4--- 5--- 6--- 7--- 8---
 _ EF20      2032           1   1  EF20 ____ ____ ____ ____ ____ ____ ____
 _ EF21      2032                  EF21 ____ ____ ____ ____ ____ ____ ____
 _ EF30      2032           1   1  EF30 ____ ____ ____ ____ ____ ____ ____
 _ EF31      2032           1   1  EF31 ____ ____ ____ ____ ____ ____ ____
 _ EF40      2032           1   1  EF40 ____ ____ ____ ____ ____ ____ ____
 _ EF41      2032           1   1  EF41 ____ ____ ____ ____ ____ ____ ____
 _ EF5A      2032           1   1  EF5A ____ ____ ____ ____ ____ ____ ____
 _ EF5B      2032           1   1  EF5B ____ ____ ____ ____ ____ ____ ____
 _ EF5C      2032           1   1  EF5C ____ ____ ____ ____ ____ ____ ____
 _ EF5D      2032           1   1  EF5D ____ ____ ____ ____ ____ ____ ____
 _ EF60      2032           1   1  EF60 ____ ____ ____ ____ ____ ____ ____
 _ EF61      2032           1   1  EF61 ____ ____ ____ ____ ____ ____ ____
 _ EF71      2032           1   1  EF71 ____ ____ ____ ____ ____ ____ ____
  F1=Help      F2=Split     F3=Exit      F4=Prompt    F5=Reset     F7=Backward
  F8=Forward   F9=Swap     F10=Actions  F11=Add      F12=Cancel

        MA                                           >                   >  04/015
```

Figure 7-44   I/O Device List

So that the CUP devices on each of the FICON VSANs within each switch can talk to this system, we also complete the CU and I/O device definitions in the same way.

7.2.6 Conclusion and Input\Output Configuration Program data set

We have completed the HCD definitions that were used in the building of the sample network that was used as the basis for the examples in this book. The definitions were activated by using HCD before the switch environment was built. Throughout the rest of this book, we bring various devices online and start workloads to help illustrate how things look in a live example.

The Input/Output Configuration Program (IOCP) data set that is generated by HCD for the input is shown in Figure 7-45.

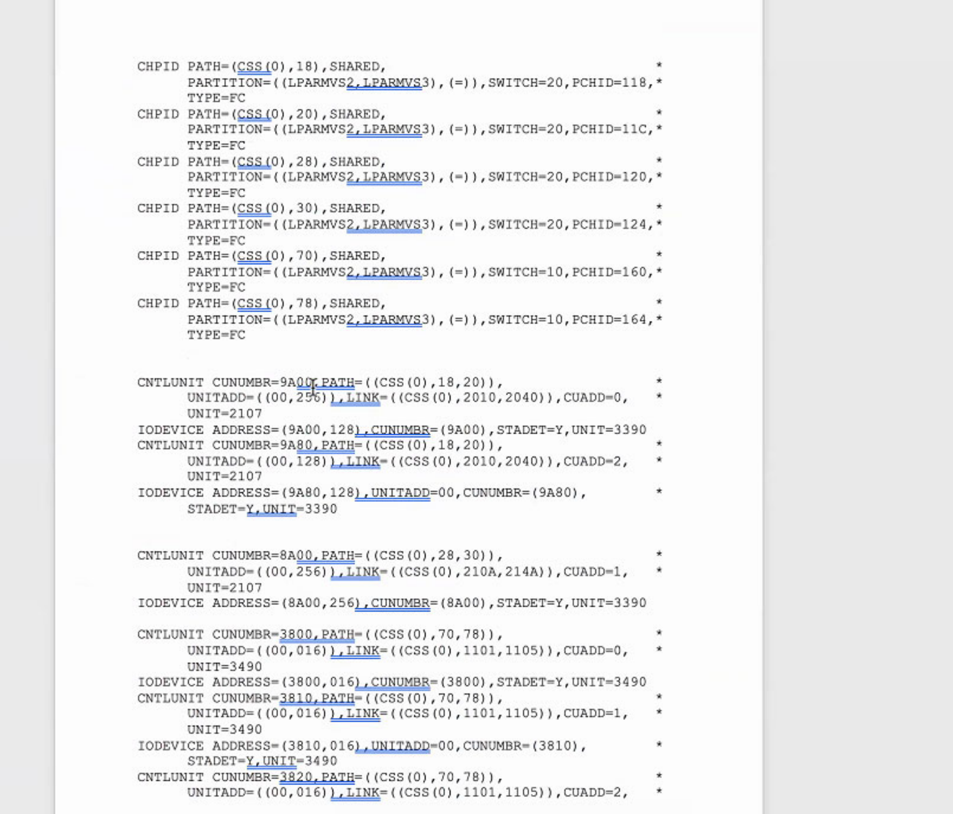

```
CHPID PATH=(CSS(0),18),SHARED,                                *
      PARTITION=((LPARMVS2,LPARMVS3),(=)),SWITCH=20,PCHID=118,*
      TYPE=FC
CHPID PATH=(CSS(0),20),SHARED,                                *
      PARTITION=((LPARMVS2,LPARMVS3),(=)),SWITCH=20,PCHID=11C,*
      TYPE=FC
CHPID PATH=(CSS(0),28),SHARED,                                *
      PARTITION=((LPARMVS2,LPARMVS3),(=)),SWITCH=20,PCHID=120,*
      TYPE=FC
CHPID PATH=(CSS(0),30),SHARED,                                *
      PARTITION=((LPARMVS2,LPARMVS3),(=)),SWITCH=20,PCHID=124,*
      TYPE=FC
CHPID PATH=(CSS(0),70),SHARED,                                *
      PARTITION=((LPARMVS2,LPARMVS3),(=)),SWITCH=10,PCHID=160,*
      TYPE=FC
CHPID PATH=(CSS(0),78),SHARED,                                *
      PARTITION=((LPARMVS2,LPARMVS3),(=)),SWITCH=10,PCHID=164,*
      TYPE=FC


CNTLUNIT CUNUMBR=9A00,PATH=((CSS(0),18,20)),                  *
      UNITADD=((00,256)),LINK=((CSS(0),2010,2040)),CUADD=0,   *
      UNIT=2107
IODEVICE ADDRESS=(9A00,128),CUNUMBR=(9A00),STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=9A80,PATH=((CSS(0),18,20)),                  *
      UNITADD=((00,128)),LINK=((CSS(0),2010,2040)),CUADD=2,   *
      UNIT=2107
IODEVICE ADDRESS=(9A80,128),UNITADD=00,CUNUMBR=(9A80),        *
      STADET=Y,UNIT=3390


CNTLUNIT CUNUMBR=8A00,PATH=((CSS(0),28,30)),                  *
      UNITADD=((00,256)),LINK=((CSS(0),210A,214A)),CUADD=1,   *
      UNIT=2107
IODEVICE ADDRESS=(8A00,256),CUNUMBR=(8A00),STADET=Y,UNIT=3390

CNTLUNIT CUNUMBR=3800,PATH=((CSS(0),70,78)),                  *
      UNITADD=((00,016)),LINK=((CSS(0),1101,1105)),CUADD=0,   *
      UNIT=3490
IODEVICE ADDRESS=(3800,016),CUNUMBR=(3800),STADET=Y,UNIT=3490
CNTLUNIT CUNUMBR=3810,PATH=((CSS(0),70,78)),                  *
      UNITADD=((00,016)),LINK=((CSS(0),1101,1105)),CUADD=1,   *
      UNIT=3490
IODEVICE ADDRESS=(3810,016),UNITADD=00,CUNUMBR=(3810),        *
      STADET=Y,UNIT=3490
CNTLUNIT CUNUMBR=3820,PATH=((CSS(0),70,78)),                  *
      UNITADD=((00,016)),LINK=((CSS(0),1101,1105)),CUADD=2,   *
```

Figure 7-45 Input/Output Configuration Program (IOCP) data set

7.2.7 Features

In this section, we show how to enable the features that are needed for the mainframe environment.

Select **Topology**, double-click **SAN384C-6**, and then select **Show more details**, as shown in Figure 7-46.

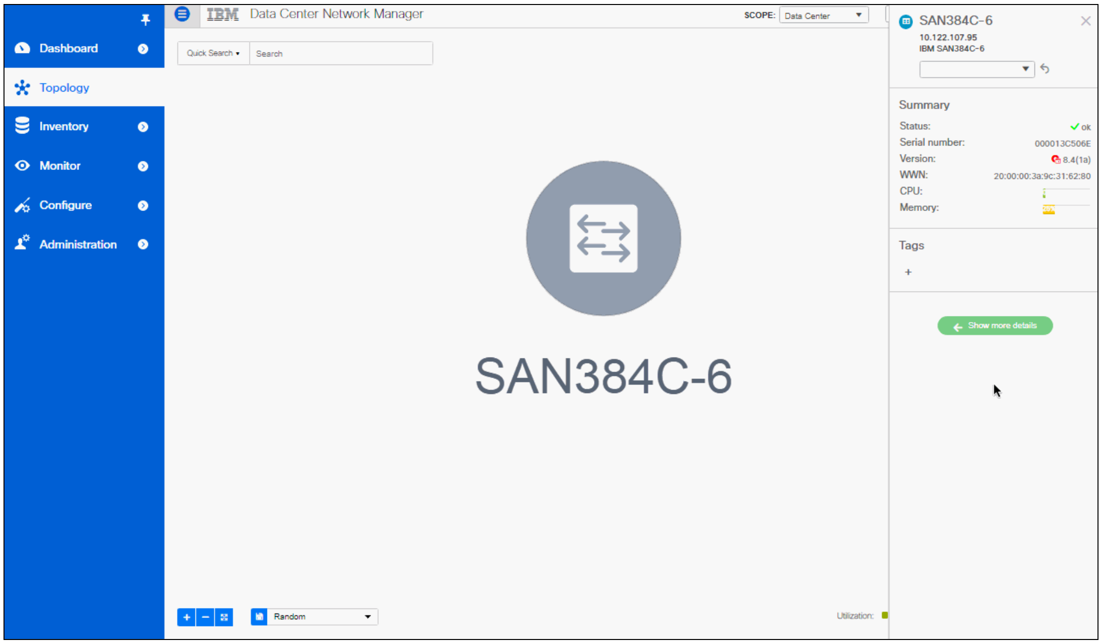

SAN384C-6
10.122.107.95
IBM SAN384C-6

Summary

Status: ✓ ok
Serial number: 000013C506E
Version: 8.4(1a)
WWN: 20:00:00:3a:9c:31:62:80
CPU:
Memory: 25%

Tags

Figure 7-46 Show more details window

Figure 7-47 shows the default System Info view tab for the IBM Storage Networking SAN384C-6 switch. Click the **Features** tab to view the active features.

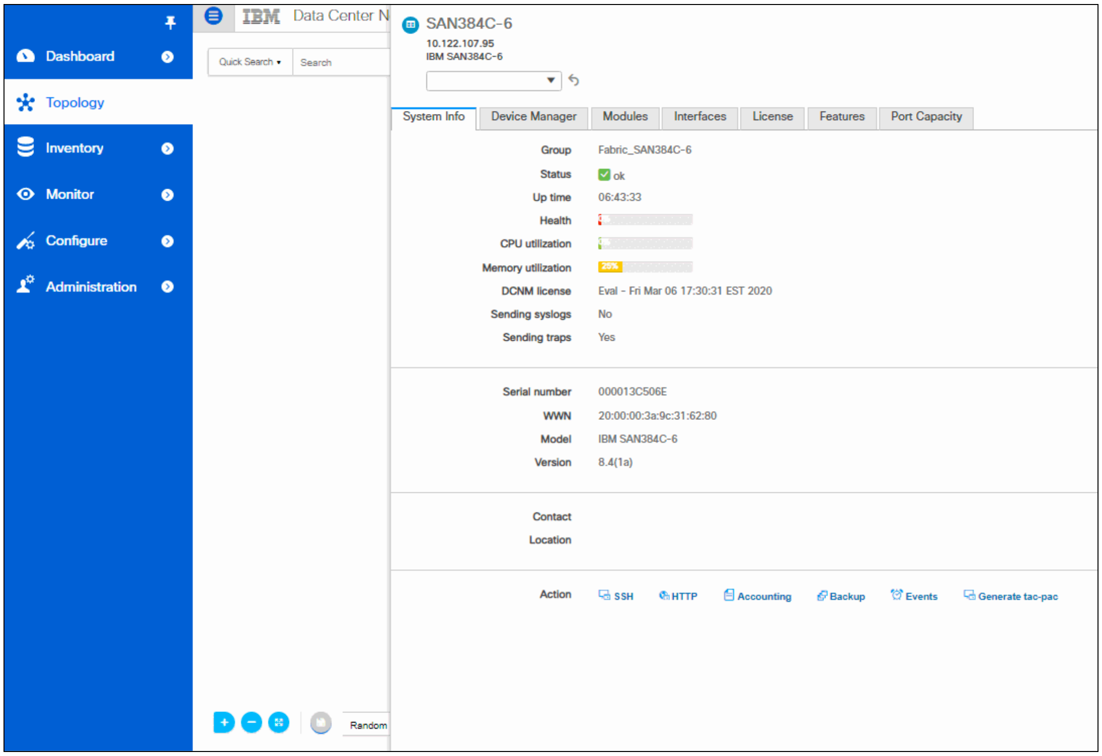

SAN384C-6
10.122.107.95
IBM SAN384C-6

Group Fabric_SAN384C-6
Status ✓ ok
Up time 06:43:33
Health
CPU utilization
Memory utilization 25%
DCNM license Eval - Fri Mar 06 17:30:31 EST 2020
Sending syslogs No
Sending traps Yes

Serial number 000013C506E
WWN 20:00:00:3a:9c:31:62:80
Model IBM SAN384C-6
Version 8.4(1a)

Contact
Location

Action

Figure 7-47 System Info

Figure 7-48 shows the enabled features on the IBM Storage Networking SAN384C-6 switch. To enable more features that are required for the FICON environment, click the **Device Manager** tab.

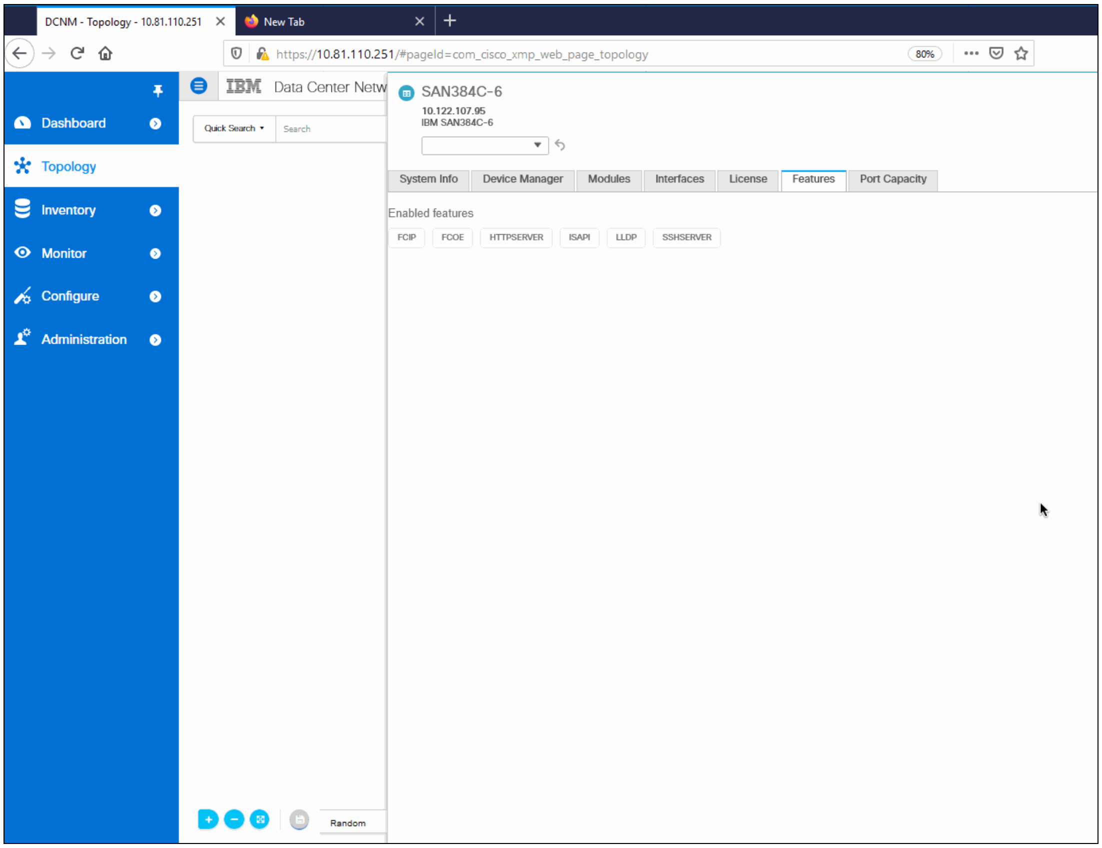

Figure 7-48 Enabled Features

Figure 7-49 shows the Device Manager tab. Select **Admin** → **Feature Control**. This window is used to enable and disable features.

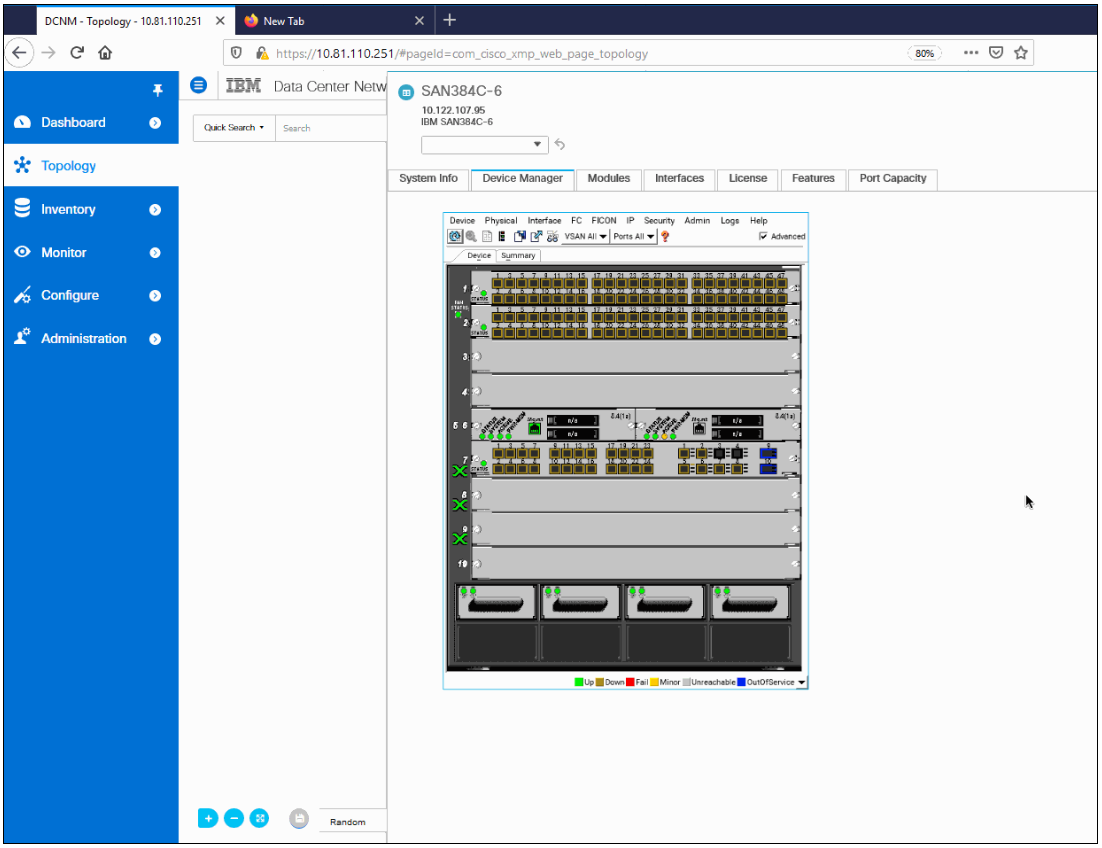

Figure 7-49 Device Manager tab

Figure 7-50 shows the Features Control tab with fabric-binding disabled. We find the fabric-binding feature under the Action column and change it to **Enable**.

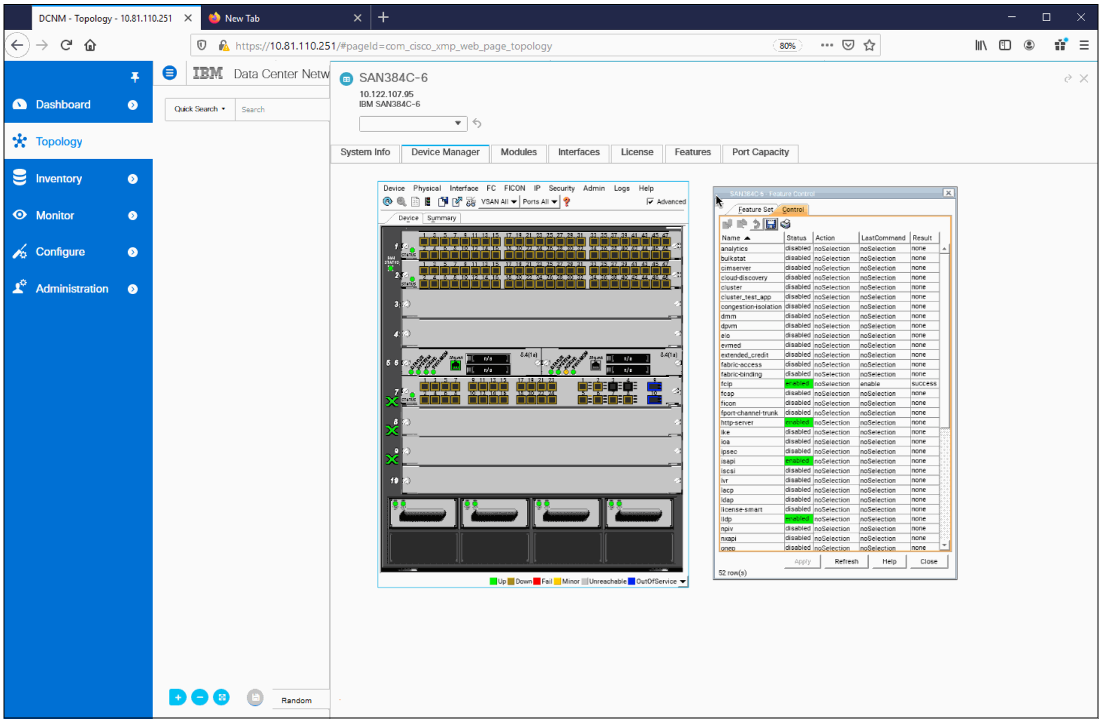

SAN384C-6
10.122.107.95
IBM SAN384C-6

SAN384C-6 - Feature Control

Name	Status	Action	LastCommand	Result
analytics	disabled	noSelection	noSelection	none
bulkstat	disabled	noSelection	noSelection	none
cimserver	disabled	noSelection	noSelection	none
cloud-discovery	disabled	noSelection	noSelection	none
cluster	disabled	noSelection	noSelection	none
cluster_test_app	disabled	noSelection	noSelection	none
congestion-isolation	disabled	noSelection	noSelection	none
dmm	disabled	noSelection	noSelection	none
dpvm	disabled	noSelection	noSelection	none
eio	disabled	noSelection	noSelection	none
evmed	disabled	noSelection	noSelection	none
extended_credit	disabled	noSelection	noSelection	none
fabric-access	disabled	noSelection	noSelection	none
fabric-binding	disabled	noSelection	noSelection	none
fcip	enabled	noSelection	enable	success
fcsp	disabled	noSelection	noSelection	none
ficon	disabled	noSelection	noSelection	none
fport-channel-trunk	disabled	noSelection	noSelection	none
http-server	enabled	noSelection	noSelection	none
ike	disabled	noSelection	noSelection	none
ioa	disabled	noSelection	noSelection	none
ipsec	disabled	noSelection	noSelection	none
isapi	enabled	noSelection	noSelection	none
iscsi	disabled	noSelection	noSelection	none
ivr	disabled	noSelection	noSelection	none
lacp	disabled	noSelection	noSelection	none
ldap	disabled	noSelection	noSelection	none
license-smart	disabled	noSelection	noSelection	none
lldp	enabled	noSelection	noSelection	none
npiv	disabled	noSelection	noSelection	none
nxapi	disabled	noSelection	noSelection	none
onep	disabled	noSelection	noSelection	none

Apply Refresh Help Close

52 row(s)

Figure 7-50 Fabric Binding feature disabled

Figure 7-51 shows that fabric-binding (highlighted in blue) was changed to enable. Click **Apply** to activate the feature.

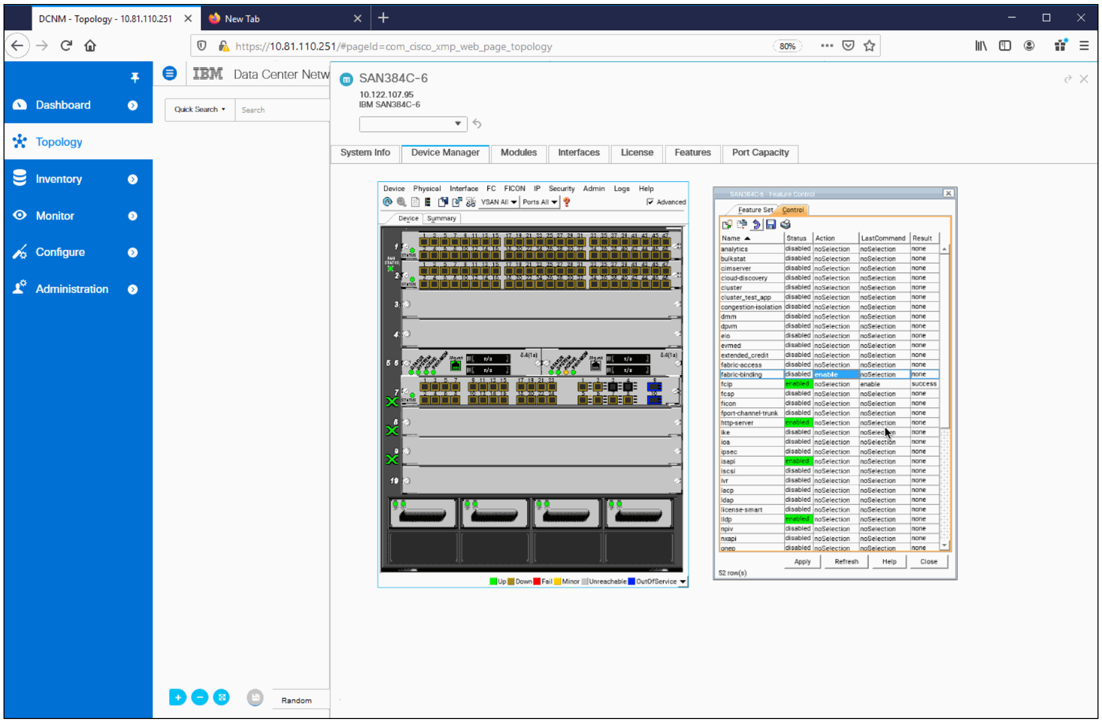

SAN384C-6
10.122.107.95
IBM SAN384C-6

SAN384C-6 - Feature Control

Name	Status	Action	LastCommand	Result
analytics	disabled	noSelection	noSelection	none
bulkstat	disabled	noSelection	noSelection	none
cimserver	disabled	noSelection	noSelection	none
cloud-discovery	disabled	noSelection	noSelection	none
cluster	disabled	noSelection	noSelection	none
cluster_test_app	disabled	noSelection	noSelection	none
congestion-isolation	disabled	noSelection	noSelection	none
dmm	disabled	noSelection	noSelection	none
dpvm	disabled	noSelection	noSelection	none
eio	disabled	noSelection	noSelection	none
evmed	disabled	noSelection	noSelection	none
extended_credit	disabled	noSelection	noSelection	none
fabric-access	disabled	noSelection	noSelection	none
fabric-binding	disabled	enable	noSelection	none
fcip	enabled	noSelection	enable	success
fcsp	disabled	noSelection	noSelection	none
ficon	disabled	noSelection	noSelection	none
fport-channel-trunk	disabled	noSelection	noSelection	none
http-server	enabled	noSelection	noSelection	none
ike	disabled	noSelection	noSelection	none
ioa	disabled	noSelection	noSelection	none
ipsec	disabled	noSelection	noSelection	none
isapi	enabled	noSelection	noSelection	none
iscsi	disabled	noSelection	noSelection	none
ivr	disabled	noSelection	noSelection	none
lacp	disabled	noSelection	noSelection	none
ldap	disabled	noSelection	noSelection	none
license-smart	disabled	noSelection	noSelection	none
lldp	enabled	noSelection	noSelection	none
npiv	disabled	noSelection	noSelection	none
nxapi	disabled	noSelection	noSelection	none
onep	disabled	noSelection	noSelection	none

Apply Refresh Help Close

52 row(s)

Figure 7-51 Fabric Binding enabled

Figure 7-52 shows that the fabric-binding status is enabled and the results column as Success.

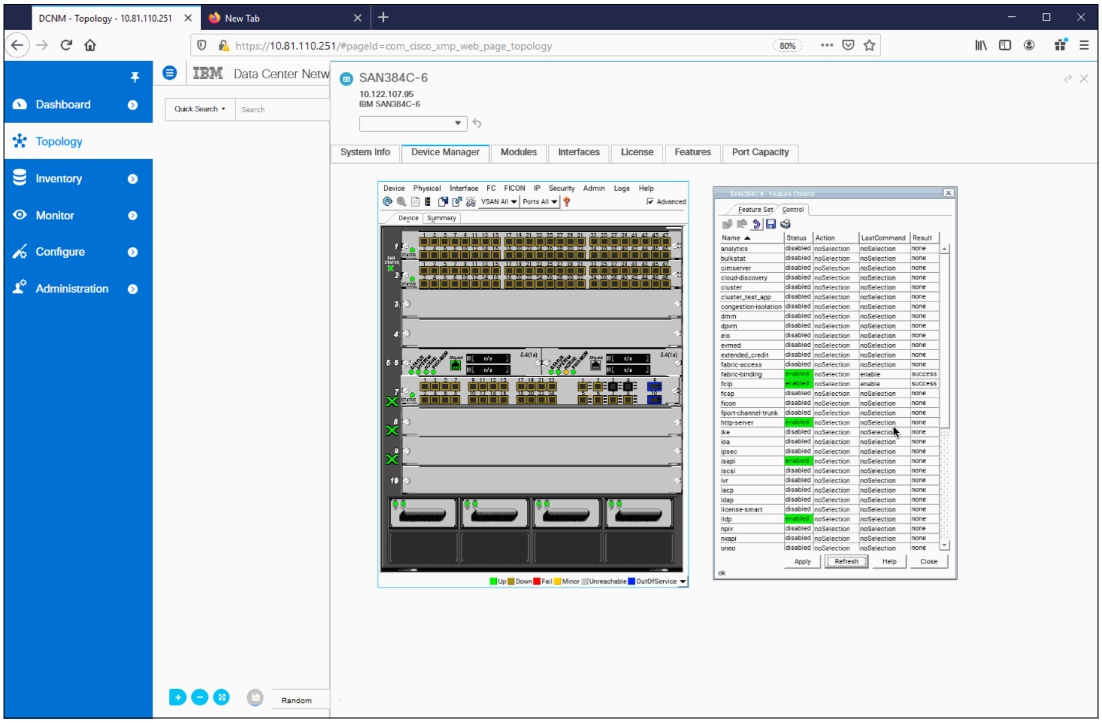

SAN384C-6
10.122.107.95
IBM SAN384C-6

Name	Status	Action	LastCommand	Result
analytics	disabled	noSelection	noSelection	none
bulkstat	disabled	noSelection	noSelection	none
cimserver	disabled	noSelection	noSelection	none
cloud-discovery	disabled	noSelection	noSelection	none
cluster	disabled	noSelection	noSelection	none
cluster_test_app	disabled	noSelection	noSelection	none
congestion-isolation	disabled	noSelection	noSelection	none
dmm	disabled	noSelection	noSelection	none
dpvm	disabled	noSelection	noSelection	none
eio	disabled	noSelection	noSelection	none
evmed	disabled	noSelection	noSelection	none
extended_credit	disabled	noSelection	noSelection	none
fabric-access	disabled	noSelection	noSelection	none
fabric-binding	enabled	noSelection	enable	success
fcip	enabled	noSelection	enable	success
fcsp	disabled	noSelection	noSelection	none
ficon	disabled	noSelection	noSelection	none
fport-channel-trunk	disabled	noSelection	noSelection	none
http-server	enabled	noSelection	noSelection	none
ike	disabled	noSelection	noSelection	none
ioa	disabled	noSelection	noSelection	none
ipsec	disabled	noSelection	noSelection	none
isapi	enabled	noSelection	noSelection	none
iscsi	disabled	noSelection	noSelection	none
ivr	disabled	noSelection	noSelection	none
lacp	disabled	noSelection	noSelection	none
ldap	disabled	noSelection	noSelection	none
license-smart	disabled	noSelection	noSelection	none
lldp	enabled	noSelection	noSelection	none
npiv	disabled	noSelection	noSelection	none
nxapi	disabled	noSelection	noSelection	none
onep	disabled	noSelection	noSelection	none

Figure 7-52 Fabric Binding enabled successfully

Now that we have enabled the Fabric Binding feature, we repeat the same steps for the FICON feature, as shown in Figure 7-53.

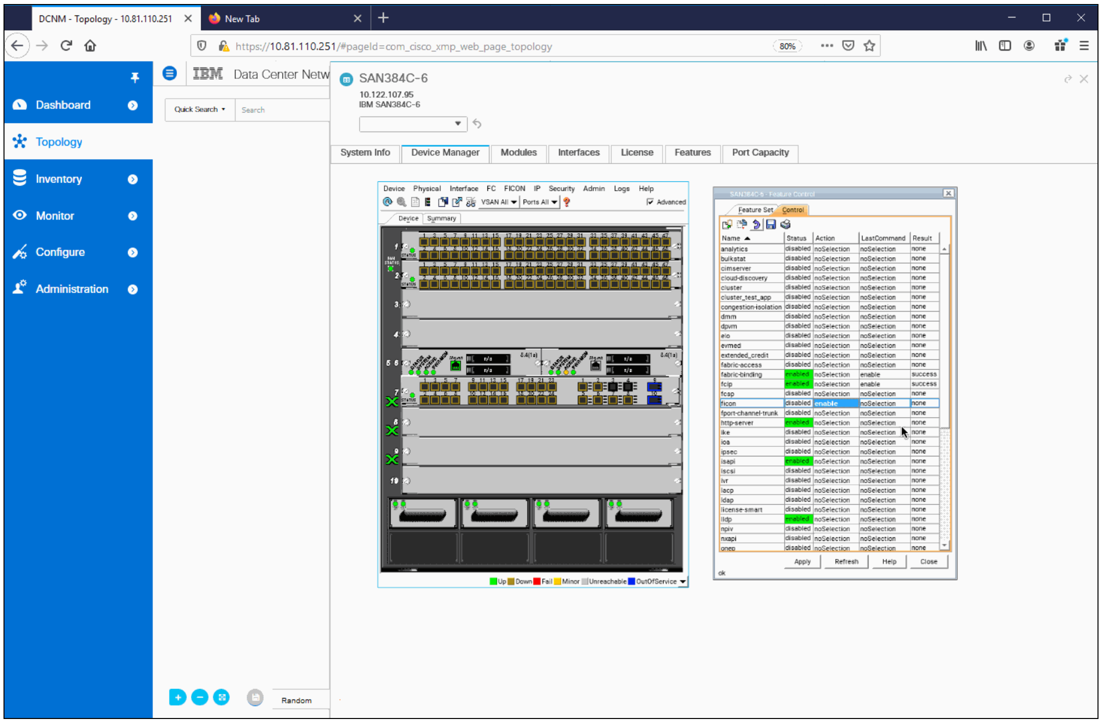

SAN384C-6
10.122.107.95
IBM SAN384C-6

Name	Status	Action	LastCommand	Result
analytics	disabled	noSelection	noSelection	none
bulkstat	disabled	noSelection	noSelection	none
cimserver	disabled	noSelection	noSelection	none
cloud-discovery	disabled	noSelection	noSelection	none
cluster	disabled	noSelection	noSelection	none
cluster_test_app	disabled	noSelection	noSelection	none
congestion-isolation	disabled	noSelection	noSelection	none
dmm	disabled	noSelection	noSelection	none
dpvm	disabled	noSelection	noSelection	none
eio	disabled	noSelection	noSelection	none
evmed	disabled	noSelection	noSelection	none
extended_credit	disabled	noSelection	noSelection	none
fabric-access	disabled	noSelection	noSelection	none
fabric-binding	enabled	noSelection	enable	success
fcip	enabled	noSelection	enable	success
fcsp	disabled	noSelection	noSelection	none
ficon	disabled	enable	noSelection	none
fport-channel-trunk	disabled	noSelection	noSelection	none
http-server	enabled	noSelection	noSelection	none
ike	disabled	noSelection	noSelection	none
ioa	disabled	noSelection	noSelection	none
ipsec	disabled	noSelection	noSelection	none
isapi	enabled	noSelection	noSelection	none
iscsi	disabled	noSelection	noSelection	none
ivr	disabled	noSelection	noSelection	none
lacp	disabled	noSelection	noSelection	none
ldap	disabled	noSelection	noSelection	none
license-smart	disabled	noSelection	noSelection	none
lldp	enabled	noSelection	noSelection	none
npiv	disabled	noSelection	noSelection	none
nxapi	disabled	noSelection	noSelection	none
onep	disabled	noSelection	noSelection	none

Figure 7-53 FICON feature disabled

Figure 7-54 shows that the FICON feature was enabled successfully.

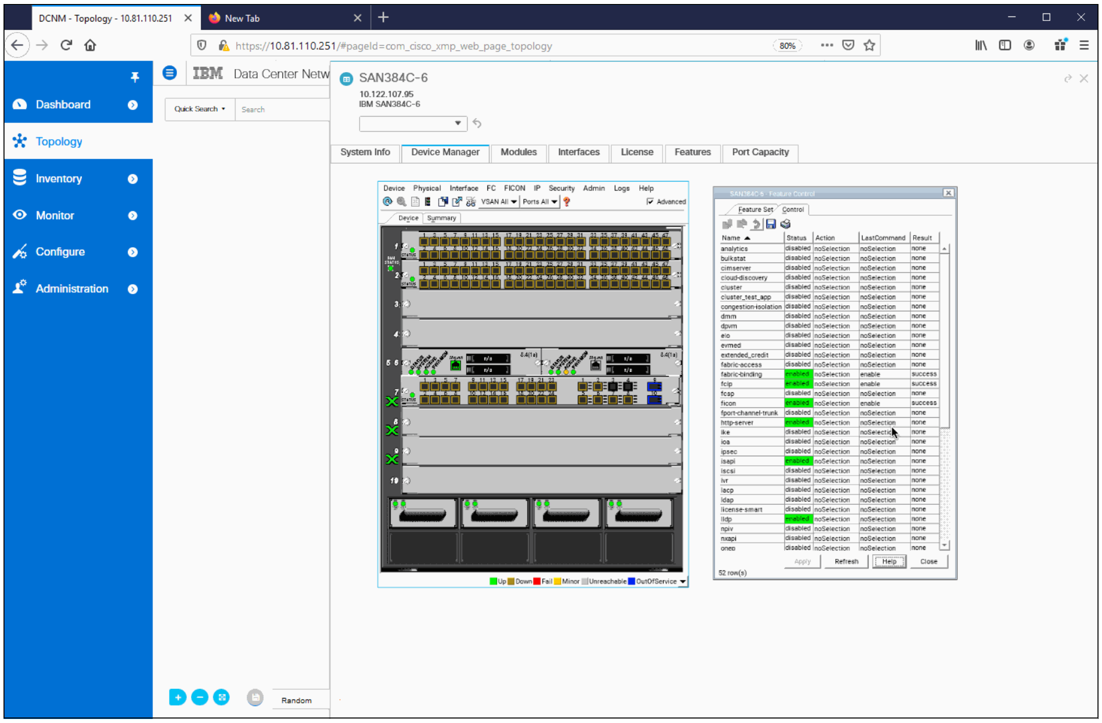

Name	Status	Action	LastCommand	Result
analytics	disabled	noSelection	noSelection	none
bulkstat	disabled	noSelection	noSelection	none
cimserver	disabled	noSelection	noSelection	none
cloud-discovery	disabled	noSelection	noSelection	none
cluster	disabled	noSelection	noSelection	none
cluster_test_app	disabled	noSelection	noSelection	none
congestion-isolation	disabled	noSelection	noSelection	none
dmm	disabled	noSelection	noSelection	none
dpvm	disabled	noSelection	noSelection	none
elo	disabled	noSelection	noSelection	none
evmed	disabled	noSelection	noSelection	none
extended_credit	disabled	noSelection	noSelection	none
fabric-access	disabled	noSelection	noSelection	none
fabric-binding	enabled	noSelection	enable	success
fcip	enabled	noSelection	enable	success
fcsp	disabled	noSelection	noSelection	none
ficon	enabled	noSelection	enable	success
fport-channel-trunk	disabled	noSelection	noSelection	none
http-server	enabled	noSelection	noSelection	none
ike	disabled	noSelection	noSelection	none
ioa	disabled	noSelection	noSelection	none
ipsec	disabled	noSelection	noSelection	none
isapi	enabled	noSelection	noSelection	none
iscsi	disabled	noSelection	noSelection	none
ivr	disabled	noSelection	noSelection	none
lacp	disabled	noSelection	noSelection	none
ldap	disabled	noSelection	noSelection	none
license-smart	disabled	noSelection	noSelection	none
lldp	enabled	noSelection	noSelection	none
npiv	disabled	noSelection	noSelection	none
nxapi	disabled	noSelection	noSelection	none
onep	disabled	noSelection	noSelection	none

Figure 7-54 FICON feature enabled

Enabling features by using the CLI

In this section, we show how to enable the features that are needed for the mainframe environment by using the CLI, as shown in Example 7-1.

Example 7-1 Enabling features by using the CLI

```
SAN384C-6#
SAN384C-6# show feature | incl enabled
http-server           1         enabled
isapi                 1         enabled
lldp                  1         enabled
sshServer             1         enabled
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# feature fabric-binding
SAN384C-6(config)#
SAN384C-6(config)# feature ficon
SAN384C-6(config)#
SAN384C-6(config)# end
SAN384C-6# show feature | include enabled
fabric-binding        1         enabled
ficon                 1         enabled
http-server           1         enabled
isapi                 1         enabled
lldp                  1         enabled
sshServer             1         enabled
SAN384C-6#
```

7.2.8 Configuring FICON Port Addresses

When setting up FICON, you first must configure the FICON Port Addressing that will be used for the switch. The FICON architecture allows for 256 FICON ports that are based on a single-byte port address. The values 0x00 - 0xFD are valid FICON addresses. 0xFE is reserved for the FICON CUP devices, and 0xFF is reserved. Because the IBM c-type switches support more than 256 ports, if you use all the ports in the switch, some port addressees must be allocated to different VSANs because the mainframe does not support duplicate addresses.

The IBM c-type switch allows up to eight FICON VSANs. In addition, all the ports have virtual port addresses.

There are two types of FICON Port Addresses:

- Physical: The interfaces are used to connect a FICON device.
- Logical: The interfaces are used for port channels and FCIP interfaces.

Note: There is no affinity between the physical location of a port within the switch and the FICON Port Address. Any port address can be at any physical location of the switch if there is not a duplicate port address in any single VSAN.

For this configuration step, we start with DM, as shown in Figure 7-55.

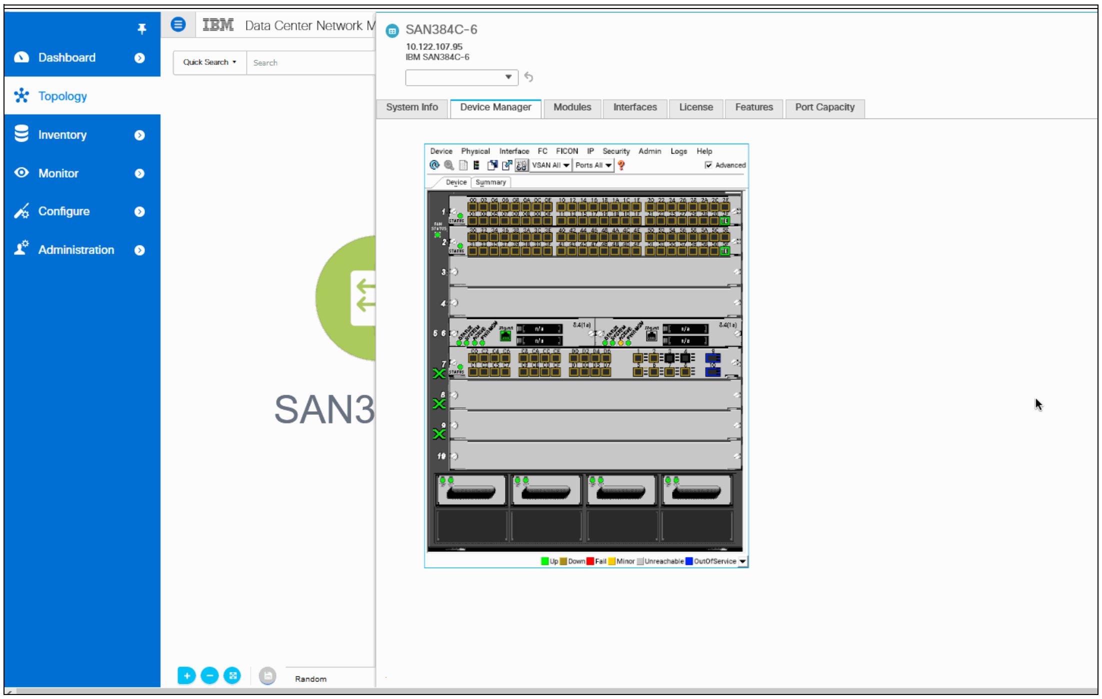

Figure 7-55 Device Manager

Note: With the FICON sunglasses button in DM, you can toggle between the Standard view and the FICON view. The Standard view numbers the ports as they are physically on the hardware line card. The FICON view numbers the ports with the FICON Port Addresses that are used for FICON routing.

Figure 7-56 shows a side-by-side representation of the Standard and FICON views for ports on an IBM Storage Networking SAN384C-6 switch.

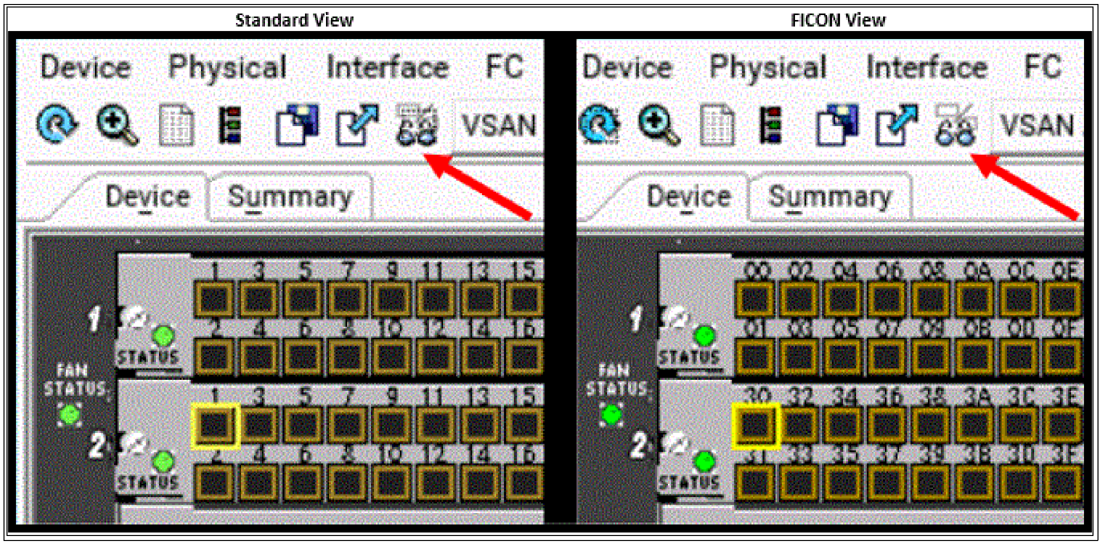

Figure 7-56 FICON sunglasses

To start the configuration of the FICON Port Address layout, select **FICON** → **Port Numbers**, as shown in Figure 7-57. By default, each c-type switch allocates 48 physical FICON Port Addresses per slot up to 0xEF. FICON Port Addresses can be arranged in any order to meet the requirements of the environment.

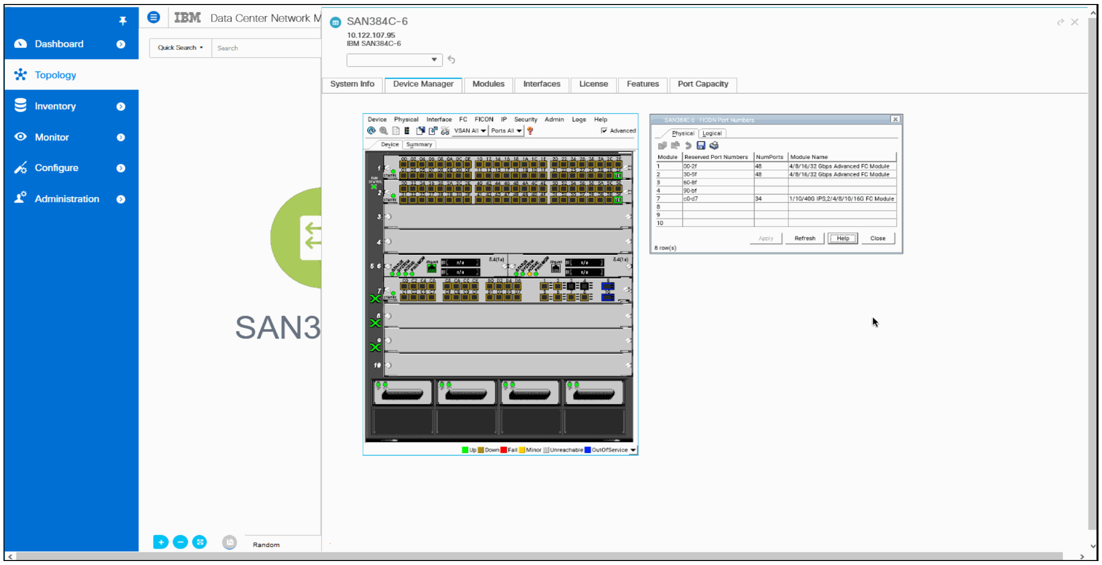

SAN384C-6 - FICON Port Numbers

Physical | Logical

Module	Reserved Port Numbers	NumPorts	Module Name
1	00-2f	48	4/8/16/32 Gbps Advanced FC Module
2	30-5f	48	4/8/16/32 Gbps Advanced FC Module
3	60-8f		
4	90-bf		
7	c0-d7	34	1/10/40G IPS,2/4/8/10/16G FC Module
8			
9			
10			

Apply | Refresh | Help | Close

8 row(s)

Figure 7-57 FICON Port Numbers

Figure 7-58 on page 293 shows FICON Port Numbers Logical tab. By default, the switch allocates the last 14 FICON Port Addresses for the logical pool. In a case where the switch is not using ISLs, the logical port addresses can be removed and reallocated as physical port addresses. Because our lab environment has several FCIP links and port channels, we use the default for logical port addresses.

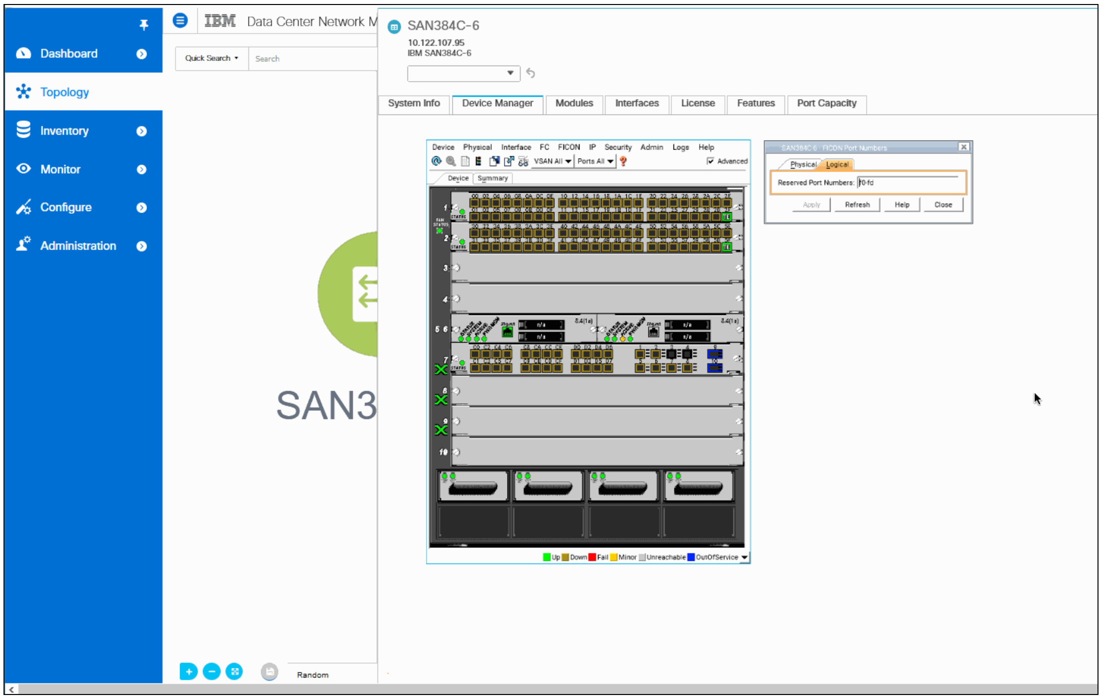

Reserved Port Numbers: f0-fd

Figure 7-58 Logical port addresses

7.2.9 Creating virtual storage area networks

In this section, we show how to create a VSAN, which provides traffic isolation, dedicated fabric services, and support for up to 256 VSANs in a switch. Users can configure VSAN IDs in the range of 2 to 4093, with the default being VSAN 1 along with an additional isolated VSAN (VSAN 4094).

> **Important:** VSAN 1 should never be used with FICON or open systems devices. New custom VSANs should be created when configuring environments.

To create a VSAN, launch DCNM and select **Topology**, select a switch, and select **Show more details** → **Device Manager**. Within Device Manager, select **FC** → **VSANs**, as shown in Figure 7-59.

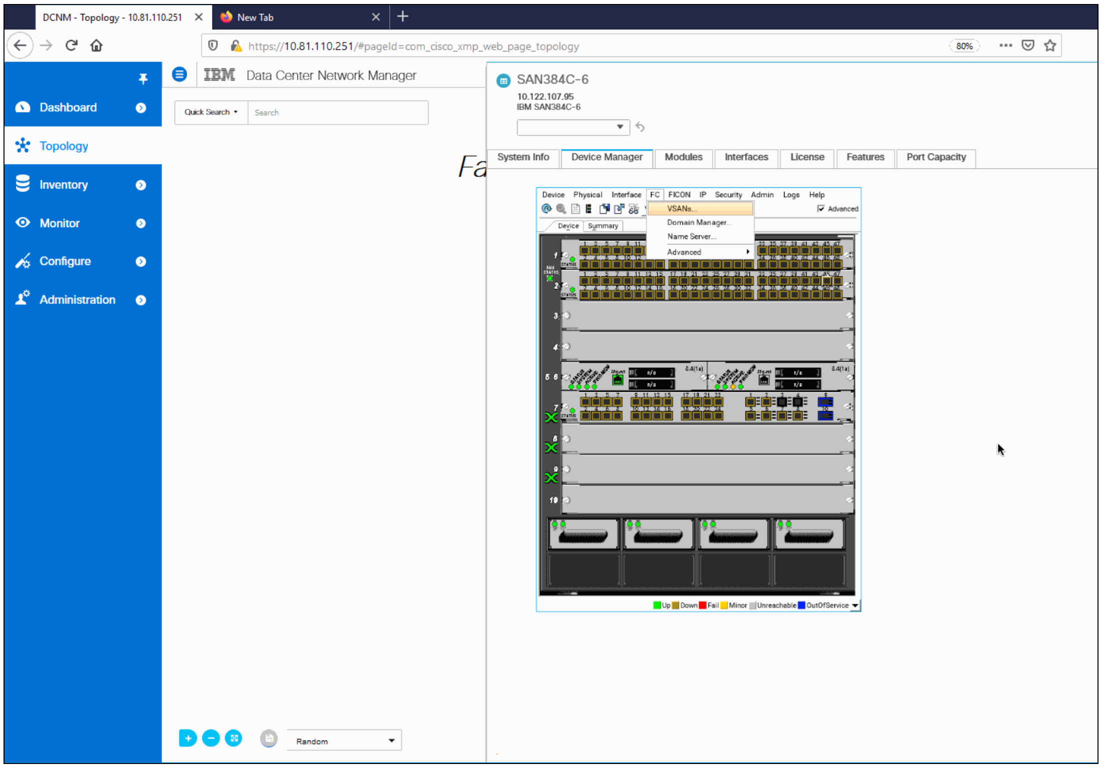

Figure 7-59 Device Manager

Figure 7-60 on page 295 shows the VSAN window. Click **Create** to enter the VSAN information.

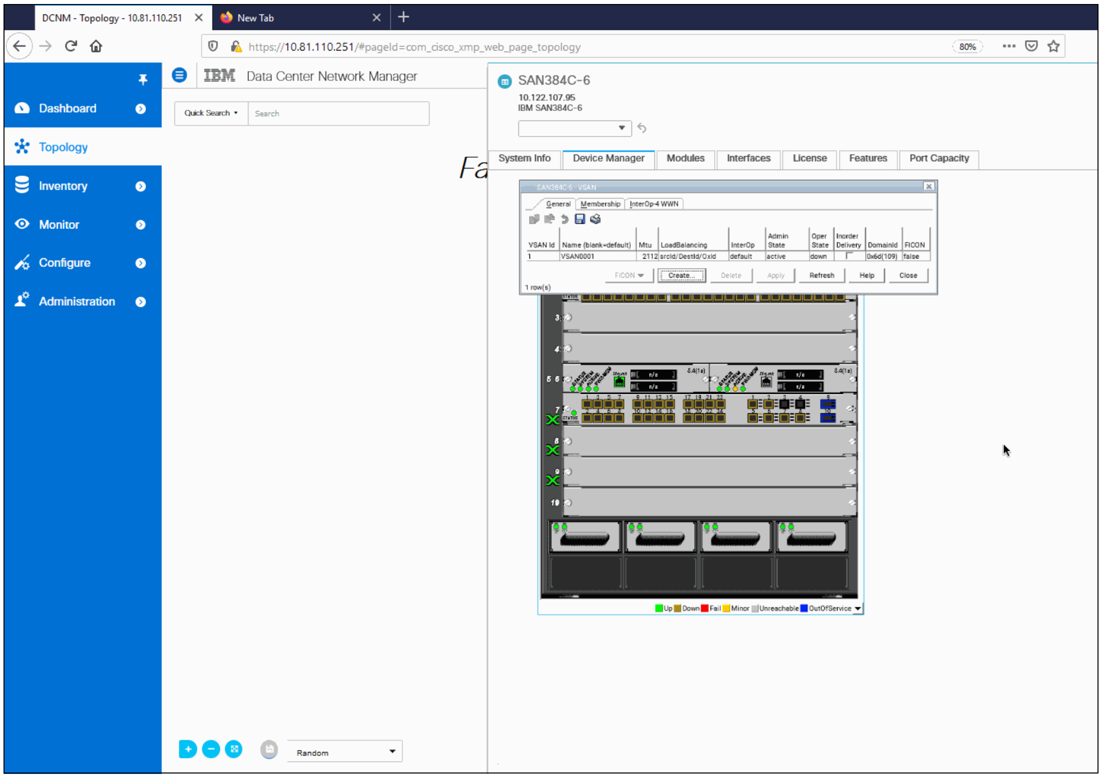

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false

Figure 7-60 VSAN window

Figure 7-61 shows the Create VSAN window. To define a FICON VSAN, select **FICON**, which changes the defaults to match the FICON VSAN characteristics.

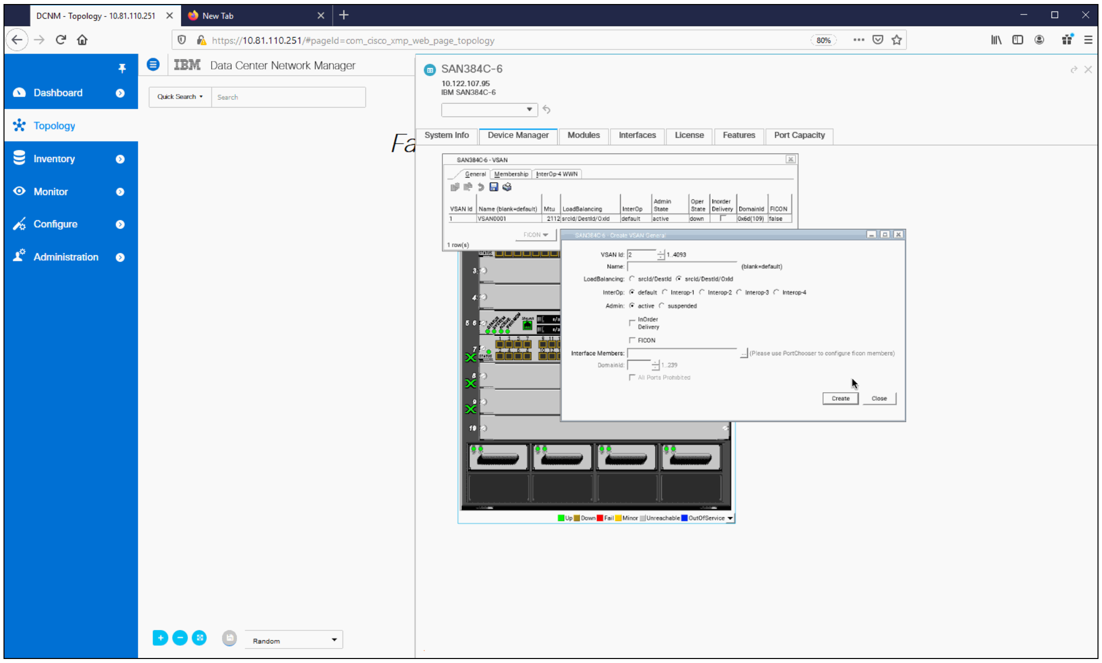

VSAN Id: 2 1..4093

Name: (blank=default)

LoadBalancing: ○ srcId/DestId ● srcId/DestId/OxId

InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4

Admin: ● active ○ suspended

□ InOrder Delivery

□ FICON

Interface Members: ... (Please use PortChooser to configure ficon members)

DomainId: 1..239

□ All Ports Prohibited

Figure 7-61 Create VSAN window

Figure 7-62 shows the default FICON parameters. In our example, we code the domain ID in hex by preceding it with a 0x. The domain ID is required to match the switch ID in the mainframe hardware configuration. In our sample environment, this ID is 0x20.

Provide the VSAN ID, name, and the domain ID, and click **Create**.

> **Note:** The VSAN ID and the domain ID are not related.

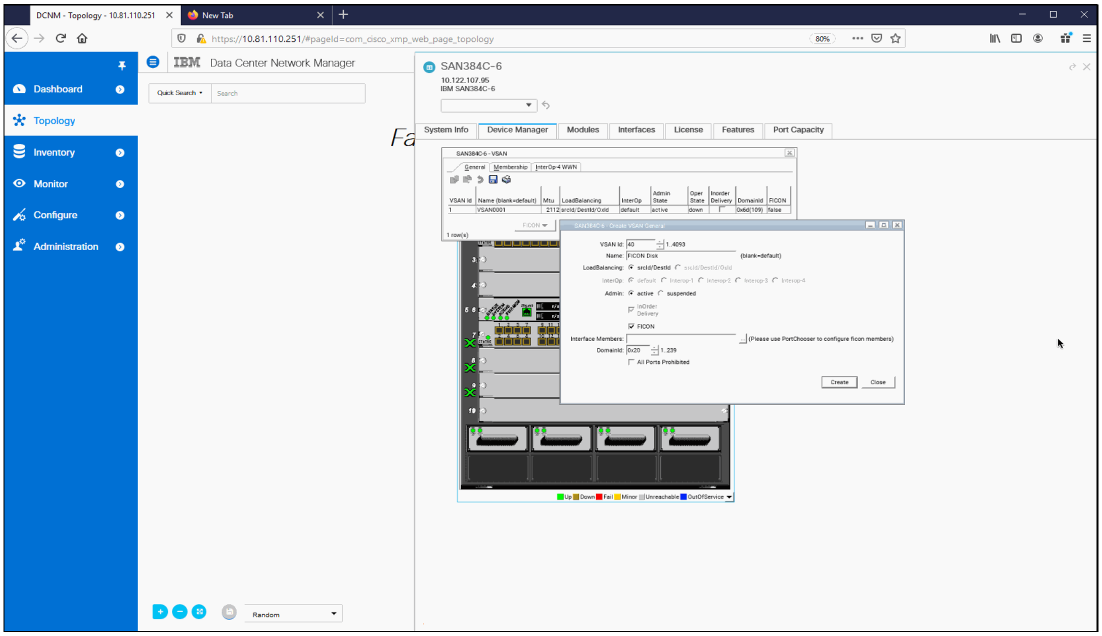

SAN384C-6 - VSAN

General | Membership | InterOp-4 WWN

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false

FICON

1 row(s)

SAN384C-6 - Create VSAN General

VSAN Id: 40 1..4093

Name: FICON Disk (blank=default)

LoadBalancing: ● srcId/DestId ○ srcId/DestId/OxId

InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4

Admin: ● active ○ suspended

☑ InOrder Delivery

☑ FICON

Interface Members: ... (Please use PortChooser to configure ficon members)

DomainId: 0x20 1..239

□ All Ports Prohibited

Create Close

Figure 7-62 Creating a VSAN

In Figure 7-63 on page 297, we validate that all required FICON parameters will be applied upon VSAN creation. This action is disruptive only when a VSAN is under modification. Click **Yes** to continue.

> **Note:** The VSAN disruption message refers to the VSAN under modification. Other VSANs are not impacted.

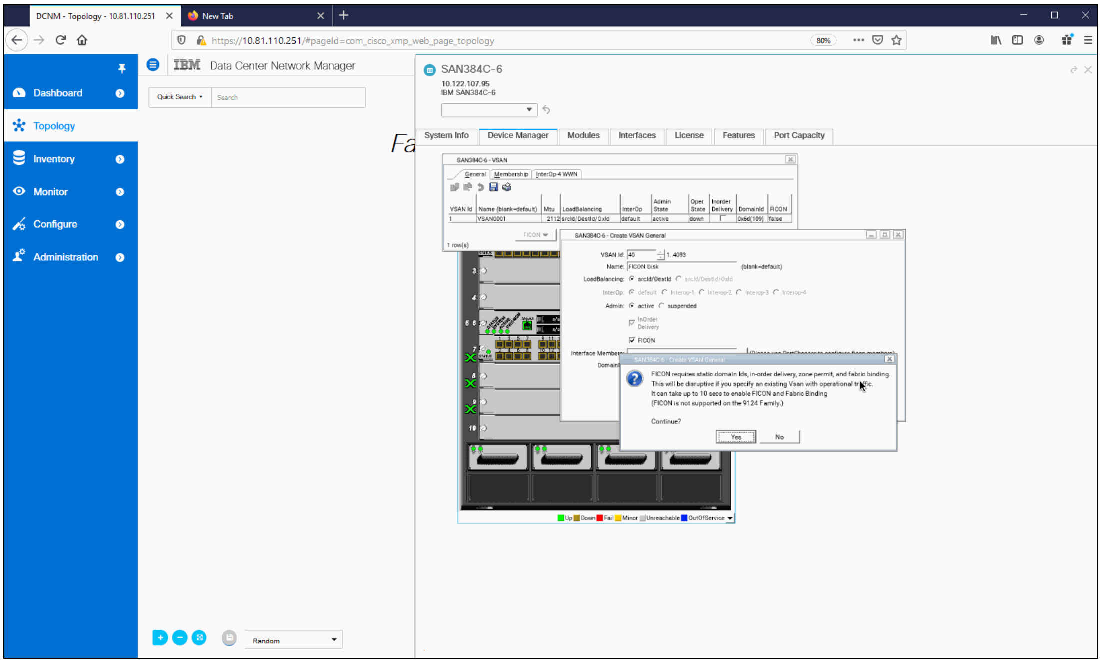

SAN384C-6

10.122.107.95
IBM SAN384C-6

SAN384C-6 - VSAN

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false

1 row(s)

SAN384C-6 - Create VSAN General

VSAN Id: 40 1..4093
Name: FICON Disk (blank=default)
LoadBalancing: ● srcId/DestId ○ srcId/DestId/OxId
InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4
Admin: ● active ○ suspended
☑ InOrder Delivery
☑ FICON

SAN384C-6 - Create VSAN General

FICON requires static domain Ids, in-order delivery, zone permit, and fabric binding.
This will be disruptive if you specify an existing Vsan with operational traffic.
It can take up to 10 secs to enable FICON and Fabric Binding
(FICON is not supported on the 9124 Family.)

Continue?

Yes No

Figure 7-63 VSAN confirmation

Figure 7-64 shows that the VSAN was created successfully. We can use the same window to create more FICON VSANs.

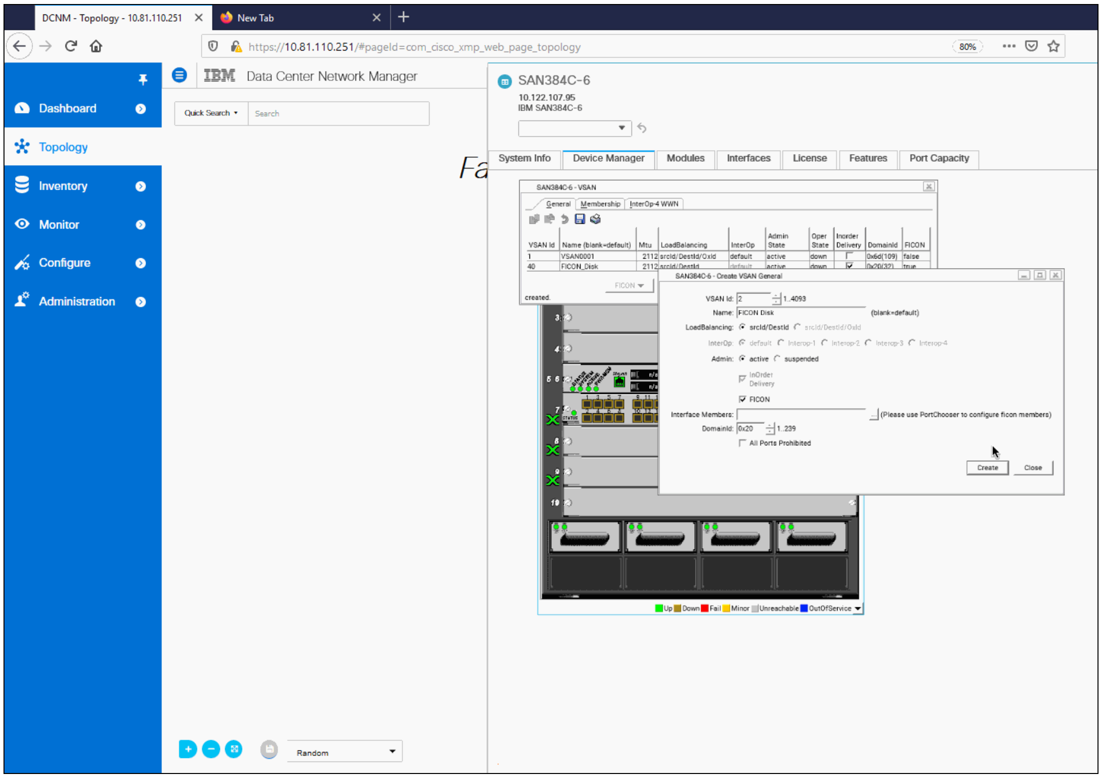

SAN384C-6

10.122.107.95
IBM SAN384C-6

SAN384C-6 - VSAN

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false
40	FICON_Disk	2112	srcId/DestId	default	active	down	☑	0x20(32)	true

created.

SAN384C-6 - Create VSAN General

VSAN Id: 2 1..4093
Name: FICON Disk (blank=default)
LoadBalancing: ● srcId/DestId ○ srcId/DestId/OxId
InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4
Admin: ● active ○ suspended
☑ InOrder Delivery
☑ FICON
Interface Members: ... (Please use PortChooser to configure ficon members)
DomainId: 0x20 1..239
□ All Ports Prohibited

Create Close

Figure 7-64 VSAN created

Figure 7-65 shows the creation of an additional VSAN for the lab environment.

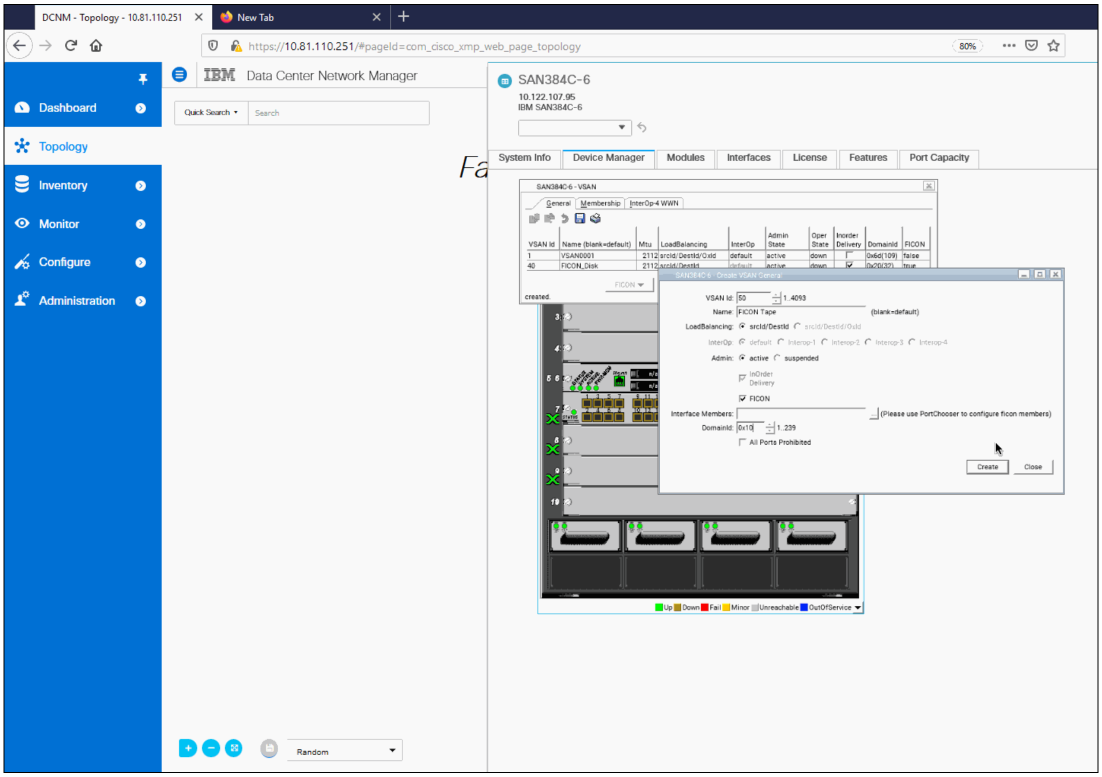

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false
40	FICON_Disk	2112	srcId/DestId	default	active	down	☑	0x20(32)	true

VSAN Id: 50 1..4093

Name: FICON Tape (blank=default)

LoadBalancing: ● srcId/DestId ○ srcId/DestId/OxId

InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4

Admin: ● active ○ suspended

☑ InOrder Delivery

☑ FICON

Interface Members: ... (Please use PortChooser to configure ficon members)

DomainId: 0x10 1..239

□ All Ports Prohibited

Create Close

Figure 7-65 VSAN created

Repeat the steps to create additional VSANs, as shown in Figure 7-66.

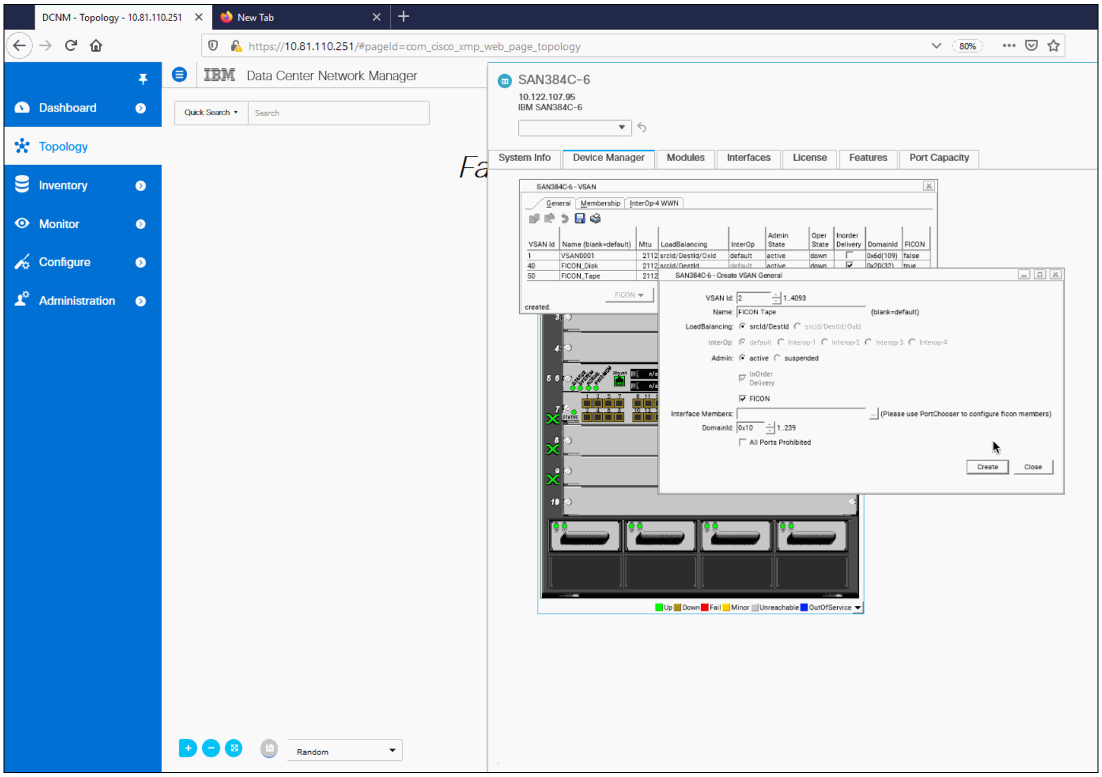

SAN384C-6 - VSAN

General | Membership | InterOp-4 WWN

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false
40	FICON_Disk	2112	srcId/DestId	default	active	down	☑	0x20(32)	true
50	FICON_Tape	2112							

FICON

created.

SAN384C-6 - Create VSAN General

VSAN Id: 2 1..4093

Name: FICON Tape (blank=default)

LoadBalancing: ● srcId/DestId ○ srcId/DestId/OxId

InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4

Admin: ● active ○ suspended

☑ InOrder Delivery

☑ FICON

Interface Members: ... (Please use PortChooser to configure ficon members)

DomainId: 0x10 1..239

□ All Ports Prohibited

Create Close

Figure 7-66 VSANs created

Figure 7-67 shows two new FICON VSANs that we created for our lab environment. Both VSANs were created with source and destination (SRCID-DESTID) based load balancing.

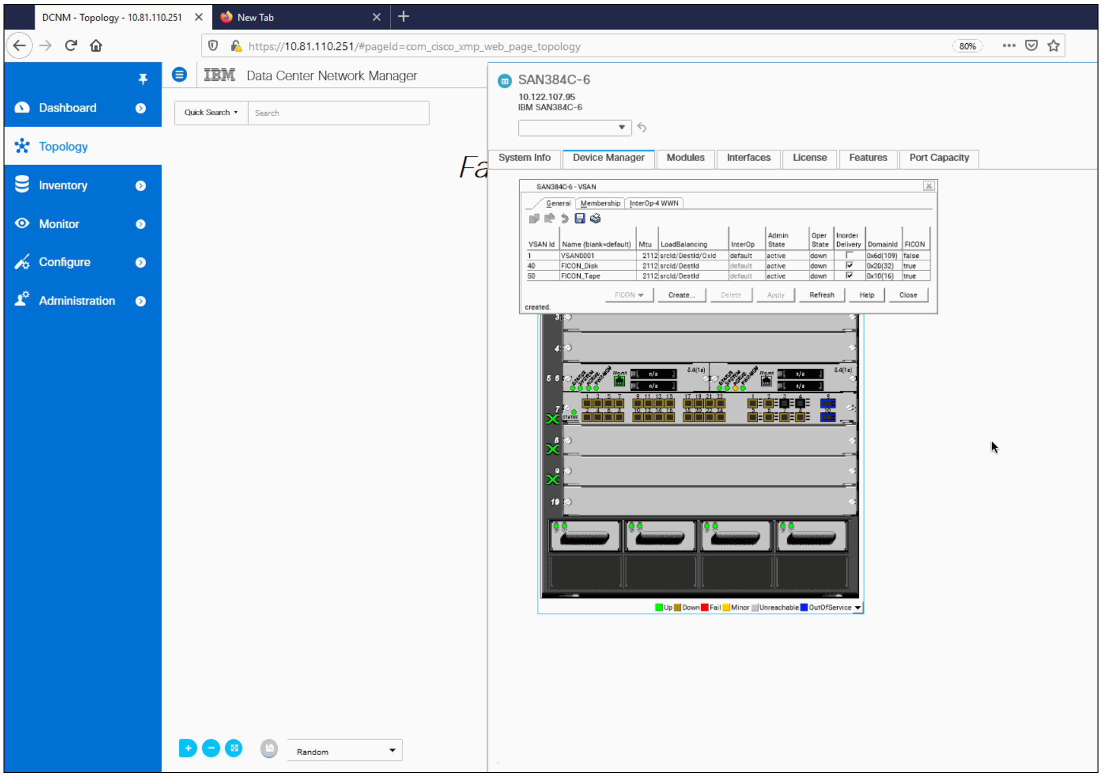

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false
40	FICON_Disk	2112	srcId/DestId	default	active	down	☑	0x20(32)	true
50	FICON_Tape	2112	srcId/DestId	default	active	down	☑	0x10(16)	true

created.

Figure 7-67 VSANs created

Figure 7-68 shows the creation of an open systems VSAN. For open systems VSANs, the domain ID is usually determined dynamically by the switch. Provide the VSAN ID and name, and then click **Create**.

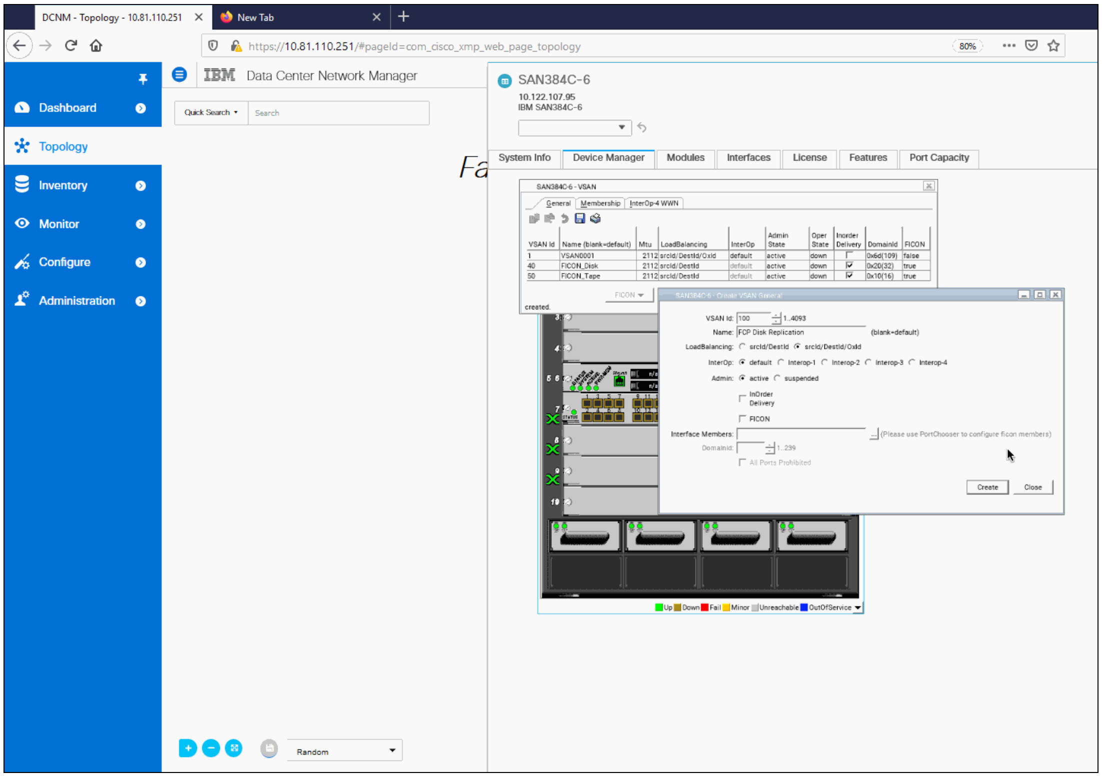

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false
40	FICON_Disk	2112	srcId/DestId	default	active	down	☑	0x20(32)	true
50	FICON_Tape	2112	srcId/DestId	default	active	down	☑	0x10(16)	true

created.

SAN384C-6 - Create VSAN General

VSAN Id: 100 1..4093
Name: FCP Disk Replication (blank=default)
LoadBalancing: ○ srcId/DestId ● srcId/DestId/OxId
InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4
Admin: ● active ○ suspended
□ InOrder Delivery
□ FICON
Interface Members: ... (Please use PortChooser to configure ficon members)
DomainId: 1..239
□ All Ports Prohibited

Create Close

Figure 7-68 Open systems VSAN creation

Figure 7-69 shows that the open systems VSAN was created with exchange-based load balancing and In Order Delivery (IOD) turned off. The statuses of the FICON VSANs are shown as `true` in the right column. Now, all three VSANs for the lab environment are created.

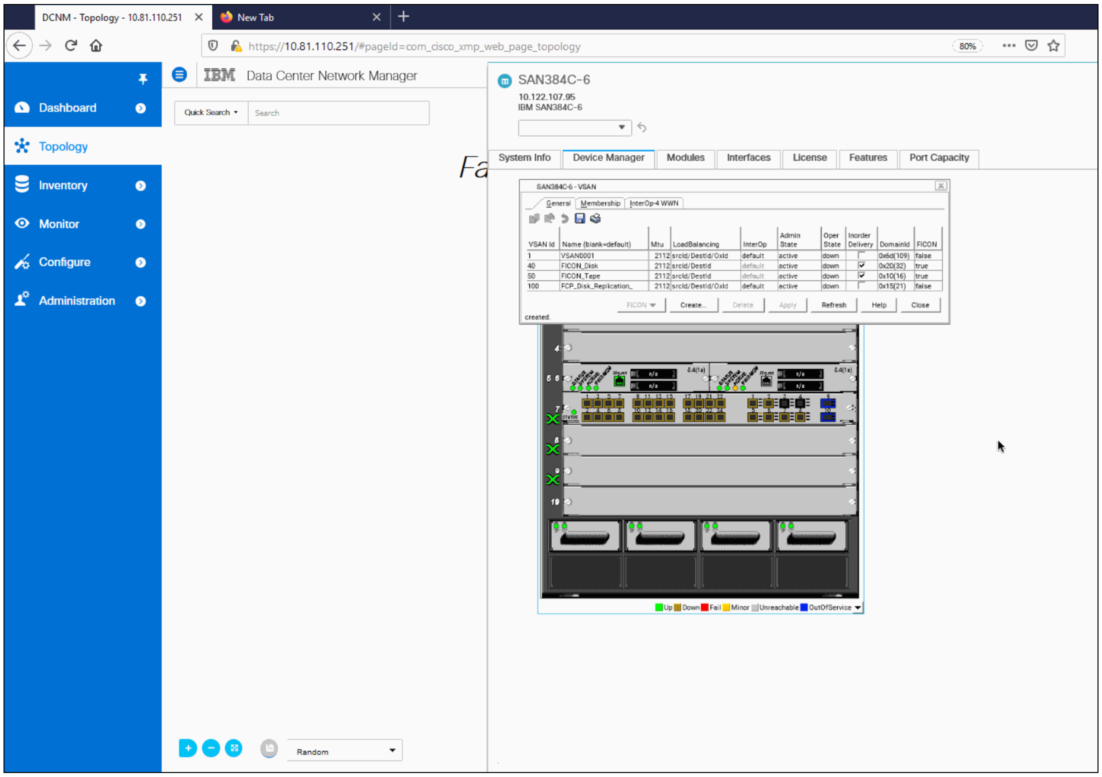

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	□	0x6d(109)	false
40	FICON_Disk	2112	srcId/DestId	default	active	down	☑	0x20(32)	true
50	FICON_Tape	2112	srcId/DestId	default	active	down	☑	0x10(16)	true
100	FCP_Disk_Replication_	2112	srcId/DestId/OxId	default	active	down	□	0x15(21)	false

created.

Figure 7-69 Open systems VSAN created

Figure 7-70 on page 303 shows the VSAN creation for our lab environment on the other switch. The VSAN IDs are the same, but the domain IDs are different on this switch because the VSAN IDs are forming a fabric with the two switches based on the unique domain IDs in that fabric.

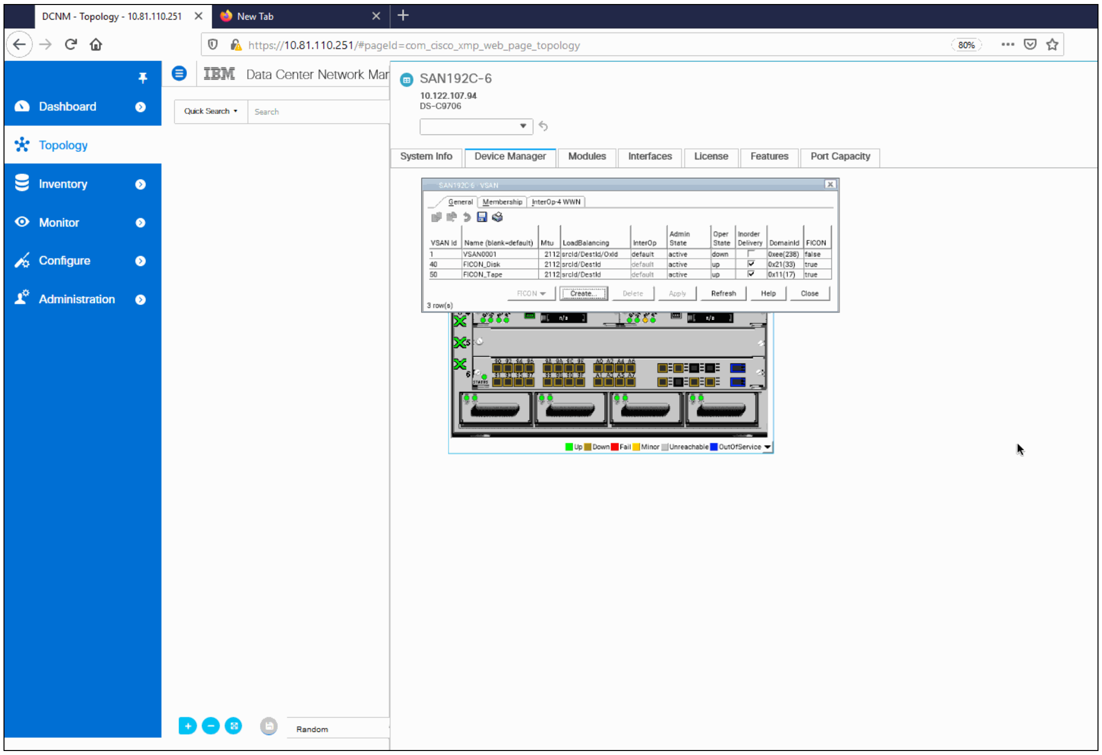

SAN192C-6

10.122.107.94
DS-C9706

VSAN Id	Name (blank=default)	Mtu	LoadBalancing	InterOp	Admin State	Oper State	Inorder Delivery	DomainId	FICON
1	VSAN0001	2112	srcId/DestId/OxId	default	active	down	☐	0xee(238)	false
40	FICON_Disk	2112	srcId/DestId	default	active	up	☑	0x21(33)	true
50	FICON_Tape	2112	srcId/DestId	default	active	up	☑	0x11(17)	true

3 row(s)

Figure 7-70 VSANs on the other switch

Configuring VSANs by using the CLI

In this section, we show how to create a VSAN by using the CLI, which provides traffic isolation, dedicated fabric services, and support for up to 256 VSANs in a switch, as shown in Example 7-2. VSAN IDs range from 2 to 4093 with the default being VSAN 1 along with an additional isolated VSAN (VSAN 4094).

> **Important:** Never use VSAN 1 for FICON or open systems devices. As a best practice, create custom VSANs when configuring these types of environments.

Example 7-2 Configuring VSANs by using the CLI

```
SAN384C-6#
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# vsan database
SAN384C-6(config-vsan-db)# vsan 40 name "FICON_Disk" loadbalancing src-dst-id
SAN384C-6(config-vsan-db)#
SAN384C-6(config-vsan-db)# show vsan 40
vsan 40 information
         name:FICON_Disk  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id
         operational state:down

SAN384C-6(config-vsan-db)# fcdroplatency network 500 vsan 40
SAN384C-6(config)# in-order-guarantee vsan 40
SAN384C-6(config)# fcdomain domain 32 static vsan 40
SAN384C-6(config)# zone default-zone permit vsan 40
```

```
SAN384C-6(config)#
SAN384C-6(config)# fabric-binding database vsan 40
SAN384C-6(config-fabric-binding)# fabric-binding activate vsan 40 force
SAN384C-6(config)#
SAN384C-6(config)# ficon vsan 40
SAN384C-6(config-ficon)#
SAN384C-6(config-ficon)# show vsan
vsan 1 information
         name:VSAN0001  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:down

vsan 40 information
         name:FICON_Disk  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id
         operational state:up

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

SAN384C-6(config-ficon)# show ficon vsan 40

FICON information for VSAN 40
  FICON is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Disabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Enabled
  Loadbalance is srcid-dstid
  Number of implemented ports are 254
  Key Counter is 0
  FCID last byte is 0(0)
  Serial number is 02.003A9C316288
  Date/Time is same as system time (Thu Mar 5 05:20:58.698909 2020)
  Device Allegiance not locked
  Codepage is us-canada
  Saved configuration files
    IPL

SAN384C-6(config-ficon)# end
Performing fast copy of configurationdone.
SAN384C-6#
```

7.2.10 Configuring FICON fabric security (Fabric Binding)

Before we can define the links between the switches, we must configure permissions for which switches are allowed to communicate with each other. This task is done on a per VSAN basis. The process defines a database that is called the Fabric Binding database, which must be the same on every switch in the VSAN. In this database, we define each switch and its associated domain ID that is allowed to pass traffic within the VSAN. The switch worldwide name (sWWN) is used to uniquely identify each switch in the VSAN.

Starting at Device Manager, select **Security** → **Fabric Binding** for the IBM Storage Networking SAN192C-6 switch, as shown in Figure 7-71.

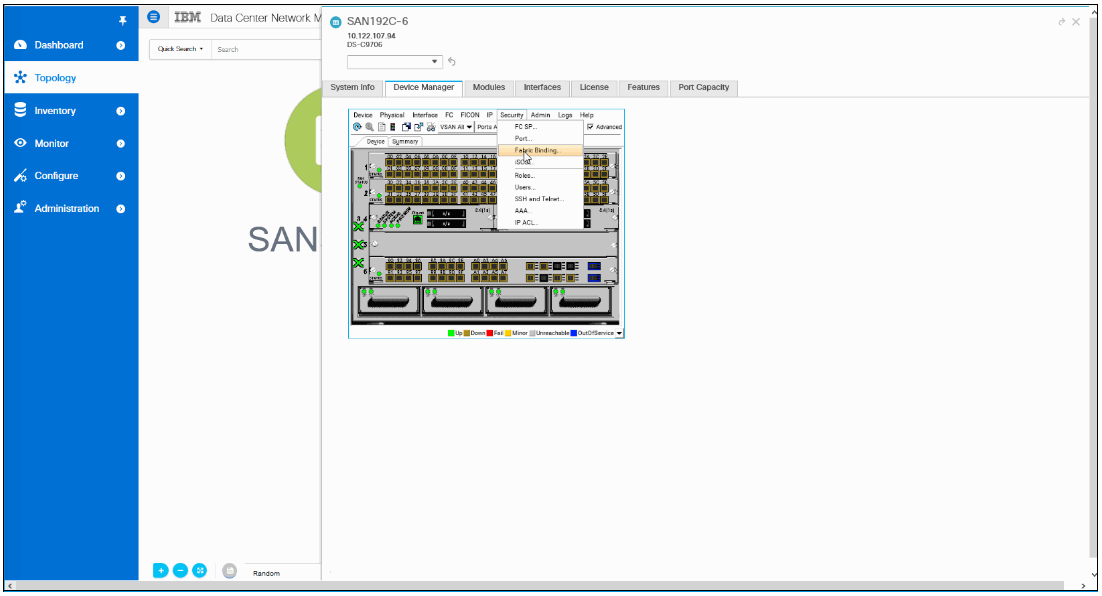

Figure 7-71 Fabric Binding

Figure 7-72 shows the Fabric Binding window. On the Config Database tab, you see the Local Switch worldwide name (WWN) for the IBM Storage Networking SAN192C-6 switch. Click **Create**.

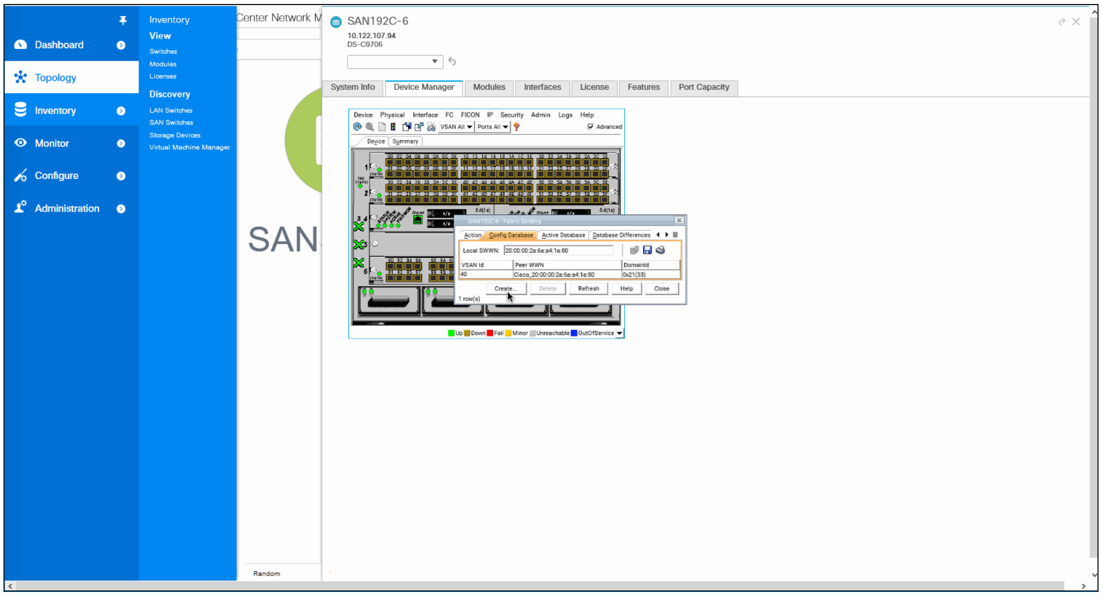

Local SWWN: 20:00:00:2a:6a:a4:1a:80

VSAN Id	Peer WWN	DomainId
40	Cisco_20:00:00:2a:6a:a4:1a:80	0x21(33)

1 row(s)

Figure 7-72 Config Databases tab

Figure 7-73 shows the Create Fabric Binding Config Database window. To populate this window, we open the Fabric Binding window on the IBM Storage Networking SAN384C-6 peer switch so that we can copy its local sWWN.

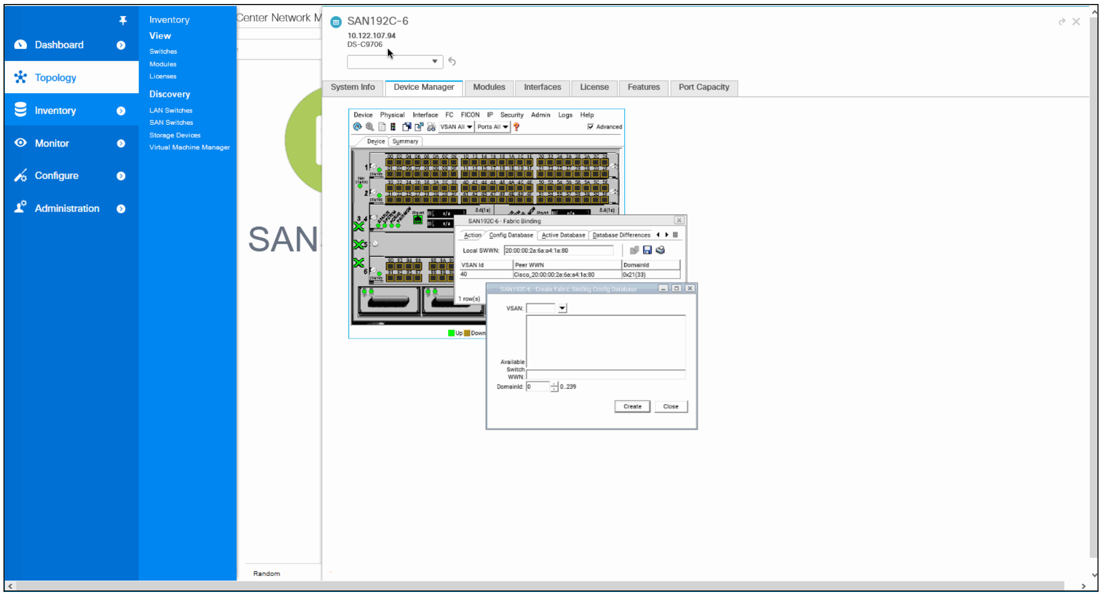

Local SWWN: 20:00:00:2a:6a:a4:1a:80

VSAN Id	Peer WWN	DomainId
40	Cisco_20:00:00:2a:6a:a4:1a:80	0x21(33)

1 row(s)

DomainId: 0

Figure 7-73 Create Fabric Binding Config Database

Figure 7-74 on page 307 shows the Fabric Binding Database window for the IBM Storage Networking SAN384C-6 switch. Capture the local sWWN information, which will be used to create the Database entry on the IBM Storage Networking SAN192C-6 switch.

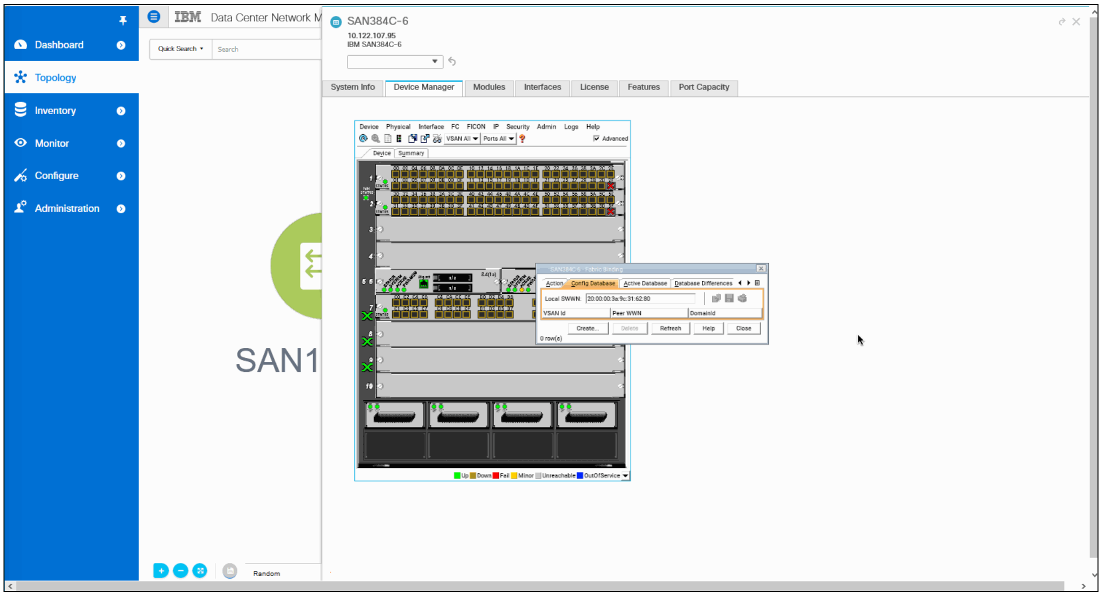

Local SWWN: 20:00:00:3a:9c:31:62:80

VSAN Id	Peer WWN	DomainId

0 row(s)

Figure 7-74 Fabric Binding creation

Select **VSAN 40** from the drop-down menu, enter the domain ID 0x20, and paste the local sWWN for the IBM Storage Networking SAN384C-6 peer switch that we captured. Click **Create**, as shown in Figure 7-75.

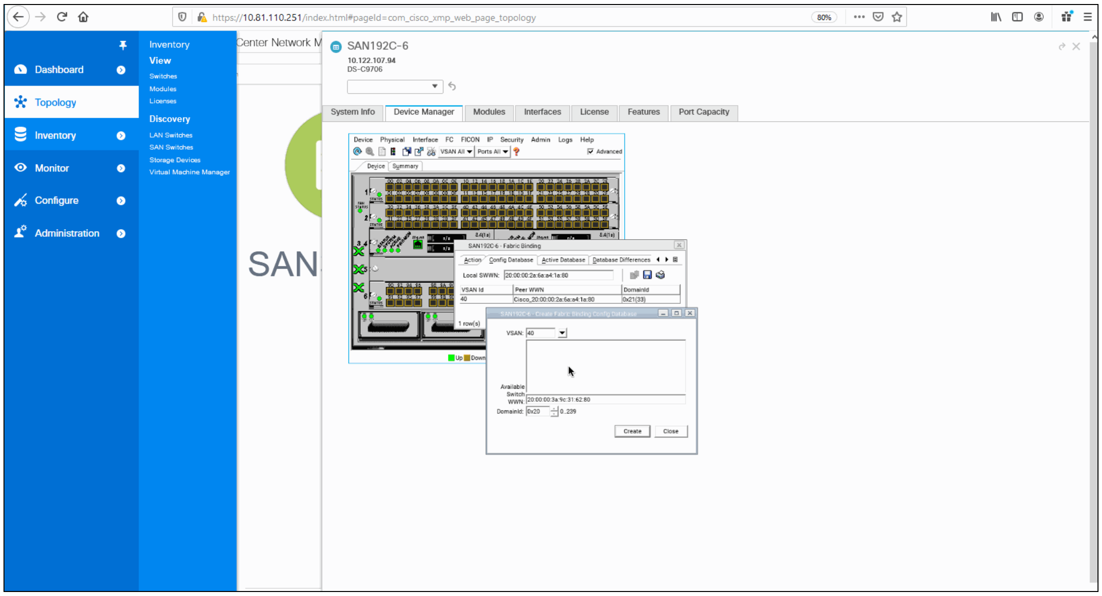

Local SWWN: 20:00:00:2a:6a:a4:1a:80

VSAN Id	Peer WWN	DomainId
40	Cisco_20:00:00:2a:6a:a4:1a:80	0x21(33)

1 row(s)

VSAN: 40

Available Switch WWN: 20:00:00:3a:9c:31:62:80

DomainId: 0x20 0..239

Figure 7-75 Fabric Binding database creation

Figure 7-76 shows the entries for local and remote switches in VSAN 40. Now, we capture the local sWWN information that will be used to create the Database entry on the IBM Storage Networking SAN384C-6 peer switch.

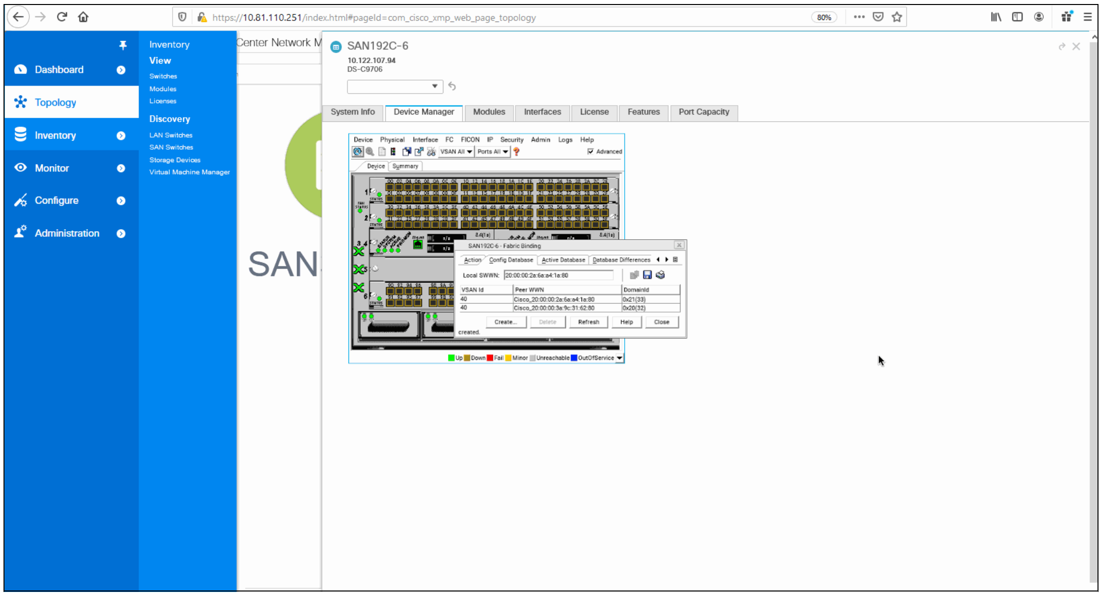

SAN192C-6 - Fabric Binding

Action | Config Database | Active Database | Database Differences

Local SWWN: 20:00:00:2a:6a:a4:1a:80

VSAN Id	Peer WWN	DomainId
40	Cisco_20:00:00:2a:6a:a4:1a:80	0x21(33)
40	Cisco_20:00:00:3a:9c:31:62:80	0x20(32)

Create... | Delete | Refresh | Help | Close

created.

Figure 7-76   Database entry created

Select **VSAN 40** from the drop-down menu, enter the domain ID 0x21, and paste the local sWWN for the IBM Storage Networking SAN192C-6 peer switch. Click **Create**, as shown in Figure 7-77.

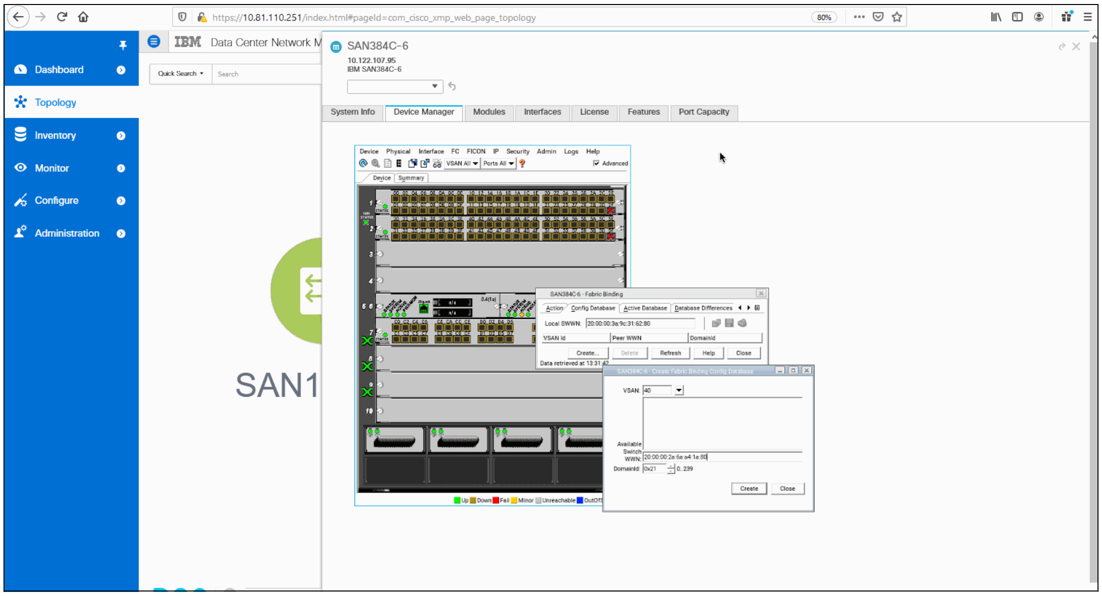

SAN384C-6 - Fabric Binding

Action | Config Database | Active Database | Database Differences

Local SWWN: 20:00:00:3a:9c:31:62:80

VSAN Id	Peer WWN	DomainId

Create... | Delete | Refresh | Help | Close

Data retrieved at 13:31:42

SAN384C-6 - Create Fabric Binding Config Database

VSAN: 40

Available Switch WWN: 20:00:00:2a:6a:a4:1a:80

DomainId: 0x21 0..239

Create | Close

Figure 7-77   Fabric Binding database creation

Figure 7-78 on page 309 shows the entries for local and remote switches in VSAN 40.

> **Important:** It is critical to validate that the Fabric Binding database is the same on each switch. If there is a mismatch, the FICON VSAN will not be allowed to become active between the switches.

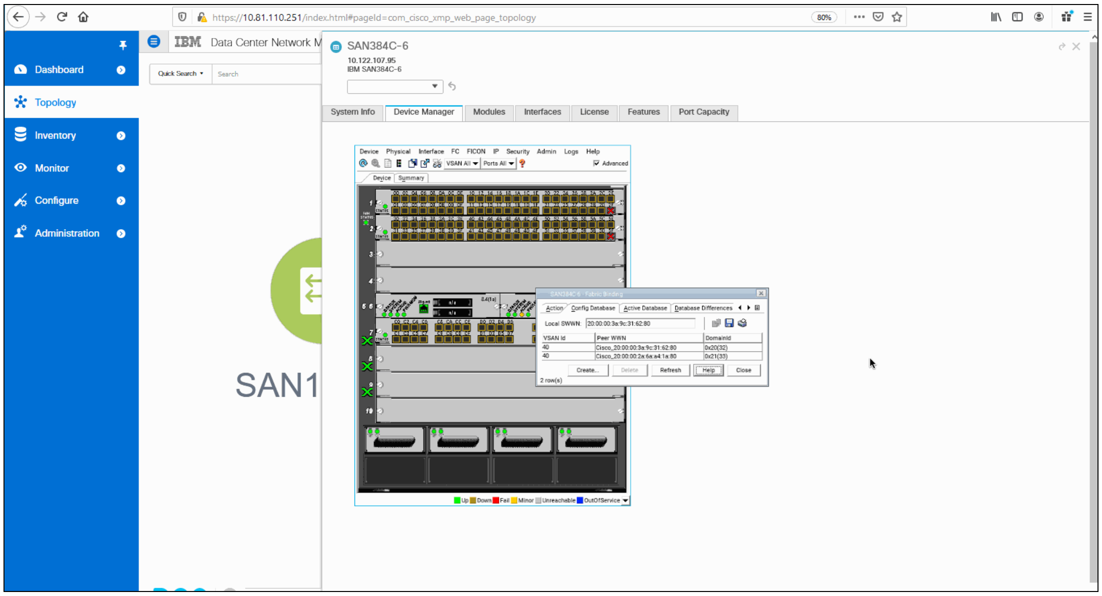

Local SWWN: 20:00:00:3a:9c:31:62:80

VSAN Id	Peer WWN	DomainId
40	Cisco_20:00:00:3a:9c:31:62:80	0x20(32)
40	Cisco_20:00:00:2a:6a:a4:1a:80	0x21(33)

2 row(s)

Figure 7-78   Database entry created

On the **Action** tab, select **ForceActivate** from the drop-down menu for VSAN 40, as shown in Figure 7-79. This action activates the Fabric Binding database. Click **Apply**. This step must be repeated on the IBM Storage Networking SAN284C-6 peer switch.

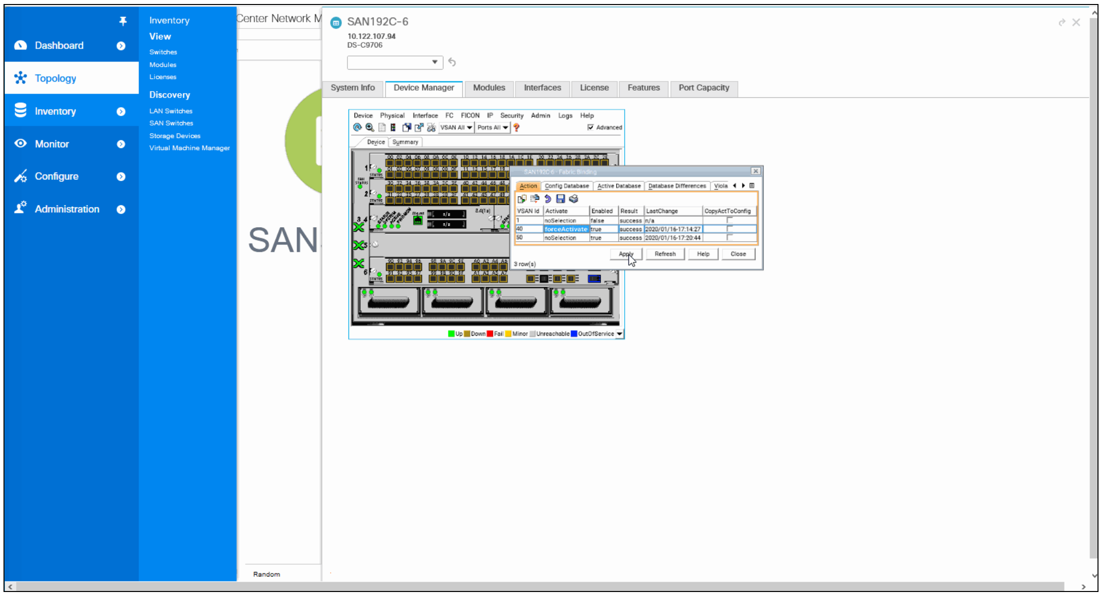

VSAN Id	Activate	Enabled	Result	LastChange	CopyActToConfig
1	noSelection	false	success	n/a	□
40	forceActivate	true	success	2020/01/16-17:14:27	□
50	noSelection	true	success	2020/01/16-17:20:44	□

3 row(s)

Figure 7-79   Database activation

Figure 7-80 shows that the Fabric Binding database was activated successfully.

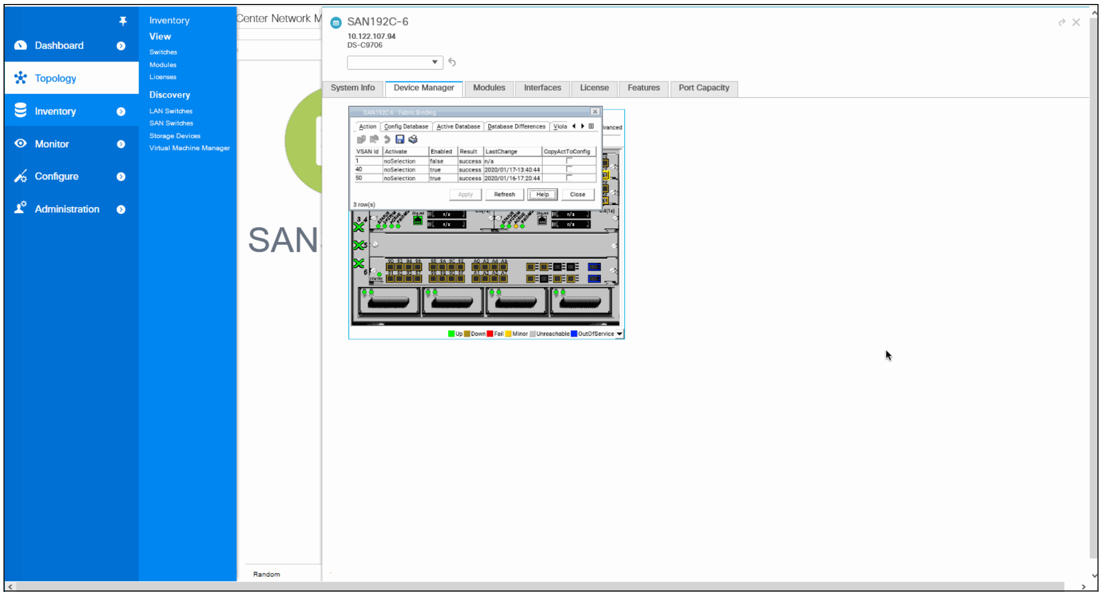

VSAN Id	Activate	Enabled	Result	LastChange	CopyActToConfig
1	noSelection	false	success	n/a	□
40	noSelection	true	success	2020/01/17-13:40:44	□
50	noSelection	true	success	2020/01/16-17:20:44	□

3 row(s)

Figure 7-80 Database activated

Configuring Fabric Binding by using the CLI

Before defining the links between switches, we must configure permissions for which switches are allowed to communicate with each other on a per VSAN basis. We define each switch and its associated domain ID that is allowed to pass traffic within the VSAN by using the CLI, as shown in Example 7-3.

Example 7-3 Configuring Fabric Binding by using the CLI

```
SAN384C-6#
SAN384C-6# show wwn switch
Switch WWN is 20:00:00:3a:9c:31:62:80
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# fabric-binding database vsan 40
SAN384C-6(config-fabric-binding)# swwn 20:00:00:3a:9c:31:62:80 domain 0x20
SAN384C-6(config-fabric-binding)# swwn 20:00:00:2a:6a:a4:1a:80 domain 0x21
SAN384C-6(config-fabric-binding)#
SAN384C-6(config-fabric-binding)# fabric-binding activate vsan 40 force
SAN384C-6(config)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6# show fabric-binding database vsan 40
--------------------------------------------------
Vsan   Logging-in Switch WWN     Domain-id
--------------------------------------------------
40     20:00:00:3a:9c:31:62:80    0x20(32) [Local]
40     20:00:00:2a:6a:a4:1a:80    0x21(33)
[Total 2 entries]
SAN384C-6#
SAN384C-6# show fabric-binding database active vsan 40
--------------------------------------------------
Vsan   Logging-in Switch WWN     Domain-id
--------------------------------------------------
```

```
40     20:00:00:3a:9c:31:62:80    0x20(32) [Local]
40     20:00:00:2a:6a:a4:1a:80    0x21(33)
[Total 2 entries]
SAN384C-6#

SAN192C-6#
SAN192C-6# show wwn switch
Switch WWN is 20:00:00:2a:6a:a4:1a:80
SAN192C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# fabric-binding database vsan 40
SAN192C-6(config-fabric-binding)# swwn 20:00:00:2a:6a:a4:1a:80 domain 0x21
SAN192C-6(config-fabric-binding)# swwn 20:00:00:3a:9c:31:62:80 domain 0x20
SAN192C-6(config-fabric-binding)#
SAN192C-6(config-fabric-binding)# fabric-binding activate vsan 40 force
SAN192C-6(config)# end
Performing fast copy of configurationdone.
SAN192C-6# show fabric-binding database vsan 40
--------------------------------------------------
Vsan   Logging-in Switch WWN     Domain-id
--------------------------------------------------
40     20:00:00:2a:6a:a4:1a:80    0x21(33) [Local]
40     20:00:00:3a:9c:31:62:80    0x20(32)
[Total 2 entries]
SAN192C-6# show fabric-binding database active vsan 40
--------------------------------------------------
Vsan   Logging-in Switch WWN     Domain-id
--------------------------------------------------
40     20:00:00:2a:6a:a4:1a:80    0x21(33) [Local]
40     20:00:00:3a:9c:31:62:80    0x20(32)
[Total 2 entries]
SAN192C-6#
```

7.2.11 Configuring Inter-Switch Links by using the DCNM

FICON links are ISLs that transport FC control and data frames between switches.

Figure 7-81 shows the Topology view of the two lab switches with no FICON ISLs.

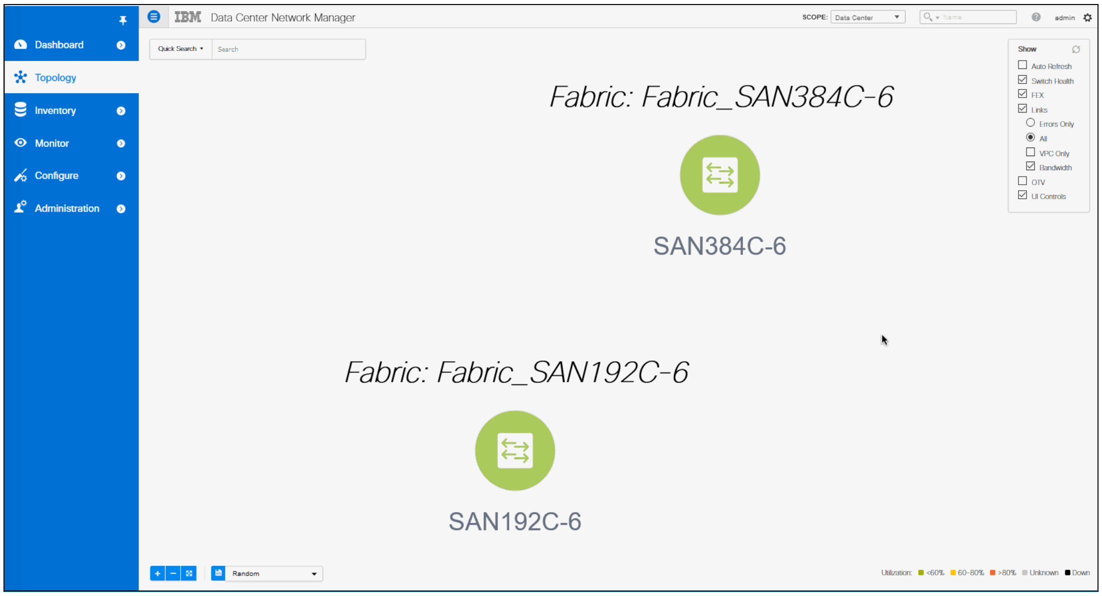

Figure 7-81 Topology view

Using Device Manager, as shown in Figure 7-82, double-click fc1/48 on the IBM Storage Networking SAN192C-6 switch to open the Interface window.

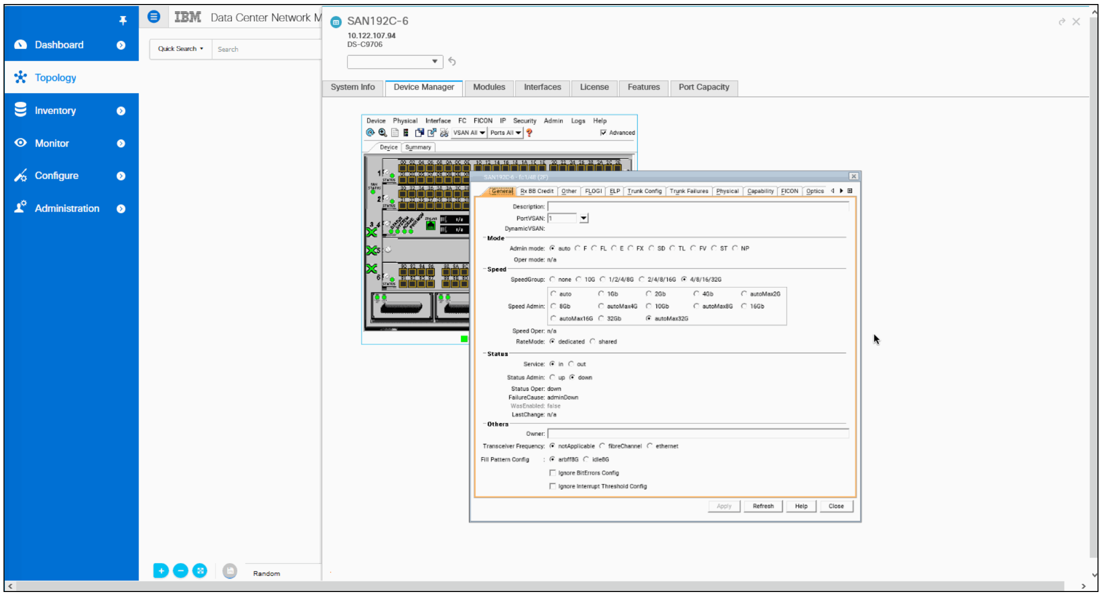

Description:
PortVSAN: 1
DynamicVSAN:

Mode
Admin mode: ● auto ○ F ○ FL ○ E ○ FX ○ SD ○ TL ○ FV ○ ST ○ NP
Oper mode: n/a

Speed
SpeedGroup: ○ none ○ 10G ○ 1/2/4/8G ○ 2/4/8/16G ● 4/8/16/32G
Speed Admin: ○ auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb ○ autoMax8G ○ 16Gb ○ autoMax16G ○ 32Gb ● autoMax32G
Speed Oper: n/a
RateMode: ● dedicated ○ shared

Status
Service: ● in ○ out
Status Admin: ○ up ● down
Status Oper: down
FailureCause: adminDown
WasEnabled: false
LastChange: n/a

Others
Owner:
Transceiver Frequency: ● notApplicable ○ fibreChannel ○ ethernet
Fill Pattern Config : ● arbff8G ○ idle8G
□ Ignore BitErrors Config
□ Ignore Interrupt Threshold Config

Figure 7-82 Interface window

When configuring an ISL interface, it is a best practice to provide a description that references both sides of the link. Change the **Mode** to E, **Speed** to 32 GbE, and **Admin Status** to Up, and click **Apply**, as shown in Figure 7-83 on page 313.

Note: It is not recommended to use auto-negotiate when configuring the ISL speed.

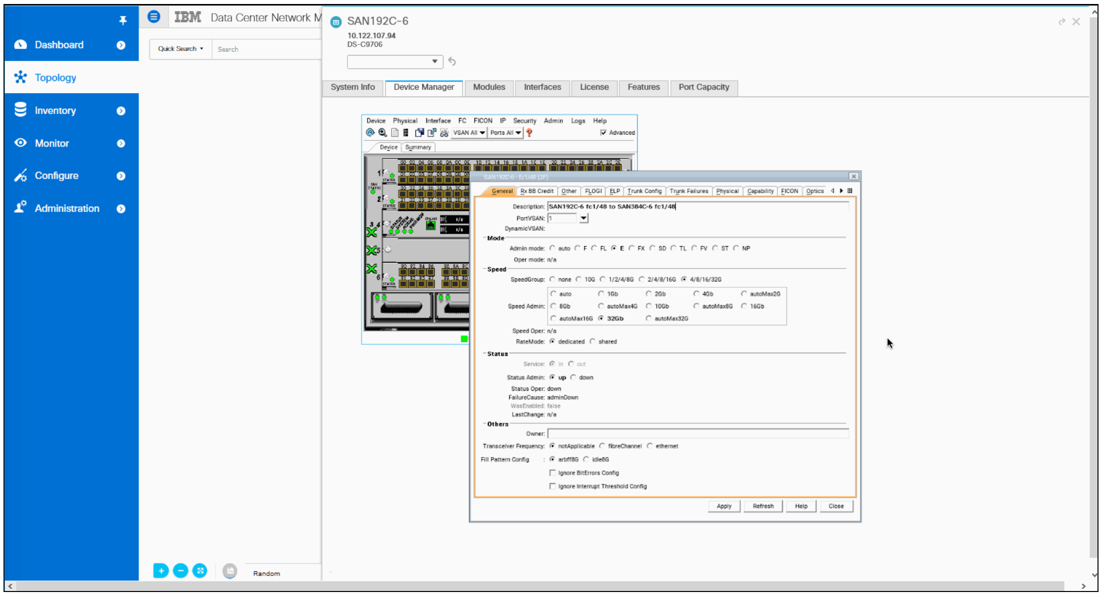

SAN192C-6 - fc1/48 (2F)

General | Rx BB Credit | Other | FLOGI | ELP | Trunk Config | Trunk Failures | Physical | Capability | FICON | Optics

Description: SAN192C-6 fc1/48 to SAN384C-6 fc1/48
PortVSAN: 1
DynamicVSAN:

Mode
Admin mode: ○ auto ○ F ○ FL ● E ○ FX ○ SD ○ TL ○ FV ○ ST ○ NP
Oper mode: n/a

Speed
SpeedGroup: ○ none ○ 10G ○ 1/2/4/8G ○ 2/4/8/16G ● 4/8/16/32G
Speed Admin: ○ auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb ○ autoMax8G ○ 16Gb ○ autoMax16G ● 32Gb ○ autoMax32G
Speed Oper: n/a
RateMode: ● dedicated ○ shared

Status
Service: ● in ○ out
Status Admin: ● up ○ down
Status Oper: down
FailureCause: adminDown
WasEnabled: false
LastChange: n/a

Others
Owner:
Transceiver Frequency: ● notApplicable ○ fibreChannel ○ ethernet
Fill Pattern Config : ● arbff8G ○ idle8G
□ Ignore BitErrors Config
□ Ignore Interrupt Threshold Config

Apply | Refresh | Help | Close

Figure 7-83 Defining ISL characteristics

We select the **Trunk Config** tab and populate the allowed VSANs for this ISL, in this case VSANs 1 and 40, as shown in Figure 7-84. Click **Apply**.

Note: As a best practice, trunk VSAN 1 in addition to the applicable FICON VSANs on all FICON ISLs.

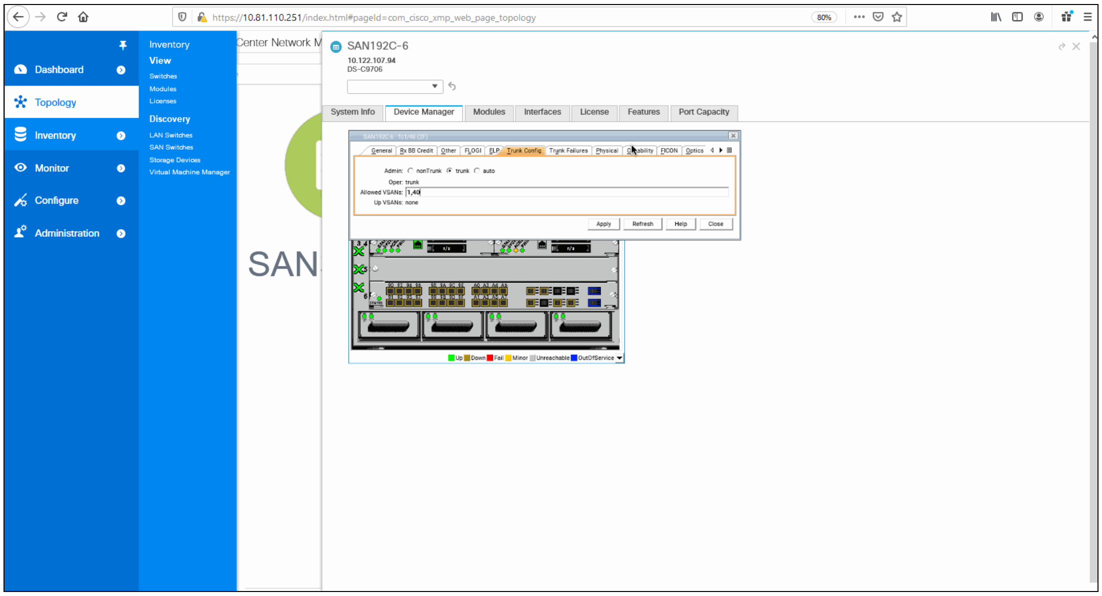

SAN192C-6 - fc1/48 (2F)

General | Rx BB Credit | Other | FLOGI | ELP | Trunk Config | Trunk Failures | Physical | Capability | FICON | Optics

Admin: ○ nonTrunk ● trunk ○ auto
Oper: trunk
Allowed VSANs: 1,40
Up VSANs: none

Apply | Refresh | Help | Close

Figure 7-84 ISL Trunk Config tab

Figure 7-85 shows how we configure the other end of the ISL on the IBM Storage Networking SAN384C-6 switch by using the same parameters as before. Click **Apply** and **Close**.

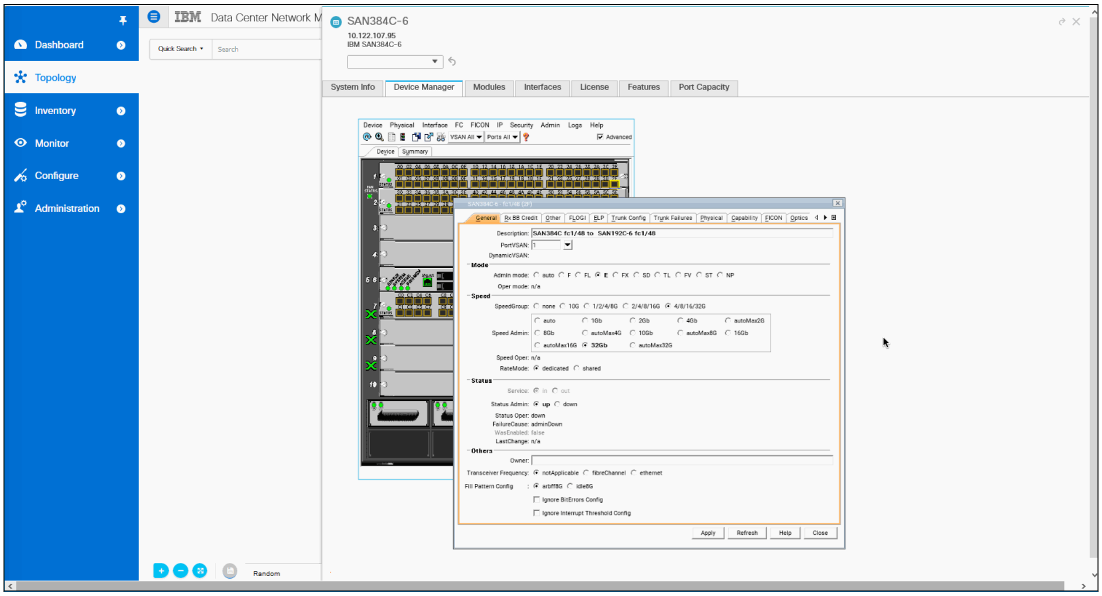

SAN384C-6 - fc1/48 (2F)

Description: SAN384C fc1/48 to SAN192C-6 fc1/48

PortVSAN: 1

DynamicVSAN:

Mode

Admin mode: ○ auto ○ F ○ FL ● E ○ FX ○ SD ○ TL ○ FV ○ ST ○ NP

Oper mode: n/a

Speed

SpeedGroup: ○ none ○ 10G ○ 1/2/4/8G ○ 2/4/8/16G ● 4/8/16/32G

Speed Admin: ○ auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb ○ autoMax8G ○ 16Gb ○ autoMax16G ● 32Gb ○ autoMax32G

Speed Oper: n/a

RateMode: ● dedicated ○ shared

Status

Service: ● in ○ out

Status Admin: ● up ○ down

Status Oper: down

FailureCause: adminDown

WasEnabled: false

LastChange: n/a

Others

Owner:

Transceiver Frequency: ● notApplicable ○ fibreChannel ○ ethernet

Fill Pattern Config : ● arbff8G ○ idle8G

□ Ignore BitErrors Config

□ Ignore Interrupt Threshold Config

Figure 7-85 Defining ISL characteristics

Now, we create a second ISL by using ports fc2/48 on each switch in the same manner. Figure 7-86 shows the ISLs coming online. When the port is double-clicked, you can see that VSAN 40 is UP.

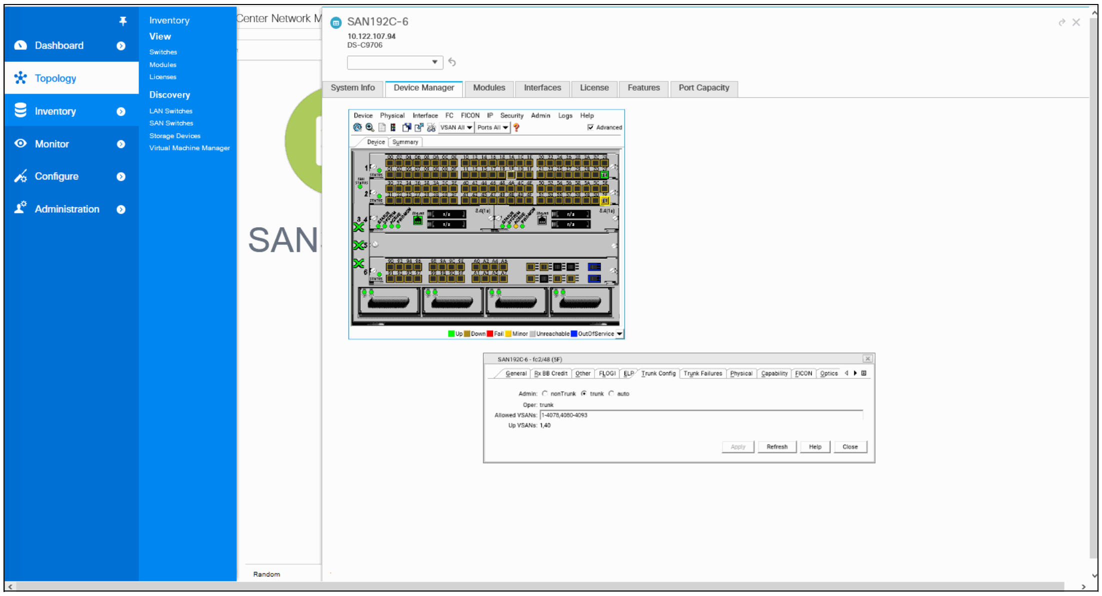

SAN192C-6 - fc2/48 (5F)

Admin: ○ nonTrunk ● trunk ○ auto

Oper: trunk

Allowed VSANs: 1-4078,4080-4093

Up VSANs: 1,40

Figure 7-86 VSAN 40 online

Figure 7-87 on page 315 shows the Topology view of the two ISLs between our lab switches.

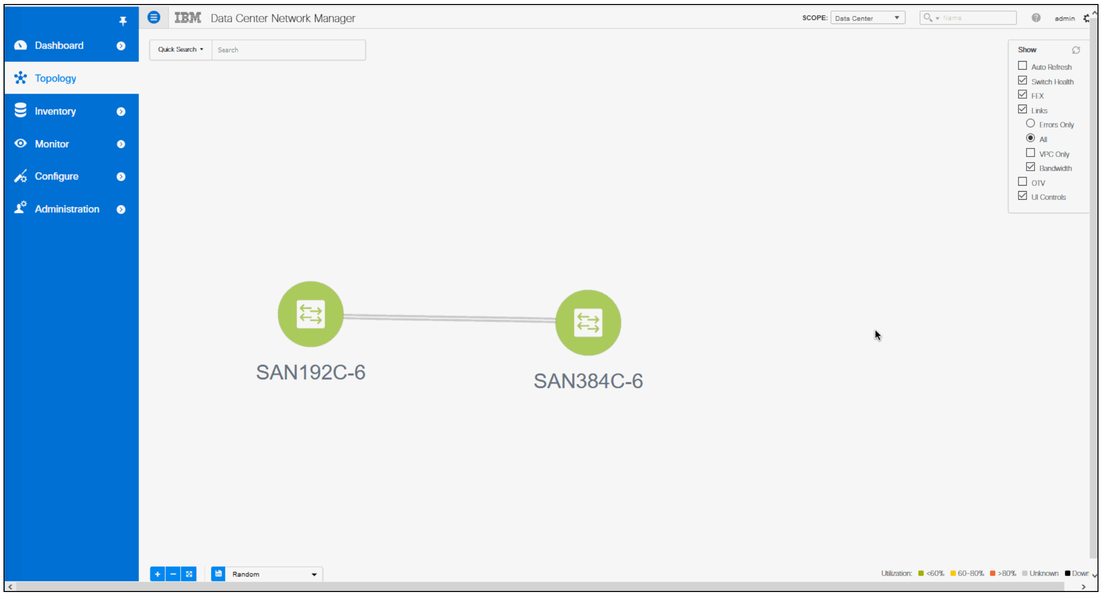

Figure 7-87 ISLs online

Configuring ISLs by using the CLI

Example 7-4 shows how to configure ISLs by using the CLI.

Example 7-4 Configuring ISLs by using the CLI

```
SAN384C-6#
SAN384C-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# interface fc1/48
SAN384C-6(config-if)# switchport mode E
SAN384C-6(config-if)# switchport speed 32000
SAN384C-6(config-if)#
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)#
SAN384C-6(config-if)# switchport trunk allowed vsan add 40
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6#
SAN384C-6# show interface fc1/48
fc1/48 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:30:00:3a:9c:31:62:80
    Peer port WWN is 20:30:00:2a:6a:a4:1a:80
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Admin Speed is 32 Gbps
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 500
    Receive B2B Credit is 500
    B2B State Change: Admin(on), Oper(up), Negotiated Value(14)
    Receive data field Size is 2112
    Beacon is turned off
    fec is enabled by default
    Logical type is core
    Trunk vsans (admin allowed and active) (1,40)
    Trunk vsans (up)                       (1,40)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 544 bits/sec,68 bytes/sec, 1 frames/sec
    5 minutes output rate 544 bits/sec,68 bytes/sec, 1 frames/sec
      1167 frames input,101488 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      1166 frames output,99680 bytes
        0 discards,0 errors
      2 input OLS,3  LRR,2 NOS,0 loop inits
      3 output OLS,1 LRR, 0 NOS, 0 loop inits
    500 receive B2B credit remaining
    500 transmit B2B credit remaining
    500 low priority transmit B2B credit remaining
    Interface last changed at Thu Mar  5 06:07:40 2020

    Last clearing of "show interface" counters :  never


SAN384C-6# show run interface fc1/48

!Command: show running-config interface fc1/48
!Running configuration last done at: Thu Mar  5 06:07:59 2020
!Time: Thu Mar  5 06:09:38 2020

version 8.4(1a)

interface fc1/48
  switchport speed 32000
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 40
  no shutdown

SAN384C-6#

<mdb other switch>

SAN192C-6#
SAN192C-6#
SAN192C-6# conf t
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# interface fc1/48
SAN192C-6(config-if)# switchport mode E
SAN192C-6(config-if)# switchport speed 32000
SAN192C-6(config-if)#
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)#
SAN192C-6(config-if)# switchport trunk allowed vsan add 40
SAN192C-6(config-if)#
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)# end
Performing fast copy of configurationdone.
SAN192C-6#
SAN192C-6# show interface fc1/48
fc1/48 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:30:00:2a:6a:a4:1a:80
    Peer port WWN is 20:30:00:3a:9c:31:62:80
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Admin Speed is 32 Gbps
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 500
    Receive B2B Credit is 500
    B2B State Change: Admin(on), Oper(up), Negotiated Value(14)
    Receive data field Size is 2112
    Beacon is turned off
    fec is enabled by default
    Logical type is core
    Trunk vsans (admin allowed and active) (1,40)
    Trunk vsans (up)                       (1,40)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 416 bits/sec,52 bytes/sec, 2 frames/sec
    5 minutes output rate 416 bits/sec,52 bytes/sec, 2 frames/sec
      1574 frames input,131460 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      1576 frames output,134700 bytes
        0 discards,0 errors
      1 input OLS,1  LRR,0 NOS,0 loop inits
      2 output OLS,3 LRR, 1 NOS, 0 loop inits
    500 receive B2B credit remaining
    500 transmit B2B credit remaining
    500 low priority transmit B2B credit remaining
    Interface last changed at Thu Mar  5 06:09:05 2020
    Last clearing of "show interface" counters :  never

SAN192C-6#
SAN192C-6# show run interface fc1/48

!Command: show running-config interface fc1/48
!Running configuration last done at: Thu Mar  5 06:11:33 2020
!Time: Thu Mar  5 06:18:29 2020

version 8.4(1a)

interface fc1/48
  switchport speed 32000
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 40
  no shutdown

SAN192C-6#
```

Then we add the second ISL:

```
SAN384C-6#
SAN384C-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# interface fc2/48
SAN384C-6(config-if)# switchport mode E
SAN384C-6(config-if)# switchport speed 32000
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)# switchport trunk allowed vsan add 40
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6# show topology vsan 40

FC Topology for VSAN 40 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
           fc1/48  0x21(33)           fc1/48  10.122.107.94(SAN192C-6)
           fc2/48  0x21(33)           fc2/48  10.122.107.94(SAN192C-6)
SAN384C-6#


SAN192C-6#
SAN192C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# interface fc2/48
SAN192C-6(config-if)# switchport mode E
```

```
SAN192C-6(config-if)# switchport speed 32000
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)# switchport trunk allowed vsan add 40
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)# end
'Performing fast copy of configurationdone.
SAN192C-6#
SAN192C-6# show topology vsan 40

FC Topology for VSAN 40 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
           fc1/48  0x20(32)           fc1/48  10.122.107.95(SAN384C-6)
           fc2/48  0x20(32)           fc2/48  10.122.107.95(SAN384C-6)
SAN192C-6#
```

7.2.12 Configuring Fibre Channel port channels by using the DCNM

To take advantage of the ability to aggregate bandwidth, enhance load balancing, and provide link-level redundancy, we will create a port channel out of the ISLs that were defined in 7.2.11, “Configuring Inter-Switch Links by using the DCNM” on page 312.

To configure a Port Channel by using the DCNM, click **Configure**, and under **SAN**, click **Port Channel**, as shown in Figure 7-88. Click **Create New Port Channel**.

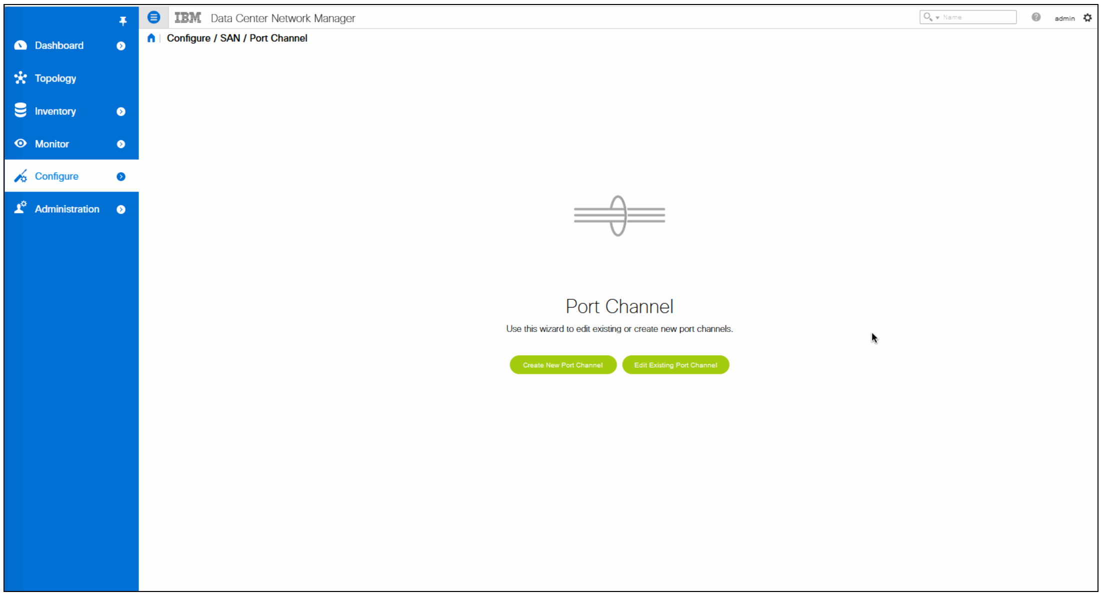

Use this wizard to edit existing or create new port channels.

Figure 7-88 Port channel wizard

Select the fabric and switch pair, as shown in Figure 7-89, that will be used to create this port channel.

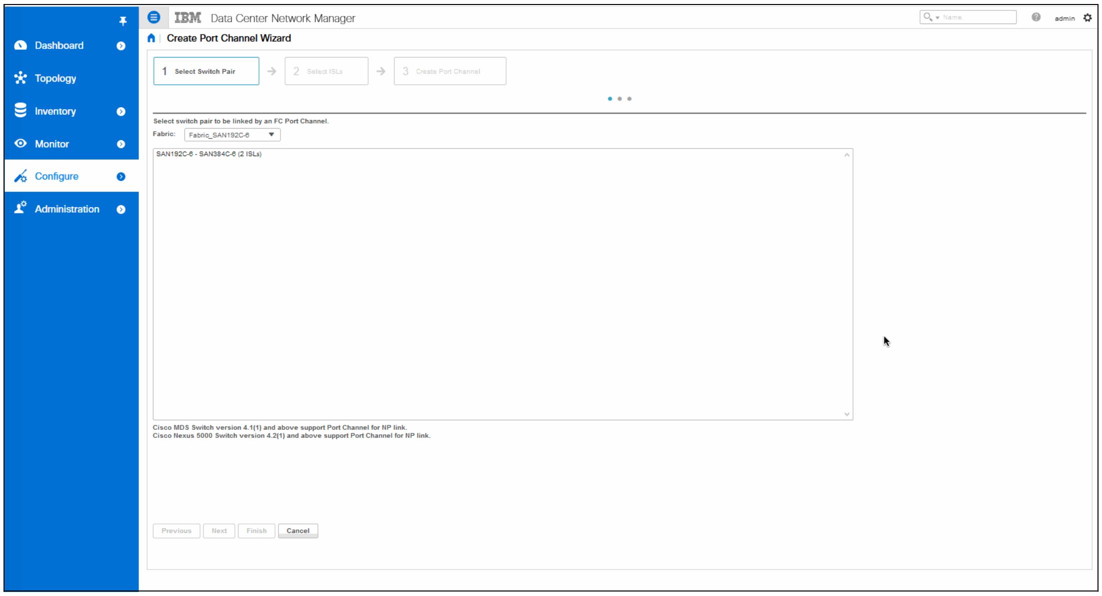

Select switch pair to be linked by an FC Port Channel.

Fabric: Fabric_SAN192C-6

SAN192C-6 - SAN384C-6 (2 ISLs)

Cisco MDS Switch version 4.1(1) and above support Port Channel for NP link.
Cisco Nexus 5000 Switch version 4.2(1) and above support Port Channel for NP link.

Figure 7-89 Selecting the switch pair

In our lab environment, there is a single fabric with one pair of connected switches. We have selected two switches, IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6, as shown in Figure 7-90. Click **Next**.

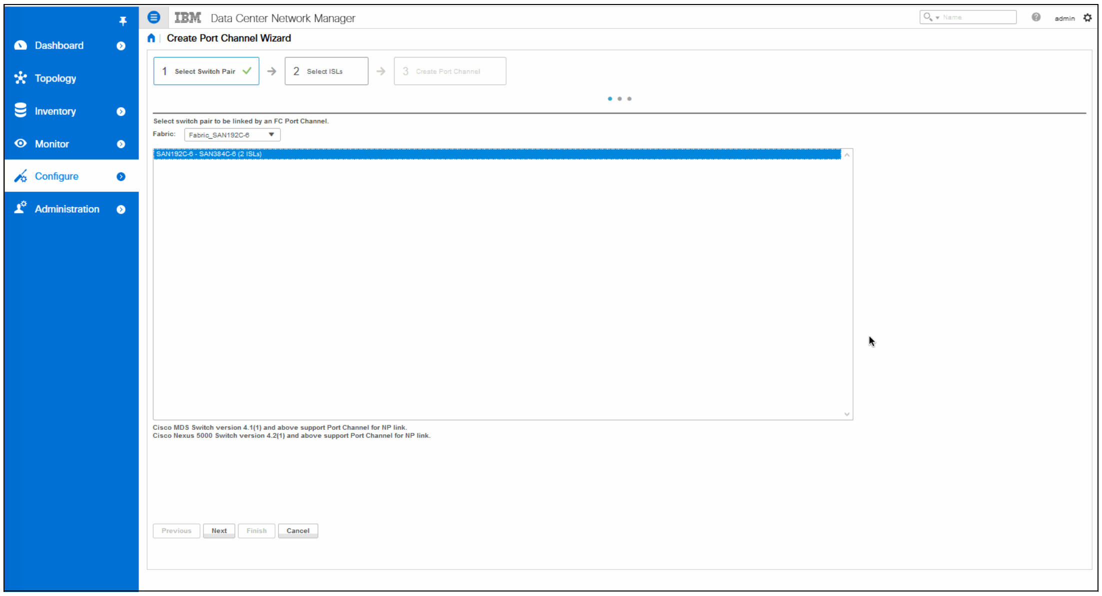

Select switch pair to be linked by an FC Port Channel.

Fabric: Fabric_SAN192C-6

SAN192C-6 - SAN384C-6 (2 ISLs)

Cisco MDS Switch version 4.1(1) and above support Port Channel for NP link.
Cisco Nexus 5000 Switch version 4.2(1) and above support Port Channel for NP link.

Figure 7-90 Selected switches

In Figure 7-91 on page 321, we leave the ISLs that are listed under Selected for the port channel being created. Click **Next**.

Note: When you enter the Port Channel Create window, all available ISLs that are configured in your fabric are listed under Selected. Any ISLs that do not need to be defined in this port channel should be moved to Available. These ISLs can be used to create additional port channels.

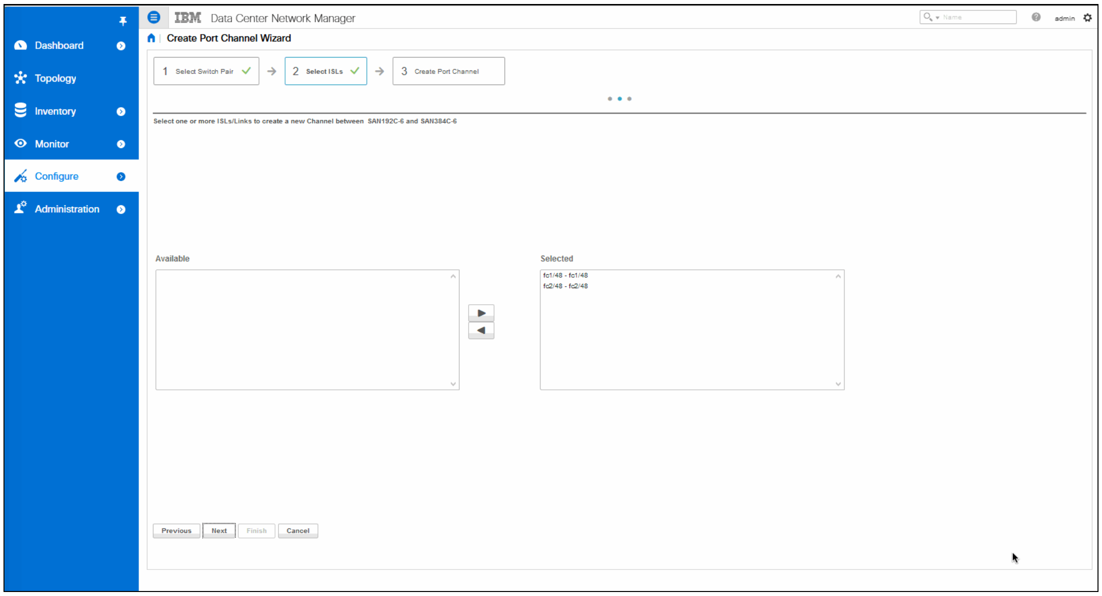

Select one or more ISLs/Links to create a new Channel between SAN192C-6 and SAN384C-6

Selected:
fc1/48 - fc1/48
fc2/48 - fc2/48

Figure 7-91 ISL selection

Figure 7-92 shows the attributes of the port channel.

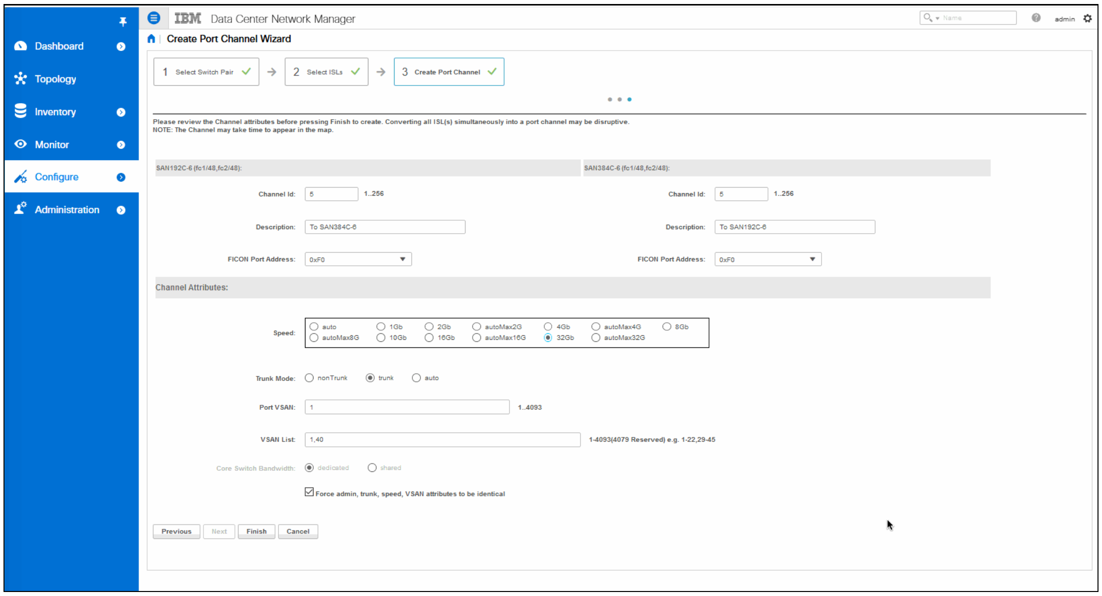

Please review the Channel attributes before pressing Finish to create. Converting all ISL(s) simultaneously into a port channel may be disruptive.
NOTE: The Channel may take time to appear in the map.

SAN192C-6 (fc1/48,fc2/48):

Channel Id: 5 1..256
Description: To SAN384C-6
FICON Port Address: 0xF0

SAN384C-6 (fc1/48,fc2/48):

Channel Id: 5 1..256
Description: To SAN192C-6
FICON Port Address: 0xF0

Channel Attributes:

Speed: ○ auto ○ 1Gb ○ 2Gb ○ autoMax2G ○ 4Gb ○ autoMax4G ○ 8Gb ○ autoMax8G ○ 10Gb ○ 16Gb ○ autoMax16G ● 32Gb ○ autoMax32G

Trunk Mode: ○ nonTrunk ● trunk ○ auto

Port VSAN: 1 1..4093

VSAN List: 1,40 1-4093(4079 Reserved) e.g. 1-22,29-45

Core Switch Bandwidth: ● dedicated ○ shared

☑ Force admin, trunk, speed, VSAN attributes to be identical

Figure 7-92 Port Channel attributes

Define the following key attributes:

- **Channel ID:** As a best practice, use the same value on both switches for the port channel when possible.
- **Description:** Reference the source and destination of the port channel.

- **FICON Port Address:** Only applicable when FICON is enabled on the switch. Must be configured when there is a FICON VSAN communicating on this port channel. As a best practice, use the same value on both switches for the port channel when possible. The value of this attribute must be taken from the pool of logical FICON Port Addresses.
- **Speed:** Auto-negotiation is not recommended when configuring port channels, so set a fixed speed.
- **Trunk Mode:** As a best practice, use **Trunk**.
- **Port VSAN:** As a best practice, use the default value of 1.
- **VSAN List:** Should have 1 and the value of any VSANs that require access to the port channel data path.
- Select the **Force admin, Speed, Trunk, and VSAN attributes** checkbox to ensure that all ISLs members in the port channel are identical.

To proceed, click **Finish**.

Figure 7-93 shows that converting ISLs to port channels can be disruptive. Click **Yes** to create the port channel.

Important: In a production environment, converting ISLs into a port channel can be disruptive if there is active traffic across ISLs. Redundant paths between switches do not mitigate this risk.

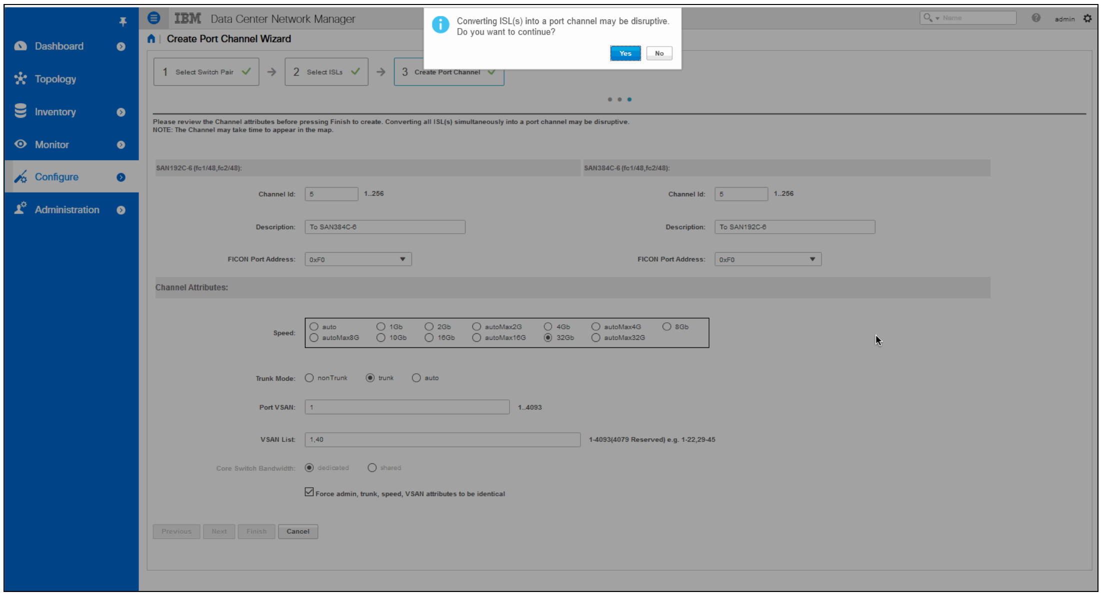

Converting ISL(s) into a port channel may be disruptive.
Do you want to continue?

Yes No

Please review the Channel attributes before pressing Finish to create. Converting all ISL(s) simultaneously into a port channel may be disruptive.
NOTE: The Channel may take time to appear in the map.

SAN192C-6 (fc1/48,fc2/48):

Channel Id: 5 1..256
Description: To SAN384C-6
FICON Port Address: 0xF0

SAN384C-6 (fc1/48,fc2/48):

Channel Id: 5 1..256
Description: To SAN192C-6
FICON Port Address: 0xF0

Channel Attributes:

Speed: ○ auto ○ 1Gb ○ 2Gb ○ autoMax2G ○ 4Gb ○ autoMax4G ○ 8Gb ○ autoMax8G ○ 10Gb ○ 16Gb ○ autoMax16G ● 32Gb ○ autoMax32G

Trunk Mode: ○ nonTrunk ● trunk ○ auto

Port VSAN: 1 1..4093

VSAN List: 1,40 1-4093(4079 Reserved) e.g. 1-22,29-45

Core Switch Bandwidth: ● dedicated ○ shared

☑ Force admin, trunk, speed, VSAN attributes to be identical

Previous Next Finish Cancel

Figure 7-93 Confirming the port channel creation

Figure 7-94 on page 323 shows that port channel configuration was applied successfully.

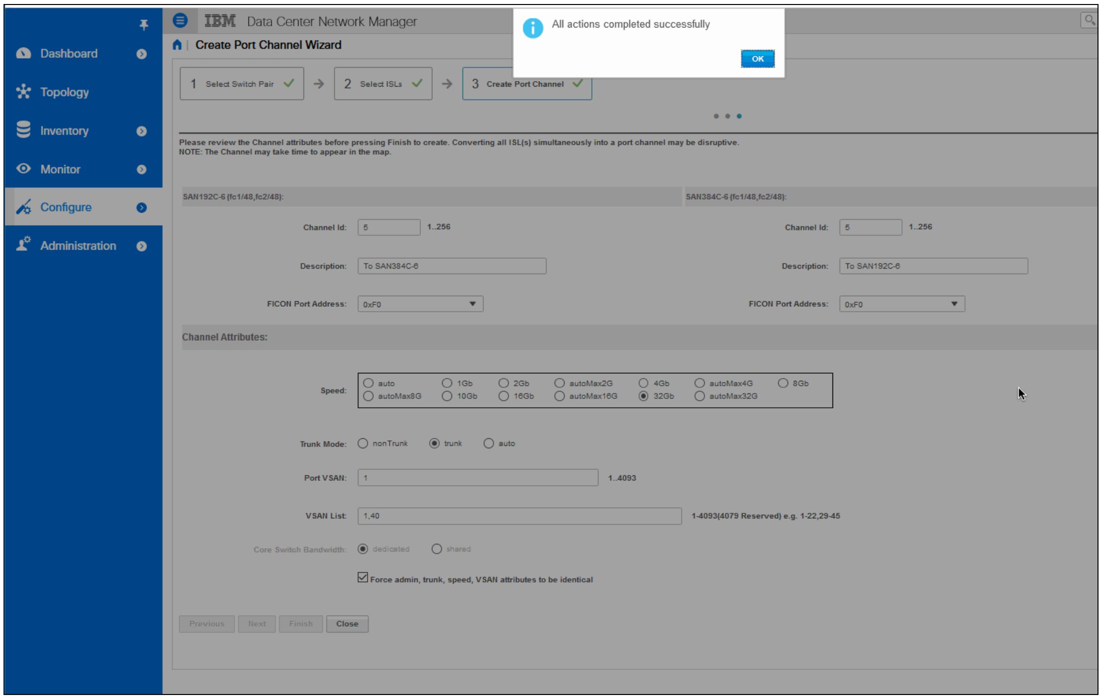

All actions completed successfully

Please review the Channel attributes before pressing Finish to create. Converting all ISL(s) simultaneously into a port channel may be disruptive.
NOTE: The Channel may take time to appear in the map.

SAN192C-6 (fc1/48,fc2/48):

Channel Id: 5 1..256

Description: To SAN384C-6

FICON Port Address: 0xF0

SAN384C-6 (fc1/48,fc2/48):

Channel Id: 5 1..256

Description: To SAN192C-6

FICON Port Address: 0xF0

Channel Attributes:

Speed: ○ auto ○ 1Gb ○ 2Gb ○ autoMax2G ○ 4Gb ○ autoMax4G ○ 8Gb ○ autoMax8G ○ 10Gb ○ 16Gb ○ autoMax16G ● 32Gb ○ autoMax32G

Trunk Mode: ○ nonTrunk ● trunk ○ auto

Port VSAN: 1 1..4093

VSAN List: 1,40 1-4093(4079 Reserved) e.g. 1-22,29-45

Core Switch Bandwidth: ● dedicated ○ shared

☑ Force admin, trunk, speed, VSAN attributes to be identical

Figure 7-94 Configuration was applied successfully

Figure 7-95 shows the Topology view and that the port channel between the IBM Storage Networking SAN384C-6 and IBM Storage Networking SAN192C-6 switches was created successfully.

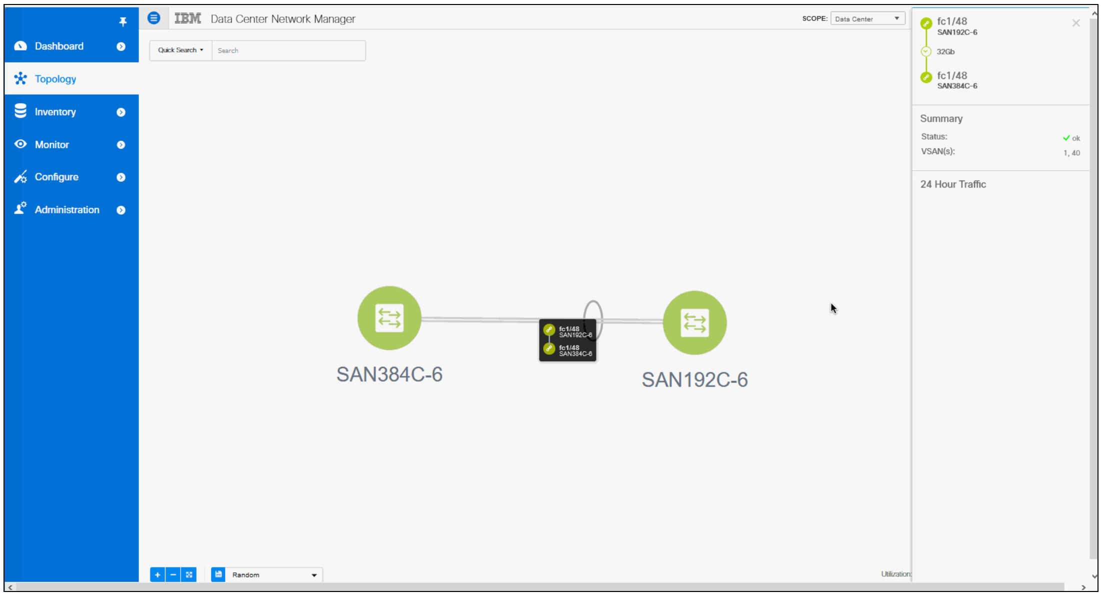

fc1/48
SAN192C-6
32Gb
fc1/48
SAN384C-6

Summary

Status: ✓ ok

VSAN(s): 1, 40

24 Hour Traffic

Figure 7-95 Port Channel Topology view

Configuring Fibre Channel port channels by using the CLI

To aggregate bandwidth, enhance load balancing and provide link level redundancy, we create a port channel out of the ISLs by using CLI, as shown in Example 7-5.

Example 7-5 Configuring a Fibre Channel port channel by using the CLI

```
SAN384C-6#
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)#
SAN384C-6(config)# int port-channel 5
SAN384C-6(config-if)# switchport mode E
SAN384C-6(config-if)# switchport speed 32000
SAN384C-6(config-if)# switchport description To SAN192C-6
SAN384C-6(config-if)#
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)# switchport trunk allowed vsan add 40
SAN384C-6(config-if)#
SAN384C-6(config-if)# ficon portnumber 0xf0
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# interface fc1/48
SAN384C-6(config-if)# channel-group 5 force
fc1/48 added to port-channel 5 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# interface fc2/48
SAN384C-6(config-if)# channel-group 5 force
fc2/48 added to port-channel 5 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6# show interface port-channel 5
port-channel5 is trunking
    Port description is To SAN192C-6
    Hardware is Fibre Channel
    Port WWN is 24:05:00:3a:9c:31:62:80
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 64 Gbps
    fec is enabled by default
    Logical type is core
    Trunk vsans (admin allowed and active) (1,40)
    Trunk vsans (up)                       (1,40)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 2016 bits/sec,252 bytes/sec, 1 frames/sec
    5 minutes output rate 2016 bits/sec,252 bytes/sec, 1 frames/sec
      142812 frames input,9817016 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      142812 frames output,7937920 bytes
        0 discards,0 errors
      16 input OLS,21  LRR,8 NOS,0 loop inits
      20 output OLS,9 LRR, 8 NOS, 0 loop inits
    Member[1] : fc1/48    [up] *
    Member[2] : fc2/48    [up]
    Interface last changed at Sat Mar  7 22:03:26 2020
SAN384C-6#
SAN384C-6# show port-channel summary
------------------------------------------------------------------------------
Interface                 Total Ports        Oper Ports        First Oper Port
------------------------------------------------------------------------------
port-channel 5                 2                  2                   fc1/48
SAN384C-6#
SAN384C-6# show topology vsan 40

FC Topology for VSAN 40 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
    port-channel5  0x21(33)    port-channel5  10.122.107.94(SAN192C-6)
SAN384C-6#


Now other switch

SAN192C-6#
SAN192C-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# int port-channel 5
SAN192C-6(config-if)# switchport mode E
SAN192C-6(config-if)# switchport speed 32000
SAN192C-6(config-if)# switchport description To SAN384C-6
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)# switchport trunk allowed vsan add 40
SAN192C-6(config-if)#
SAN192C-6(config-if)# ficon portnumber 0xf0
SAN192C-6(config-if)#
SAN192C-6(config-if)# no shut
```

```
SAN192C-6(config-if)#
SAN192C-6(config-if)# interface fc1/48
SAN192C-6(config-if)# channel-group 5 force
fc1/48 added to port-channel 5 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN192C-6(config-if)# no shut
SAN192C-6(config-if)#
SAN192C-6(config-if)#
SAN192C-6(config-if)# interface fc2/48
SAN192C-6(config-if)# channel-group 5 force
fc2/48 added to port-channel 5 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN192C-6(config-if)# no shut
SAN192C-6(config-if)#
SAN192C-6(config-if)# end
Performing fast copy of configurationdone.
SAN192C-6#
SAN192C-6# show topology vsan 40

FC Topology for VSAN 40 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
    port-channel5  0x20(32)    port-channel5  10.122.107.95(SAN384C-6)
SAN192C-6#
```

7.2.13 Configuring the FICON mainframe disk and tape interfaces

This section shows how to configure the FICON mainframe disk and tape interfaces by using DM.

Note: Some configuration steps might vary depending on your environment requirements.

In DCNM, select **Topology** → **Switch** → **Device Manager**. The General tab of the interface characteristics window opens. For this port, we will be connecting a 16 GbE mainframe channel. Start by entering the description of the CHPID, Physical Channel ID (PCHID), and Port VSAN manually or by using the drop-down menu for the available VSANs. It is a best practice that we use Forward Error Correction (FEC) for the connection for the 16 GbE mainframe channel, so we set the interface to 16 GbE, as shown in Figure 7-96 on page 327, under the Speed category. Click **Apply**.

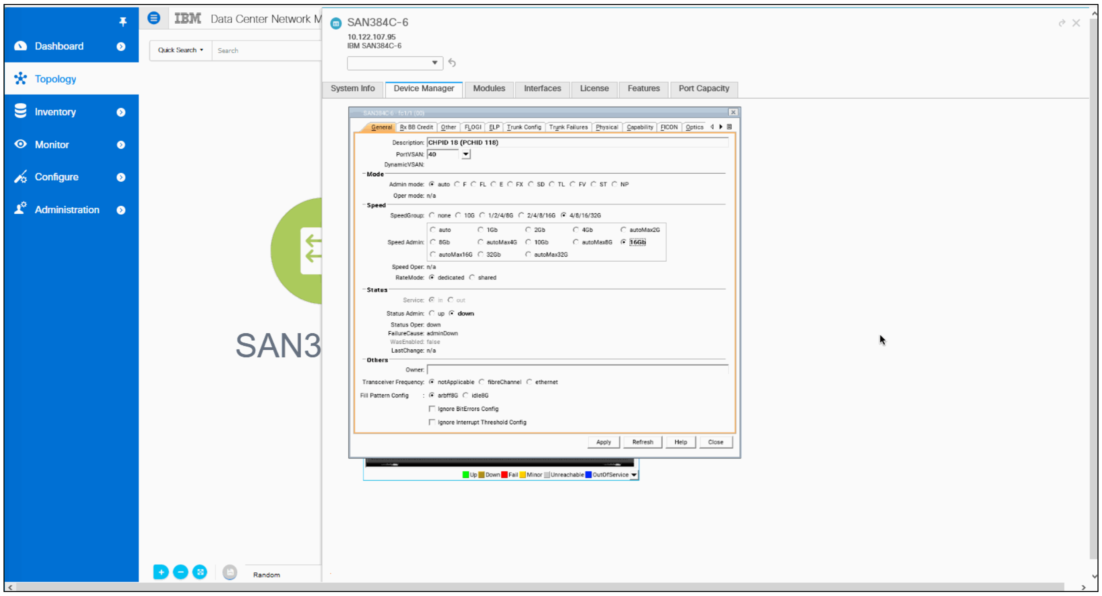

SAN384C-6
10.122.107.95
IBM SAN384C-6

SAN384C-6 - fc1/1 (00)

Description: CHPID 18 (PCHID 118)
PortVSAN: 40
DynamicVSAN:

Mode
Admin mode: ● auto ○ F ○ FL ○ E ○ FX ○ SD ○ TL ○ FV ○ ST ○ NP
Oper mode: n/a

Speed
SpeedGroup: ○ none ○ 10G ○ 1/2/4/8G ○ 2/4/8/16G ● 4/8/16/32G
Speed Admin: ○ auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb ○ autoMax8G ● 16Gb ○ autoMax16G ○ 32Gb ○ autoMax32G
Speed Oper: n/a
RateMode: ● dedicated ○ shared

Status
Service: ● in ○ out
Status Admin: ○ up ● down
Status Oper: down
FailureCause: adminDown
WasEnabled: false
LastChange: n/a

Others
Owner:
Transceiver Frequency: ● notApplicable ○ fibreChannel ○ ethernet
Fill Pattern Config : ● arbff8G ○ idle8G
□ Ignore BitErrors Config
□ Ignore Interrupt Threshold Config

Figure 7-96 General tab

Click the **Other** tab and select the **Up** dialog button for Admin FEC and Admin FEC TTS, as shown in Figure 7-97. By doing this action, the switch can perform FEC negotiation with the CHPID when it comes online. Click **Apply**.

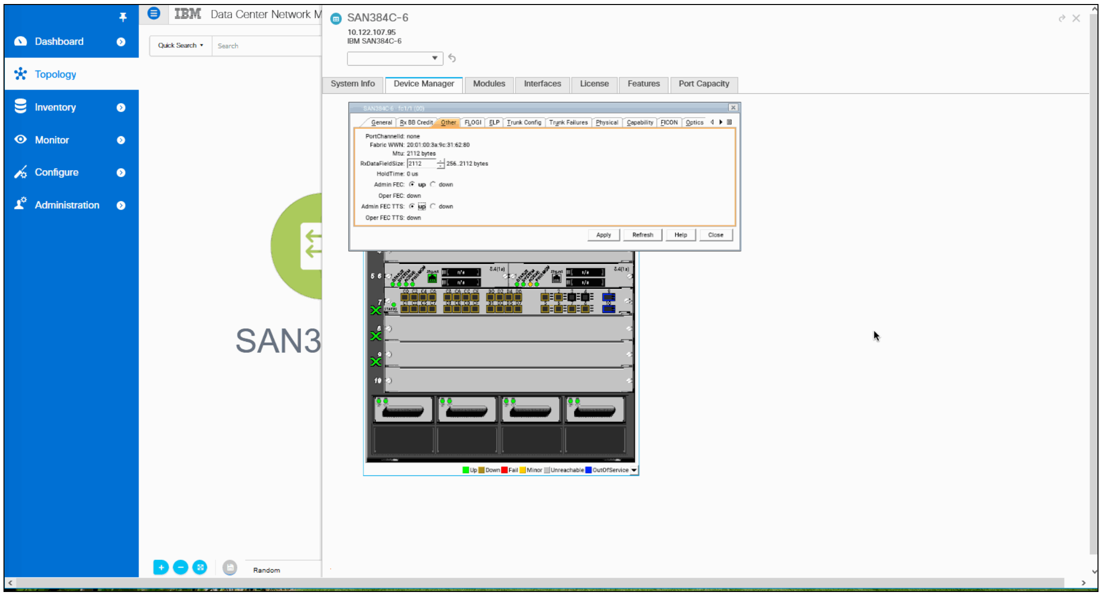

SAN384C-6
10.122.107.95
IBM SAN384C-6

SAN384C-6 - fc1/1 (00)

PortChannelId: none
Fabric WWN: 20:01:00:3a:9c:31:62:80
Mtu: 2112 bytes
RxDataFieldSize: 2112 256..2112 bytes
HoldTime: 0 us
Admin FEC: ● up ○ down
Oper FEC: down
Admin FEC TTS: ● up ○ down
Oper FEC TTS: down

Figure 7-97 FEC selection

Click the **General** tab. Under Status, select the Status Admin **Up** dialog button, as shown in Figure 7-98. This action brings the 16 GbE switch port online.

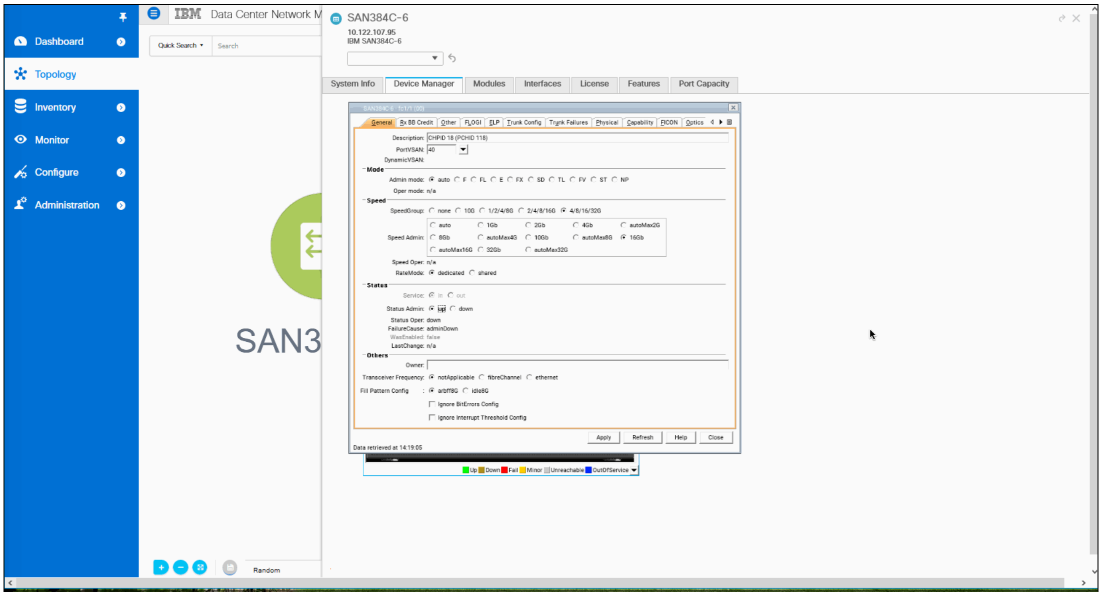

Description: CHPID 18 (PCHID 118)
PortVSAN: 40
DynamicVSAN:

Mode

Admin mode: ● auto ○ F ○ FL ○ E ○ FX ○ SD ○ TL ○ FV ○ ST ○ NP
Oper mode: n/a

Speed

SpeedGroup: ○ none ○ 10G ○ 1/2/4/8G ○ 2/4/8/16G ● 4/8/16/32G
Speed Admin: ○ auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb ○ autoMax8G ● 16Gb ○ autoMax16G ○ 32Gb ○ autoMax32G
Speed Oper: n/a
RateMode: ● dedicated ○ shared

Status

Service: ● in ○ out
Status Admin: ● up ○ down
Status Oper: down
FailureCause: adminDown
WasEnabled: false
LastChange: n/a

Others

Owner:
Transceiver Frequency: ● notApplicable ○ fibreChannel ○ ethernet
Fill Pattern Config : ● arbff8G ○ idle8G
□ Ignore BitErrors Config
□ Ignore Interrupt Threshold Config

Data retrieved at 14:19:05

Figure 7-98 General tab

From the mainframe console, we configure the CHPID as online, as shown in Figure 7-99.

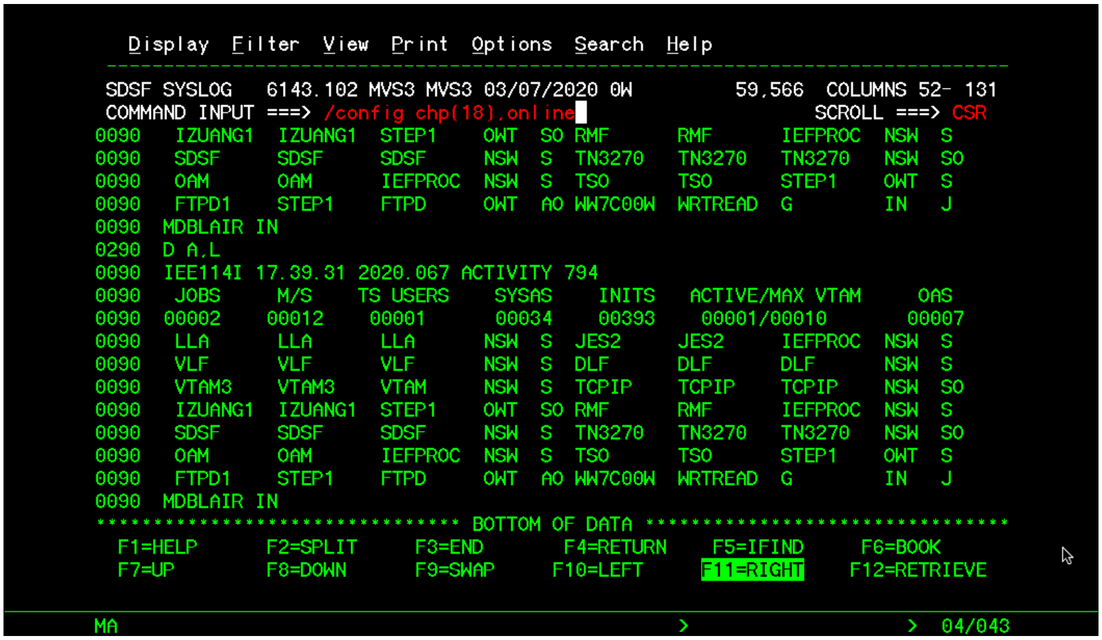

```
  Display  Filter  View  Print  Options  Search  Help
 -------------------------------------------------------------------------------
 SDSF SYSLOG   6143.102 MVS3 MVS3 03/07/2020 0W         59.566  COLUMNS 52- 131
 COMMAND INPUT ===> /config chp(18),online                     SCROLL ===> CSR
0090   IZUANG1  IZUANG1  STEP1    OWT  SO RMF      RMF      IEFPROC  NSW  S
0090   SDSF     SDSF     SDSF     NSW  S  TN3270   TN3270   TN3270   NSW  SO
0090   OAM      OAM      IEFPROC  NSW  S  TSO      TSO      STEP1    OWT  S
0090   FTPD1    STEP1    FTPD     OWT  AO WW7C00W  WRTREAD  G        IN   J
0090  MDBLAIR IN
0290  D A,L
0090  IEE114I 17.39.31 2020.067 ACTIVITY 794
0090   JOBS     M/S    TS USERS    SYSAS    INITS   ACTIVE/MAX VTAM     OAS
0090  00002    00012    00001      00034    00393    00001/00010       00007
0090   LLA      LLA      LLA      NSW  S  JES2     JES2     IEFPROC  NSW  S
0090   VLF      VLF      VLF      NSW  S  DLF      DLF      DLF      NSW  S
0090   VTAM3    VTAM3    VTAM     NSW  S  TCPIP    TCPIP    TCPIP    NSW  SO
0090   IZUANG1  IZUANG1  STEP1    OWT  SO RMF      RMF      IEFPROC  NSW  S
0090   SDSF     SDSF     SDSF     NSW  S  TN3270   TN3270   TN3270   NSW  SO
0090   OAM      OAM      IEFPROC  NSW  S  TSO      TSO      STEP1    OWT  S
0090   FTPD1    STEP1    FTPD     OWT  AO WW7C00W  WRTREAD  G        IN   J
0090  MDBLAIR IN
****************************** BOTTOM OF DATA ******************************
  F1=HELP      F2=SPLIT     F3=END       F4=RETURN    F5=IFIND     F6=BOOK
  F7=UP        F8=DOWN      F9=SWAP     F10=LEFT     F11=RIGHT    F12=RETRIEVE

MA                                             >                   >  04/043
```

Figure 7-99 Showing the CHPID as online

Figure 7-100 on page 329 shows that the mainframe CHPID is logged in at 16 GbE (based on the operational speed and status).

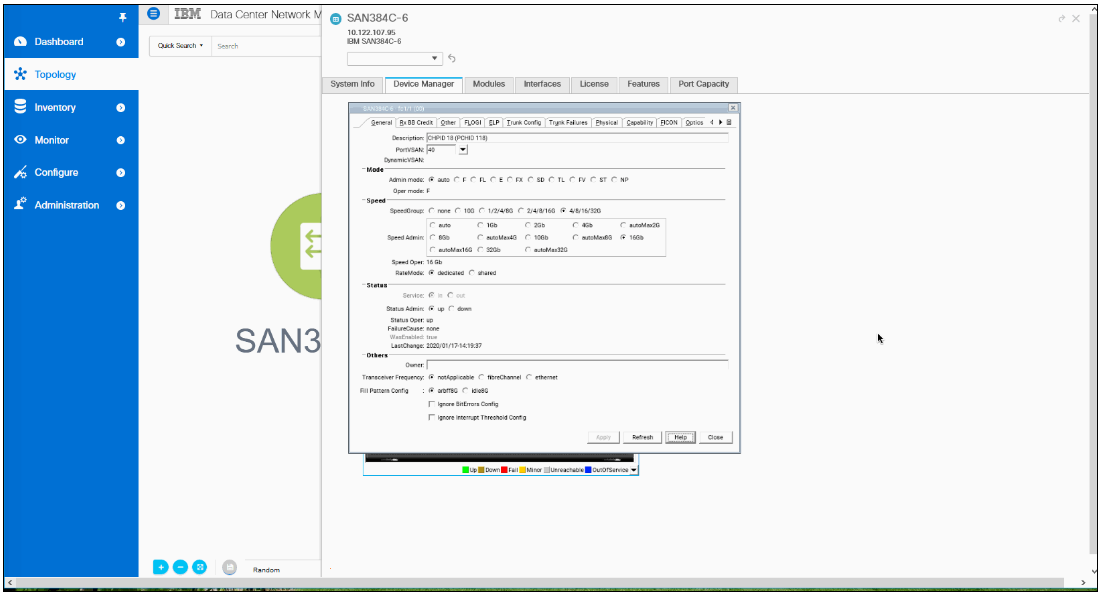

Description: CHPID 18 (PCHID 118)
PortVSAN: 40
DynamicVSAN:

Mode

Admin mode: ● auto ○ F ○ FL ○ E ○ FX ○ SD ○ TL ○ FV ○ ST ○ NP
Oper mode: F

Speed

SpeedGroup: ○ none ○ 10G ○ 1/2/4/8G ○ 2/4/8/16G ● 4/8/16/32G
Speed Admin: ○ auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb ○ autoMax8G ● 16Gb ○ autoMax16G ○ 32Gb ○ autoMax32G
Speed Oper: 16 Gb
RateMode: ● dedicated ○ shared

Status

Service: ● in ○ out
Status Admin: ● up ○ down
Status Oper: up
FailureCause: none
WasEnabled: true
LastChange: 2020/01/17-14:19:37

Others

Owner:
Transceiver Frequency: ● notApplicable ○ fibreChannel ○ ethernet
Fill Pattern Config : ● arbff8G ○ idle8G
□ Ignore BitErrors Config
□ Ignore Interrupt Threshold Config

Figure 7-100 CHPID online

To view the FICON Request Node Identification (RNID) information, click the **FICON** tab, as shown in Figure 7-101. As we can see, this node is CHPID 18 on the IBM2965 mainframe with serial number E8F77, according to the various fields in the RNID. We also can validate that FEC is operational between the mainframe and the switch. Click **Close**.

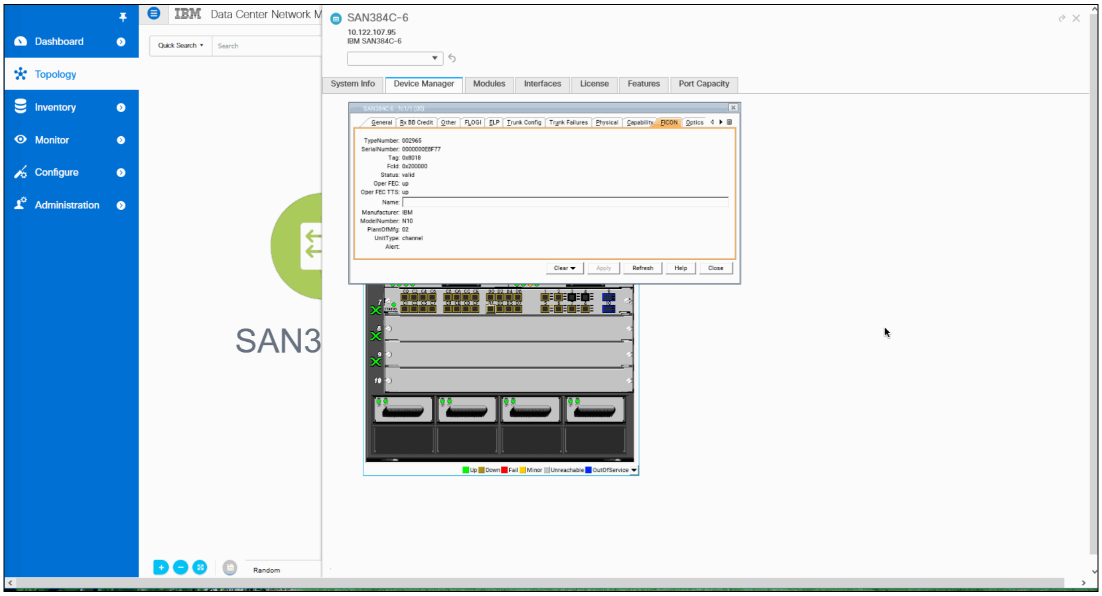

TypeNumber: 002965
SerialNumber: 0000000E8F77
Tag: 0x8018
FcId: 0x200000
Status: valid
Oper FEC: up
Oper FEC TTS: up
Name:
Manufacturer: IBM
ModelNumber: N10
PlantOfMfg: 02
UnitType: channel
Alert:

Figure 7-101 RNID

Figure 7-102 shows the CHPID online, which is indicated by the green box with the CH.

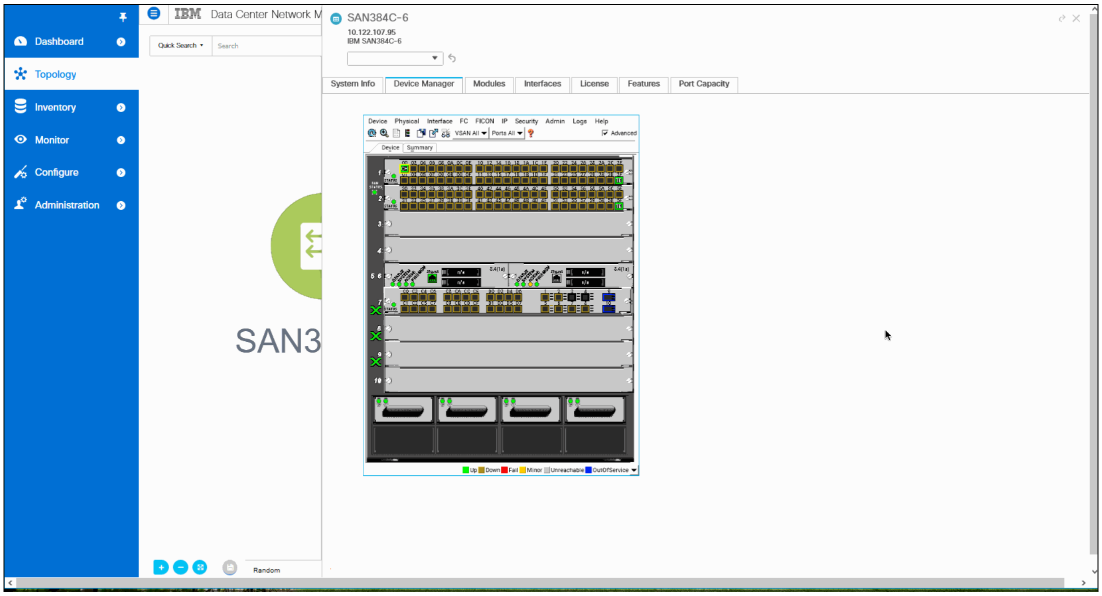

Figure 7-102 CHPID online

Figure 7-103 shows a summary of all online devices. The first interface column shows both physical interface fc1/1 on the switch and the FICON Port Address (00) as it would be viewed from the mainframe host. The RNID information for the CHPID ports is displayed in the Connected To column.

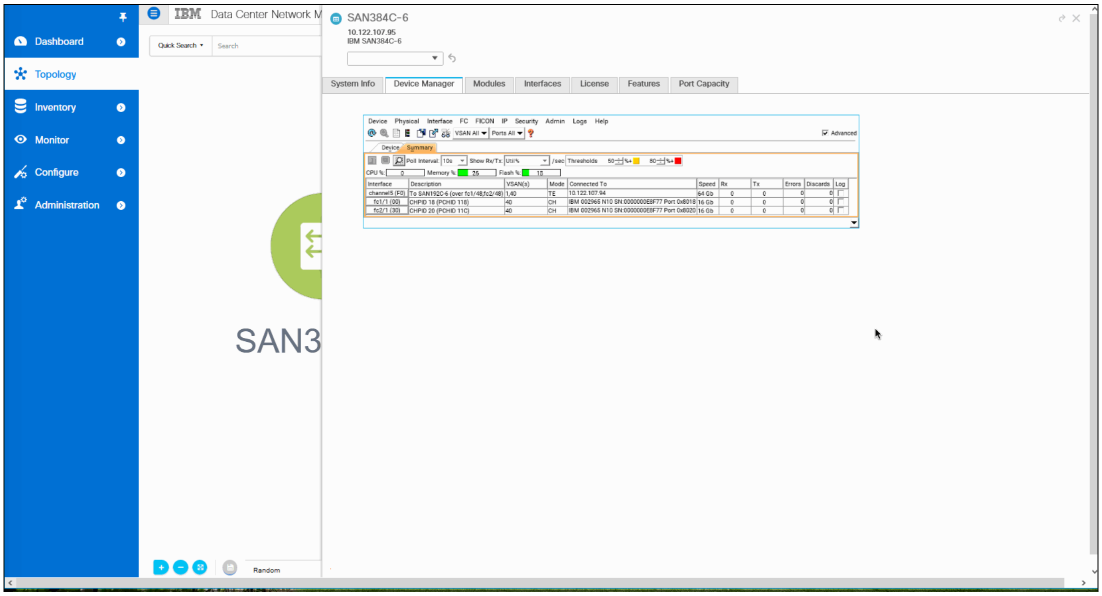

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 🟨 80 %+ 🟥

CPU %: 0 Memory %: 26 Flash %: 18

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
channel5 (F0)	To SAN192C-6 (over fc1/48,fc2/48)	1,40	TE	10.122.107.94	64 Gb	0	0	0	0	□
fc1/1 (00)	CHPID 18 (PCHID 118)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8018	16 Gb	0	0	0	0	□
fc2/1 (30)	CHPID 20 (PCHID 11C)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8020	16 Gb	0	0	0	0	□

Figure 7-103 Summary tab

We configure the IBM storage array ports at FICON Port Addresses 0x10 and 0x40 by using the same process and bring them online. Figure 7-104 on page 331 shows the IBM storage ports as being online and identified as CU ports.

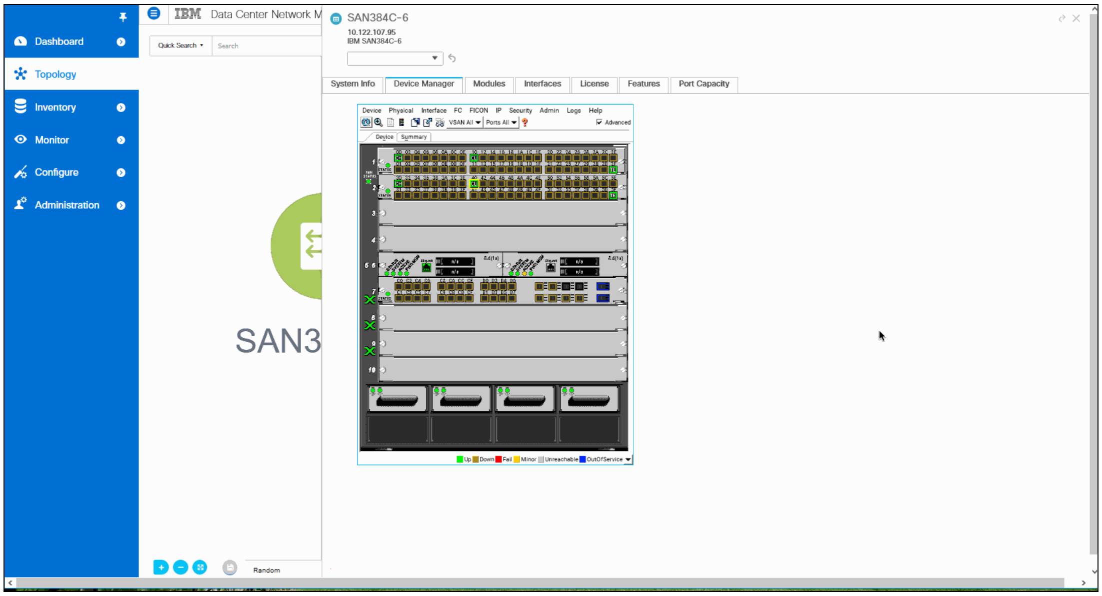

Figure 7-104 Control Unit Ports online

Figure 7-105 shows a summary of all online devices. On the IBM storage array CU ports, the FICON Port Addresses are the same as referenced in the mainframe hardware configuration.

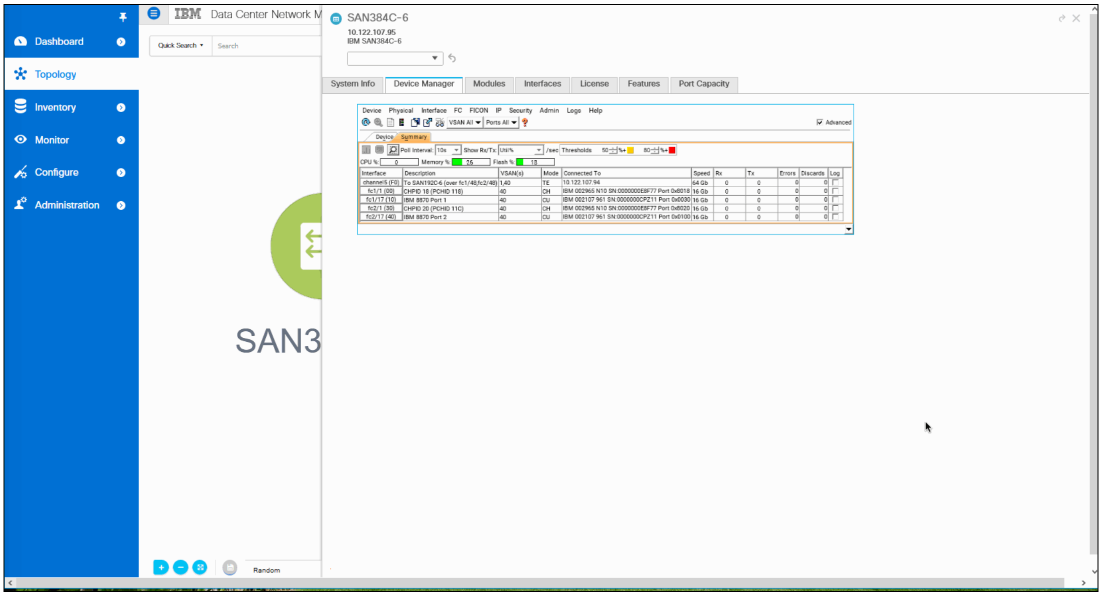

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 🟨 80 %+ 🟥

CPU %: 0 Memory %: 26 Flash %: 18

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
channel5 (F0)	To SAN192C-6 (over fc1/48,fc2/48)	1,40	TE	10.122.107.94	64 Gb	0	0	0	0	□
fc1/1 (00)	CHPID 18 (PCHID 118)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8018	16 Gb	0	0	0	0	□
fc1/17 (10)	IBM 8870 Port 1	40	CU	IBM 002107 961 SN:0000000CPZ11 Port 0x0030	16 Gb	0	0	0	0	□
fc2/1 (30)	CHPID 20 (PCHID 11C)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8020	16 Gb	0	0	0	0	□
fc2/17 (40)	IBM 8870 Port 2	40	CU	IBM 002107 961 SN:0000000CPZ11 Port 0x0100	16 Gb	0	0	0	0	□

Figure 7-105 Summary of online devices

Figure 7-106 shows the mainframe view of the device paths that are online for CHPIDs 18 and 20 so that they can communicate with the DS8870 storage frame on ports 0x10 (CHPID 18) and 0x40 (CHPID 20).

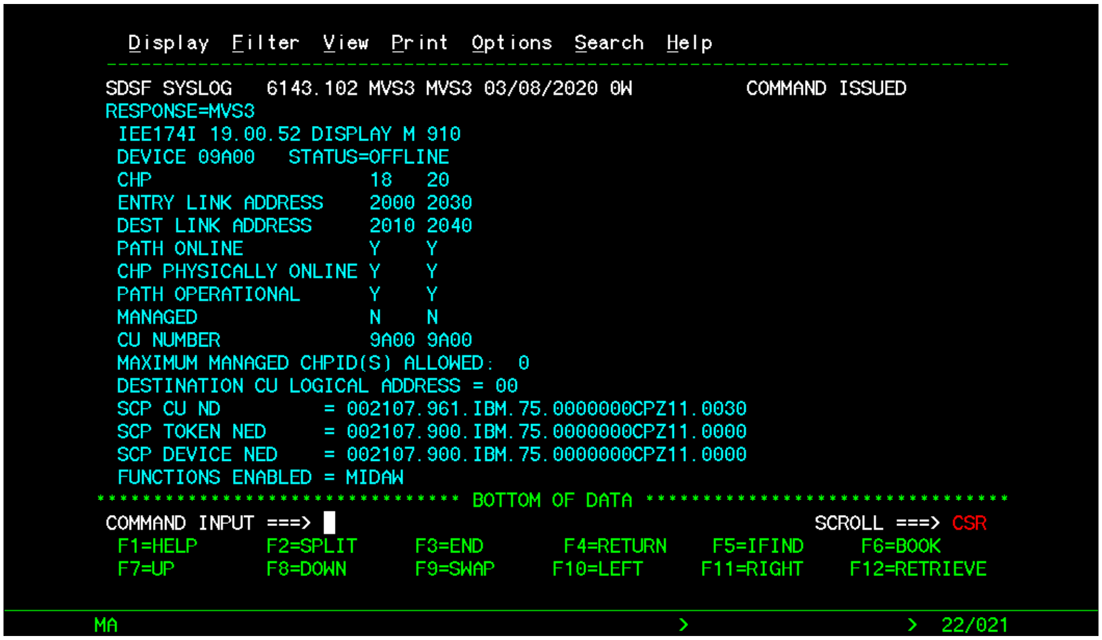

```
  Display  Filter  View  Print  Options  Search  Help
 -------------------------------------------------------------------------------
 SDSF SYSLOG   6143.102 MVS3 MVS3 03/08/2020 0W         COMMAND ISSUED
 RESPONSE=MVS3
  IEE174I 19.00.52 DISPLAY M 910
  DEVICE 09A00   STATUS=OFFLINE
  CHP                   18   20
  ENTRY LINK ADDRESS    2000 2030
  DEST LINK ADDRESS     2010 2040
  PATH ONLINE           Y    Y
  CHP PHYSICALLY ONLINE Y    Y
  PATH OPERATIONAL      Y    Y
  MANAGED               N    N
  CU NUMBER             9A00 9A00
  MAXIMUM MANAGED CHPID(S) ALLOWED:  0
  DESTINATION CU LOGICAL ADDRESS = 00
  SCP CU ND         = 002107.961.IBM.75.0000000CPZ11.0030
  SCP TOKEN NED     = 002107.900.IBM.75.0000000CPZ11.0000
  SCP DEVICE NED    = 002107.900.IBM.75.0000000CPZ11.0000
  FUNCTIONS ENABLED = MIDAW
******************************** BOTTOM OF DATA ********************************
 COMMAND INPUT ===>                                            SCROLL ===> CSR
  F1=HELP      F2=SPLIT     F3=END       F4=RETURN    F5=IFIND     F6=BOOK
  F7=UP        F8=DOWN      F9=SWAP     F10=LEFT     F11=RIGHT    F12=RETRIEVE

MA                                          >                   >  22/021
```

Figure 7-106 Mainframe view of the device paths that are online for CHPIDs 18 and 20

We define the cascaded CHPIDs for the lab environment, as shown in Figure 7-107.

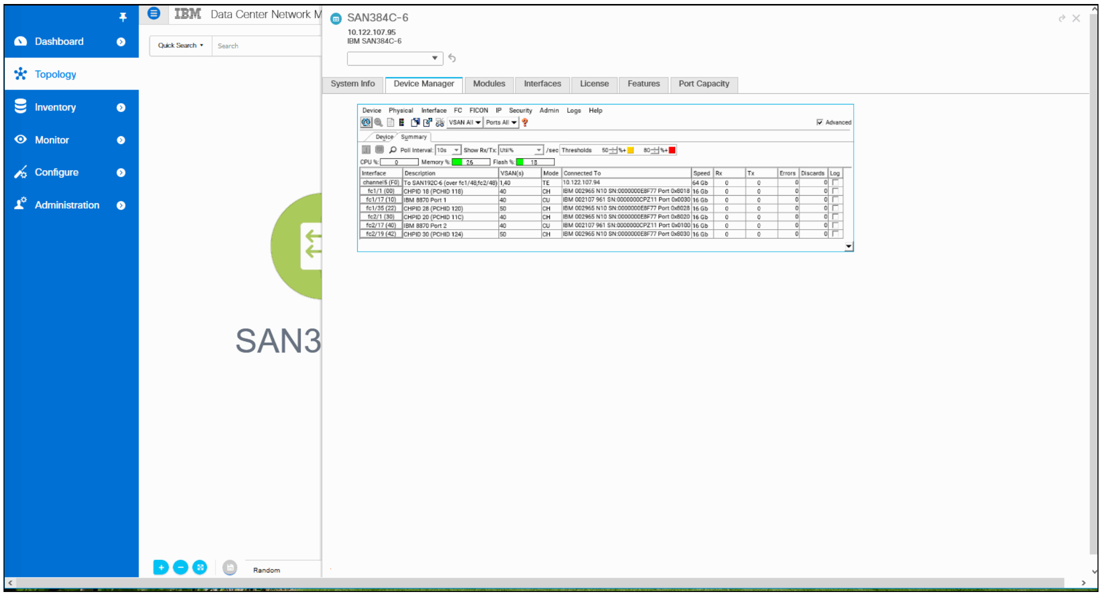

SAN384C-6
10.122.107.95
IBM SAN384C-6

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 80 %+
CPU %: 0 Memory %: 26 Flash %: 18

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
channel5 (F0)	To SAN192C-6 (over fc1/48,fc2/48)	1,40	TE	10.122.107.94	64 Gb	0	0	0	0	□
fc1/1 (00)	CHPID 18 (PCHID 118)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8018	16 Gb	0	0	0	0	□
fc1/17 (10)	IBM 8870 Port 1	40	CU	IBM 002107 961 SN:0000000CPZ11 Port 0x0030	16 Gb	0	0	0	0	□
fc1/35 (22)	CHPID 28 (PCHID 120)	50	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8028	16 Gb	0	0	0	0	□
fc2/1 (30)	CHPID 20 (PCHID 11C)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8020	16 Gb	0	0	0	0	□
fc2/17 (40)	IBM 8870 Port 2	40	CU	IBM 002107 961 SN:0000000CPZ11 Port 0x0100	16 Gb	0	0	0	0	□
fc2/19 (42)	CHPID 30 (PCHID 124)	50	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8030	16 Gb	0	0	0	0	□

Figure 7-107 Cascaded CHPIDs

On the IBM Storage Networking SAN192C-6 switch, we configure the cascaded storage array ports at FICON Port Addresses 0x0A and 0x4A and bring them online. Figure 7-108 on page 333 shows the cascaded storage ports as being online and identified as CU ports.

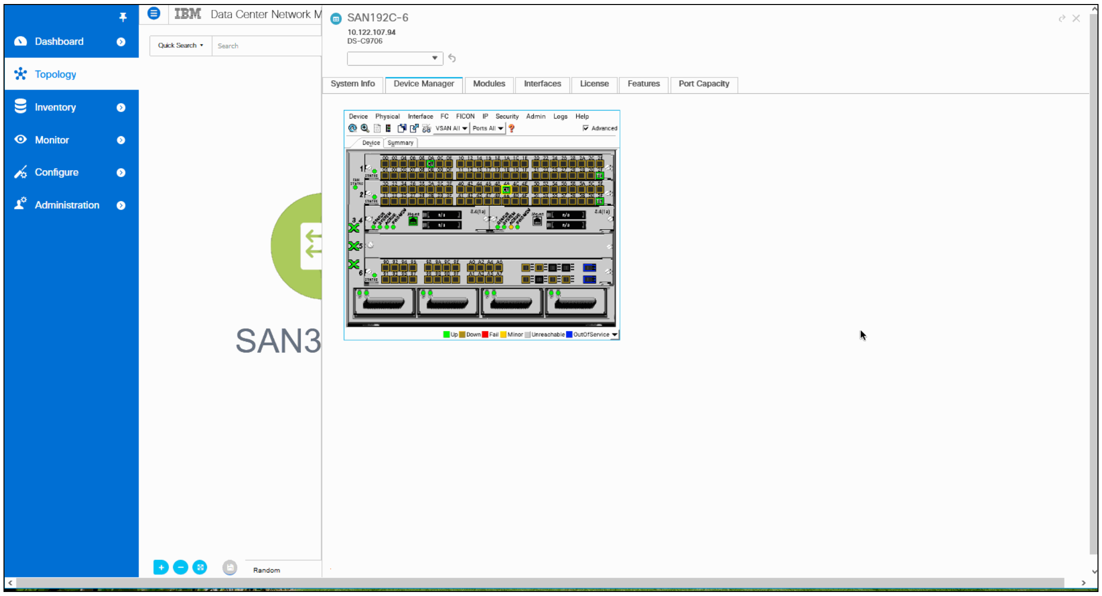

Figure 7-108 Cascaded CU ports online

Figure 7-109 shows a summary of all the online devices on the IBM Storage Networking SAN192C-6 switch.

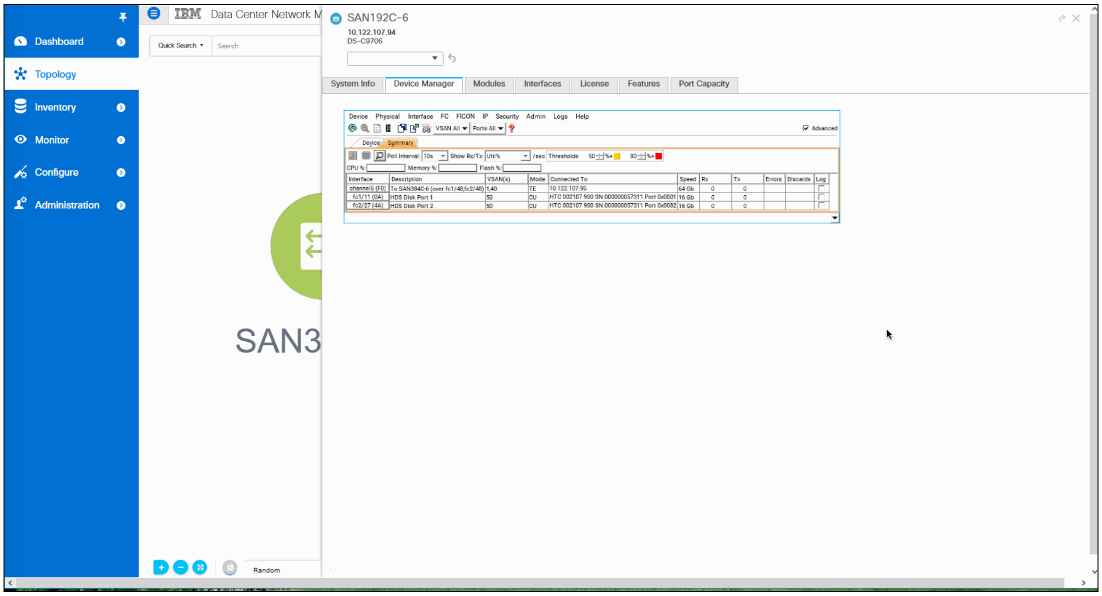

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 🟨 80 %+ 🟥

CPU %: Memory %: Flash %:

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
channel5 (F0)	To SAN384C-6 (over fc1/48,fc2/48)	1,40	TE	10.122.107.95	64 Gb	0	0			□
fc1/11 (0A)	HDS Disk Port 1	50	CU	HTC 002107 900 SN:000000057311 Port 0x0001	16 Gb	0	0			□
fc2/27 (4A)	HDS Disk Port 2	50	CU	HTC 002107 900 SN:000000057311 Port 0x0082	16 Gb	0	0			□

Figure 7-109 Online devices

Important: For security reasons, mainframe channels come online in FICON VSANs.

Figure 7-110 shows an invalid attachment for CHPID 0.20 where it was brought online before the port was in a FICON VSAN on the switch.

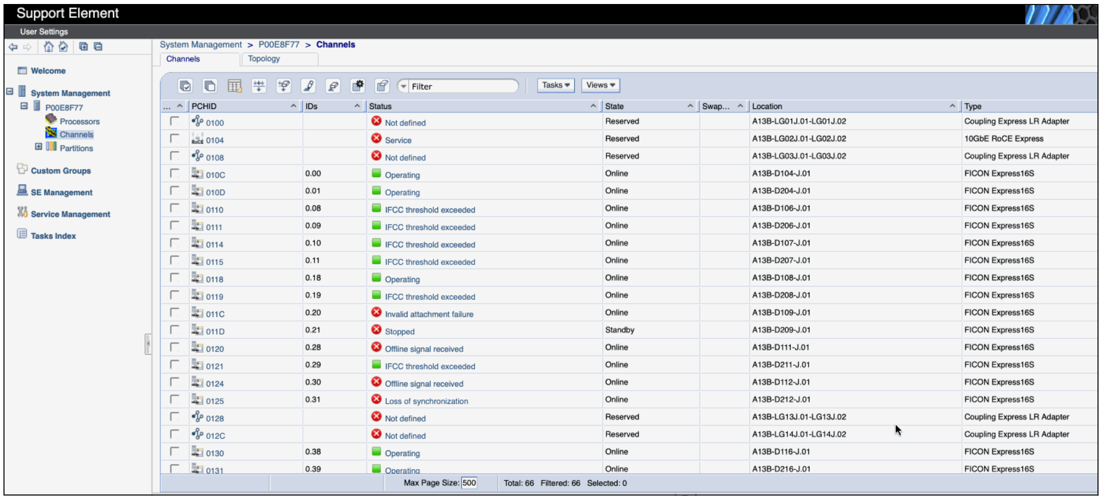

...	PCHID	IDs	Status	State	Swap...	Location	Type
□	0100		Not defined	Reserved		A13B-LG01J.01-LG01J.02	Coupling Express LR Adapter
□	0104		Service	Reserved		A13B-LG02J.01-LG02J.02	10GbE RoCE Express
□	0108		Not defined	Reserved		A13B-LG03J.01-LG03J.02	Coupling Express LR Adapter
□	010C	0.00	Operating	Online		A13B-D104-J.01	FICON Express16S
□	010D	0.01	Operating	Online		A13B-D204-J.01	FICON Express16S
□	0110	0.08	IFCC threshold exceeded	Online		A13B-D106-J.01	FICON Express16S
□	0111	0.09	IFCC threshold exceeded	Online		A13B-D206-J.01	FICON Express16S
□	0114	0.10	IFCC threshold exceeded	Online		A13B-D107-J.01	FICON Express16S
□	0115	0.11	IFCC threshold exceeded	Online		A13B-D207-J.01	FICON Express16S
□	0118	0.18	Operating	Online		A13B-D108-J.01	FICON Express16S
□	0119	0.19	IFCC threshold exceeded	Online		A13B-D208-J.01	FICON Express16S
□	011C	0.20	Invalid attachment failure	Online		A13B-D109-J.01	FICON Express16S
□	011D	0.21	Stopped	Standby		A13B-D209-J.01	FICON Express16S
□	0120	0.28	Offline signal received	Online		A13B-D111-J.01	FICON Express16S
□	0121	0.29	IFCC threshold exceeded	Online		A13B-D211-J.01	FICON Express16S
□	0124	0.30	Offline signal received	Online		A13B-D112-J.01	FICON Express16S
□	0125	0.31	Loss of synchronization	Online		A13B-D212-J.01	FICON Express16S
□	0128		Not defined	Reserved		A13B-LG13J.01-LG13J.02	Coupling Express LR Adapter
□	012C		Not defined	Reserved		A13B-LG14J.01-LG14J.02	Coupling Express LR Adapter
□	0130	0.38	Operating	Online		A13B-D116-J.01	FICON Express16S
□	0131	0.39	Operating	Online		A13B-D216-J.01	FICON Express16S

Max Page Size: 500 Total: 66 Filtered: 66 Selected: 0

Figure 7-110 Support element

Configuring FICON mainframe disk and tape interfaces by using the CLI

Example 7-6 shows how to configure the FICON mainframe disk and tape interfaces by using the CLI.

Example 7-6 Configuring the disk and tape interfaces

```
configure the 1st two channels in VSAN 40

SAN384C-6#
SAN384C-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# vsan database
SAN384C-6(config-vsan-db)# vsan 40 interface fc1/1
SAN384C-6(config-vsan-db)# exit
SAN384C-6(config)#
SAN384C-6(config)# interface fc1/1
SAN384C-6(config-if)# switchport description CHPID 18 (PCHID 118)
SAN384C-6(config-if)# switchport mode F
SAN384C-6(config-if)# switchport speed 16000
SAN384C-6(config-if)#
SAN384C-6(config-if)# switchport fec
SAN384C-6(config-if)# switchport fec tts
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)# exit
SAN384C-6(config)#
SAN384C-6(config)# vsan database
SAN384C-6(config-vsan-db)# vsan 40 interface fc2/1
SAN384C-6(config-vsan-db)# exit
SAN384C-6(config)#
SAN384C-6(config)# interface fc2/1
SAN384C-6(config-if)# switchport description CHPID 20 (PCHID 11C)
SAN384C-6(config-if)# switchport mode F
SAN384C-6(config-if)# switchport speed 16000
SAN384C-6(config-if)# switchport fec
```

```
SAN384C-6(config-if)# switchport fec tts
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)#
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6#

Configure the IBM disk ports in VSAN 40

SAN384C-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# vsan database
SAN384C-6(config-vsan-db)# vsan 40 interface fc1/17
SAN384C-6(config-vsan-db)# vsan 40 interface fc2/17
SAN384C-6(config-vsan-db)# exit
SAN384C-6(config)#
SAN384C-6(config)# interface fc1/17
SAN384C-6(config-if)# switchport description IBM 8870 Port 1
SAN384C-6(config-if)# switchport mode F
SAN384C-6(config-if)# switchport speed 16000
SAN384C-6(config-if)# switchport fec
SAN384C-6(config-if)# switchport fec tts
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)#
SAN384C-6(config-if)# interface fc2/17
SAN384C-6(config-if)# switchport description IBM 8870 Port 2
SAN384C-6(config-if)# switchport mode F
SAN384C-6(config-if)# switchport speed 16000
SAN384C-6(config-if)# switchport fec
SAN384C-6(config-if)# switchport fec tts
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)#
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6# show int fc1/17
fc1/17 is up
    Port description is IBM 8870 Port 1
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:11:00:3a:9c:31:62:80
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is F, FCID is 0x201000
    Port vsan is 40
    Admin Speed is 16 Gbps
    Operating Speed is 16 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 90
    Receive B2B Credit is 32
    B2B State Change: Admin(on), Oper(down)
    Receive data field Size is 2112
    Beacon is turned off
    admin fec state is on
    oper fec state is up
    Logical type is edge
    Peer is type 002107  model 961 manufactured by IBM, ficon tag is 0x0030
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      38 frames input,2556 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      38 frames output,2784 bytes
        0 discards,0 errors
      1 input OLS,1  LRR,2 NOS,0 loop inits
      5 output OLS,2 LRR, 0 NOS, 0 loop inits
    32 receive B2B credit remaining
    90 transmit B2B credit remaining
    90 low priority transmit B2B credit remaining
    Interface last changed at Sun Mar  8 22:49:07 2020

    Last clearing of "show interface" counters :  never

SAN384C-6#
SAN384C-6# show ficon vsan 40 portaddress 0x10
Port Address 16(0x10) is up in vsan 40
    Port number is 16(0x10), Interface is fc1/17
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 255(0xff)
    Admin port mode is F
    Port mode is F, FCID is 0x201000
    Peer is type 002107 model 961 manufactured by IBM
    Serial num is 0000000CPZ11, FICON tag is 0x0030

SAN384C-6# show int fc1/1
fc1/1 is up
    Port description is CHPID 18 (PCHID 118)
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:01:00:3a:9c:31:62:80
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is F, FCID is 0x200000
    Port vsan is 40
    Admin Speed is 16 Gbps
    Operating Speed is 16 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 90
    Receive B2B Credit is 32
    B2B State Change: Admin(on), Oper(down)
    Receive data field Size is 2112
    Beacon is turned off
    admin fec state is on
    oper fec state is up
    Logical type is edge
    Peer is type 002965  model N10 manufactured by IBM, ficon tag is 0x8018
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      2835 frames input,165420 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      4617 frames output,380120 bytes
        0 discards,0 errors
      5 input OLS,3  LRR,26 NOS,0 loop inits
      5 output OLS,1 LRR, 24 NOS, 0 loop inits
    32 receive B2B credit remaining
    90 transmit B2B credit remaining
    90 low priority transmit B2B credit remaining
    Interface last changed at Sun Mar  8 22:51:25 2020

    Last clearing of "show interface" counters :  never

SAN384C-6# show ficon vsan 40 portaddress 0x00
Port Address 0(0) is up in vsan 40
    Port number is 0(0), Interface is fc1/1
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 255(0xff)
    Admin port mode is F
    Port mode is F, FCID is 0x200000
    Peer is type 002965 model N10 manufactured by IBM
    Serial num is 0000000E8F77, FICON tag is 0x8018

SAN384C-6#
```

7.2.14 Fibre Channel over IP tunneling

FCIP is a protocol that is used to connect FC switches over a wide area network (WAN) to provide connectivity to remote geographical SAN locations. FCIP links are ISLs that transport FC control and data frames between switches by leveraging the FCIP tunneling protocol. IBM c-type switches IBM Storage Networking SAN768C-6, IBM Storage Networking SAN384C-6, and IBM Storage Networking SAN192C-6 support FCIP by using the 24/10 SAN Extension Modules with twenty-four 16 Gbps FC and eight 1/10 GbE ports, and the IP Storage (IPS) module on the IBM Storage Networking SAN50C-R switch with 40 ports of 16 Gbps FC and 10 ports of 1/10 GbE.

IP Security and Internet Key Exchange Protocols

In this section, we describe the common terms that are used with the IPsec and IKE protocols.

When configuring IPsec and IKE, two security associations (SAs) are required for outbound and inbound communication so that you can establish bidirectional communication between two participating switches to encrypt and decrypt IP packets. The security association database (SAD) stores sets of SA records.

The following information is included within the SA records:

- Security parameter index (SPI).
- IBM c-type switches that support IPsec.

- Transforms that provide data authentication and confidentiality, which are composed of an acceptable combination of security protocols, algorithms, and additional settings that can be applied to IPsec-protected traffic. When IPsec SA, negotiation takes place when the peers agree to use a particular transform set when protecting a data flow.
- A session key is used by the transform for security services.
- A lifetime counter tracks creation of SAs up until the time expires. If the time expires, the SAs are no longer operational, but can be automatically rekeyed if needed.
- Tunnel mode and Transport modes are the only two modes of operation that are generally available for IPsec. IPsec supports only tunnel mode, which provides secure communication paths between two switches on a configured FCIP link.
- Anti-replay is a security service where the receiver can reject old or duplicate packets to protect itself against replay attacks. IPsec provides this optional service by using a sequence number that is combined with using data authentication.
- Has data authentication, integrity, confidentiality, and the origin from which the data was sent.
- Perfect forward secrecy (PFS) is a cryptographic characteristic that is associated with a secret shared value.
- Security Policy Database (SPD) is a list of policies that decides whether a packet needs processing in clear text or if it should be dropped. IPsec SPD is part of the crypto map user configuration, and IKE SPD is configured by users.

IPsec

IPsec provides data confidentiality, data integrity, and data authentication between two participating switches. IPsec provides IP layer security services that protect one or more data flows between a pair of switches that are connected over an FCIP tunnel. IPsec in combination with IKE generates encryption and authentication keys. IPsec provides security for transmission at the network layer to protect authenticated IP packets between switches.

IPsec protects data that is transmitted across public networks from observation, modification, and spoofing, which allows virtual private networks (VPNs), intranets, extranets, and remote users access.

The Encapsulating Security Payload (ESP) protocol, which is a member of the IPsec suite, encapsulates the data to be protected and provides data privacy services, data authentication, and optional anti-replay services.

IKE negotiates IPsec SAs and generates keys for switches by using the IPsec feature and allows you to refresh IPsec SAs, which provides dynamic authentication of peers and anti-replay services while supporting a manageable and scalable IPsec configuration.

> **Note:** When implementing IPsec and IKE, each GbE interface on the IBM Storage Networking SAN50C-R switch and the 24/10-Port SAN Extension Module must be configured in its own IP subnet to ensure that the IPsec tunnel works.

7.2.15 Using the FCIP wizard in DCNM

When configuring the FCIP feature, you must explicitly enable all the switches in the fabric. The FCIP feature is disabled by default. The configuration and verification command suite is available only when FCIP is enabled on a switch.

Important: The ENTERPRISE_PKG license is required to configure IPsec on the IBM Storage Networking SAN50C-R switch or the 24/10-Port SAN Extension Module.

Note: A best practice for configuring an IP route is to configure a static route to each GbE interface by designating a subnet mask of 255.255.255.255 when you add the route.

To create and manage FCIP links with DCNM 11, use the FCIP wizard. Make sure that you can ping the GbE interfaces from local and remote switches to verify connectivity.

Before implementing FCIP, complete the following steps:

1. Check Fabric Health.
2. Back up the switch configuration.
3. Document the FCIP and IPsec configuration.
4. Perform a post-backup of the switch configuration.

Important: If you encounter a problem with your configuration, do not automatically restore the FCIP and IPsec configuration backup because it restores the entire switch configuration, which might impact the existing FCIP links that function properly. For help, contact IBM Support.

Example 7-7 shows how to back up the switch configuration

Example 7-7 Backing up the running switch configuration

```
switch#copy running config to startup-config - Fabric X.
switch#copy running config to tftp server - Fabric X.
switch#
```

To create FCIP links by using the FCIP wizard, complete the following steps:

1. Go to the Welcome window.
2. Select the switch pairs.
3. Specify the IP Address/Route.
4. Specify the Tunnel Properties.
5. Create the FCIP ISL.

7.2.16 Configuring FCIP links per IPS port by using the DCNM

This section shows how to configure a single FCIP link on an IPS port by using the Cisco DCNM wizard.

In the DCNM GUI, select **Configure** → **FCIP** to access the FCIP wizard, as shown in Figure 7-111.

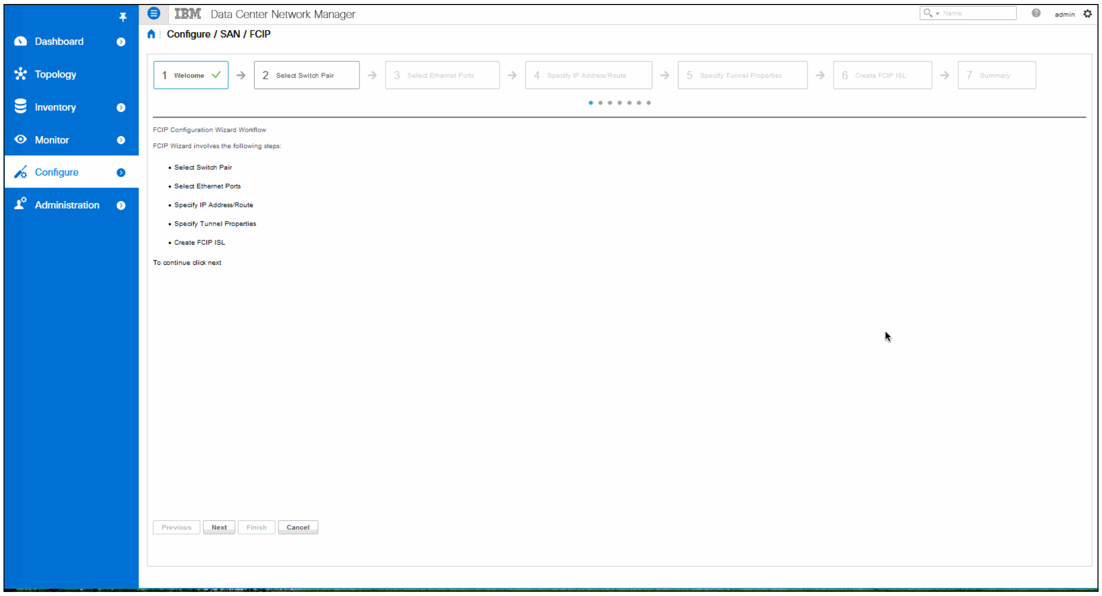

FCIP Configuration Wizard Workflow

FCIP Wizard involves the following steps:

- Select Switch Pair
- Select Ethernet Ports
- Specify IP Address/Route
- Specify Tunnel Properties
- Create FCIP ISL

To continue click next

Figure 7-111 DCNM FCIP wizard

We select the two switch end points that we will use to create our FCIP tunnel, as shown in Figure 7-112.

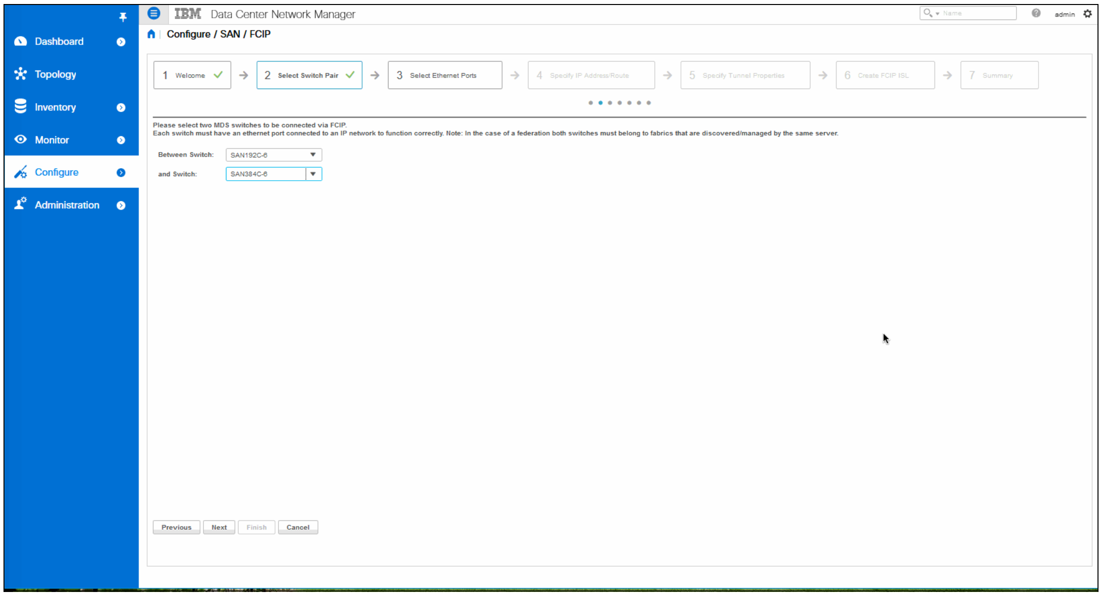

Please select two MDS switches to be connected via FCIP.
Each switch must have an ethernet port connected to an IP network to function correctly. Note: In the case of a federation both switches must belong to fabrics that are discovered/managed by the same server.

Between Switch: SAN192C-6

and Switch: SAN384C-6

Figure 7-112 Switch pairs

Figure 7-113 on page 341 shows the selection of which 10 GbE IPS port will be used for the new FCIP tunnel interface that will provide physical connectivity between our two end-point switches, IBM Storage Networking SAN384C-6 and IBM Storage Networking SAN192C-6, by using the 24/10 SAN extension module. We also select **(Jumbo Frames)**. Click **Next**.

Note: As a best practice, use Jumbo Frames whenever the network allows it because they increase the performance of the FCIP tunnel. By default, Ethernet has a 1500-byte maximum transmission unit (MTU) size, which means that each FC frame must be segmented to send, and be reassembled upon receipt, which can add extra latency.

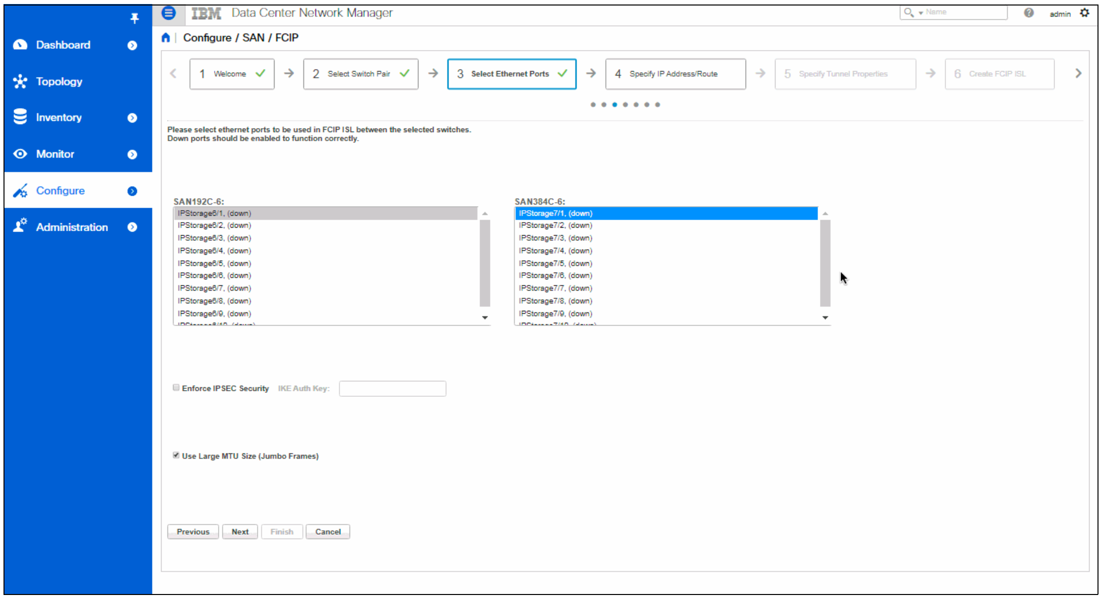

Please select ethernet ports to be used in FCIP ISL between the selected switches.
Down ports should be enabled to function correctly.

SAN192C-6:	SAN384C-6:
IPStorage6/1, (down)	IPStorage7/1, (down)
IPStorage6/2, (down)	IPStorage7/2, (down)
IPStorage6/3, (down)	IPStorage7/3, (down)
IPStorage6/4, (down)	IPStorage7/4, (down)
IPStorage6/5, (down)	IPStorage7/5, (down)
IPStorage6/6, (down)	IPStorage7/6, (down)
IPStorage6/7, (down)	IPStorage7/7, (down)
IPStorage6/8, (down)	IPStorage7/8, (down)
IPStorage6/9, (down)	IPStorage7/9, (down)

□ Enforce IPSEC Security IKE Auth Key:

☑ Use Large MTU Size (Jumbo Frames)

Figure 7-113 Selecting Ethernet ports

As shown in Figure 7-114, we provide the IP addresses for each of the IP storage ports on each end point. Routes are not needed in this example because the IP addresses are in the same subnet. Click **Next**.

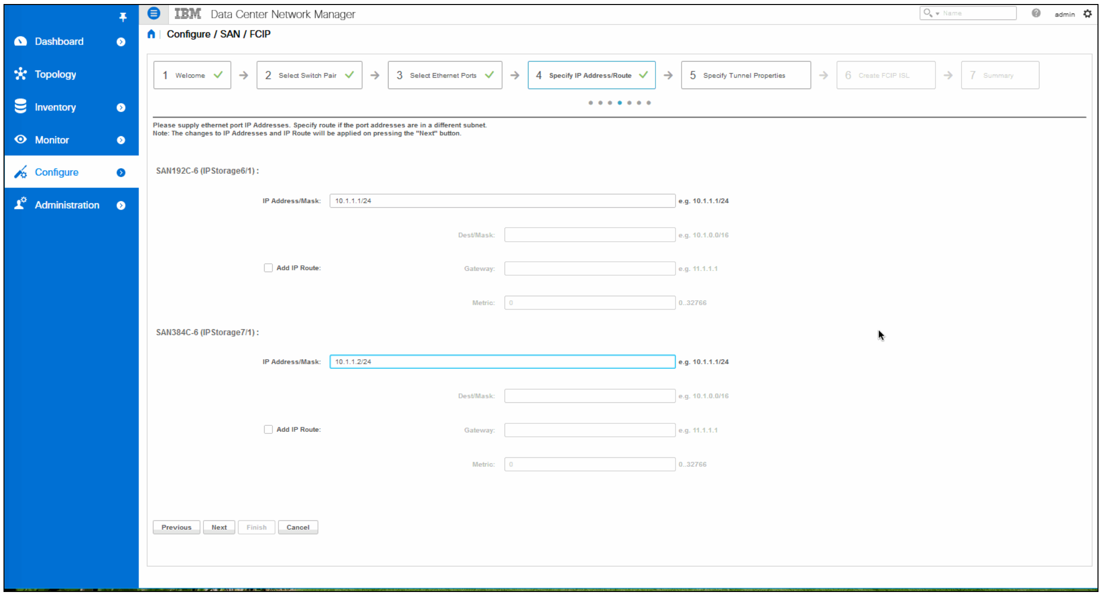

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN192C-6 (IPStorage6/1) :

IP Address/Mask: 10.1.1.1/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

SAN384C-6 (IPStorage7/1) :

IP Address/Mask: 10.1.1.2/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

Figure 7-114 Specifying the IP addresses

Click **Yes** to continue, as shown in Figure 7-115.

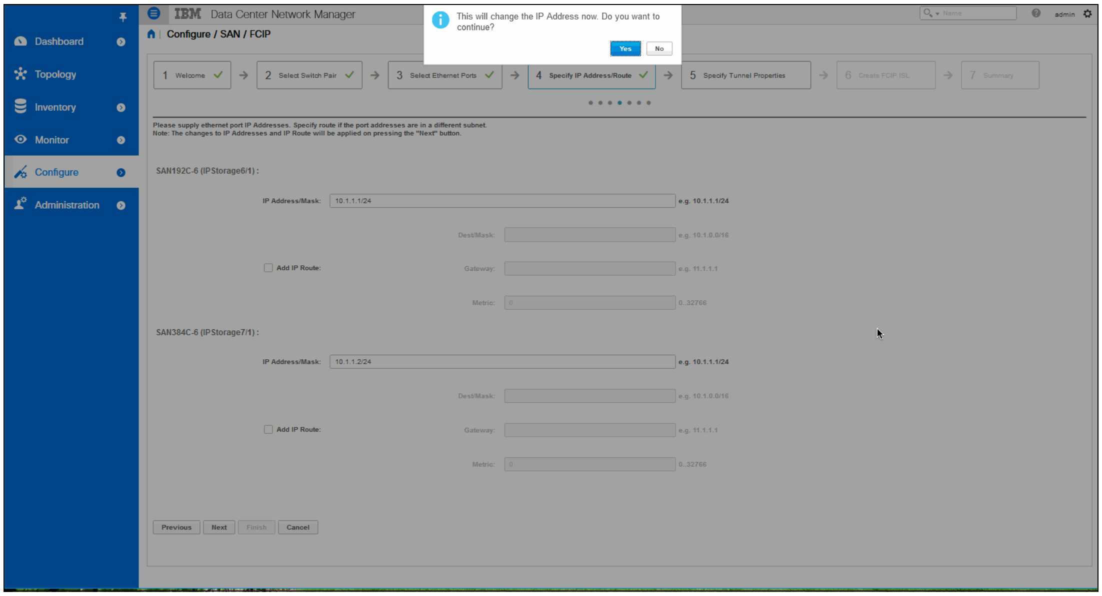

This will change the IP Address now. Do you want to continue?

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN192C-6 (IPStorage6/1) :

IP Address/Mask: 10.1.1.1/24 e.g. 10.1.1.1/24
Dest/Mask: e.g. 10.1.0.0/16
□ Add IP Route: Gateway: e.g. 11.1.1.1
Metric: 0 0..32766

SAN384C-6 (IPStorage7/1) :

IP Address/Mask: 10.1.1.2/24 e.g. 10.1.1.1/24
Dest/Mask: e.g. 10.1.0.0/16
□ Add IP Route: Gateway: e.g. 11.1.1.1
Metric: 0 0..32766

Figure 7-115 IP addresses

In the tunnel properties window, we select **Measure RTT** (round-trip time), as shown in Figure 7-116. Click **Close** to continue.

Important: Measure RTT is used to test the network connection and provide the time that it takes for a packet to cross the network and return an acknowledgment.

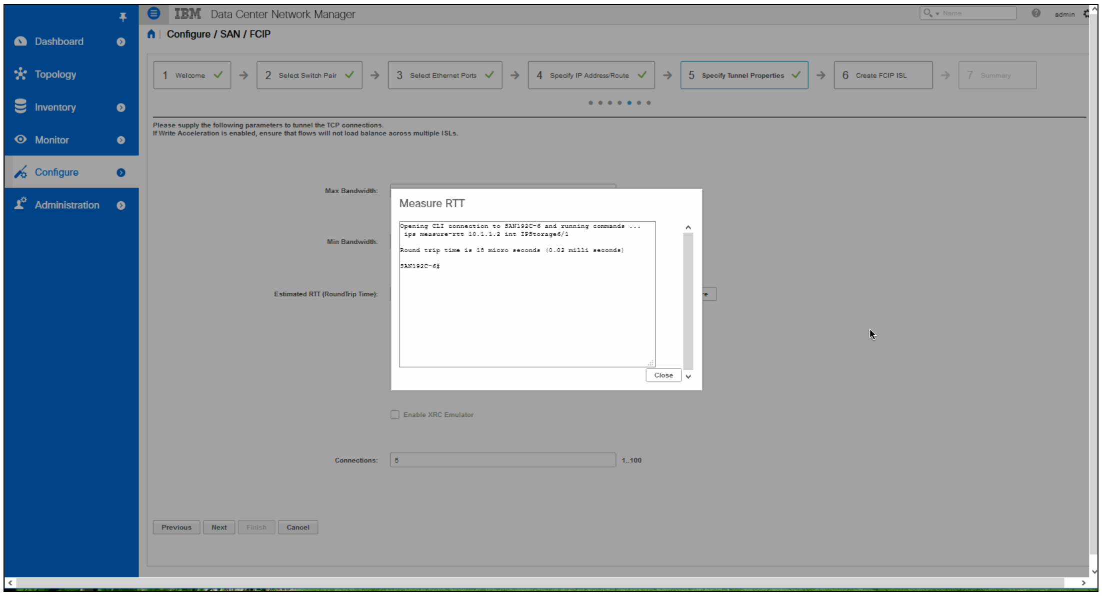

Please supply the following parameters to tunnel the TCP connections.
If Write Acceleration is enabled, ensure that flows will not load balance across multiple ISLs.

Max Bandwidth:

Min Bandwidth:

Estimated RTT (RoundTrip Time):

Measure RTT

```
Opening CLI connection to SAN192C-6 and running commands ...
 ips measure-rtt 10.1.1.2 int IPStorage6/1

Round trip time is 18 micro seconds (0.02 milli seconds)

SAN192C-6#
```

□ Enable XRC Emulator

Connections: 5 1..100

Figure 7-116 Measure RTT window

In Figure 7-117, we provide the Minimum and Maximum bandwidth for the FCIP tunnel and the round-trip time (RTT) for the FCIP tunnel. In this lab example, we defined a 10 GbE dedicated network bandwidth. The maximum is set at 10 GbE, and the minimum is set at 95% of the maximum. Click **Next** to continue.

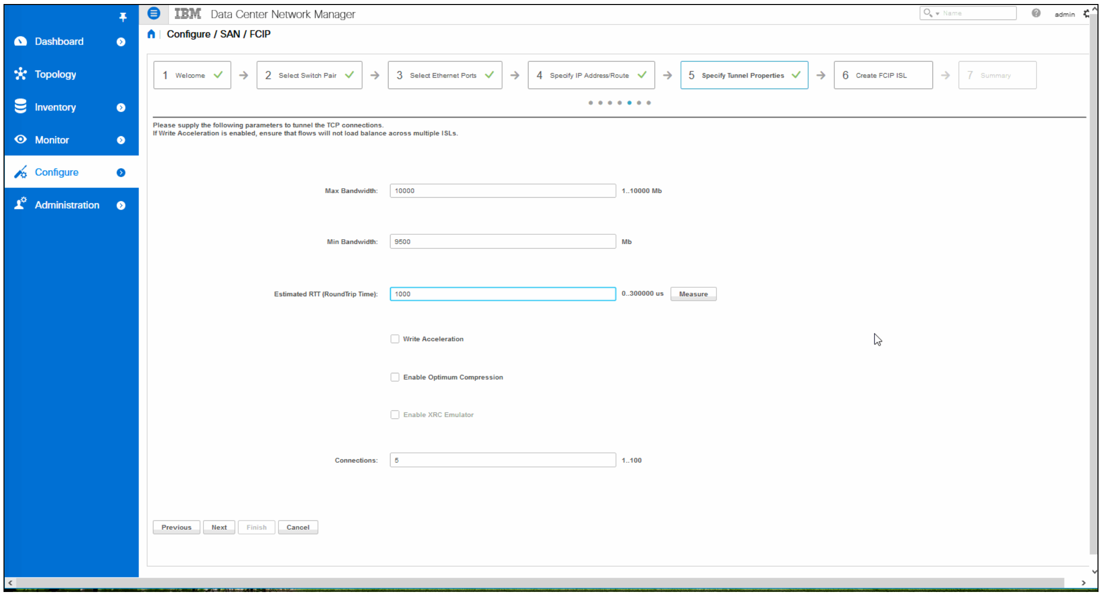

Please supply the following parameters to tunnel the TCP connections.
If Write Acceleration is enabled, ensure that flows will not load balance across multiple ISLs.

Max Bandwidth: 10000 1..10000 Mb

Min Bandwidth: 9500 Mb

Estimated RTT (RoundTrip Time): 1000 0..300000 us Measure

□ Write Acceleration

□ Enable Optimum Compression

□ Enable XRC Emulator

Connections: 5 1..100

Figure 7-117 Tunnel properties

In Figure 7-118, we specify the final parameters for the FCIP link configuration:

- **Profile ID:** 100, which provides detailed information about the local IP address and TCP parameters.
- **Tunnel ID:** 100, which is used to create the name of the new FCIP interface, fcip100.
- **FICON Port Address:** 0xF1, which is only applicable when FICON is enabled on the switch. The FICON Port Address must be configured when there is a FICON VSAN communicating on this FCIP tunnel. As a best practice, use the same value on both switches for the FCIP tunnel when possible. The value of this attribute must be taken from the pool of logical FICON Port Addresses.
- **Trunk Mode:** As a best practice, use **Trunk**.
- **VSAN List:** Should be 1 and the value of any VSANs that require access to the FCIP tunnel. In this example, we use VSAN 50.

To proceed, click **Next**.

> **Note:** As a best practice, keep the Profile and Tunnel IDs the same on both switches when creating an FCIP configuration. In addition, when FCIP links carry FICON traffic, the FICON Port Addresses should be the same.

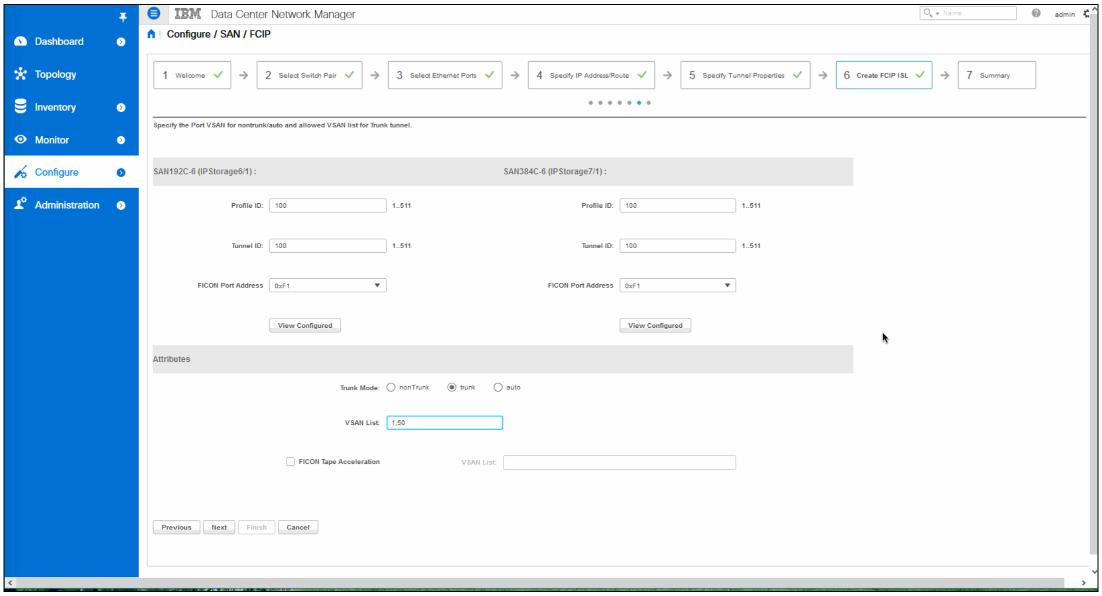

Specify the Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.

SAN192C-6 (IPStorage6/1) :

Profile ID: 100 1..511

Tunnel ID: 100 1..511

FICON Port Address 0xF1

SAN384C-6 (IPStorage7/1) :

Profile ID: 100 1..511

Tunnel ID: 100 1..511

FICON Port Address 0xF1

Attributes

Trunk Mode: ○ nonTrunk ● trunk ○ auto

VSAN List: 1,50

□ FICON Tape Acceleration VSAN List:

Figure 7-118 Creating an FCIP ISL

Figure 7-119 on page 345 shows a summary of the configuration that will be applied to create the single FCIP link. To proceed, click **Finish**.

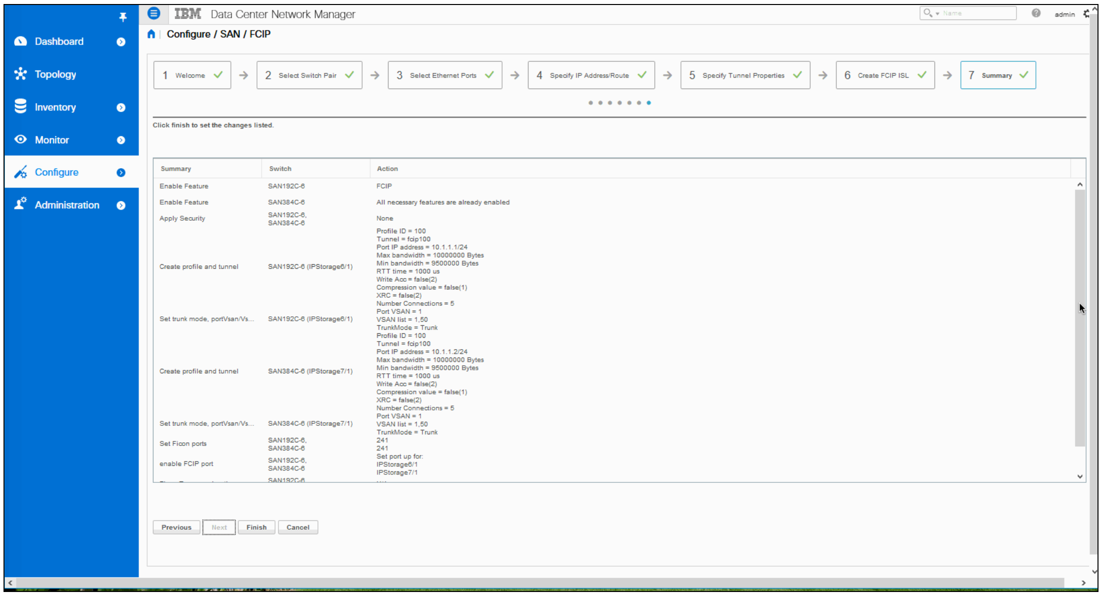

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN192C-6	FCIP
Enable Feature	SAN384C-6	All necessary features are already enabled
Apply Security	SAN192C-6, SAN384C-6	None
Create profile and tunnel	SAN192C-6 (IPStorage6/1)	Profile ID = 100 Tunnel = fcip100 Port IP address = 10.1.1.1/24 Max bandwidth = 10000000 Bytes Min bandwidth = 9500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN192C-6 (IPStorage6/1)	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/1)	Profile ID = 100 Tunnel = fcip100 Port IP address = 10.1.1.2/24 Max bandwidth = 10000000 Bytes Min bandwidth = 9500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/1)	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Set Ficon ports	SAN192C-6, SAN384C-6	241 241
enable FCIP port	SAN192C-6, SAN384C-6	Set port up for: IPStorage6/1 IPStorage7/1
[illegible]	SAN192C-6	[illegible]

Figure 7-119 Summary review

Figure 7-120 shows that the configuration successfully completed. To proceed, click **OK** and then **Close**.

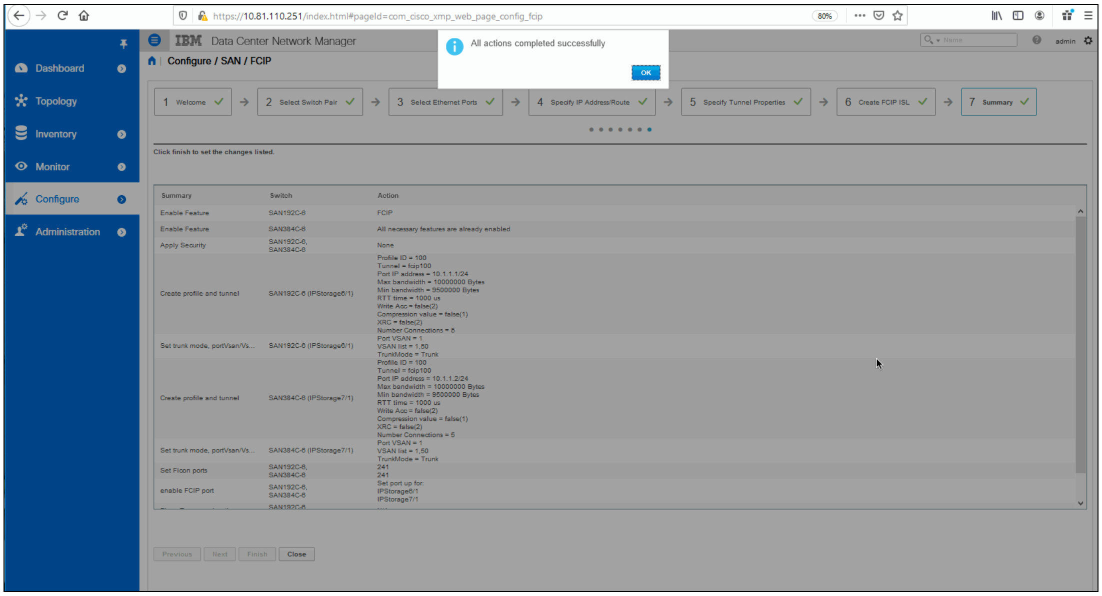

All actions completed successfully

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN192C-6	FCIP
Enable Feature	SAN384C-6	All necessary features are already enabled
Apply Security	SAN192C-6, SAN384C-6	None
Create profile and tunnel	SAN192C-6 (IPStorage6/1)	Profile ID = 100 Tunnel = fcip100 Port IP address = 10.1.1.1/24 Max bandwidth = 10000000 Bytes Min bandwidth = 9500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN192C-6 (IPStorage6/1)	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/1)	Profile ID = 100 Tunnel = fcip100 Port IP address = 10.1.1.2/24 Max bandwidth = 10000000 Bytes Min bandwidth = 9500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/1)	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Set Ficon ports	SAN192C-6, SAN384C-6	241 241
enable FCIP port	SAN192C-6, SAN384C-6	Set port up for: IPStorage6/1 IPStorage7/1
[illegible]	SAN192C-6	[illegible]

Figure 7-120 Successfully completed

Figure 7-121 shows the new FCIP link from the IBM Storage Networking SAN192C-6 switch (fcip100) to the IBM Storage Networking SAN384C-6 (fcip100) switch.

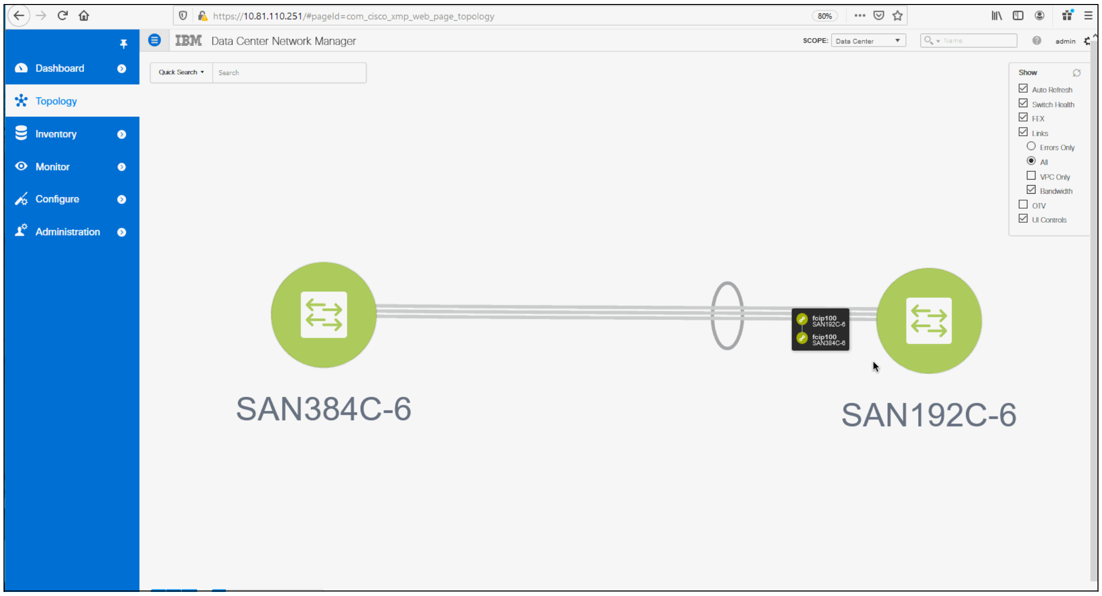

Figure 7-121 Showing the FCIP link

Configuring the FCIP link on the second IPS port by using the CLI

Example 7-8 shows how to configure an FCIP link on the second IPS port by using the CLI.

Example 7-8 Configuring a single FCIP link on the second IPS port by using the CLI.

```
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# feature fcip
SAN384C-6(config)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6# show feature | incl fcip
fcip                  1        enabled
SAN384C-6#

SAN384C-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# interface IPStorage7/2
SAN384C-6(config-if)# switchport mtu 2300
SAN384C-6(config-if)# ip address 10.1.2.2 255.255.255.0
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)#
SAN384C-6(config-if)# fcip profile 110
SAN384C-6(config-profile)# ip address 10.1.2.2
SAN384C-6(config-profile)#
SAN384C-6(config-profile)# tcp max-bandwidth-mbps 10000
min-available-bandwidth-mbps 9500  round-trip-time-ms 1
SAN384C-6(config-profile)#
SAN384C-6(config-profile)# interface fcip110
SAN384C-6(config-if)# use-profile 110
SAN384C-6(config-if)#
SAN384C-6(config-if)# peer-info ipaddr 10.1.2.1
```

```
SAN384C-6(config-if)#
SAN384C-6(config-if)# tcp-connections 5
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)# switchport trunk allowed vsan add 50
SAN384C-6(config-if)#
SAN384C-6(config-if)# ficon portnumber 0xf2
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shutdown
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6# show fcip summary

-------------------------------------------------------------------------------
Tun prof    IPS-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
-------------------------------------------------------------------------------
100 100  IPS7/1     10.1.1.1         TRNK  Y N N  N   N   10000M/9500M  1000
110 110  IPS7/2     10.1.2.1         TRNK  Y N N  N   N   10000M/9500M  1000

SAN384C-6# show int fcip110
fcip110 is trunking
    Hardware is IPStorage
    Port WWN is 21:9e:00:3a:9c:31:62:80
    Peer port WWN is 21:5e:00:2a:6a:a4:1a:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Operating Speed is 10000 Mbps
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    Interface last changed at Mon Mar  9 00:23:06 2020

    Using Profile id 110  (interface IPStorage7/2)
    Peer Information
      Peer Internet address is 10.1.2.1 and port is 3225
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    FICON Load Balancer configured off for all vsans
    FICON Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 5
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      5 Active TCP connections
        12 Attempts for active connections, 0 close of connections
```

```
       Path MTU 2300 bytes
       Current retransmission timeout is 200 ms
       Current Send Buffer Size: 149580 KB, Requested Send Buffer Size: 125000 KB
       CWM Burst Size: 50 KB
CONN<0>
   Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65534
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 310 KB, Maximum: 310 KB, Scale: 7
     Congestion window: Current: 279 KB, Slow start threshold: 285 KB
     Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
     Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
   Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65533
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 58 KB, Maximum: 58 KB, Scale: 7
     Congestion window: Current: 52 KB, Slow start threshold: 275 KB
     Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
     Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
   Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65531
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 58 KB, Maximum: 58 KB, Scale: 7
     Congestion window: Current: 52 KB, Slow start threshold: 275 KB
     Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
     Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
   Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65529
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 58 KB, Maximum: 58 KB, Scale: 7
     Congestion window: Current: 52 KB, Slow start threshold: 275 KB
     Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
     Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
   Control connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65527
   TCP Parameters
     Advertized window: Current: 8182 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 50 KB, Maximum: 50 KB, Scale: 7
     Congestion window: Current: 50 KB, Slow start threshold: 275 KB
     Measured RTT : 12 us Min RTT: 13 us Max RTT: 13 us
```

```
      Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
    5 minutes input rate 248 bits/sec, 31 bytes/sec, 0 frames/sec
    5 minutes output rate 216 bits/sec, 27 bytes/sec, 0 frames/sec
      82 frames input, 9764 bytes
         82 Class F frames input, 9764 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      83 frames output, 8504 bytes
         83 Class F frames output, 8504 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames

SAN384C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
          fcip100  0x11(17)          fcip100  10.122.107.94(SAN192C-6)
          fcip110  0x11(17)          fcip110  10.122.107.94(SAN192C-6)
SAN384C-6#
SAN384C-6# show ficon vsan 50 portaddress 0xF2
Port Address 242(0xf2) is up in vsan 50
    Port number is 242(0xf2), Interface is fcip110
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 255(0xff)
    Admin port mode is auto
    Port mode is TE
    Peer is type OMDS9K model 706 manufactured by CSC
    Serial num is 002A6AA41A82, FICON tag is 0x00F2

And on the other switch

SAN192C-6#
SAN192C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# interface IPStorage6/2
SAN192C-6(config-if)# switchport mtu 2300
SAN192C-6(config-if)# ip address 10.1.2.1 255.255.255.0
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)#
SAN192C-6(config-if)# fcip profile 110
SAN192C-6(config-profile)# ip address 10.1.2.1
SAN192C-6(config-profile)#
SAN192C-6(config-profile)# tcp max-bandwidth-mbps 10000
min-available-bandwidth-mbps 9500  round-trip-time-ms 1
SAN192C-6(config-profile)# interface fcip110
SAN192C-6(config-if)# use-profile 110
SAN192C-6(config-if)# peer-info ipaddr 10.1.2.2
SAN192C-6(config-if)# tcp-connections 5
```

```
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)# switchport trunk allowed vsan add 50
SAN192C-6(config-if)# ficon portnumber 0xf2
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)#
SAN192C-6(config-if)# end
Performing fast copy of configurationdone.
SAN192C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
          fcip100  0x10(16)          fcip100  10.122.107.95(SAN384C-6)
          fcip110  0x10(16)          fcip110  10.122.107.95(SAN384C-6)
SAN192C-6#
```

7.2.17 Configuring an FCIP Port Channel by using the DCNM

This section shows how to configure an FCIP Port Channel by using the SAN switches IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6, and the fcip100, fcip110 ISLs, by using the DCNM.

Select the **Create New Port Channel**, as shown in Figure 7-122.

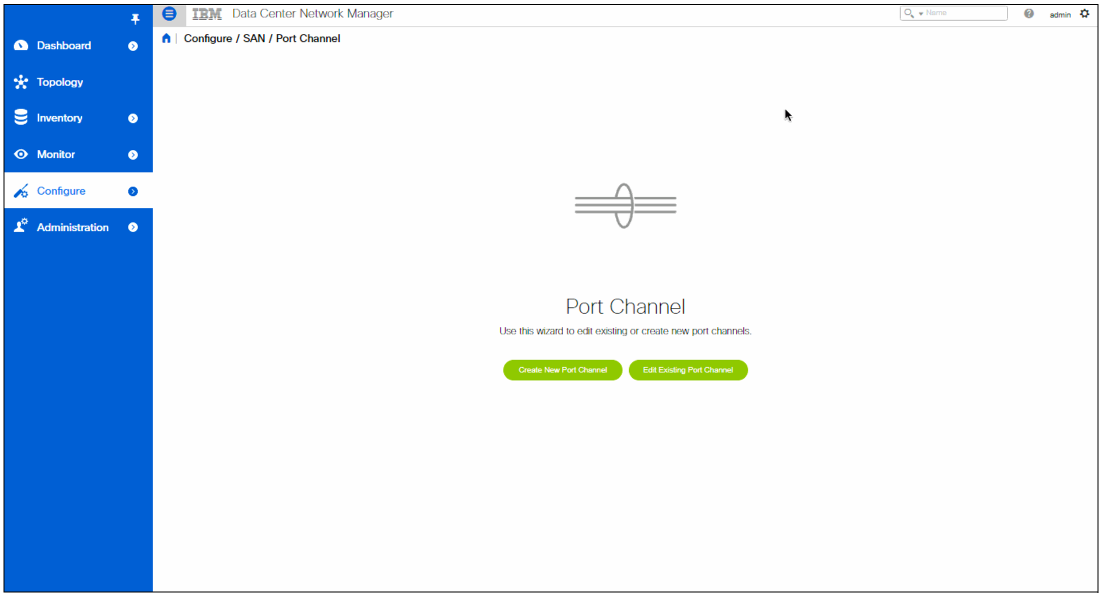

Configure / SAN / Port Channel

Port Channel

Use this wizard to edit existing or create new port channels.

Figure 7-122 Port Channel wizard

Select the switch pairs IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 to participate in the Port Channel creation by using previously created ISLs, as shown in Figure 7-123 on page 351.

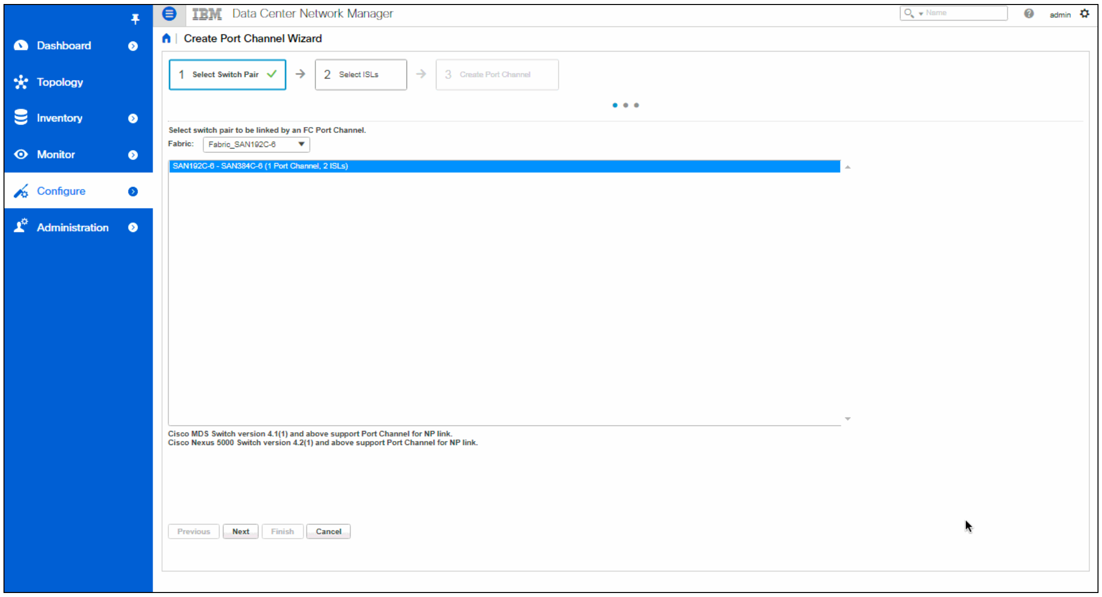

Select switch pair to be linked by an FC Port Channel.

Fabric: Fabric_SAN192C-6

SAN192C-6 - SAN384C-6 (1 Port Channel, 2 ISLs)

Cisco MDS Switch version 4.1(1) and above support Port Channel for NP link.
Cisco Nexus 5000 Switch version 4.2(1) and above support Port Channel for NP link.

Figure 7-123 Selecting a switch pair

Select fcip 100 and fcip 110 to provide ISL redundancy within the FCIP Port Channel, as shown in Figure 7-124.

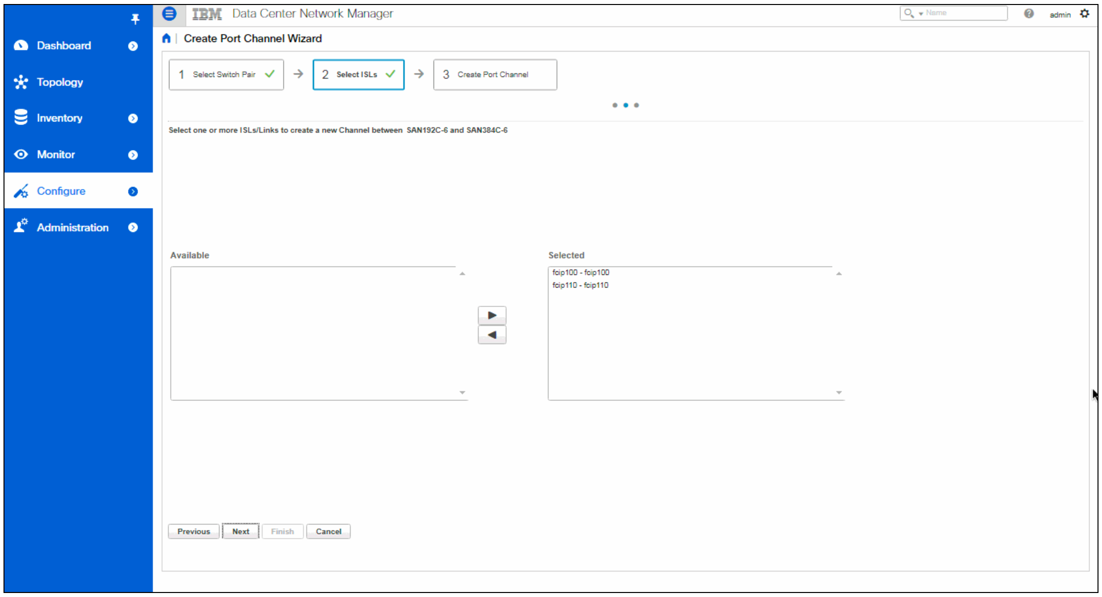

Select one or more ISLs/Links to create a new Channel between SAN192C-6 and SAN384C-6

Selected:
fcip100 - fcip100
fcip110 - fcip110

Figure 7-124 Selecting the ISLs

Review the Port Channel attributes for the selected switches, as shown in Figure 7-125. Click **Finish** to continue.

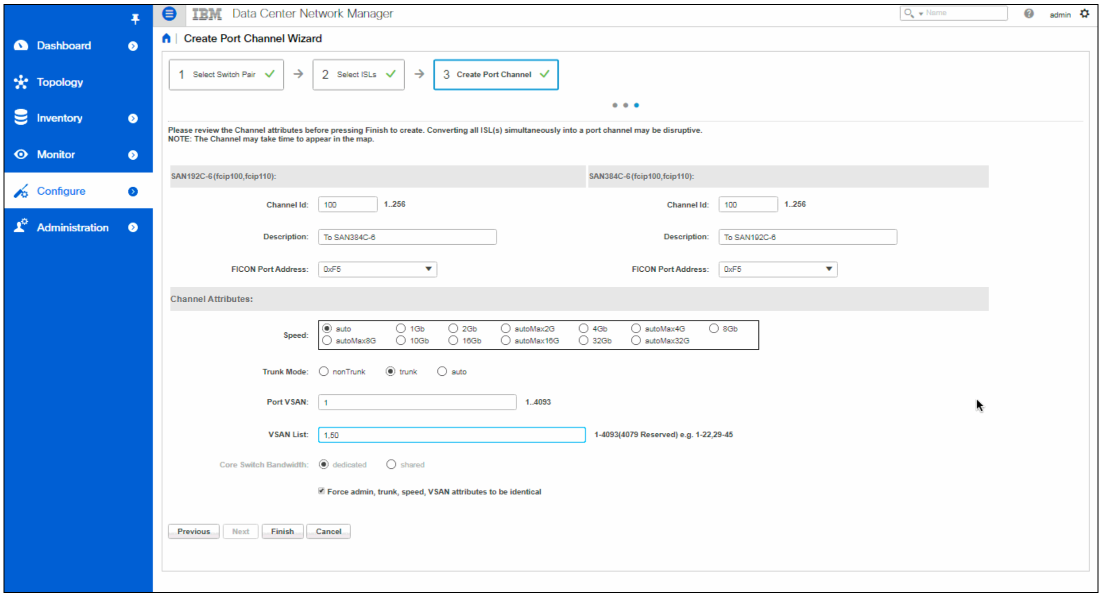

Please review the Channel attributes before pressing Finish to create. Converting all ISL(s) simultaneously into a port channel may be disruptive.
NOTE: The Channel may take time to appear in the map.

SAN192C-6(fcip100,fcip110):

Channel Id: 100 1..256

Description: To SAN384C-6

FICON Port Address: 0xF5

SAN384C-6(fcip100,fcip110):

Channel Id: 100 1..256

Description: To SAN192C-6

FICON Port Address: 0xF5

Channel Attributes:

Speed: ● auto ○ 1Gb ○ 2Gb ○ autoMax2G ○ 4Gb ○ autoMax4G ○ 8Gb ○ autoMax8G ○ 10Gb ○ 16Gb ○ autoMax16G ○ 32Gb ○ autoMax32G

Trunk Mode: ○ nonTrunk ● trunk ○ auto

Port VSAN: 1 1..4093

VSAN List: 1,50 1-4093(4079 Reserved) e.g. 1-22,29-45

Core Switch Bandwidth: ● dedicated ○ shared

☑ Force admin, trunk, speed, VSAN attributes to be identical

Figure 7-125 Creating the Port Channel

Figure 7-126 shows that Port Channel creation completed successfully. Click **Close**.

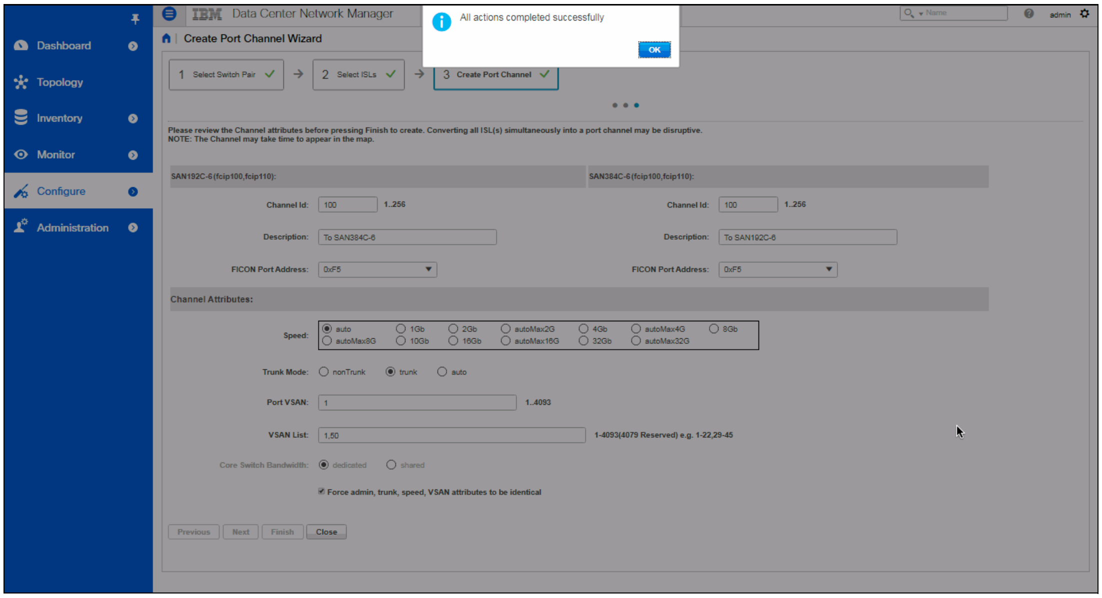

All actions completed successfully

Figure 7-126 Port Channel created successfully

Configuring FCIP Port Channels by using the CLI

Example 7-9 on page 353 shows how to configure an FCIP Port Channel on the IBM Storage Networking SAN384C-6 primary switch by using the CLI.

Example 7-9 Configuring FCIP Port Channels on the IBM Storage Networking SAN384C-6 primary switch

```
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# interface port-channel 100
SAN384C-6(config-if)# ficon portnumber 0xF5
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk on
ly the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)# switchport trunk allowed vsan add 50
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)# interface fcip 100
SAN384C-6(config-if)# channel-group 100 force
fcip100 added to port-channel 100 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN384C-6(config-if)# int fcip 110
SAN384C-6(config-if)# channel-group 100 force
fcip110 added to port-channel 100 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN384C-6(config-if)# Performing fast copy of configurationdone.

SAN384C-6# show interface port-channel 100
port-channel100 is trunking
    Hardware is IPStorage
    Port WWN is 24:64:00:3a:9c:31:62:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 20 Gbps
    Logical type is core
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 424 bits/sec, 53 bytes/sec, 0 frames/sec
    5 minutes output rate 368 bits/sec, 46 bytes/sec, 0 frames/sec
      1452 frames input, 167716 bytes
         1452 Class F frames input, 167716 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      1467 frames output, 154764 bytes
         1467 Class F frames output, 154764 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
    Member[1] : fcip100   [up] *
    Member[2] : fcip110   [up]

SAN384C-6# show interface fcip 100
fcip100 is trunking
    Hardware is IPStorage
    Port WWN is 21:9a:00:3a:9c:31:62:80
    Peer port WWN is 21:5a:00:2a:6a:a4:1a:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Operating Speed is 10000 Mbps
    Belongs to port-channel100
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    Interface last changed at Fri Jan 17 20:51:40 2020

    Using Profile id 100  (interface IPStorage7/1)
    Peer Information
      Peer Internet address is 10.1.1.1 and port is 3225
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    FICON Load Balancer configured off for all vsans
    FICON Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 5
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      5 Active TCP connections
        25 Attempts for active connections, 3 close of connections
        Path MTU 2300 bytes
        Current retransmission timeout is 200 ms
        Current Send Buffer Size: 149580 KB, Requested Send Buffer Size: 125000
KB
        CWM Burst Size: 50 KB
 CONN<0>
    Data connection: Local 10.1.1.2:3225, Remote 10.1.1.1:65498
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 310 KB, Maximum: 310 KB, Scale: 7
      Congestion window: Current: 279 KB, Slow start threshold: 294 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<1>
    Data connection: Local 10.1.1.2:3225, Remote 10.1.1.1:65497
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 57 KB, Maximum: 57 KB, Scale: 7
      Congestion window: Current: 51 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
```

```
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<2>
    Data connection: Local 10.1.1.2:3225, Remote 10.1.1.1:65495
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 58 KB, Maximum: 58 KB, Scale: 7
      Congestion window: Current: 52 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<3>
    Data connection: Local 10.1.1.2:3225, Remote 10.1.1.1:65493
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 57 KB, Maximum: 57 KB, Scale: 7
      Congestion window: Current: 51 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<4>
    Control connection: Local 10.1.1.2:3225, Remote 10.1.1.1:65491
    TCP Parameters
      Advertized window: Current: 8178 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 50 KB, Maximum: 50 KB, Scale: 7
      Congestion window: Current: 50 KB, Slow start threshold: 275 KB
      Measured RTT : 11 us Min RTT: 12 us Max RTT: 13 us
      Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
    5 minutes input rate 352 bits/sec, 44 bytes/sec, 0 frames/sec
    5 minutes output rate 312 bits/sec, 39 bytes/sec, 0 frames/sec
      963 frames input, 111700 bytes
         963 Class F frames input, 111700 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      976 frames output, 104976 bytes
         976 Class F frames output, 104976 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames

SAN384C-6# show interface fcip 110
fcip110 is trunking
    Hardware is IPStorage
    Port WWN is 21:9e:00:3a:9c:31:62:80
    Peer port WWN is 21:5e:00:2a:6a:a4:1a:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Operating Speed is 10000 Mbps
    Belongs to port-channel100
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    Interface last changed at Fri Jan 17 20:52:06 2020

    Using Profile id 110  (interface IPStorage7/2)
    Peer Information
      Peer Internet address is 10.1.2.1 and port is 3225
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    FICON Load Balancer configured off for all vsans
    FICON Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 5
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      5 Active TCP connections
        30 Attempts for active connections, 4 close of connections
        Path MTU 2300 bytes
        Current retransmission timeout is 200 ms
        Current Send Buffer Size: 149580 KB, Requested Send Buffer Size: 125000
KB
        CWM Burst Size: 50 KB
 CONN<0>
    Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65496
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 310 KB, Maximum: 310 KB, Scale: 7
      Congestion window: Current: 279 KB, Slow start threshold: 288 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<1>
    Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65495
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 57 KB, Maximum: 57 KB, Scale: 7
      Congestion window: Current: 51 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<2>
    Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65493
    TCP Parameters
```

```
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 58 KB, Maximum: 58 KB, Scale: 7
      Congestion window: Current: 52 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<3>
    Data connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65491
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 58 KB, Maximum: 58 KB, Scale: 7
      Congestion window: Current: 52 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<4>
    Control connection: Local 10.1.2.2:3225, Remote 10.1.2.1:65489
    TCP Parameters
      Advertized window: Current: 8189 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 56 KB, Maximum: 56 KB, Scale: 7
      Congestion window: Current: 50 KB, Slow start threshold: 275 KB
      Measured RTT : 500000 us Min RTT: 13 us Max RTT: 0 us
      Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
    5 minutes input rate 72 bits/sec, 9 bytes/sec, 0 frames/sec
    5 minutes output rate 56 bits/sec, 7 bytes/sec, 0 frames/sec
      497 frames input, 56752 bytes
         497 Class F frames input, 56752 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      499 frames output, 50524 bytes
         499 Class F frames output, 50524 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames

SAN384C-6#
SAN384C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
  port-channel100  0x11(17)  port-channel100  10.122.107.94(SAN192C-6)
SAN384C-6#
```

Configuring a Port Channel on the IBM Storage Networking SAN192C-6 partner switch by using the CLI

Example 7-10 shows how to configure a Port Channel on the IBM Storage Networking SAN192C-6 partner switch by using the CLI.

Example 7-10 Configuring a Port Channel on the IBM Storage Networking SAN192C-6 partner switch by using the CLI

```
SAN192C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# interface port-channel 100
SAN192C-6(config-if)# switchport tr
trunk                  trunk-max-npiv-limit
SAN192C-6(config-if)# ficon portnumber 0xF5
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk on
ly the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)# switchport trunk
trunk                  trunk-max-npiv-limit
SAN192C-6(config-if)# switchport trunk allowed vsan add 50
SAN192C-6(config-if)# no shut
SAN192C-6(config-if)# int fcip 100
SAN192C-6(config-if)# channel-group 100 force
fcip100 added to port-channel 100 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN192C-6(config-if)# interface fcip 110
SAN192C-6(config-if)# channel-group 100 force
fcip110 added to port-channel 100 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SAN192C-6(config-if)# Performing fast copy of configurationdone.

SAN192C-6# show in
in-order-guarantee    incompatibility-all   inventory
inactive-if-config    install
incompatibility       interface
SAN192C-6# show interface port-channel 100
port-channel100 is trunking
    Hardware is IPStorage
    Port WWN is 24:64:00:2a:6a:a4:1a:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 20 Gbps
    Logical type is core
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 368 bits/sec, 46 bytes/sec, 0 frames/sec
    5 minutes output rate 424 bits/sec, 53 bytes/sec, 0 frames/sec
      1455 frames input, 152316 bytes
```

```
         1455 Class F frames input, 152316 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      1472 frames output, 170900 bytes
         1472 Class F frames output, 170900 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
    Member[1] : fcip100   [up] *
    Member[2] : fcip110   [up]

SAN192C-6# show interface fcip 100
fcip100 is trunking
    Hardware is IPStorage
    Port WWN is 21:5a:00:2a:6a:a4:1a:80
    Peer port WWN is 21:9a:00:3a:9c:31:62:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Operating Speed is 10000 Mbps
    Belongs to port-channel100
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    Interface last changed at Fri Jan 17 20:52:40 2020

    Using Profile id 100  (interface IPStorage6/1)
    Peer Information
      Peer Internet address is 10.1.1.2 and port is 3225
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    FICON Load Balancer configured off for all vsans
    FICON Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 5
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      5 Active TCP connections
        25 Attempts for active connections, 6 close of connections
        Path MTU 2300 bytes
        Current retransmission timeout is 200 ms
        Current Send Buffer Size: 125310 KB, Requested Send Buffer Size: 125000
KB
        CWM Burst Size: 50 KB
 CONN<0>
    Data connection: Local 10.1.1.1:65498, Remote 10.1.1.2:3225
    TCP Parameters
      Advertized window: Current: 310 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
      Congestion window: Current: 254 KB, Slow start threshold: 3450 KB
```

```
     Measured RTT : 500000 us Min RTT: 920 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
   Data connection: Local 10.1.1.1:65497, Remote 10.1.1.2:3225
   TCP Parameters
     Advertized window: Current: 57 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
     Congestion window: Current: 14 KB, Slow start threshold: 3546 KB
     Measured RTT : 500000 us Min RTT: 36 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
   Data connection: Local 10.1.1.1:65495, Remote 10.1.1.2:3225
   TCP Parameters
     Advertized window: Current: 58 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
     Congestion window: Current: 14 KB, Slow start threshold: 3546 KB
     Measured RTT : 500000 us Min RTT: 39 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
   Data connection: Local 10.1.1.1:65493, Remote 10.1.1.2:3225
   TCP Parameters
     Advertized window: Current: 57 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
     Congestion window: Current: 14 KB, Slow start threshold: 3546 KB
     Measured RTT : 500000 us Min RTT: 36 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
   Control connection: Local 10.1.1.1:65491, Remote 10.1.1.2:3225
   TCP Parameters
     Advertized window: Current: 50 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8178 KB, Maximum: 8178 KB, Scale: 7
     Congestion window: Current: 21 KB, Slow start threshold: 3543 KB
     Measured RTT : 12 us Min RTT: 12 us Max RTT: 14 us
     Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
   5 minutes input rate 312 bits/sec, 39 bytes/sec, 0 frames/sec
   5 minutes output rate 352 bits/sec, 44 bytes/sec, 0 frames/sec
     964 frames input, 103200 bytes
        964 Class F frames input, 103200 bytes
        0 Class 2/3 frames input, 0 bytes
        0 Reass frames
```

```
         0 Error frames timestamp error 0
      967 frames output, 112740 bytes
         967 Class F frames output, 112740 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames

SAN192C-6# show interface fcip 110
fcip110 is trunking
    Hardware is IPStorage
    Port WWN is 21:5e:00:2a:6a:a4:1a:80
    Peer port WWN is 21:9e:00:3a:9c:31:62:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Operating Speed is 10000 Mbps
    Belongs to port-channel100
    Trunk vsans (admin allowed and active) (1,50)
    Trunk vsans (up)                       (1,50)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    Interface last changed at Fri Jan 17 20:53:06 2020

    Using Profile id 110  (interface IPStorage6/2)
    Peer Information
      Peer Internet address is 10.1.2.2 and port is 3225
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    FICON Load Balancer configured off for all vsans
    FICON Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 5
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      5 Active TCP connections
        27 Attempts for active connections, 5 close of connections
        Path MTU 2300 bytes
        Current retransmission timeout is 200 ms
        Current Send Buffer Size: 125310 KB, Requested Send Buffer Size: 125000
KB
        CWM Burst Size: 50 KB
 CONN<0>
    Data connection: Local 10.1.2.1:65496, Remote 10.1.2.2:3225
    TCP Parameters
      Advertized window: Current: 310 KB, Maximum: 24580 KB, Scale: 7
      Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
      Congestion window: Current: 254 KB, Slow start threshold: 2657 KB
      Measured RTT : 500000 us Min RTT: 921 us Max RTT: 0 us
      Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
```

```
CONN<1>
   Data connection: Local 10.1.2.1:65495, Remote 10.1.2.2:3225
   TCP Parameters
     Advertized window: Current: 57 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
     Congestion window: Current: 14 KB, Slow start threshold: 2729 KB
     Measured RTT : 500000 us Min RTT: 37 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
   Data connection: Local 10.1.2.1:65493, Remote 10.1.2.2:3225
   TCP Parameters
     Advertized window: Current: 58 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
     Congestion window: Current: 14 KB, Slow start threshold: 2728 KB
     Measured RTT : 500000 us Min RTT: 39 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
   Data connection: Local 10.1.2.1:65491, Remote 10.1.2.2:3225
   TCP Parameters
     Advertized window: Current: 58 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
     Congestion window: Current: 14 KB, Slow start threshold: 2728 KB
     Measured RTT : 500000 us Min RTT: 40 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
   Control connection: Local 10.1.2.1:65489, Remote 10.1.2.2:3225
   TCP Parameters
     Advertized window: Current: 56 KB, Maximum: 24580 KB, Scale: 7
     Peer receive window: Current: 8189 KB, Maximum: 8189 KB, Scale: 7
     Congestion window: Current: 15 KB, Slow start threshold: 2728 KB
     Measured RTT : 500000 us Min RTT: 12 us Max RTT: 0 us
     Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
   5 minutes input rate 56 bits/sec, 7 bytes/sec, 0 frames/sec
   5 minutes output rate 72 bits/sec, 9 bytes/sec, 0 frames/sec
     491 frames input, 49116 bytes
        491 Class F frames input, 49116 bytes
        0 Class 2/3 frames input, 0 bytes
        0 Reass frames
        0 Error frames timestamp error 0
     505 frames output, 58160 bytes
        505 Class F frames output, 58160 bytes
        0 Class 2/3 frames output, 0 bytes
        0 Error frames
```

```
Member[1] : fcip100    [up] *
Member[2] : fcip110    [up]

SAN192C-6#
SAN192C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
  port-channel100  0x10(16)  port-channel100  10.122.107.95(SAN384C-6)
SAN192C-6#
```

7.2.18 Configuring FCIP with VLAN sub-interfaces

In this section, we show how to configure multiple FCIP interfaces per IPS port by using VLAN sub-interfaces.

To accomplish this task, we must first configure multiple VLANs on the physical IPS interface, which in turn creates the VLAN sub-interface. Then, we create the FCIP interfaces, which are tied to these VLAN sub-interfaces. When using VLAN sub-interfaces, you must match the VLANs that are on the FCIP IPS interfaces with VLANs that are configured on the Ethernet switch that is physically attached to the IBM c-type switch.

On the IBM Storage Networking SAN384C-6 device tab, we double-click the IPStorage 7/3 port, which shows the default MTU of 1500, as shown in Figure 7-127.

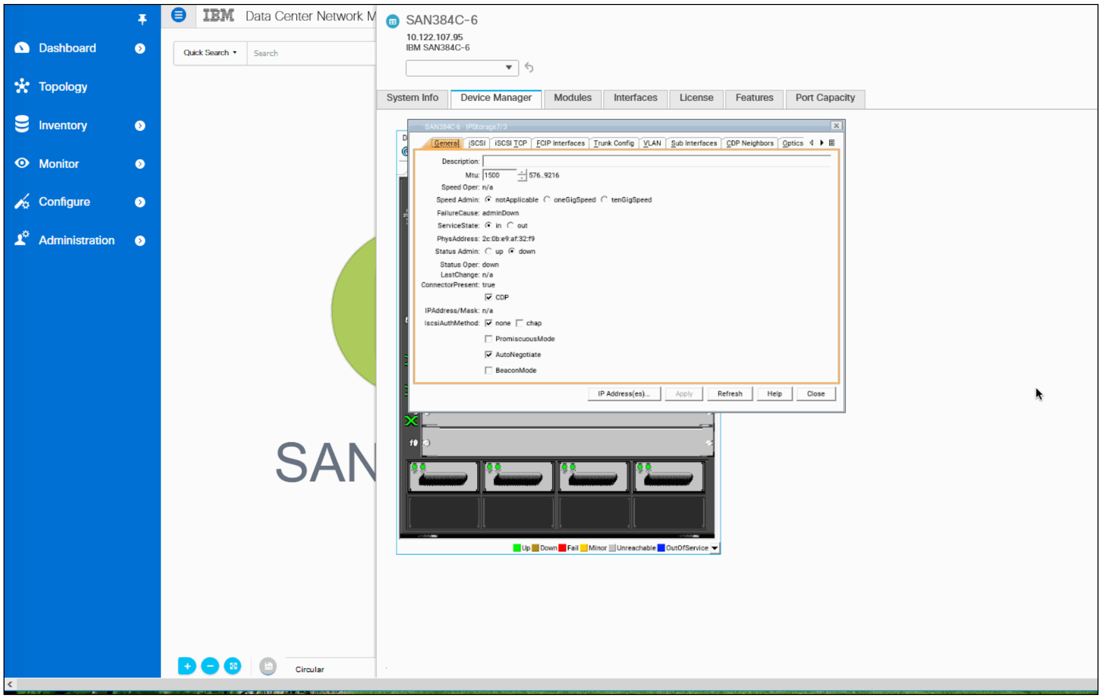
SAN384C-6
10.122.107.95
IBM SAN384C-6

SAN384C-6 · IPStorage7/3

Description:
Mtu: 1500 576..9216
Speed Oper: n/a
Speed Admin: ● notApplicable ○ oneGigSpeed ○ tenGigSpeed
FailureCause: adminDown
ServiceState: ● in ○ out
PhysAddress: 2c:0b:e9:af:32:f9
Status Admin: ○ up ● down
Status Oper: down
LastChange: n/a
ConnectorPresent: true
☑ CDP
IPAddress/Mask: n/a
IscsiAuthMethod: ☑ none □ chap
□ PromiscuousMode
☑ AutoNegotiate
□ BeaconMode

Figure 7-127 Default MTU of 1500

Figure 7-128 shows changing the MTU to 2500, which prevents the switch from segmenting frames.

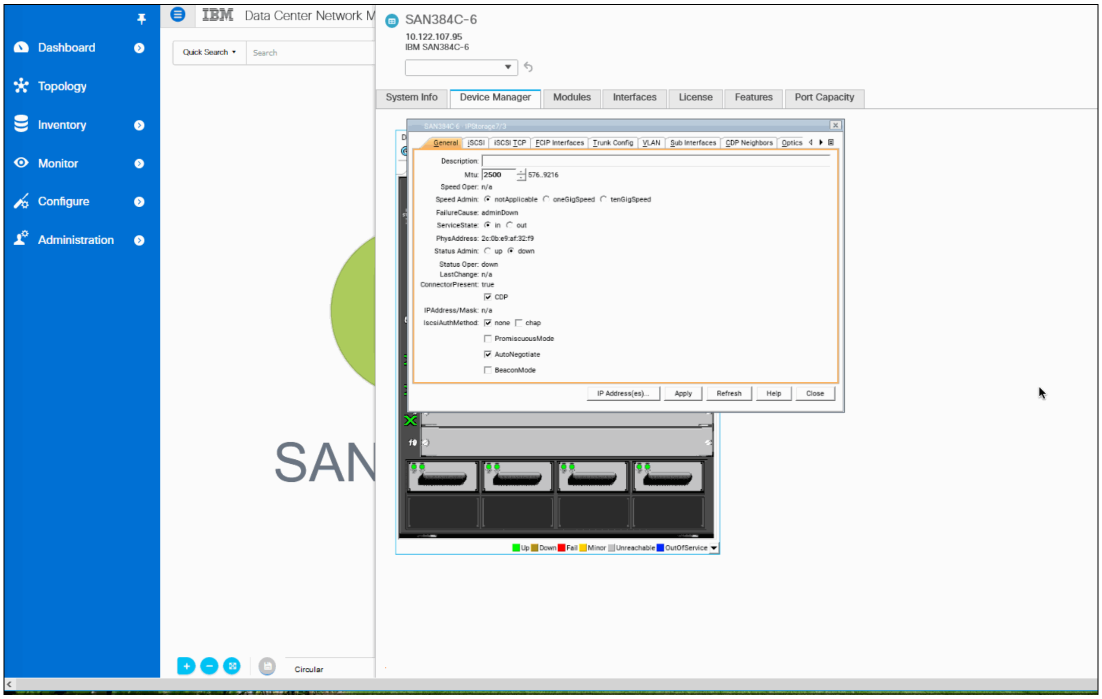

Description:
Mtu: 2500 576..9216
Speed Oper: n/a
Speed Admin: ● notApplicable ○ oneGigSpeed ○ tenGigSpeed
FailureCause: adminDown
ServiceState: ● in ○ out
PhysAddress: 2c:0b:e9:af:32:f9
Status Admin: ○ up ● down
Status Oper: down
LastChange: n/a
ConnectorPresent: true
☑ CDP
IPAddress/Mask: n/a
IscsiAuthMethod: ☑ none □ chap
□ PromiscuousMode
☑ AutoNegotiate
□ BeaconMode

Figure 7-128 Changing the MTU to 2500

In Figure 7-129, we select the **VLAN** tab to add VLANs 1000 and 1010, which match our Ethernet switch configuration to create a VLAN trunk between the IBM Storage Networking SAN384C-6 switch and the Ethernet switch. Click **Apply**.

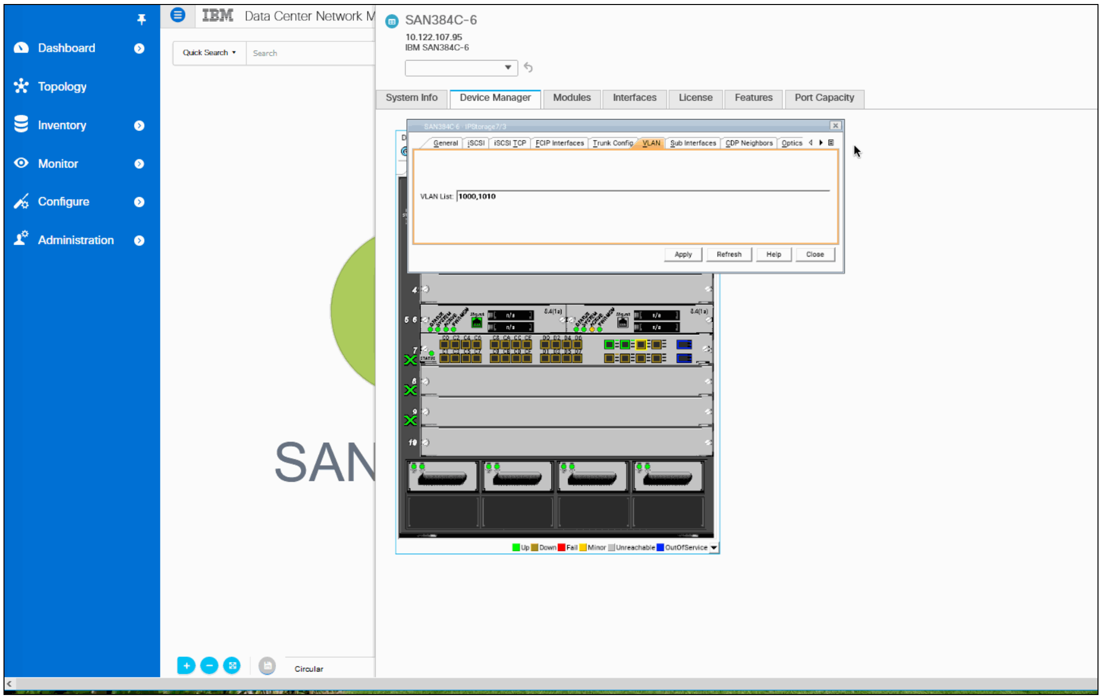

VLAN List: 1000,1010

Figure 7-129 VLAN tab

In Figure 7-130, we select the **Sub Interfaces** tab to validate the creation of the new sub-interfaces and change the MTU size to 2500. Click **Apply** and **Close**.

> **Note:** The naming convention for the "IPStorage 7/3.1000" sub interface starts with the IPS interface name followed by a period and then the VLAN number.

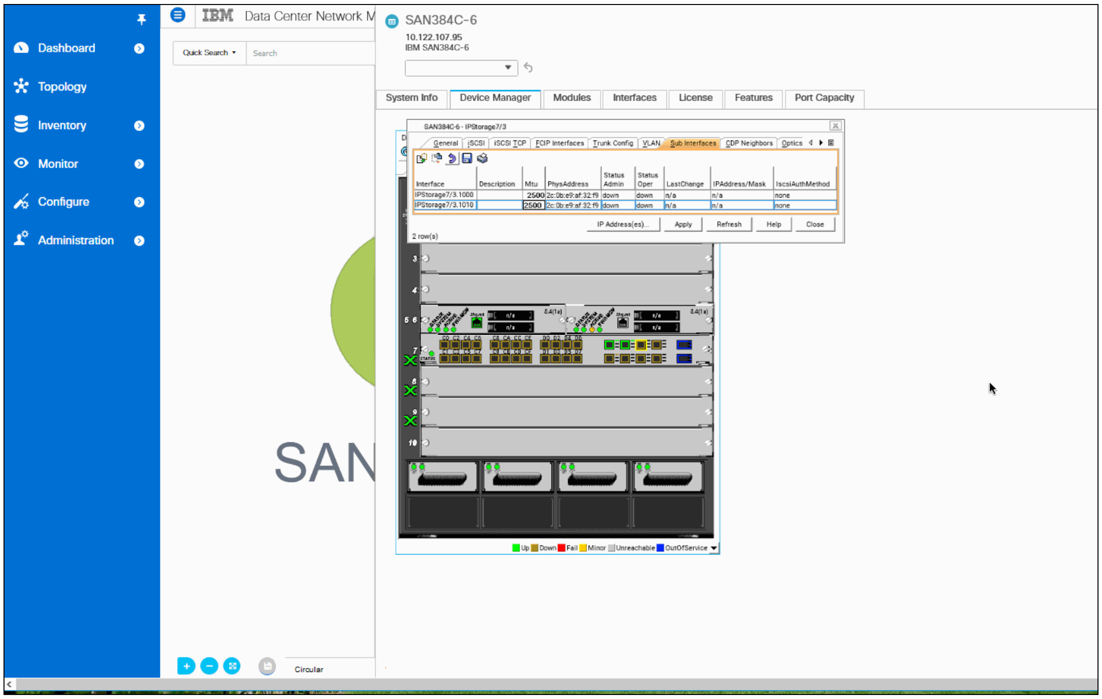

Interface	Description	Mtu	PhysAddress	Status Admin	Status Oper	LastChange	IPAddress/Mask	IscsiAuthMethod
IPStorage7/3.1000		2500	2c:0b:e9:af:32:f9	down	down	n/a	n/a	none
IPStorage7/3.1010		2500	2c:0b:e9:af:32:f9	down	down	n/a	n/a	none

Figure 7-130 Sub Interfaces tab

Figure 7-131 shows the FCIP wizard.

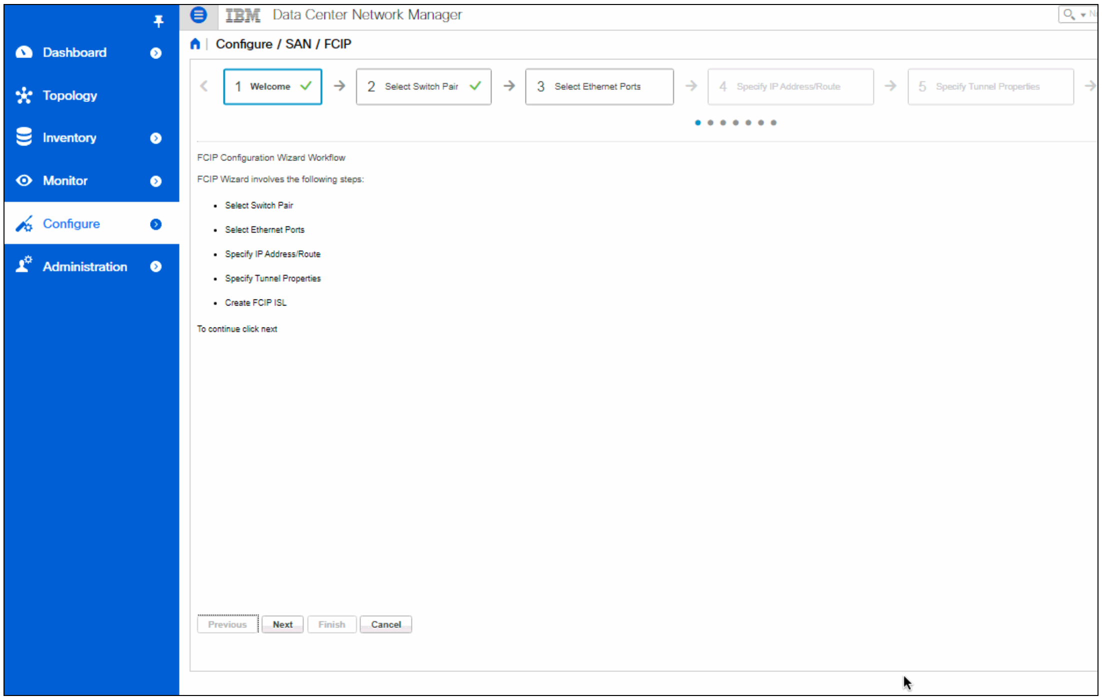

FCIP Configuration Wizard Workflow

FCIP Wizard involves the following steps:

- Select Switch Pair
- Select Ethernet Ports
- Specify IP Address/Route
- Specify Tunnel Properties
- Create FCIP ISL

To continue click next

Figure 7-131 DCNM FCIP wizard

We select the two switch end points that we will use to create our FCIP tunnel, as shown in Figure 7-132.

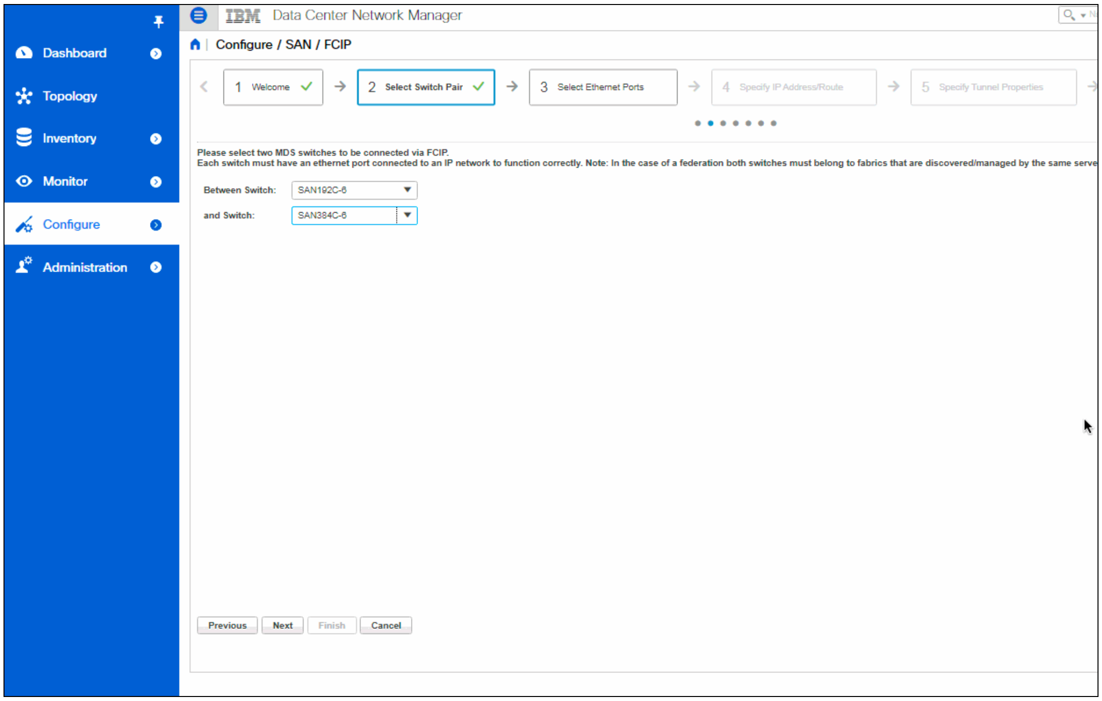

Please select two MDS switches to be connected via FCIP.
Each switch must have an ethernet port connected to an IP network to function correctly. Note: In the case of a federation both switches must belong to fabrics that are discovered/managed by the same serve

Between Switch: SAN192C-6

and Switch: SAN384C-6

Figure 7-132 Switch pair end points

We select the VLAN sub-interfaces that were just created, which will be used for the new FCIP tunnels, as shown in Figure 7-133. These sub-interfaces provide physical connectivity between the IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6 end-point switches. We also select **(Jumbo Frames)**. Click **Next**.

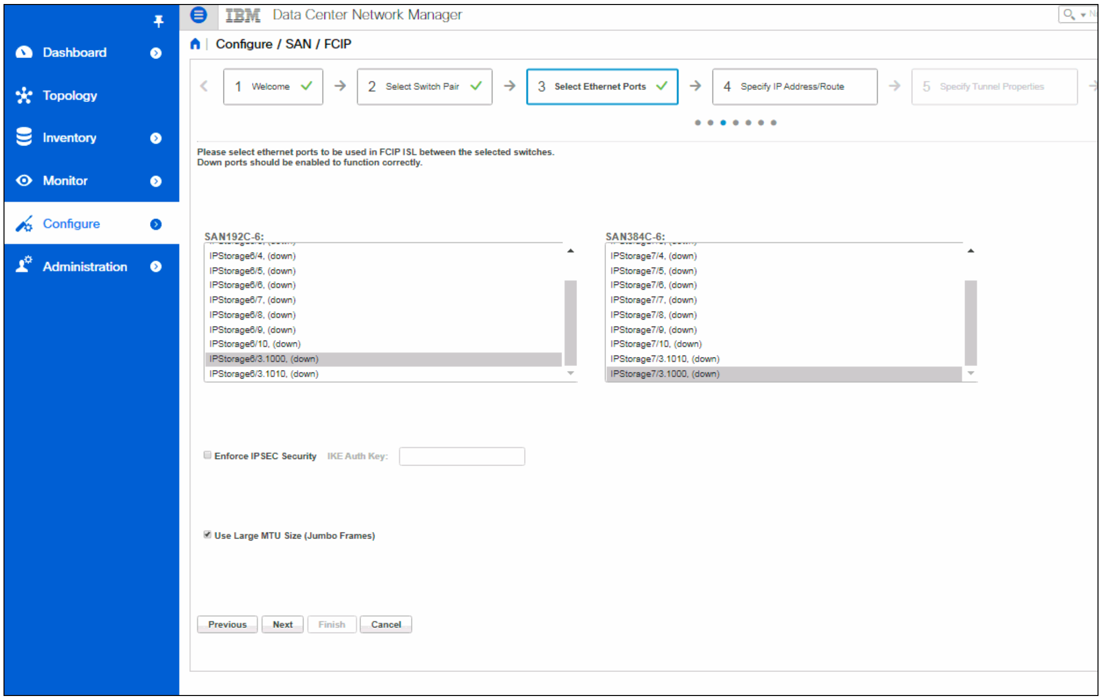

Please select ethernet ports to be used in FCIP ISL between the selected switches.
Down ports should be enabled to function correctly.

SAN192C-6:

IPStorage6/4, (down)
IPStorage6/5, (down)
IPStorage6/6, (down)
IPStorage6/7, (down)
IPStorage6/8, (down)
IPStorage6/9, (down)
IPStorage6/10, (down)
IPStorage6/3.1000, (down)
IPStorage6/3.1010, (down)

SAN384C-6:

IPStorage7/4, (down)
IPStorage7/5, (down)
IPStorage7/6, (down)
IPStorage7/7, (down)
IPStorage7/8, (down)
IPStorage7/9, (down)
IPStorage7/10, (down)
IPStorage7/3.1010, (down)
IPStorage7/3.1000, (down)

☐ Enforce IPSEC Security IKE Auth Key:

☑ Use Large MTU Size (Jumbo Frames)

Figure 7-133 Selecting the VLAN sub-interfaces

We provide the IP addresses for each of the VLAN sub-interface ports on each switch end point, as shown in Figure 7-134. Routes are not needed in this example because the IP addresses are in the same subnet. Click **Next**.

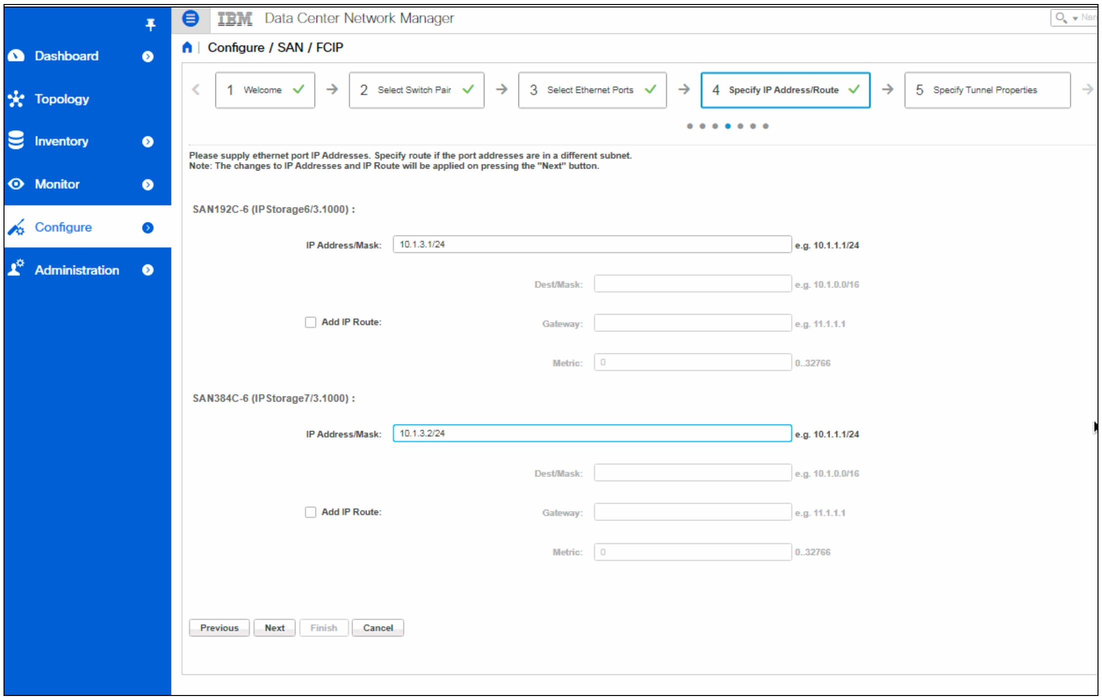

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN192C-6 (IPStorage6/3.1000) :

IP Address/Mask: 10.1.3.1/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

SAN384C-6 (IPStorage7/3.1000) :

IP Address/Mask: 10.1.3.2/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

Figure 7-134 Specifying the IP addresses

Click **Yes** to continue, as shown in Figure 7-135.

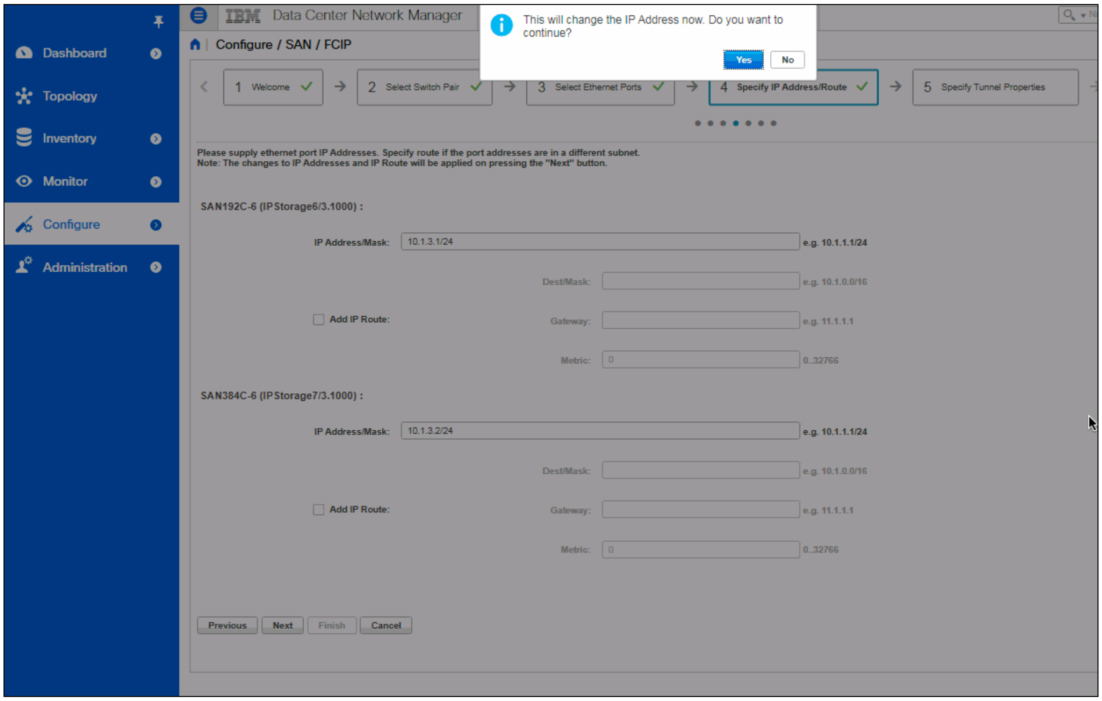

This will change the IP Address now. Do you want to continue?

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN192C-6 (IPStorage6/3.1000) :

IP Address/Mask: 10.1.3.1/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

SAN384C-6 (IPStorage7/3.1000) :

IP Address/Mask: 10.1.3.2/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

Figure 7-135 IP addresses

We provide the maximum and minimum bandwidth for the FCIP tunnel and the RTT for the FCIP tunnel, as shown in Figure 7-136. In this example, we configure a 10 GbE IPS interface and limit the tunnel bandwidth to a maximum of 5 GbE and minimum of 4.5 GbE to dedicate 50% of the 10 GbE bandwidth for this sub-interface. Click **Next** to continue.

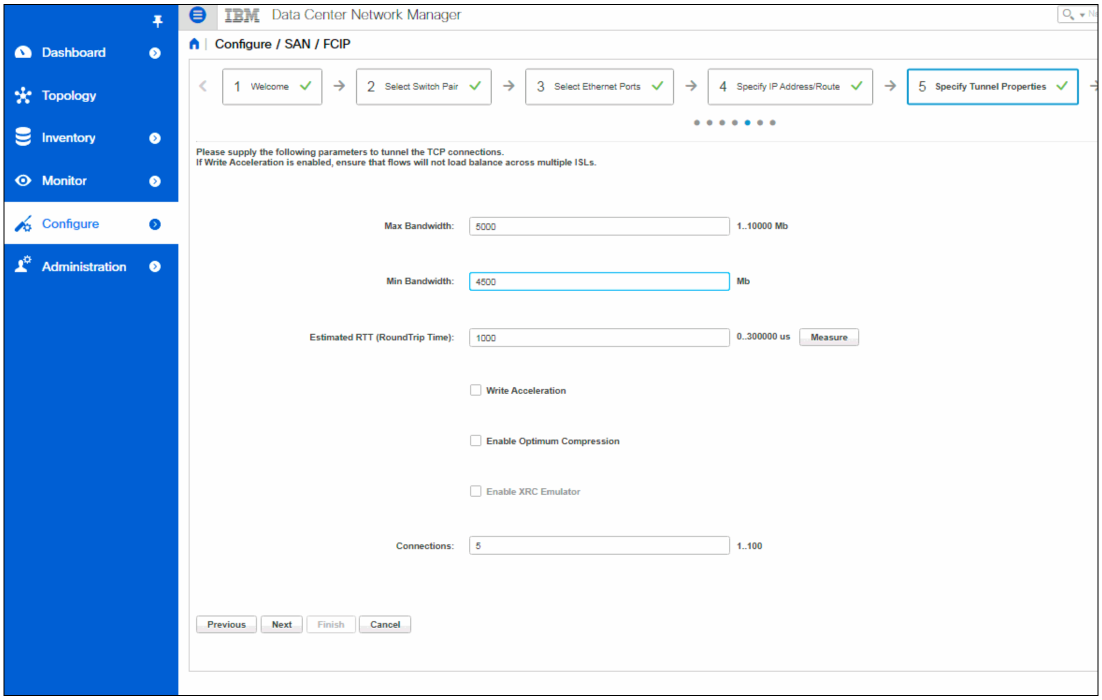

Please supply the following parameters to tunnel the TCP connections.
If Write Acceleration is enabled, ensure that flows will not load balance across multiple ISLs.

Max Bandwidth: 5000 1..10000 Mb

Min Bandwidth: 4500 Mb

Estimated RTT (RoundTrip Time): 1000 0..300000 us Measure

□ Write Acceleration

□ Enable Optimum Compression

□ Enable XRC Emulator

Connections: 5 1..100

Figure 7-136 Specifying the tunnel properties

We specify the final parameters for the FCIP link configuration, as shown in Figure 7-137:

- **Profile ID:** 120, which provides detailed information about the local IP address and TCP parameters.
- **Tunnel ID:** 120, which is used to create the name of the new FCIP interface fcip120.
- **FICON Port Address:** 0xF3, which is only applicable when FICON is enabled on the switch. The FICON Port Address must be configured when there is a FICON VSAN communicating on this FCIP tunnel. As a best practice, use the same value on both switches for the FCIP tunnel when possible. The value of this attribute must be taken from the pool of logical FICON Port Addresses.
- **Trunk Mode:** As a best practice, use **Trunk**.
- **VSAN List:** Should have 1 and the value of any VSANs that require access to the FCIP tunnel. In this example, it is **VSAN 50**.

To proceed, click **Next**.

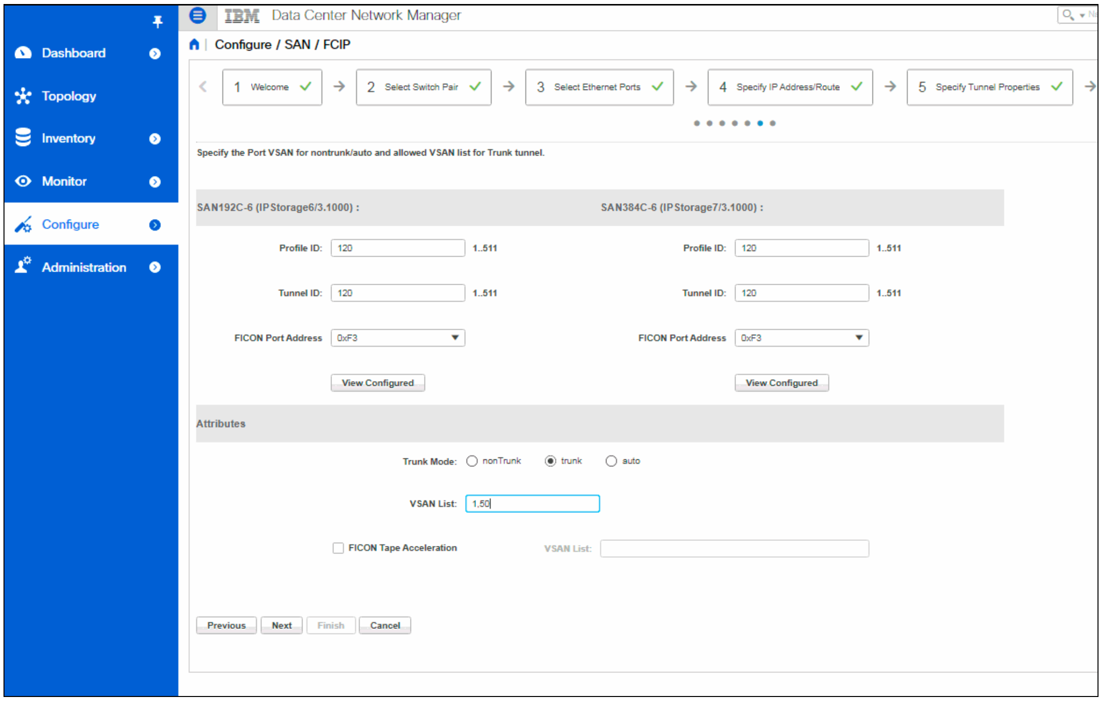

Specify the Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.

SAN192C-6 (IPStorage6/3.1000) :

Profile ID: 120 1..511

Tunnel ID: 120 1..511

FICON Port Address 0xF3

SAN384C-6 (IPStorage7/3.1000) :

Profile ID: 120 1..511

Tunnel ID: 120 1..511

FICON Port Address 0xF3

Attributes

Trunk Mode: ○ nonTrunk ● trunk ○ auto

VSAN List: 1,50

□ FICON Tape Acceleration VSAN List:

Figure 7-137 Creating an FCIP ISL

Select **View configured** → **Profiles**, and then validate which profiles and TCP ports are in use, as shown in Figure 7-138 on page 371. These values should be unique per FCIP tunnel creation.

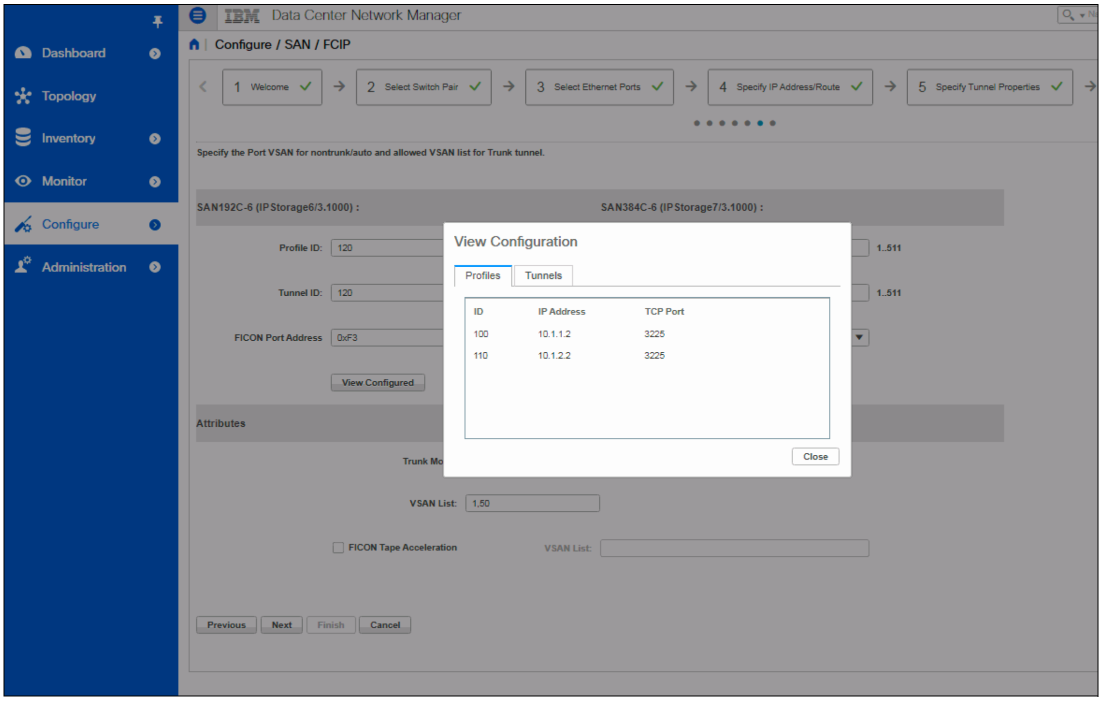

Specify the Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.

SAN192C-6 (IPStorage6/3.1000) :

SAN384C-6 (IPStorage7/3.1000) :

Profile ID: 120 1..511

Tunnel ID: 120 1..511

FICON Port Address 0xF3

VSAN List: 1,50

□ FICON Tape Acceleration VSAN List:

ID	IP Address	TCP Port
100	10.1.1.2	3225
110	10.1.2.2	3225

Figure 7-138 Viewing the configuration profiles

Select **View configured** → **Tunnels**, and then validate which profiles, tunnels, and IP addresses were created and are in use, as shown in Figure 7-139. These values should be unique per FCIP tunnel creation.

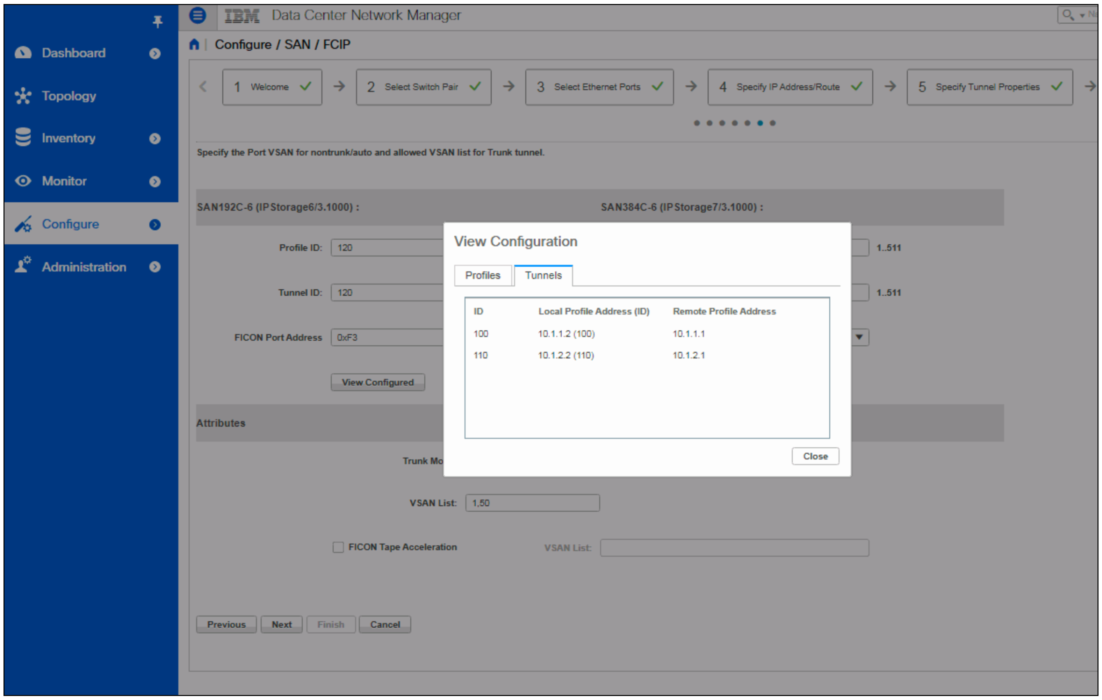

Specify the Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.

SAN192C-6 (IPStorage6/3.1000) :

SAN384C-6 (IPStorage7/3.1000) :

Profile ID: 120 1..511

Tunnel ID: 120 1..511

FICON Port Address 0xF3

VSAN List: 1,50

□ FICON Tape Acceleration VSAN List:

ID	Local Profile Address (ID)	Remote Profile Address
100	10.1.1.2 (100)	10.1.1.1
110	10.1.2.2 (110)	10.1.2.1

Figure 7-139 Viewing the configuration tunnels

Figure 7-140 shows a summary of the configuration that will be applied to create the single FCIP link. To proceed, click **Finish**.

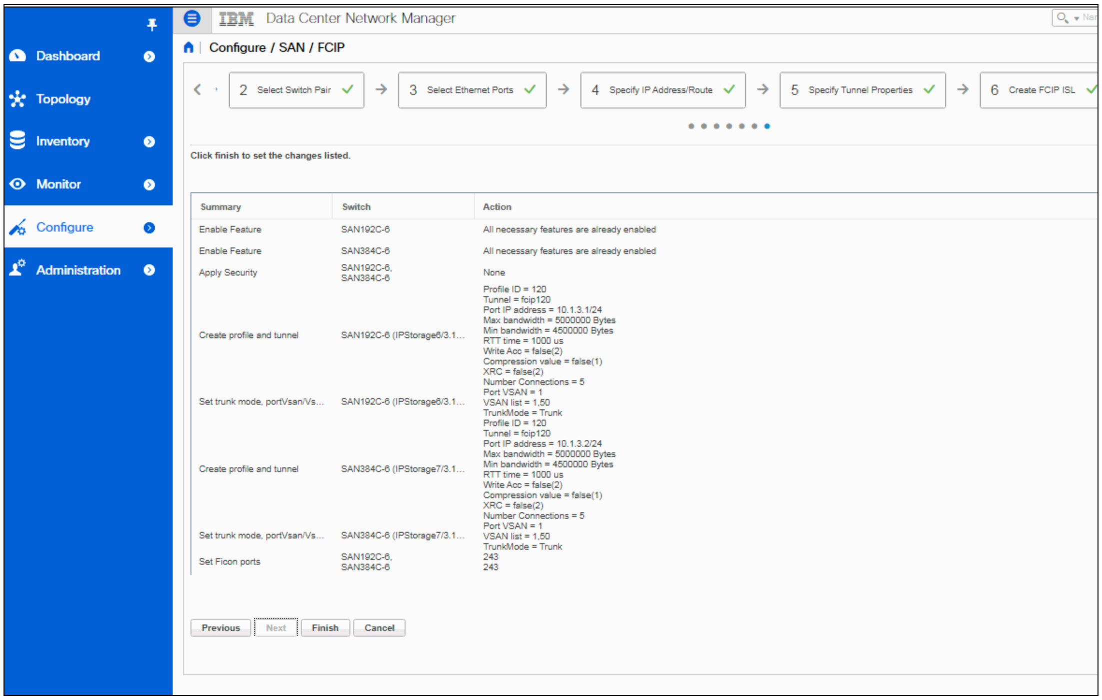

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN192C-6	All necessary features are already enabled
Enable Feature	SAN384C-6	All necessary features are already enabled
Apply Security	SAN192C-6, SAN384C-6	None
Create profile and tunnel	SAN192C-6 (IPStorage6/3.1...	Profile ID = 120 Tunnel = fcip120 Port IP address = 10.1.3.1/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN192C-6 (IPStorage6/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/3.1...	Profile ID = 120 Tunnel = fcip120 Port IP address = 10.1.3.2/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Set Ficon ports	SAN192C-6, SAN384C-6	243 243

Figure 7-140 Summary review

Figure 7-141 shows that the configuration successfully completed. To proceed, click **OK** and then **Close**.

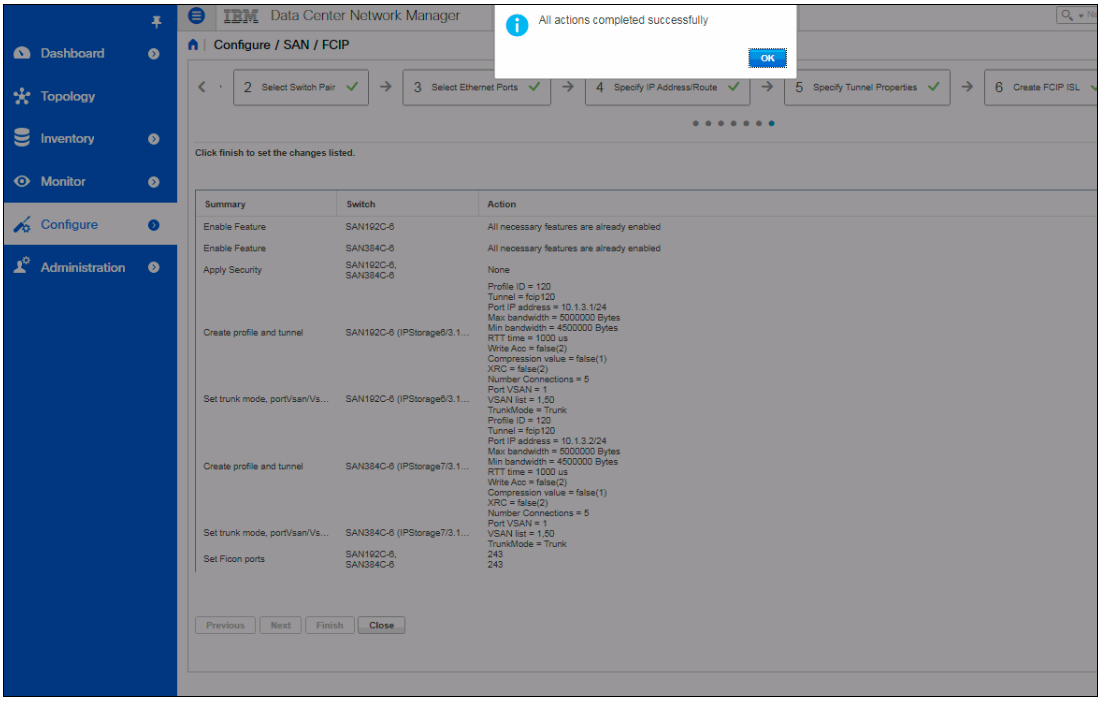

All actions completed successfully

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN192C-6	All necessary features are already enabled
Enable Feature	SAN384C-6	All necessary features are already enabled
Apply Security	SAN192C-6, SAN384C-6	None
Create profile and tunnel	SAN192C-6 (IPStorage6/3.1...	Profile ID = 120 Tunnel = fcip120 Port IP address = 10.1.3.1/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN192C-6 (IPStorage6/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/3.1...	Profile ID = 120 Tunnel = fcip120 Port IP address = 10.1.3.2/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = false(1) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Set Ficon ports	SAN192C-6, SAN384C-6	243 243

Figure 7-141 Completed successfully

Now, we configure our second FICON FCIP Link on the same physical 10 GbE Ethernet IP interface. This link will be encrypted and compressed to protect data in transit between a primary production site and an alternative location by using public transport networks.

We change the MTU to 2500, as shown in Figure 7-142.

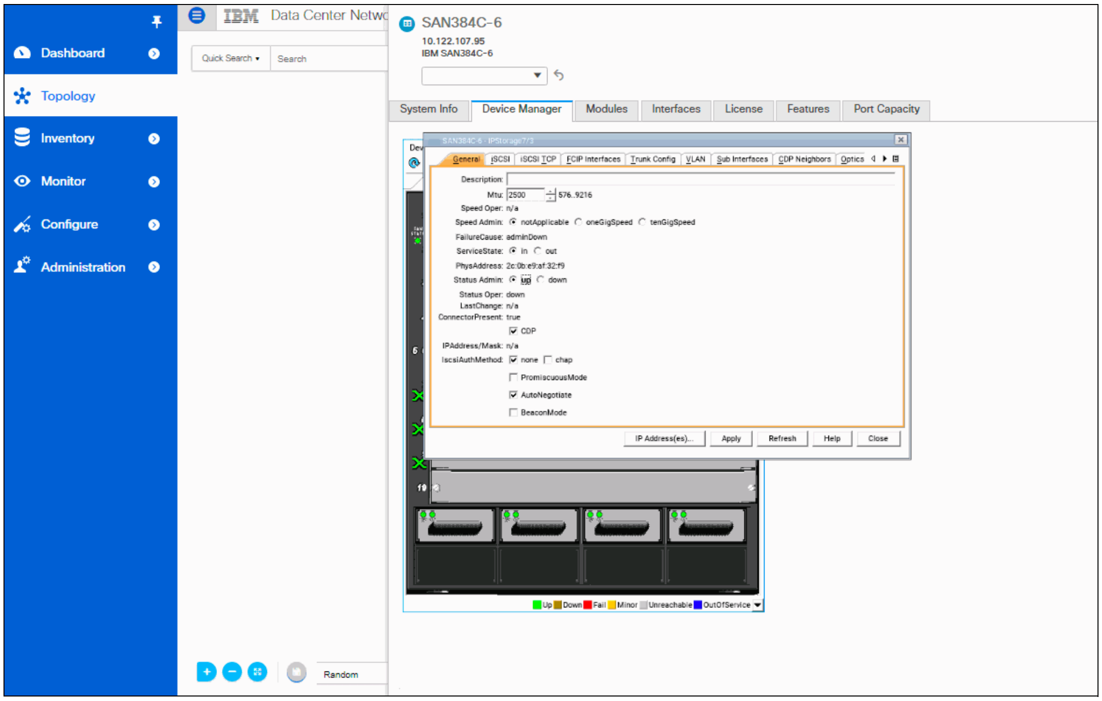

Description:
Mtu: 2500 576..9216
Speed Oper: n/a
Speed Admin: ● notApplicable ○ oneGigSpeed ○ tenGigSpeed
FailureCause: adminDown
ServiceState: ● in ○ out
PhysAddress: 2c:0b:e9:af:32:f9
Status Admin: ● up ○ down
Status Oper: down
LastChange: n/a
ConnectorPresent: true
☑ CDP
IPAddress/Mask: n/a
IscsiAuthMethod: ☑ none □ chap
□ PromiscuousMode
☑ AutoNegotiate
□ BeaconMode

Figure 7-142 Changing the MTU to 2500

We select the **Sub Interfaces** tab and change the MTU size to 2500, as shown in Figure 7-143. Click **Apply** and **Close**.

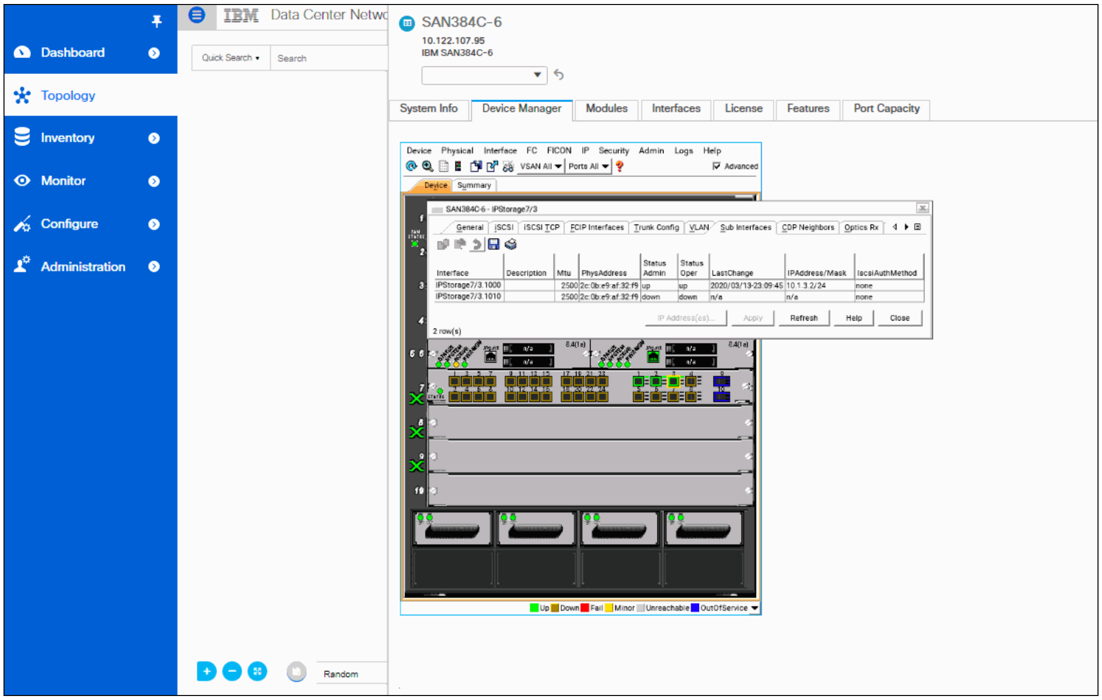

Interface	Description	Mtu	PhysAddress	Status Admin	Status Oper	LastChange	IPAddress/Mask	IscsiAuthMethod
IPStorage7/3.1000		2500	2c:0b:e9:af:32:f9	up	up	2020/03/13-23:09:45	10.1.3.2/24	none
IPStorage7/3.1010		2500	2c:0b:e9:af:32:f9	down	down	n/a	n/a	none

2 row(s)

Figure 7-143 Sub Interfaces tab

We select the **Enforce IPSEC Security** checkbox to enable encryption and input the IKE authentication key on the second VLAN sub-interface, as shown in Figure 7-144 on page 375. This setting provides physical encrypted connectivity between our two end point switches IBM Storage Networking SAN192C-6 and IBM Storage Networking SAN384C-6. We also select **(Jumbo Frames)**. Click **Next**.

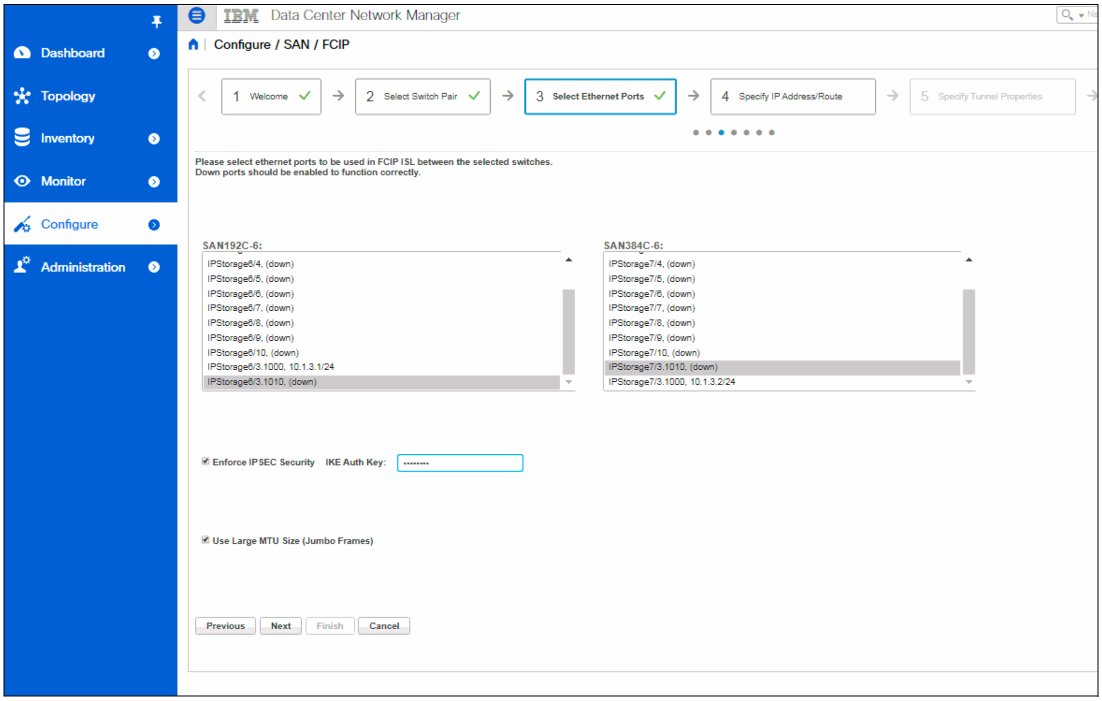

Please select ethernet ports to be used in FCIP ISL between the selected switches.
Down ports should be enabled to function correctly.

SAN192C-6:

- IPStorage6/4, (down)
- IPStorage6/5, (down)
- IPStorage6/6, (down)
- IPStorage6/7, (down)
- IPStorage6/8, (down)
- IPStorage6/9, (down)
- IPStorage6/10, (down)
- IPStorage6/3.1000, 10.1.3.1/24
- IPStorage6/3.1010, (down)

SAN384C-6:

- IPStorage7/4, (down)
- IPStorage7/5, (down)
- IPStorage7/6, (down)
- IPStorage7/7, (down)
- IPStorage7/8, (down)
- IPStorage7/9, (down)
- IPStorage7/10, (down)
- IPStorage7/3.1010, (down)
- IPStorage7/3.1000, 10.1.3.2/24

☑ Enforce IPSEC Security IKE Auth Key: ••••••••

☑ Use Large MTU Size (Jumbo Frames)

Figure 7-144 Selecting the VLAN sub-interfaces

We provide the IP addresses for each of the VLAN sub-interface ports on each switch end point, as shown in Figure 7-145. Click **Next**.

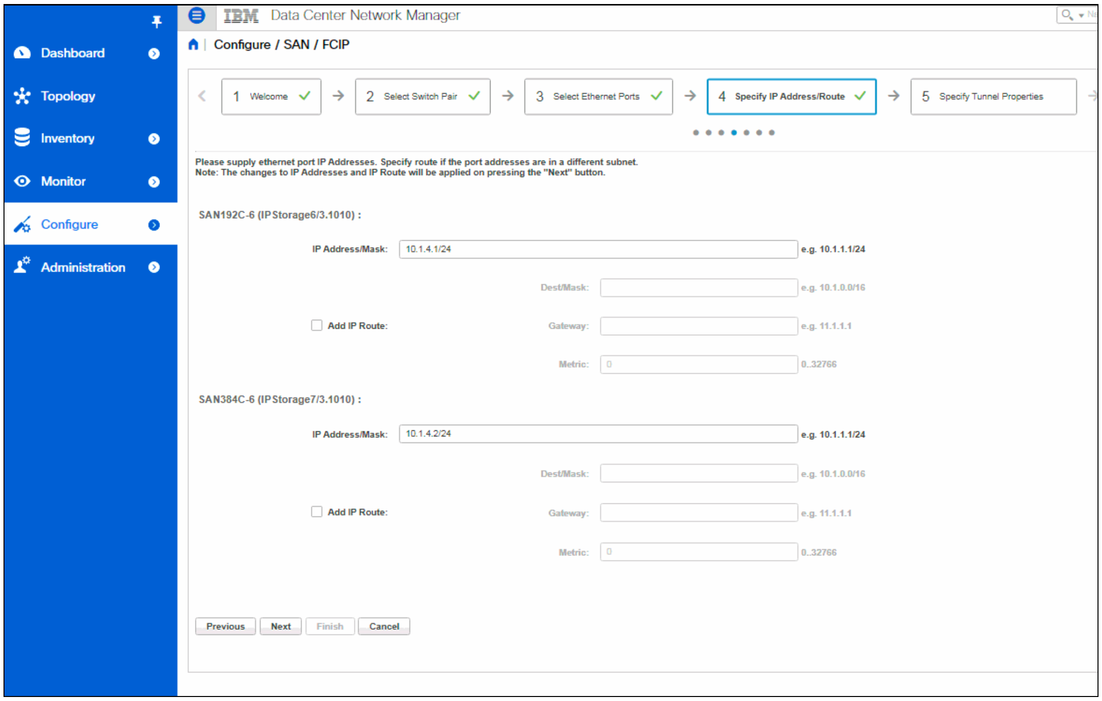

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN192C-6 (IPStorage6/3.1010) :

IP Address/Mask: 10.1.4.1/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

SAN384C-6 (IPStorage7/3.1010) :

IP Address/Mask: 10.1.4.2/24 e.g. 10.1.1.1/24

Dest/Mask: e.g. 10.1.0.0/16

□ Add IP Route: Gateway: e.g. 11.1.1.1

Metric: 0 0..32766

Figure 7-145 Selecting the IP addresses

We select the **Enable Optimum Compression** checkbox, as shown in Figure 7-146. In addition, we provide the maximum and minimum bandwidth and the RTT for the FCIP tunnel. Click **Next** to continue.

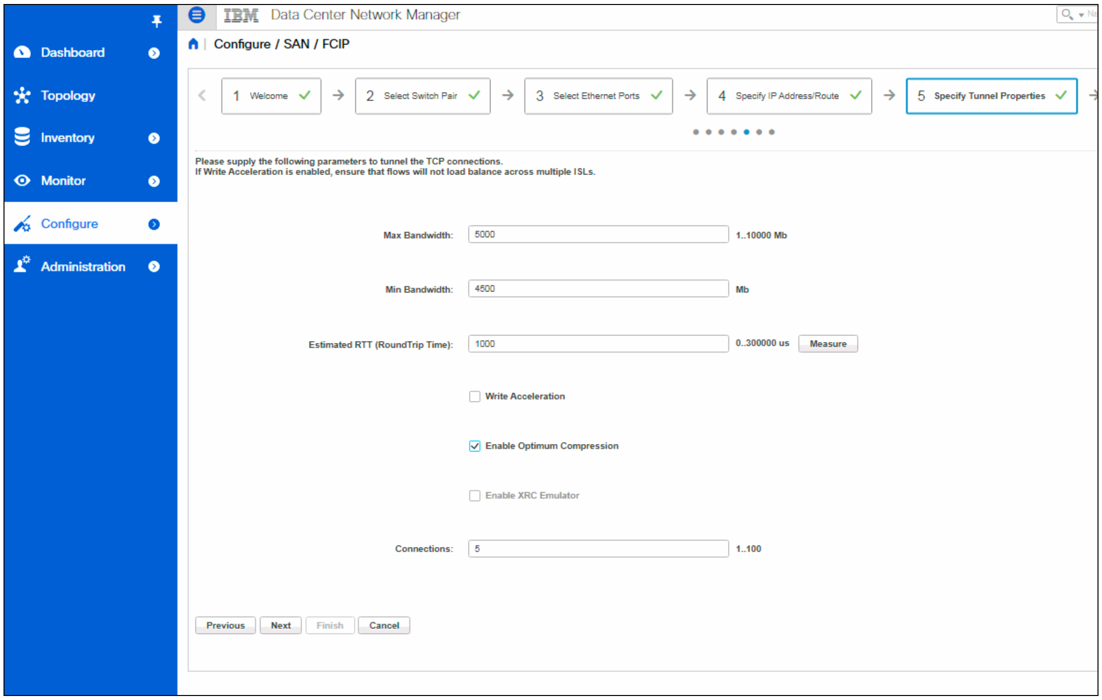

Please supply the following parameters to tunnel the TCP connections.
If Write Acceleration is enabled, ensure that flows will not load balance across multiple ISLs.

Max Bandwidth: 5000 1..10000 Mb

Min Bandwidth: 4500 Mb

Estimated RTT (RoundTrip Time): 1000 0..300000 us

□ Write Acceleration

☑ Enable Optimum Compression

□ Enable XRC Emulator

Connections: 5 1..100

Figure 7-146 Specifying the tunnel properties

We specify the final parameters for the second FCIP link configuration, as shown in Figure 7-147 on page 377:

- **Profile ID:** 130, which provides detailed information about the local IP address and TCP parameters.
- **Tunnel ID:** 130, which is used to create the name of the new FCIP interface fcip130.
- **FICON Port Address:** 0xF4, which is only applicable when FICON is enabled on the switch. The FICON Port Address must be configured when there is a FICON VSAN communicating on this FCIP tunnel. As a best practice, use the same value on both switches for the FCIP tunnel when possible. The value of this attribute must be taken from the pool of logical FICON Port Addresses.
- **Trunk Mode:** As a best practice, use **Trunk**.
- **VSAN List:** Should have 1 and the value of any VSANs that require access to the FCIP tunnel. In this example, it is **VSAN 50**.

To proceed, click **Next**.

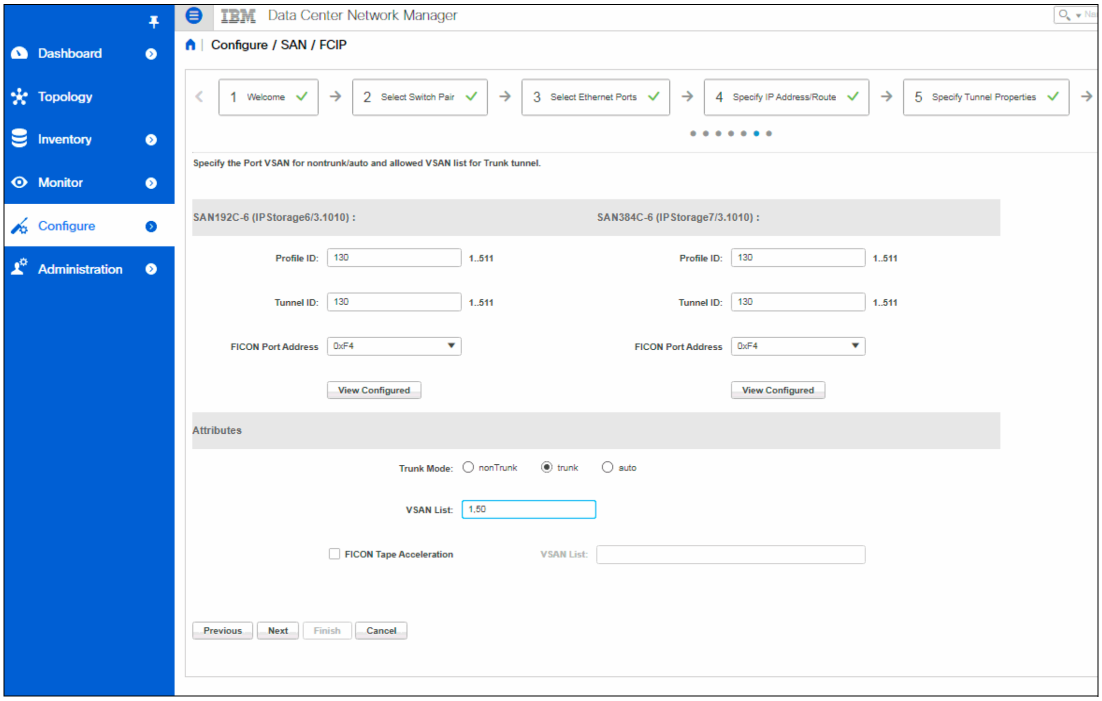

Specify the Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.

SAN192C-6 (IPStorage6/3.1010) :			SAN384C-6 (IPStorage7/3.1010) :		
Profile ID:	130	1..511	Profile ID:	130	1..511
Tunnel ID:	130	1..511	Tunnel ID:	130	1..511
FICON Port Address	0xF4		FICON Port Address	0xF4	
	View Configured			View Configured	

Attributes

Trunk Mode: ○ nonTrunk ● trunk ○ auto

VSAN List: 1,50

□ FICON Tape Acceleration VSAN List:

Figure 7-147 Creating the FCIP ISLs

Figure 7-148 shows a summary of the configuration that will be applied to create the second FCIP link. To proceed, click **Finish**.

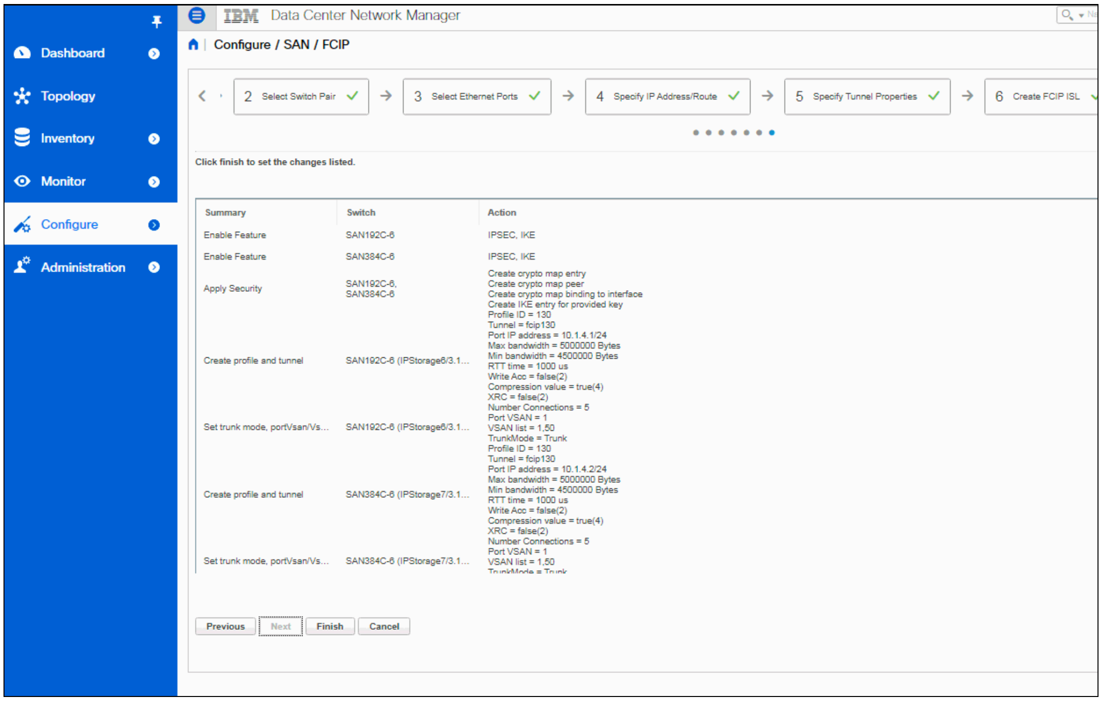

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN192C-6	IPSEC, IKE
Enable Feature	SAN384C-6	IPSEC, IKE
Apply Security	SAN192C-6, SAN384C-6	Create crypto map entry Create crypto map peer Create crypto map binding to interface Create IKE entry for provided key
Create profile and tunnel	SAN192C-6 (IPStorage6/3.1...	Profile ID = 130 Tunnel = fcip130 Port IP address = 10.1.4.1/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN192C-6 (IPStorage6/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/3.1...	Profile ID = 130 Tunnel = fcip130 Port IP address = 10.1.4.2/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk

Figure 7-148 Summary review

Figure 7-149 shows that the encryption and compression successfully completed. To proceed, click **OK** and then **Close**.

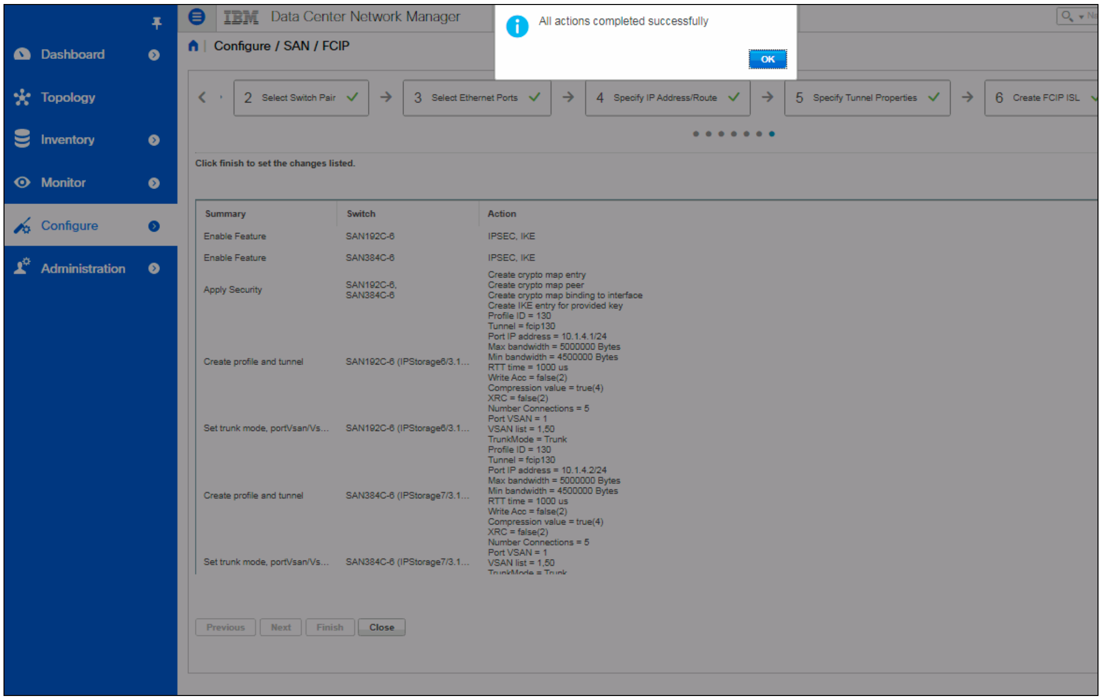

All actions completed successfully

OK

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN192C-6	IPSEC, IKE
Enable Feature	SAN384C-6	IPSEC, IKE
Apply Security	SAN192C-6, SAN384C-6	Create crypto map entry Create crypto map peer Create crypto map binding to interface Create IKE entry for provided key
Create profile and tunnel	SAN192C-6 (IPStorage6/3.1...	Profile ID = 130 Tunnel = fcip130 Port IP address = 10.1.4.1/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN192C-6 (IPStorage6/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/3.1...	Profile ID = 130 Tunnel = fcip130 Port IP address = 10.1.4.2/24 Max bandwidth = 5000000 Bytes Min bandwidth = 4500000 Bytes RTT time = 1000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/3.1...	Port VSAN = 1 VSAN list = 1,50 TrunkMode = Trunk

Figure 7-149 Completed successfully

Figure 7-150 shows the Device Manager view of the newly created FCIP 120 and FCIP 130 interfaces.

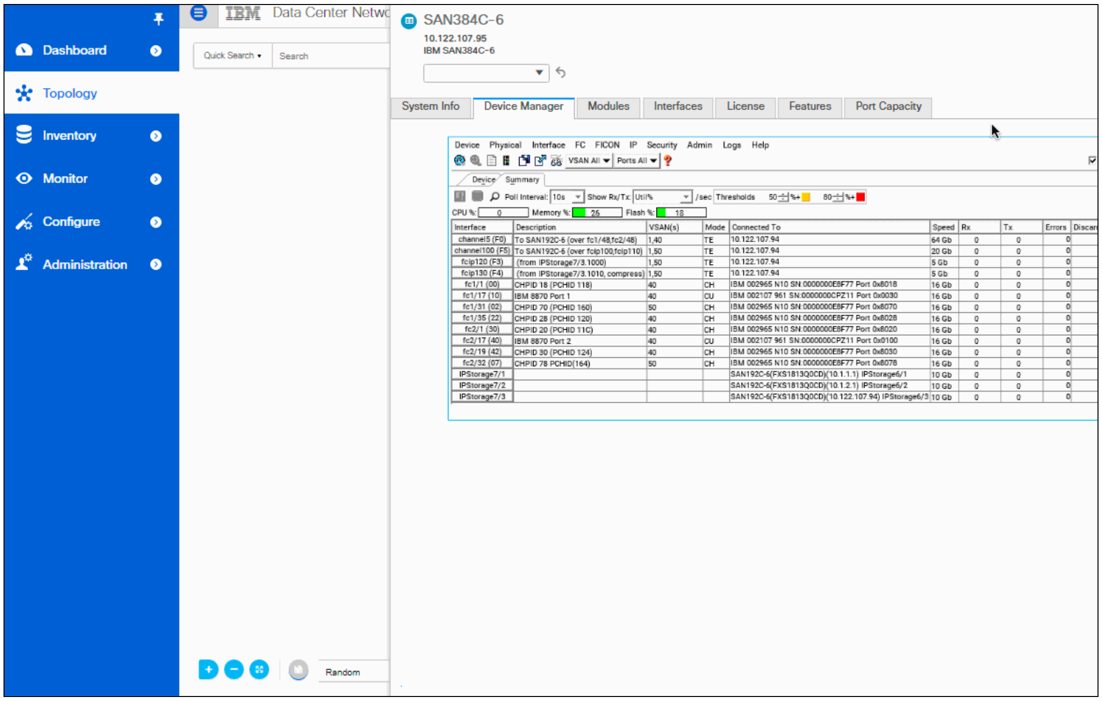

SAN384C-6

10.122.107.95
IBM SAN384C-6

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 🟨 80 %+ 🟥

CPU %: 0 Memory %: 26 Flash %: 18

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discar
channel5 (F0)	To SAN192C-6 (over fc1/48,fc2/48)	1,40	TE	10.122.107.94	64 Gb	0	0	0	
channel100 (F5)	To SAN192C-6 (over fcip100,fcip110)	1,50	TE	10.122.107.94	20 Gb	0	0	0	
fcip120 (F3)	(from IPStorage7/3.1000)	1,50	TE	10.122.107.94	5 Gb	0	0	0	
fcip130 (F4)	(from IPStorage7/3.1010, compress)	1,50	TE	10.122.107.94	5 Gb	0	0	0	
fc1/1 (00)	CHPID 18 (PCHID 118)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8018	16 Gb	0	0	0	
fc1/17 (10)	IBM 8870 Port 1	40	CU	IBM 002107 961 SN:0000000CPZ11 Port 0x0030	16 Gb	0	0	0	
fc1/31 (02)	CHPID 70 (PCHID 160)	50	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8070	16 Gb	0	0	0	
fc1/35 (22)	CHPID 28 (PCHID 120)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8028	16 Gb	0	0	0	
fc2/1 (30)	CHPID 20 (PCHID 11C)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8020	16 Gb	0	0	0	
fc2/17 (40)	IBM 8870 Port 2	40	CU	IBM 002107 961 SN:0000000CPZ11 Port 0x0100	16 Gb	0	0	0	
fc2/19 (42)	CHPID 30 (PCHID 124)	40	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8030	16 Gb	0	0	0	
fc2/32 (07)	CHPID 78 PCHID(164)	50	CH	IBM 002965 N10 SN:0000000E8F77 Port 0x8078	16 Gb	0	0	0	
IPStorage7/1				SAN192C-6(FXS1813Q0CD)(10.1.1.1) IPStorage6/1	10 Gb	0	0	0	
IPStorage7/2				SAN192C-6(FXS1813Q0CD)(10.1.2.1) IPStorage6/2	10 Gb	0	0	0	
IPStorage7/3				SAN192C-6(FXS1813Q0CD)(10.122.107.94) IPStorage6/3	10 Gb	0	0	0	

Figure 7-150 Device Manager view of the new FCIP interfaces

Configuring VLAN sub-interfaces with FCIP by using the CLI

Example 7-11 shows how to configure FCIP with VLAN sub-interfaces on the first sub-interface for the IBM Storage Networking SAN384C-6 and IBM Storage Networking SAN192C-6 switches without encryption or compression enabled by using the CLI.

Example 7-11 Configuring the VLAN sub-interfaces with FCIP by using the CLI

```
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# interface IPStorage7/3
SAN384C-6(config-if)# switchport mtu 2500
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# interface IPStorage 7/3.1000
SAN384C-6(config-if)# ip address 10.1.3.2 255.255.255.0
SAN384C-6(config-if)# switchport mtu 2500
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# fcip profile 120
SAN384C-6(config-profile)# ip address 10.1.3.2
SAN384C-6(config-profile)# tcp max-bandwidth-mbps 5000
min-available-bandwidth-mbps 4500  round-trip-time-ms 1
SAN384C-6(config-profile)#
SAN384C-6(config-profile)# interface fcip120
SAN384C-6(config-if)# use-profile 120
SAN384C-6(config-if)# peer-info ipaddr 10.1.3.1
SAN384C-6(config-if)# tcp-connections 5
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)# switchport trunk allowed vsan add 50
SAN384C-6(config-if)#
SAN384C-6(config-if)# ficon portnumber 0xf3
SAN384C-6(config-if)#
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6# show fcip summary


-------------------------------------------------------------------------------
Tun prof    IPS-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
-------------------------------------------------------------------------------
100 100  IPS7/1     10.1.1.1        TRNK  Y N N  N   N   10000M/9500M  1000
110 110  IPS7/2     10.1.2.1        TRNK  Y N N  N   N   10000M/9500M  1000
120 120  IPS7/3.1000 10.1.3.1         TRNK  Y N N  N   N   5000M/4500M  1000

SAN384C-6# show int fcip120
fcip120 is trunking
    Hardware is IPStorage
    Port WWN is 21:a2:00:3a:9c:31:62:80
    Peer port WWN is 21:62:00:2a:6a:a4:1a:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
```

```
   Port mode is TE
   Port vsan is 1
   Operating Speed is 5 Gbps
   Trunk vsans (admin allowed and active) (1,50)
   Trunk vsans (up)                       (1,50)
   Trunk vsans (isolated)                 ()
   Trunk vsans (initializing)             ()
   Interface last changed at Sun Mar 15 19:19:29 2020

   Using Profile id 120  (interface IPStorage7/3.1000)
   Peer Information
     Peer Internet address is 10.1.3.1 and port is 3225
   Write acceleration mode is configured off
   Tape acceleration mode is configured off
   Tape Accelerator flow control buffer size is automatic
   FICON XRC Accelerator is configured off
   FICON Load Balancer configured off for all vsans
   FICON Tape acceleration configured off for all vsans
   IP Compression is disabled
   Maximum number of TCP connections is 5
   QOS control code point is 0
   QOS data code point is 0
   TCP Connection Information
     5 Active TCP connections
       11 Attempts for active connections, 0 close of connections
       Path MTU 2500 bytes
       Current retransmission timeout is 200 ms
       Current Send Buffer Size: 87080 KB, Requested Send Buffer Size: 62500 KB
       CWM Burst Size: 50 KB
CONN<0>
   Data connection: Local 10.1.3.2:3225, Remote 10.1.3.1:65424
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 222 KB, Maximum: 222 KB, Scale: 6
     Congestion window: Current: 199 KB, Slow start threshold: 202 KB
     Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
     Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
   Data connection: Local 10.1.3.2:3225, Remote 10.1.3.1:65423
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 30 KB, Maximum: 30 KB, Scale: 6
     Congestion window: Current: 27 KB, Slow start threshold: 139 KB
     Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
     Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
   Data connection: Local 10.1.3.2:3225, Remote 10.1.3.1:65421
   TCP Parameters
     Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 6
```

```
      Peer receive window: Current: 31 KB, Maximum: 31 KB, Scale: 6
      Congestion window: Current: 28 KB, Slow start threshold: 139 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<3>
    Data connection: Local 10.1.3.2:3225, Remote 10.1.3.1:65419
    TCP Parameters
      Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 6
      Peer receive window: Current: 31 KB, Maximum: 31 KB, Scale: 6
      Congestion window: Current: 28 KB, Slow start threshold: 139 KB
      Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
      Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<4>
    Control connection: Local 10.1.3.2:3225, Remote 10.1.3.1:65417
    TCP Parameters
      Advertized window: Current: 4090 KB, Maximum: 24580 KB, Scale: 6
      Peer receive window: Current: 27 KB, Maximum: 27 KB, Scale: 6
      Congestion window: Current: 30 KB, Slow start threshold: 139 KB
      Measured RTT : 500000 us Min RTT: 11 us Max RTT: 0 us
      Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      72 frames input, 8796 bytes
         72 Class F frames input, 8796 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      73 frames output, 7632 bytes
         73 Class F frames output, 7632 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames

SAN384C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
  port-channel100  0x11(17)  port-channel100  10.122.107.94(SAN192C-6)
          fcip120  0x11(17)          fcip120  10.122.107.94(SAN192C-6)
SAN384C-6#
```

SAN192C-6

```
SAN192C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# interface IPStorage6/3
```

```
SAN192C-6(config-if)# switchport mtu 2500
SAN192C-6(config-if)# no shut
SAN192C-6(config-if)#
SAN192C-6(config-if)# interface IPStorage6/3.1000
SAN192C-6(config-if)# ip address 10.1.3.1 255.255.255.0
SAN192C-6(config-if)# switchport mtu 2500
SAN192C-6(config-if)# no shut
SAN192C-6(config-if)#
SAN192C-6(config-if)# fcip profile 120
SAN192C-6(config-profile)# ip address 10.1.3.1
SAN192C-6(config-profile)# tcp max-bandwidth-mbps 5000
min-available-bandwidth-mbps 4500  round-trip-time-ms 1
SAN192C-6(config-profile)#
SAN192C-6(config-profile)# interface fcip120
SAN192C-6(config-if)# use-profile 120
SAN192C-6(config-if)# peer-info ipaddr 10.1.3.2
SAN192C-6(config-if)# tcp-connections 5
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)# switchport trunk allowed vsan add 50
SAN192C-6(config-if)#
SAN192C-6(config-if)# ficon portnumber 0xf3
SAN192C-6(config-if)# no shut
SAN192C-6(config-if)# end
Performing fast copy of configurationdone.
SAN192C-6#
SAN192C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
  port-channel100  0x10(16)  port-channel100  10.122.107.95(SAN384C-6)
          fcip120  0x10(16)          fcip120  10.122.107.95(SAN384C-6)
SAN192C-6#
```

Example 7-12 shows how we configured FCIP on VLAN sub-interfaces on the second sub-interface for the IBM Storage Networking SAN384C-6 and IBM Storage Networking SAN192C-6 switches with encryption and compression enabled by using the CLI.

Example 7-12 Configuring FCIP on VLAN sub-interfaces with encryption and compression

```
SAN384C-6#
SAN384C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN384C-6(config)# feature crypto ike
SAN384C-6(config)# feature crypto ipsec
SAN384C-6(config)#
SAN384C-6(config)# interface IPStorage7/3
SAN384C-6(config-if)# switchport mtu 2500
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# interface IPStorage7/3.1010
SAN384C-6(config-if)# ip address 10.1.4.2 255.255.255.0
```

```
SAN384C-6(config-if)# switchport mtu 2500
SAN384C-6(config-if)#
SAN384C-6(config-if)# fcip profile 130
SAN384C-6(config-profile)# ip address 10.1.4.2
SAN384C-6(config-profile)# tcp max-bandwidth-mbps 5000
min-available-bandwidth-mbps 4500  round-trip-time-ms 1
SAN384C-6(config-profile)#
SAN384C-6(config-profile)# interface fcip130
SAN384C-6(config-if)# use-profile 130
SAN384C-6(config-if)# peer-info ipaddr 10.1.4.1
SAN384C-6(config-if)# tcp-connections 5
SAN384C-6(config-if)# ip-compression auto
SAN384C-6(config-if)#
SAN384C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN384C-6(config-if)# switchport trunk allowed vsan add 50
SAN384C-6(config-if)#
SAN384C-6(config-if)# ficon portnumber 0xf4
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# crypto ike domain ipsec
SAN384C-6(config-ike-ipsec)# policy 1
SAN384C-6(config-ike-ipsec-policy)# key 7 swwxoomi address 10.1.4.1
SAN384C-6(config-ike-ipsec)#
SAN384C-6(config-ike-ipsec)# crypto map domain ipsec crset-fcip130-redbook 1
SAN384C-6(config-crypto-map-ip)# set peer 10.1.4.1
SAN384C-6(config-crypto-map-ip)# match address access_list_fcip130_redbook
SAN384C-6(config-crypto-map-ip)# set transform-set ipsec_default_transform_set
SAN384C-6(config-crypto-map-ip)#
SAN384C-6(config-crypto-map-ip)# ip access-list access_list_fcip130_redbook permit
ip 10.1.4.2 0.0.0.0 10.1.4.1 0.0.0.0
SAN384C-6(config)#
SAN384C-6(config)# interface IPStorage7/3.1010
SAN384C-6(config-if)# crypto map domain ipsec crset-fcip130-redbook
SAN384C-6(config-if)# no shut
SAN384C-6(config-if)#
SAN384C-6(config-if)# end
Performing fast copy of configurationdone.
SAN384C-6#
SAN384C-6# show fcip summary

-------------------------------------------------------------------------------
Tun prof    IPS-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
-------------------------------------------------------------------------------
100 100  IPS7/1     10.1.1.1         TRNK  Y N N  N   N   10000M/9500M  1000
110 110  IPS7/2     10.1.2.1         TRNK  Y N N  N   N   10000M/9500M  1000
120 120  IPS7/3.1000 10.1.3.1         TRNK  Y N N  N   N   5000M/4500M  1000
130 130  IPS7/3.1010 10.1.4.1         TRNK  Y N N  Y   A   5000M/4500M  1000

SAN384C-6#
SAN384C-6# show int fcip130
fcip130 is trunking
```

```
   Hardware is IPStorage
   Port WWN is 21:a3:00:3a:9c:31:62:80
   Peer port WWN is 21:63:00:2a:6a:a4:1a:80
   Admin port mode is auto, trunk mode is on
   snmp link state traps are enabled
   Port mode is TE
   Port vsan is 1
   Operating Speed is 5 Gbps
   Trunk vsans (admin allowed and active) (1,50)
   Trunk vsans (up)                       (1,50)
   Trunk vsans (isolated)                 ()
   Trunk vsans (initializing)             ()
   Interface last changed at Sun Mar 15 19:58:26 2020

   Using Profile id 130  (interface IPStorage7/3.1010)
   Peer Information
     Peer Internet address is 10.1.4.1 and port is 3225
   FCIP tunnel is protected by IPsec
   Write acceleration mode is configured off
   Tape acceleration mode is configured off
   Tape Accelerator flow control buffer size is automatic
   FICON XRC Accelerator is configured off
   FICON Load Balancer configured off for all vsans
   FICON Tape acceleration configured off for all vsans
   IP Compression is enabled and set for auto
   Maximum number of TCP connections is 5
   QOS control code point is 0
   QOS data code point is 0
   TCP Connection Information
     5 Active TCP connections
       30 Attempts for active connections, 1 close of connections
       Path MTU 2400 bytes
       Current retransmission timeout is 200 ms
       Current Send Buffer Size: 62558 KB, Requested Send Buffer Size: 62500 KB
       CWM Burst Size: 50 KB
CONN<0>
   Data connection: Local 10.1.4.2:65460, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 58 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
     Congestion window: Current: 182 KB, Slow start threshold: 182 KB
     Measured RTT : 1386 us Min RTT: 500000 us Max RTT: 1386 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
   Data connection: Local 10.1.4.2:65459, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 14 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
     Congestion window: Current: 14 KB, Slow start threshold: 133 KB
     Measured RTT : 67 us Min RTT: 500000 us Max RTT: 67 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
```

```
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<2>
    Data connection: Local 10.1.4.2:65457, Remote 10.1.4.1:3225
    TCP Parameters
      Advertized window: Current: 14 KB, Maximum: 24580 KB, Scale: 6
      Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
      Congestion window: Current: 14 KB, Slow start threshold: 133 KB
      Measured RTT : 68 us Min RTT: 500000 us Max RTT: 68 us
      Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<3>
    Data connection: Local 10.1.4.2:65455, Remote 10.1.4.1:3225
    TCP Parameters
      Advertized window: Current: 14 KB, Maximum: 24580 KB, Scale: 6
      Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
      Congestion window: Current: 14 KB, Slow start threshold: 133 KB
      Measured RTT : 78 us Min RTT: 500000 us Max RTT: 78 us
      Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
 CONN<4>
    Control connection: Local 10.1.4.2:65453, Remote 10.1.4.1:3225
    TCP Parameters
      Advertized window: Current: 22 KB, Maximum: 24580 KB, Scale: 6
      Peer receive window: Current: 4089 KB, Maximum: 4095 KB, Scale: 6
      Congestion window: Current: 17 KB, Slow start threshold: 133 KB
      Measured RTT : 31 us Min RTT: 500000 us Max RTT: 74 us
      Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 152 us
    TCP Connection Rate
      Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
      Input Frames: 0/sec, Output Frames: 0/sec
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      72 frames input, 8796 bytes
         72 Class F frames input, 8796 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      73 frames output, 7632 bytes
         73 Class F frames output, 7632 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames

SAN384C-6# show int fcip130 counters
fcip130
    TCP Connection Information
      5 Active TCP connections
        30 Attempts for active connections, 1 close of connections
        Path MTU 2400 bytes
        Current retransmission timeout is 200 ms
        Current Send Buffer Size: 62600 KB, Requested Send Buffer Size: 62500 KB
```

```
       CWM Burst Size: 50 KB
CONN<0>
   Data connection: Local 10.1.4.2:65460, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 58 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
     Congestion window: Current: 182 KB, Slow start threshold: 573 KB
     Measured RTT : 500000 us Min RTT: 1386 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
   Data connection: Local 10.1.4.2:65459, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 14 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
     Congestion window: Current: 14 KB, Slow start threshold: 541 KB
     Measured RTT : 500000 us Min RTT: 67 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
   Data connection: Local 10.1.4.2:65457, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 14 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
     Congestion window: Current: 14 KB, Slow start threshold: 541 KB
     Measured RTT : 500000 us Min RTT: 68 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
   Data connection: Local 10.1.4.2:65455, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 14 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
     Congestion window: Current: 14 KB, Slow start threshold: 541 KB
     Measured RTT : 500000 us Min RTT: 78 us Max RTT: 0 us
     Round trip time: Smoothed 0 ms, Variance: 0 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
     Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
   Control connection: Local 10.1.4.2:65453, Remote 10.1.4.1:3225
   TCP Parameters
     Advertized window: Current: 26 KB, Maximum: 24580 KB, Scale: 6
     Peer receive window: Current: 4089 KB, Maximum: 4089 KB, Scale: 6
     Congestion window: Current: 17 KB, Slow start threshold: 541 KB
     Measured RTT : 33 us Min RTT: 31 us Max RTT: 37 us
     Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
   TCP Connection Rate
     Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
```

```
      Input Frames: 0/sec, Output Frames: 0/sec
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      76 frames input, 9164 bytes
         76 Class F frames input, 9164 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      77 frames output, 8000 bytes
         77 Class F frames output, 8000 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
    IP compression statistics
      9340 rxbytes
         6270 rxbytes compressed, 0 rxbytes non-compressed
         1.49 rx compression ratio
      8000 txbytes
         6023 txbytes compressed, 0 txbytes non-compressed
         1.33 tx compression ratio
    IP compression flow control statistics
      0 bytes queued for hw compression
      0 queued for hardware compression
      0 queued for hardware decompression
      0 slowed tcp flow control
      0 accelerated tcp flow control
      0 side band flow control ON
      0 side band flow control OFF
    IP compression hung statistics
      0 times compression engine hung detected
      0 jobs replayed for hardware compression
      0 jobs replayed for hardware decompression
      0 compression jobs not processed during compression engine reset
      0 compression response job not processed during compression engine reset
      0 decompression jobs not processed during decompression engine reset
      0 decompression response job not processed during decompression engine reset
SAN384C-6#
SAN384C-6# show crypto ike domain ipsec sa
Tunn Local Addr           Remote Addr          Encr  Hash  Auth Method
Lifetime
----------------------------------------------------------------------------------
--
  2  10.1.4.2[500]        10.1.4.1[500]       3des  sha1  pre-shared-key  86400
----------------------------------------------------------------------------------
--
NOTE: tunnel id ended with * indicates an IKEv1 tunnel.

SAN384C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
  port-channel100  0x11(17)  port-channel100  10.122.107.94(SAN192C-6)
          fcip120  0x11(17)          fcip120  10.122.107.94(SAN192C-6)
          fcip130  0x11(17)          fcip130  10.122.107.94(SAN192C-6)
```

```
SAN384C-6#
```

SAN192C-6

```
SAN192C-6#
SAN192C-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SAN192C-6(config)# feature crypto ike
SAN192C-6(config)# feature crypto ipsec
SAN192C-6(config)#
SAN192C-6(config)# interface IPStorage6/3
SAN192C-6(config-if)# switchport mtu 2500
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)#
SAN192C-6(config-if)# interface IPStorage6/3.1010
SAN192C-6(config-if)# ip address 10.1.4.1 255.255.255.0
SAN192C-6(config-if)# switchport mtu 2500
SAN192C-6(config-if)#
SAN192C-6(config-if)# fcip profile 130
SAN192C-6(config-profile)# ip address 10.1.4.1
SAN192C-6(config-profile)# tcp max-bandwidth-mbps 5000
min-available-bandwidth-mbps 4500  round-trip-time-ms 1
SAN192C-6(config-profile)#
SAN192C-6(config-profile)# interface fcip130
SAN192C-6(config-if)# use-profile 130
SAN192C-6(config-if)# peer-info ipaddr 10.1.4.2
SAN192C-6(config-if)# tcp-connections 5
SAN192C-6(config-if)# ip-compression auto
SAN192C-6(config-if)# switchport trunk allowed vsan 1
Warning: This command will remove all VSANs currently being trunked and trunk only
the specified VSANs.
Do you want to continue? (y/n) [n] y
SAN192C-6(config-if)# switchport trunk allowed vsan add 50
SAN192C-6(config-if)# ficon portnumber 0xf4
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)#
SAN192C-6(config-if)# crypto ike domain ipsec
SAN192C-6(config-ike-ipsec)# policy 1
SAN192C-6(config-ike-ipsec-policy)# key 7 swwxoomi address 10.1.4.2
SAN192C-6(config-ike-ipsec)#
SAN192C-6(config-ike-ipsec)# crypto map domain ipsec crset-fcip130-redbook 1
SAN192C-6(config-crypto-map-ip)# set peer 10.1.4.2
SAN192C-6(config-crypto-map-ip)# match address access_list_fcip130_redbook
SAN192C-6(config-crypto-map-ip)# set transform-set ipsec_default_transform_set
SAN192C-6(config-crypto-map-ip)#
SAN192C-6(config-crypto-map-ip)# ip access-list access_list_fcip130_redbook permit
ip 10.1.4.1 0.0.0.0 10.1.4.2 0.0.0.0
SAN192C-6(config)#
SAN192C-6(config)# interface IPStorage6/3.1010
SAN192C-6(config-if)# crypto map domain ipsec crset-fcip130-redbook
SAN192C-6(config-if)# no shutdown
SAN192C-6(config-if)# end
Performing fast copy of configurationdone.
SAN192C-6# show crypto ike domain ipsec sa
```

```
Tunn Local Addr          Remote Addr          Encr  Hash  Auth Method
Lifetime
---------------------------------------------------------------------------------
--
  4  10.1.4.1[500]       10.1.4.2[500]       3des  sha1  pre-shared-key  86400
---------------------------------------------------------------------------------
--
NOTE: tunnel id ended with * indicates an IKEv1 tunnel.
SAN192C-6# show fcip summary

-------------------------------------------------------------------------------
Tun prof    IPS-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
-------------------------------------------------------------------------------
100 100  IPS6/1     10.1.1.2         TRNK  Y N N  N   N   10000M/9500M  1000
110 110  IPS6/2     10.1.2.2         TRNK  Y N N  N   N   10000M/9500M  1000
120 120  IPS6/3.1000 10.1.3.2         TRNK  Y N N  N   N   5000M/4500M  1000
130 130  IPS6/3.1010 10.1.4.2         TRNK  Y N N  Y   A   5000M/4500M  1000

SAN192C-6# show topology vsan 50

FC Topology for VSAN 50 :
--------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
--------------------------------------------------------------------------------
  port-channel100  0x10(16)  port-channel100  10.122.107.95(SAN384C-6)
          fcip120  0x10(16)          fcip120  10.122.107.95(SAN384C-6)
          fcip130  0x10(16)          fcip130  10.122.107.95(SAN384C-6)
SAN192C-6#
```

7.2.19 Configuring FCIP for open systems VSANs

This section shows how to configure FCIP links for open systems VSANs. In this example, we configure FCIP links between the IBM Storage Networking SAN384C-6 and IBM Storage Networking SAN50C-R switches by using the DCNM.

Configuring FCIP

Start DCNM and select **Configure** → **SAN** → **FCIP**, as shown in Figure 7-151.

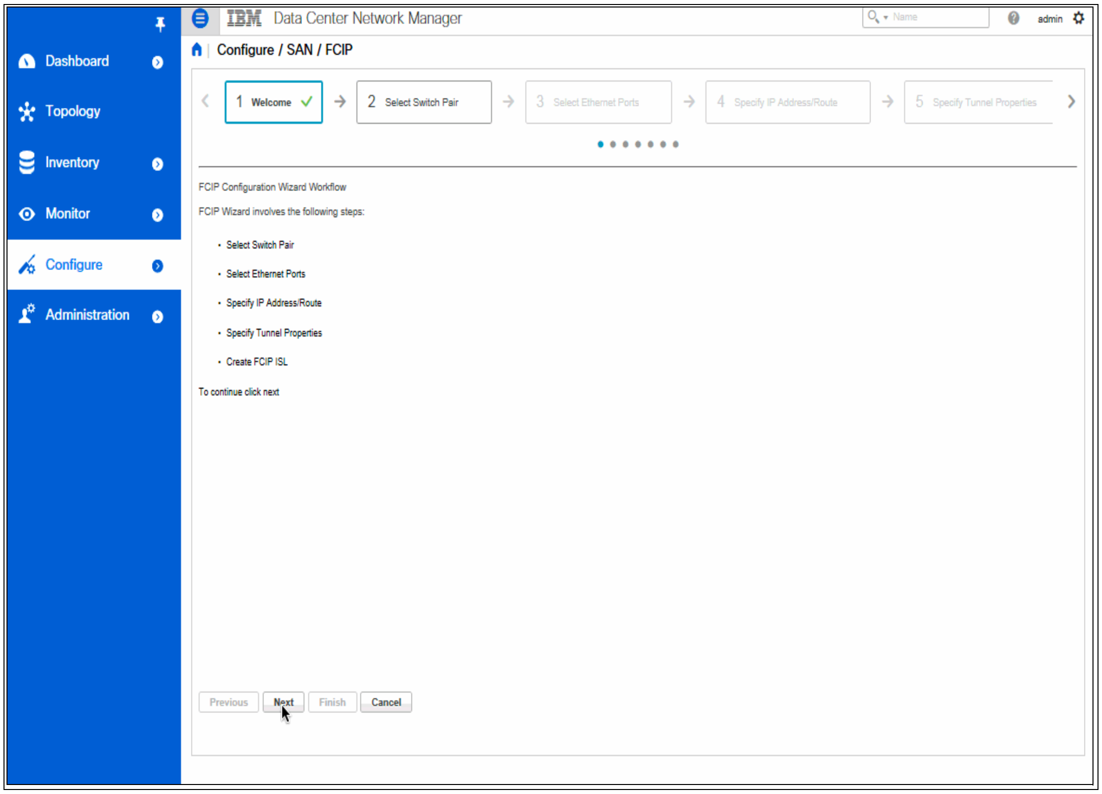

FCIP Configuration Wizard Workflow

FCIP Wizard involves the following steps:

- Select Switch Pair
- Select Ethernet Ports
- Specify IP Address/Route
- Specify Tunnel Properties
- Create FCIP ISL

To continue click next

Figure 7-151 Welcome page summary

Click **Next**. Select a switch pair, as shown in Figure 7-152.

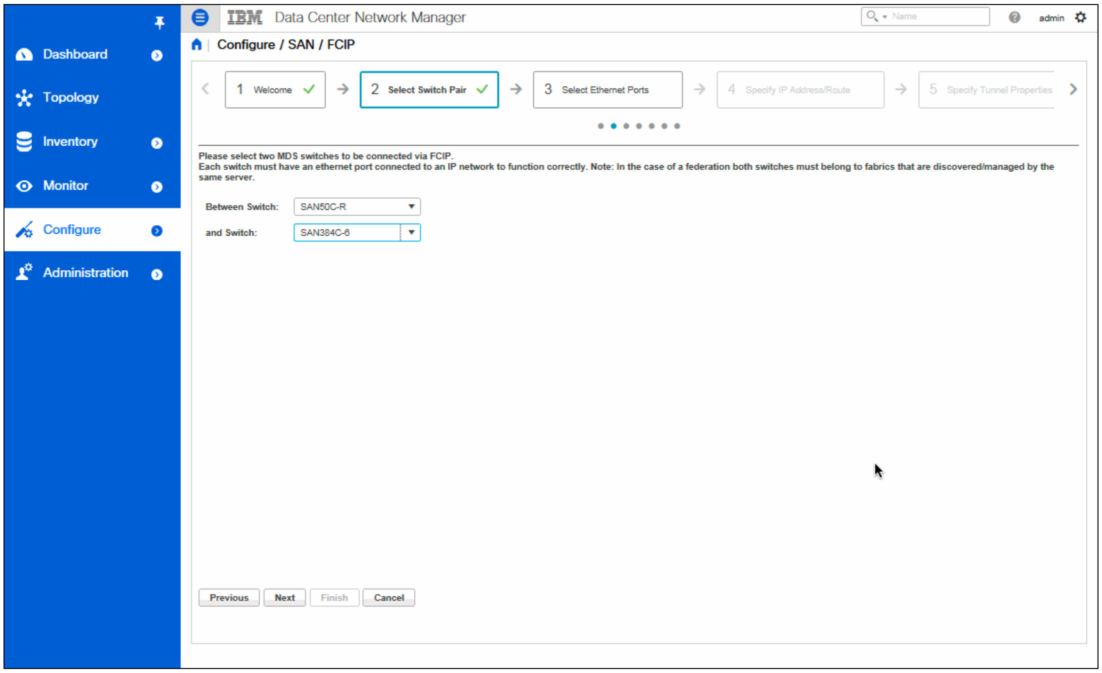

Please select two MDS switches to be connected via FCIP.
Each switch must have an ethernet port connected to an IP network to function correctly. Note: In the case of a federation both switches must belong to fabrics that are discovered/managed by the same server.

Between Switch: SAN50C-R

and Switch: SAN384C-6

Figure 7-152 Selecting a switch pair

Select **Ethernet Ports** and check the Enforce IPSEC Security check and provide an **Auth Key**, as shown in Figure 7-153. Click **Next**.

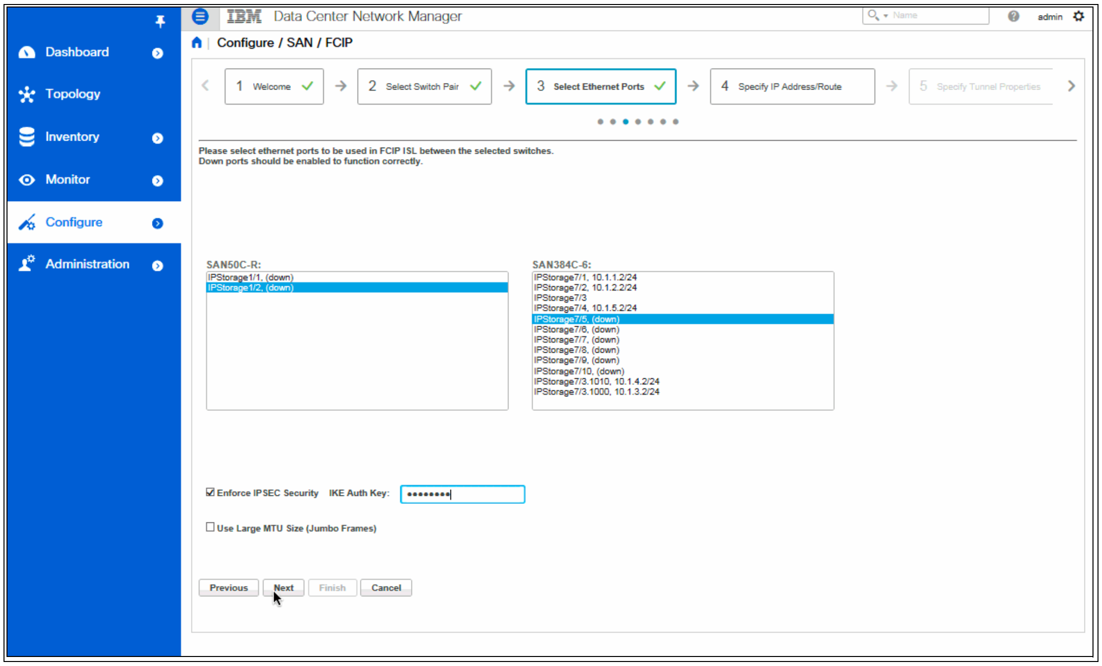

Please select ethernet ports to be used in FCIP ISL between the selected switches.
Down ports should be enabled to function correctly.

SAN50C-R:

IPStorage1/1, (down)
IPStorage1/2, (down)

SAN384C-6:

IPStorage7/1, 10.1.1.2/24
IPStorage7/2, 10.1.2.2/24
IPStorage7/3
IPStorage7/4, 10.1.5.2/24
IPStorage7/5, (down)
IPStorage7/6, (down)
IPStorage7/7, (down)
IPStorage7/8, (down)
IPStorage7/9, (down)
IPStorage7/10, (down)
IPStorage7/3.1010, 10.1.4.2/24
IPStorage7/3.1000, 10.1.3.2/24

☑ Enforce IPSEC Security IKE Auth Key: ••••••••

☐ Use Large MTU Size (Jumbo Frames)

Figure 7-153 Ethernet ports that are used for the FCIP tunnel

Routes are needed in this example because the IP addresses are not in the same subnet. Specify the **IP Address/Route**, as shown in Figure 7-154, and click **Next**.

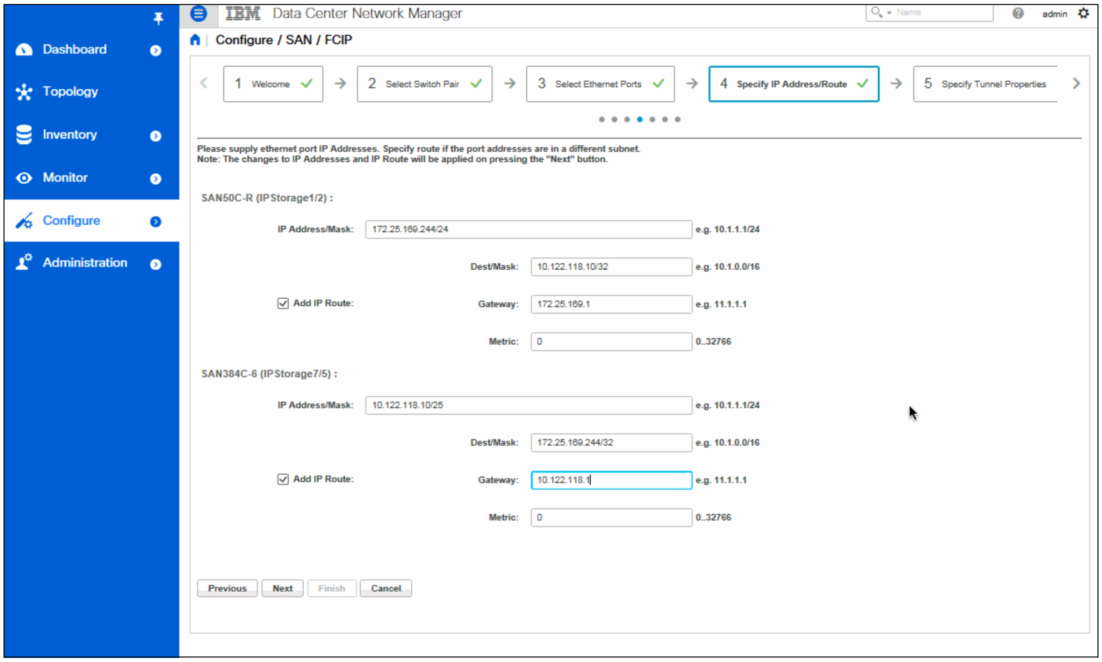

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN50C-R (IPStorage1/2) :

IP Address/Mask: 172.25.169.244/24 e.g. 10.1.1.1/24

☑ Add IP Route:

Dest/Mask: 10.122.118.10/32 e.g. 10.1.0.0/16

Gateway: 172.25.169.1 e.g. 11.1.1.1

Metric: 0 0..32766

SAN384C-6 (IPStorage7/5) :

IP Address/Mask: 10.122.118.10/25 e.g. 10.1.1.1/24

☑ Add IP Route:

Dest/Mask: 172.25.169.244/32 e.g. 10.1.0.0/16

Gateway: 10.122.118.1 e.g. 11.1.1.1

Metric: 0 0..32766

Figure 7-154 IP Address/Route

Note: As a best practice, engage your network team when using a routed network for FCIP links. A destination IP, mask, and gateway address are required for each site. This information is used to configure static routes on the IBM c-type switches.

In Figure 7-155, select **Yes** to continue.

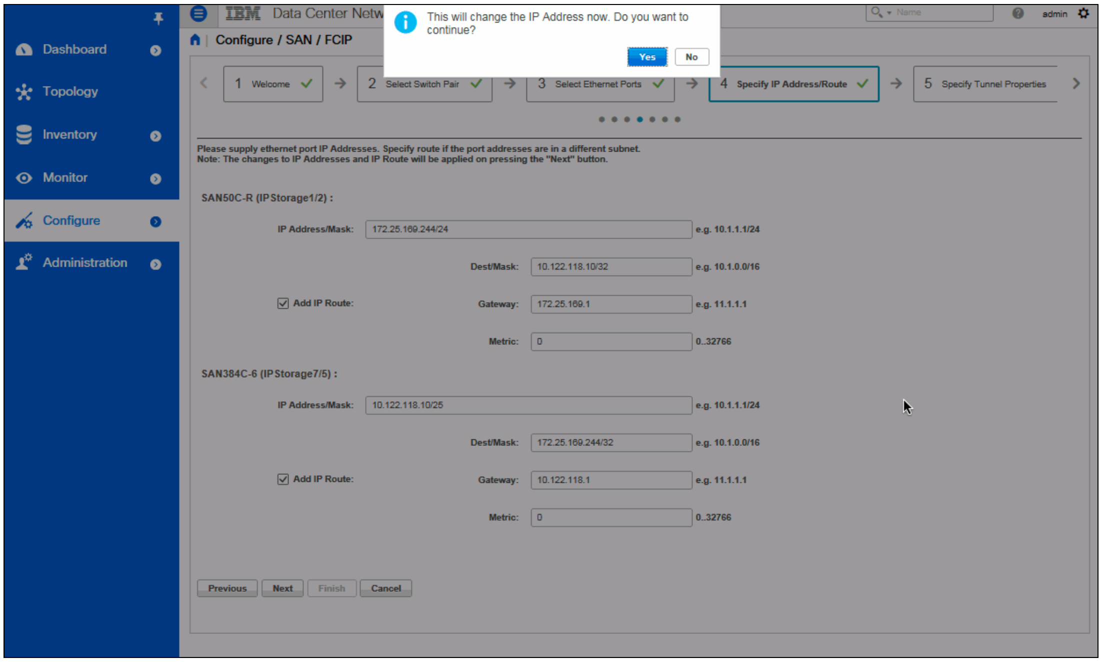

This will change the IP Address now. Do you want to continue?

Yes No

Please supply ethernet port IP Addresses. Specify route if the port addresses are in a different subnet.
Note: The changes to IP Addresses and IP Route will be applied on pressing the "Next" button.

SAN50C-R (IPStorage1/2) :

IP Address/Mask: 172.25.169.244/24 e.g. 10.1.1.1/24

☑ Add IP Route:

Dest/Mask: 10.122.118.10/32 e.g. 10.1.0.0/16

Gateway: 172.25.169.1 e.g. 11.1.1.1

Metric: 0 0..32766

SAN384C-6 (IPStorage7/5) :

IP Address/Mask: 10.122.118.10/25 e.g. 10.1.1.1/24

☑ Add IP Route:

Dest/Mask: 172.25.169.244/32 e.g. 10.1.0.0/16

Gateway: 10.122.118.1 e.g. 11.1.1.1

Metric: 0 0..32766

Previous Next Finish Cancel

Figure 7-155 Configuring the IP address and route

Specify the tunnel properties, as shown in Figure 7-156.

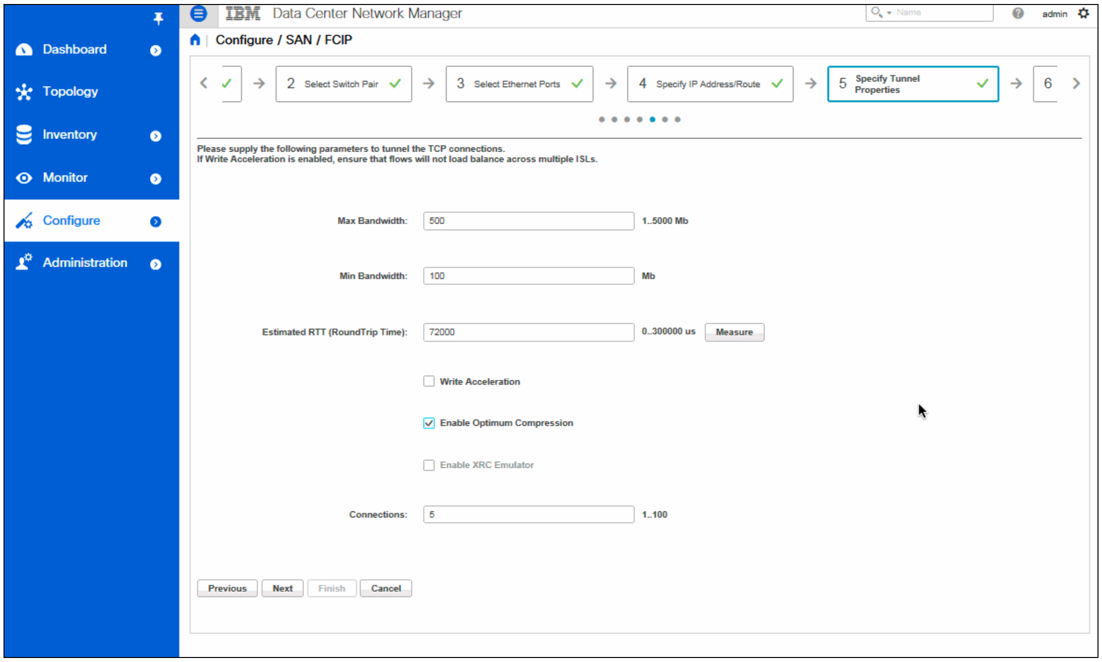

Please supply the following parameters to tunnel the TCP connections.
If Write Acceleration is enabled, ensure that flows will not load balance across multiple ISLs.

Max Bandwidth: 500 1..5000 Mb

Min Bandwidth: 100 Mb

Estimated RTT (RoundTrip Time): 72000 0..300000 us Measure

☐ Write Acceleration

☑ Enable Optimum Compression

☐ Enable XRC Emulator

Connections: 5 1..100

Previous Next Finish Cancel

Figure 7-156 Tunnel properties

Measure the estimated RTT, as shown in Figure 7-157.

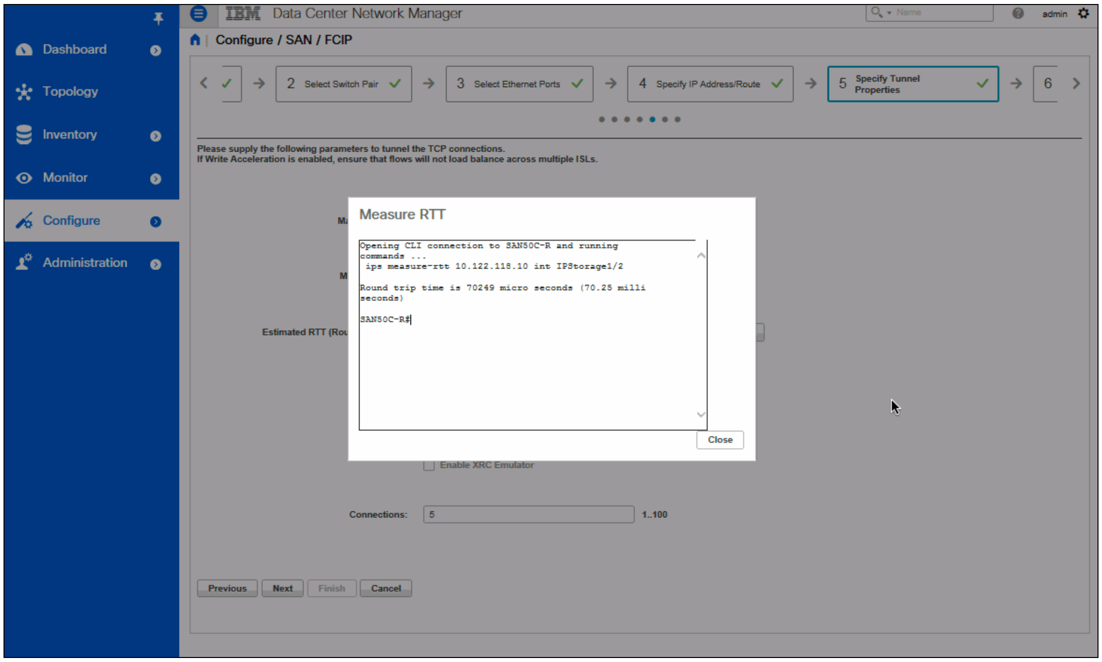

Please supply the following parameters to tunnel the TCP connections.
If Write Acceleration is enabled, ensure that flows will not load balance across multiple ISLs.

Measure RTT

```
Opening CLI connection to SAN50C-R and running
commands ...
 ips measure-rtt 10.122.118.10 int IPStorage1/2

Round trip time is 70249 micro seconds (70.25 milli
seconds)

SAN50C-R#
```

□ Enable XRC Emulator

Connections: 5 1..100

Figure 7-157 Measuring RTT

Create the FCIP ISL, as shown in Figure 7-158.

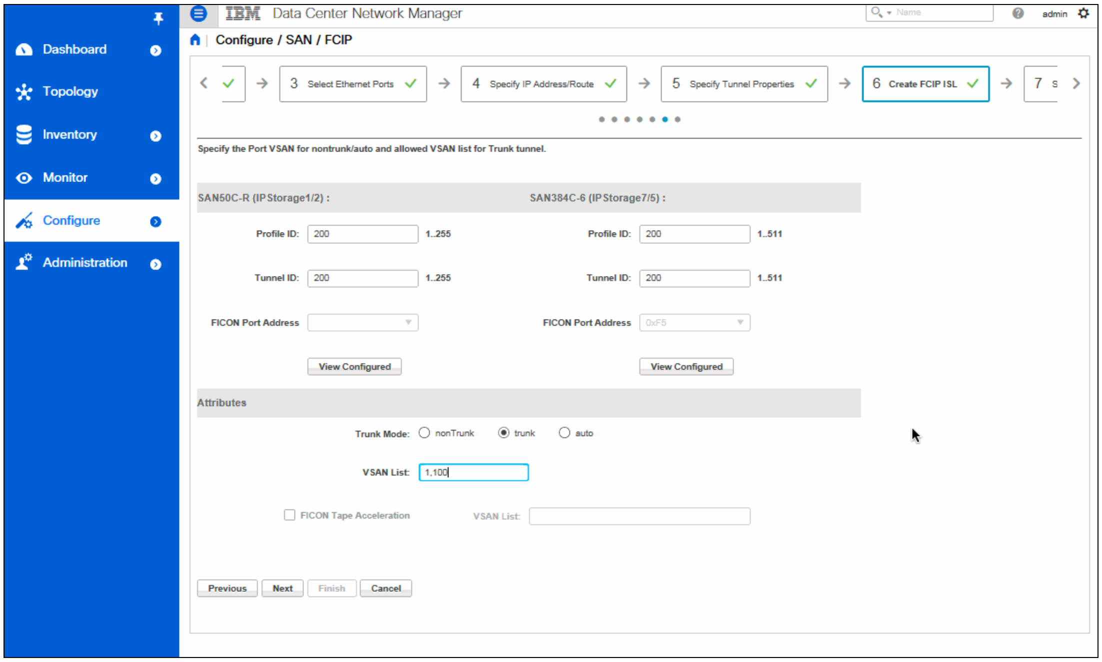

Specify the Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.

SAN50C-R (IPStorage1/2) :

Profile ID: 200 1..255

Tunnel ID: 200 1..255

FICON Port Address

SAN384C-6 (IPStorage7/5) :

Profile ID: 200 1..511

Tunnel ID: 200 1..511

FICON Port Address 0xF5

Attributes

Trunk Mode: ○ nonTrunk ● trunk ○ auto

VSAN List: 1,100

□ FICON Tape Acceleration VSAN List:

Figure 7-158 FCIP ISL configuration

Verify that routes exist if you are using different subnets, and then click **OK** to proceed, as shown Figure 7-159.

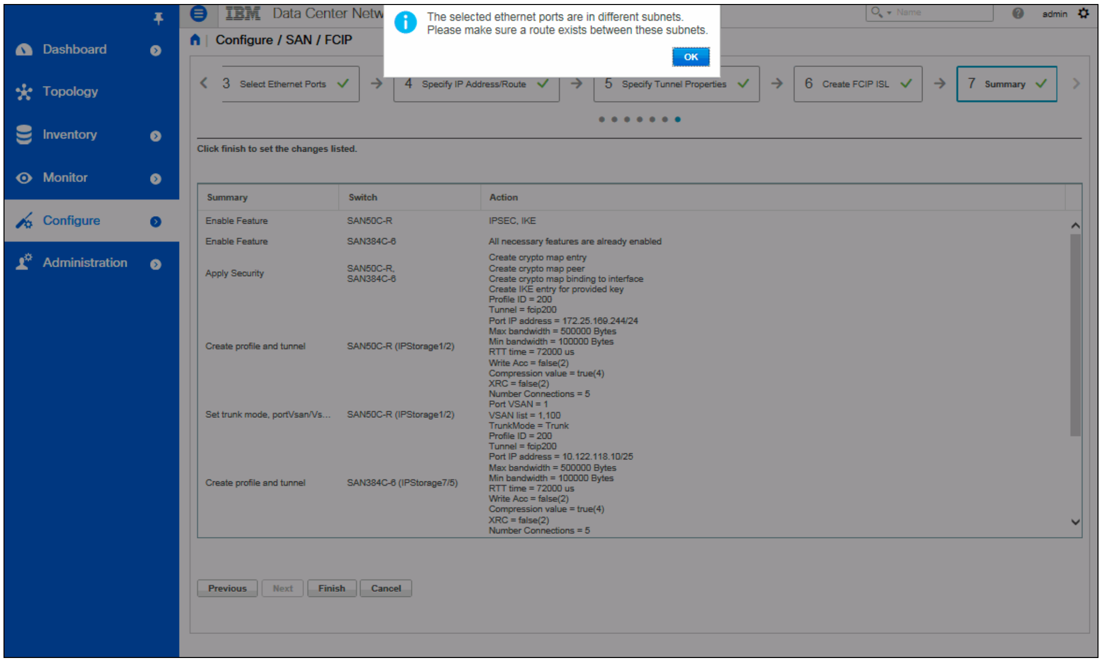

The selected ethernet ports are in different subnets.
Please make sure a route exists between these subnets.

OK

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN50C-R	IPSEC, IKE
Enable Feature	SAN384C-6	All necessary features are already enabled
Apply Security	SAN50C-R, SAN384C-6	Create crypto map entry Create crypto map peer Create crypto map binding to interface Create IKE entry for provided key
Create profile and tunnel	SAN50C-R (IPStorage1/2)	Profile ID = 200 Tunnel = fcip200 Port IP address = 172.25.169.244/24 Max bandwidth = 500000 Bytes Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN50C-R (IPStorage1/2)	Port VSAN = 1 VSAN list = 1,100 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/5)	Profile ID = 200 Tunnel = fcip200 Port IP address = 10.122.118.10/25 Max bandwidth = 500000 Bytes Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5

Figure 7-159 Verifying that routes exist

Figure 7-160 shows the FCIP configuration summary.

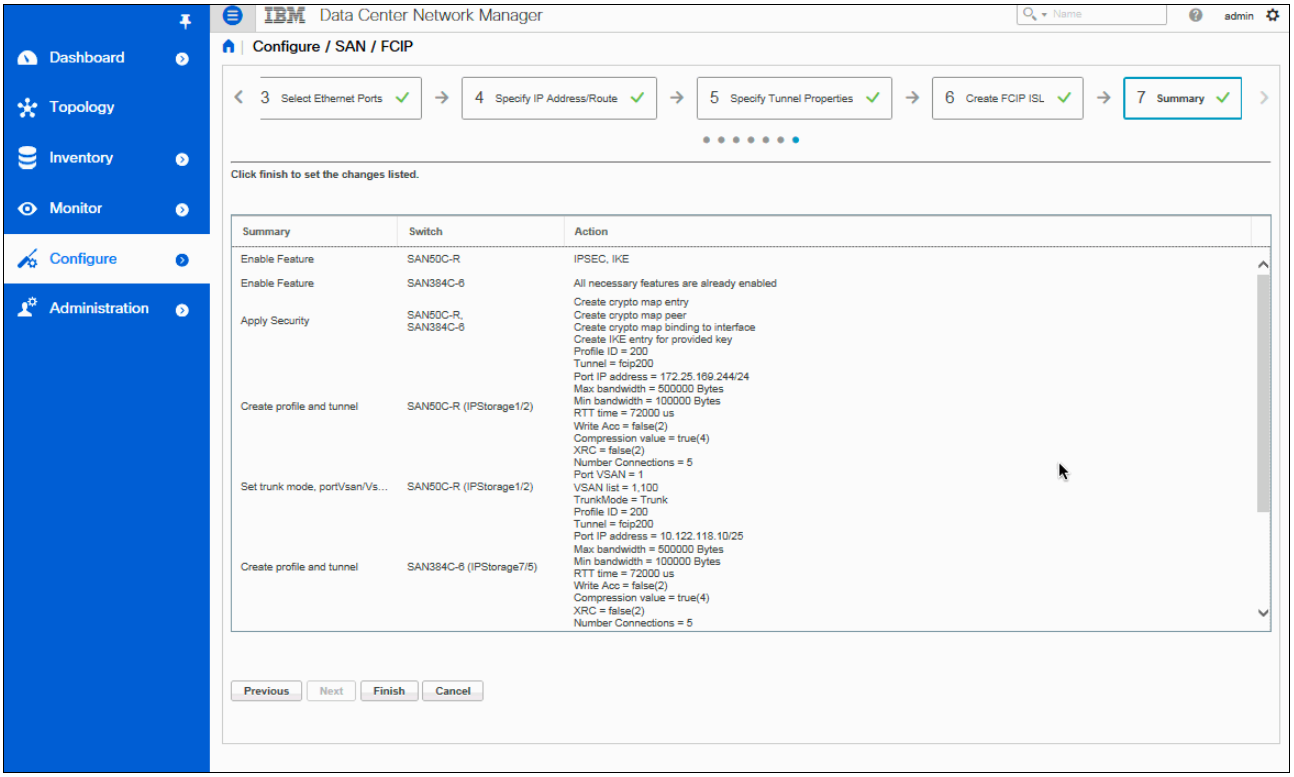

Click finish to set the changes listed.

Summary	Switch	Action
Enable Feature	SAN50C-R	IPSEC, IKE
Enable Feature	SAN384C-6	All necessary features are already enabled
Apply Security	SAN50C-R, SAN384C-6	Create crypto map entry Create crypto map peer Create crypto map binding to interface Create IKE entry for provided key
Create profile and tunnel	SAN50C-R (IPStorage1/2)	Profile ID = 200 Tunnel = fcip200 Port IP address = 172.25.169.244/24 Max bandwidth = 500000 Bytes Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN50C-R (IPStorage1/2)	Port VSAN = 1 VSAN list = 1,100 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/5)	Profile ID = 200 Tunnel = fcip200 Port IP address = 10.122.118.10/25 Max bandwidth = 500000 Bytes Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5

Figure 7-160 FCIP configuration summary

Figure 7-161 on page 395 shows the FCIP configuration summary continued.

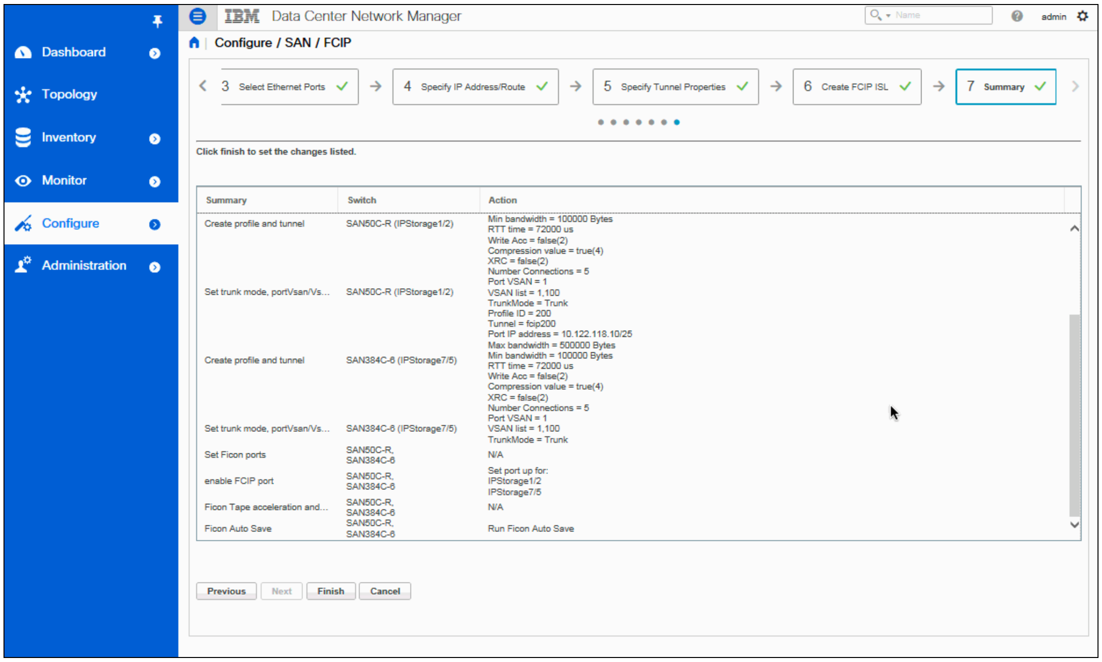

Click finish to set the changes listed.

Summary	Switch	Action
Create profile and tunnel	SAN50C-R (IPStorage1/2)	Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN50C-R (IPStorage1/2)	Port VSAN = 1 VSAN list = 1,100 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/5)	Profile ID = 200 Tunnel = fcip200 Port IP address = 10.122.118.10/25 Max bandwidth = 500000 Bytes Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/5)	Port VSAN = 1 VSAN list = 1,100 TrunkMode = Trunk
Set Ficon ports	SAN50C-R, SAN384C-6	N/A
enable FCIP port	SAN50C-R, SAN384C-6	Set port up for: IPStorage1/2 IPStorage7/5
Ficon Tape acceleration and...	SAN50C-R, SAN384C-6	N/A
Ficon Auto Save	SAN50C-R, SAN384C-6	Run Ficon Auto Save

Figure 7-161 FCIP configuration summary (cont.)

Click **OK** and then **Close** to apply the configuration settings, as shown in Figure 7-162.

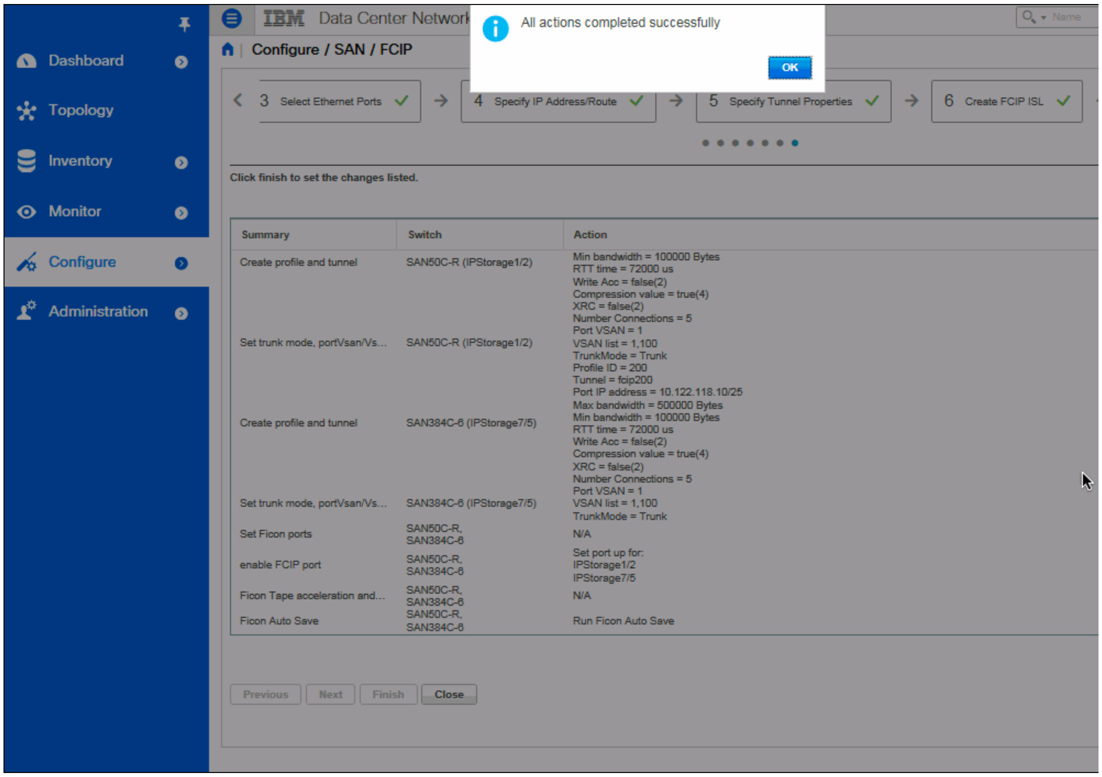

All actions completed successfully

Click finish to set the changes listed.

Summary	Switch	Action
Create profile and tunnel	SAN50C-R (IPStorage1/2)	Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN50C-R (IPStorage1/2)	Port VSAN = 1 VSAN list = 1,100 TrunkMode = Trunk
Create profile and tunnel	SAN384C-6 (IPStorage7/5)	Profile ID = 200 Tunnel = fcip200 Port IP address = 10.122.118.10/25 Max bandwidth = 500000 Bytes Min bandwidth = 100000 Bytes RTT time = 72000 us Write Acc = false(2) Compression value = true(4) XRC = false(2) Number Connections = 5
Set trunk mode, portVsan/Vs...	SAN384C-6 (IPStorage7/5)	Port VSAN = 1 VSAN list = 1,100 TrunkMode = Trunk
Set Ficon ports	SAN50C-R, SAN384C-6	N/A
enable FCIP port	SAN50C-R, SAN384C-6	Set port up for: IPStorage1/2 IPStorage7/5
Ficon Tape acceleration and...	SAN50C-R, SAN384C-6	N/A
Ficon Auto Save	SAN50C-R, SAN384C-6	Run Ficon Auto Save

Figure 7-162 Applying the FCIP configuration

Verifying that the FCIP link status is up and IPsec security is configured

Note: Log in to both local and remote switches when validating that the FCIP links are up and IPsec security is configured.

Example 7-13 shows that the FCIP link status is up.

Example 7-13 FCIP link status

```
SAN50C-R# show fcip summary

-------------------------------------------------------------------------------
Tun prof    IPS-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
-------------------------------------------------------------------------------
200 200  IPS1/2     10.122.118.10    TRNK  Y N N  Y   A     500M/100M  72000

SAN50C-R#
```

Example 7-14 shows the IPsec and IKE security configuration for FCIP tunnels.

Example 7-14 FCIP tunnels

```
SAN50C-R# show crypto ike domain ipsec sa
Tunn Local Addr          Remote Addr         Encr  Hash  Auth Method
Lifetime
--------------------------------------------------------------------------------
--
  2  172.25.169.244[500]  10.122.118.10[500]   3des  sha1  pre-shared-key  86400
--------------------------------------------------------------------------------
--
NOTE: tunnel id ended with * indicates an IKEv1 tunnel.

SAN50C-R#
```

8

IBM Storage Networking c-type operations

The IBM c-type switches are robust and reliable. They have proven excellent availability in the field and should give you years of trouble-free service. However, to realize these benefits, you must monitor the switches, perform necessary maintenance such as upgrade system code as necessary, and diagnose and resolve any outstanding warning or error conditions. Failure to maintain efficient day-to-day operations and any required troubleshooting might affect the IBM c-type switches and the devices that are either directly or indirectly connected to them, such as, storage arrays, tape libraries, mainframe servers, and open system servers.

This chapter provides an overview of useful tasks that you might want to consider regarding your day-to-day operations and any troubleshooting that might be required in your environment.

The following topics are covered in this chapter:

- Performance monitoring tools
- Backing up a switch configuration
- Call Home
- Read Diagnostic Parameters
- Port beacon and location LED
- On-board Failure Logging
- FICON Director Activity Report

8.1 Performance monitoring tools

This section provides insight into some standard and advanced performance monitoring features and tools that are used with IBM c-type storage area network (SAN) switches and directors that can be leveraged to aid in SAN fabric resiliency. The performance features and tools that are described in this section are only a subset of all the tools that are available.

To view performance information in your SAN environment, IBM c-type switches use Data Center Network Manager (DCNM) as the standard base tool for performance monitoring. In tandem with Device Manager (DM), both tools can provide several mechanisms that you can use to monitor and view real-time, light-weight, and high-level historical data for IBM c-type switch performance and troubleshooting. Data can be graphed over time to provide a real-time insight into the performance of the port, such as the following items:

- Real-time SAN Inter-Switch Link (ISL) statistics
- SAN modules, ports, and a host of additional SAN elements
- The entire SAN fabric health
- Ingress and egress Fibre Channel (FC) traffic errors
- Class 2 traffic that shows buffer-to-buffer (B2B) and end-to-end credit flow control statistics
- Checking for oversubscription
- Threshold monitoring
- RX and TX utilization percentages
- Link failures, InvalidCrcs, InvalidTxWaitCounts, and Sync Losses
- IBM Fibre Connection (FICON) data fabrics

Real-time performance statistics allow administrators to configure custom polling interval settings for statistical data collection that can help you troubleshoot IBM c-type SAN fabric issues. The results can be displayed in the DM user interface.

DM is used for monitoring and configuring ports on the IBM c-type Family switches. When gathering DM statistics, you can configure selective polling intervals to monitor the performance of your SAN environment and troubleshoot any potential problems that exceed specified thresholds.

A polling interval can be set at 1 hour and 30 minutes or as low as 10 seconds. The results that you can view are as follows:

- Absolute value or Value per second
- Minimum or maximum value per second

There are two types of performance views that DM provides:

- The Device view tab, which you can use to configure the monitor option per port
- The Summary tab

To configure these settings, you must first log in to DM, as shown in Figure 8-1 on page 399.

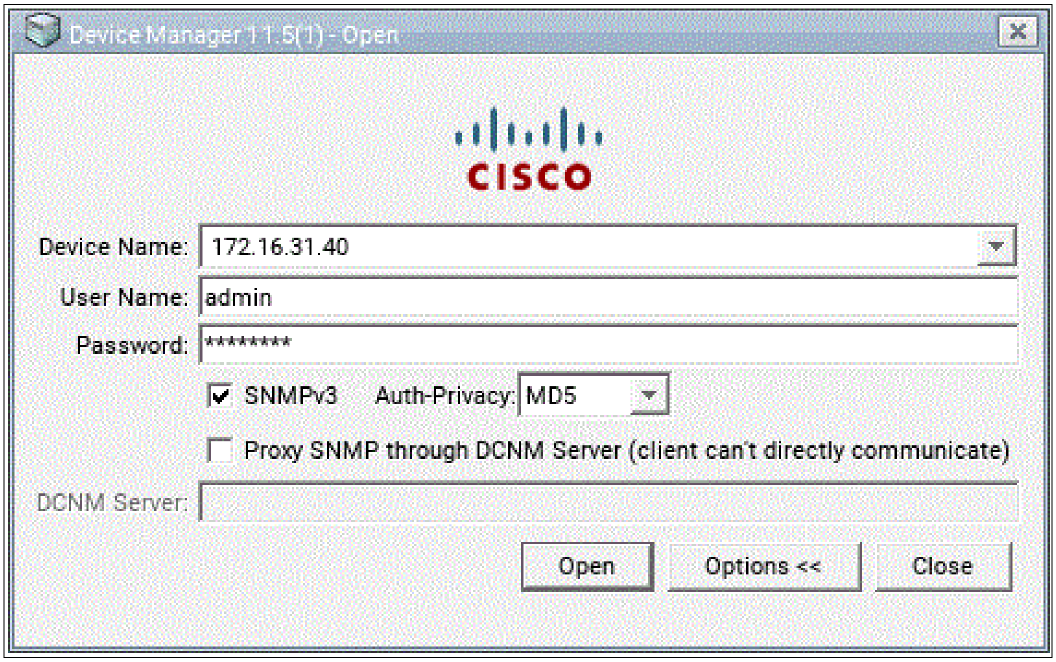

Device Name: 172.16.31.40

User Name: admin

Password: ********

☑ SNMPv3 Auth-Privacy: MD5

□ Proxy SNMP through DCNM Server (client can't directly communicate)

DCNM Server:

Figure 8-1 Device Manager login

The per port monitoring option provides many statistics. We select the **Device** tab view, right-click fc1/1, and select **MONITOR** to view the real-time monitor dialog box, as shown in Figure 8-2.

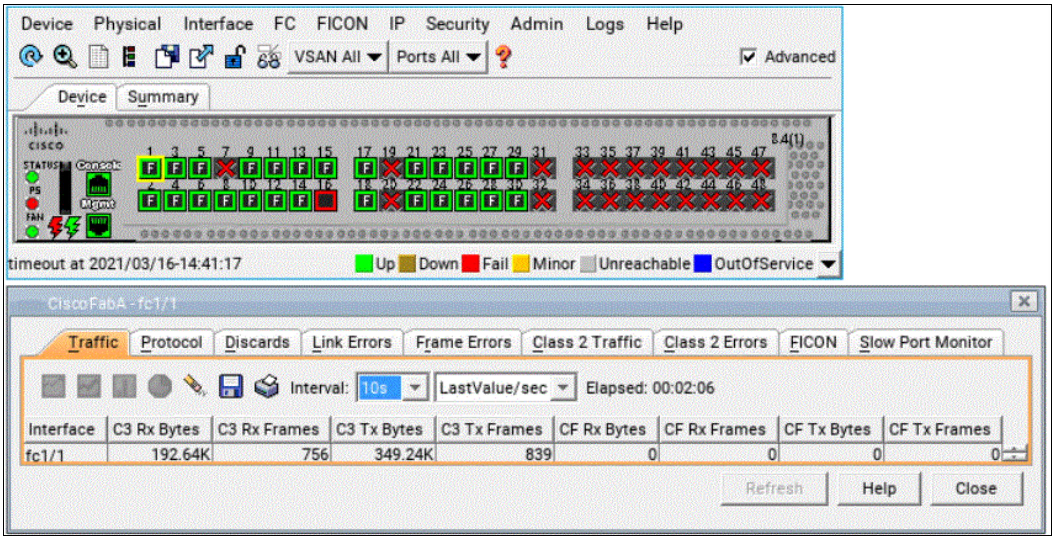

timeout at 2021/03/16-14:41:17

Interval: 10s LastValue/sec Elapsed: 00:02:06

Interface	C3 Rx Bytes	C3 Rx Frames	C3 Tx Bytes	C3 Tx Frames	CF Rx Bytes	CF Rx Frames	CF Tx Bytes	CF Tx Frames
fc1/1	192.64K	756	349.24K	839	0	0	0	0

Figure 8-2 Traffic Monitor view

The Summary view shows the active connected devices, port speed, and an option to configure parameters, as shown in Figure 8-3.

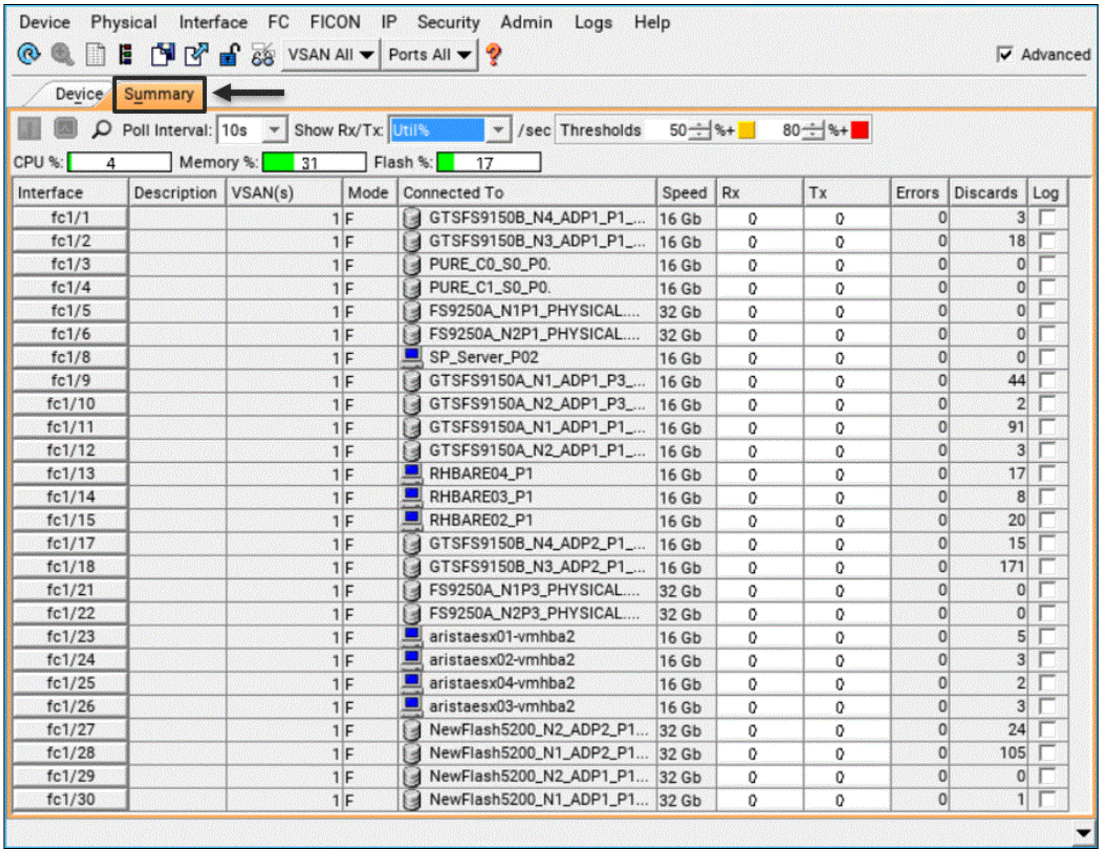

☑ Advanced

Poll Interval: 10s　Show Rx/Tx: Util% /sec　Thresholds 50 %+ 🟨 80 %+ 🟥

CPU %: 4　Memory %: 31　Flash %: 17

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
fc1/1		1	F	GTSFS9150B_N4_ADP1_P1_...	16 Gb	0	0	0	3	□
fc1/2		1	F	GTSFS9150B_N3_ADP1_P1_...	16 Gb	0	0	0	18	□
fc1/3		1	F	PURE_C0_S0_P0.	16 Gb	0	0	0	0	□
fc1/4		1	F	PURE_C1_S0_P0.	16 Gb	0	0	0	0	□
fc1/5		1	F	FS9250A_N1P1_PHYSICAL....	32 Gb	0	0	0	0	□
fc1/6		1	F	FS9250A_N2P1_PHYSICAL....	32 Gb	0	0	0	0	□
fc1/8		1	F	SP_Server_P02	16 Gb	0	0	0	0	□
fc1/9		1	F	GTSFS9150A_N1_ADP1_P3_...	16 Gb	0	0	0	44	□
fc1/10		1	F	GTSFS9150A_N2_ADP1_P3_...	16 Gb	0	0	0	2	□
fc1/11		1	F	GTSFS9150A_N1_ADP1_P1_...	16 Gb	0	0	0	91	□
fc1/12		1	F	GTSFS9150A_N2_ADP1_P1_...	16 Gb	0	0	0	3	□
fc1/13		1	F	RHBARE04_P1	16 Gb	0	0	0	17	□
fc1/14		1	F	RHBARE03_P1	16 Gb	0	0	0	8	□
fc1/15		1	F	RHBARE02_P1	16 Gb	0	0	0	20	□
fc1/17		1	F	GTSFS9150B_N4_ADP2_P1_...	16 Gb	0	0	0	15	□
fc1/18		1	F	GTSFS9150B_N3_ADP2_P1_...	16 Gb	0	0	0	171	□
fc1/21		1	F	FS9250A_N1P3_PHYSICAL....	32 Gb	0	0	0	0	□
fc1/22		1	F	FS9250A_N2P3_PHYSICAL....	32 Gb	0	0	0	0	□
fc1/23		1	F	aristaesx01-vmhba2	16 Gb	0	0	0	5	□
fc1/24		1	F	aristaesx02-vmhba2	16 Gb	0	0	0	3	□
fc1/25		1	F	aristaesx04-vmhba2	16 Gb	0	0	0	2	□
fc1/26		1	F	aristaesx03-vmhba2	16 Gb	0	0	0	3	□
fc1/27		1	F	NewFlash5200_N2_ADP2_P1...	32 Gb	0	0	0	24	□
fc1/28		1	F	NewFlash5200_N1_ADP2_P1...	32 Gb	0	0	0	105	□
fc1/29		1	F	NewFlash5200_N2_ADP1_P1...	32 Gb	0	0	0	0	□
fc1/30		1	F	NewFlash5200_N1_ADP1_P1...	32 Gb	0	0	0	1	□

Figure 8-3 Summary view tab

When deciding how you want your data to be interpreted, be sure to set the required polling intervals, Rx/Tx, and Thresholds settings, as shown in Figure 8-4.

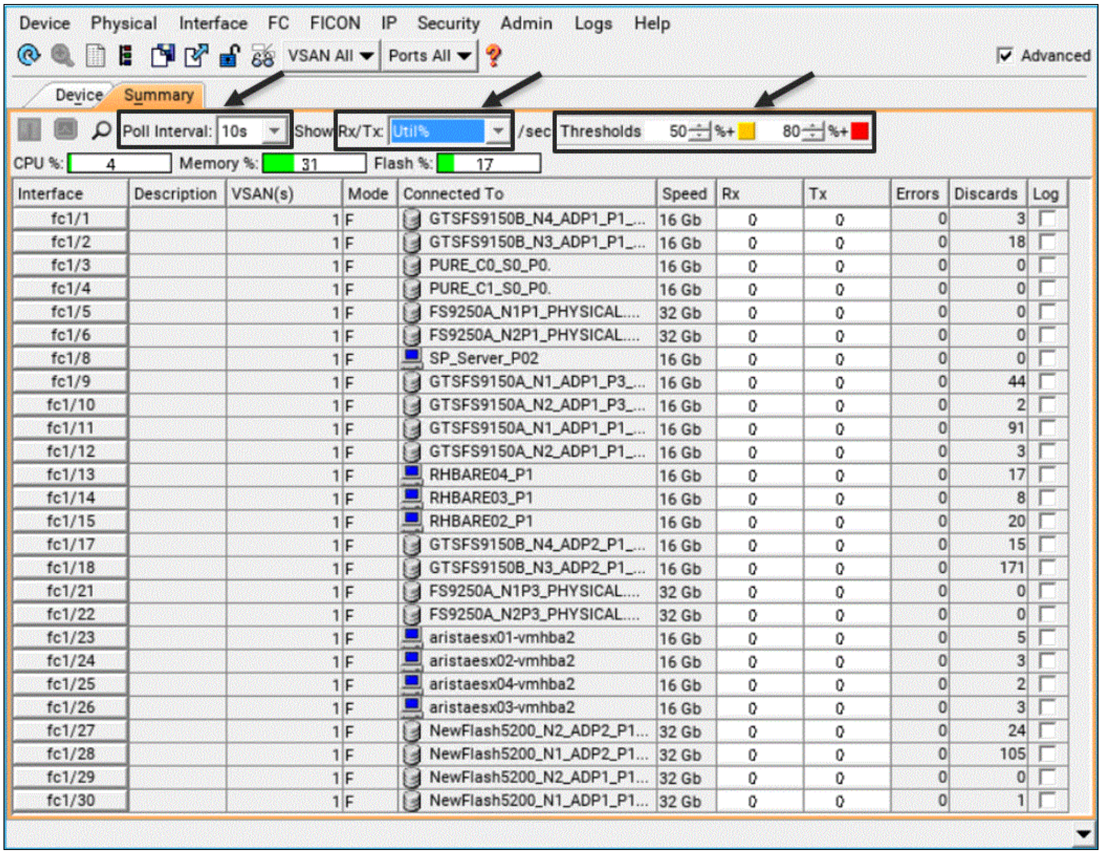
Device Physical Interface FC FICON IP Security Admin Logs Help
VSAN All
Ports All
☑ Advanced
Device
Summary

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 🟨 80 %+ 🟥

CPU %: 4 Memory %: 31 Flash %: 17

Interface	Description	VSAN(s)	Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
fc1/1		1	F	GTSFS9150B_N4_ADP1_P1_...	16 Gb	0	0	0	3	☐
fc1/2		1	F	GTSFS9150B_N3_ADP1_P1_...	16 Gb	0	0	0	18	☐
fc1/3		1	F	PURE_C0_S0_P0.	16 Gb	0	0	0	0	☐
fc1/4		1	F	PURE_C1_S0_P0.	16 Gb	0	0	0	0	☐
fc1/5		1	F	FS9250A_N1P1_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/6		1	F	FS9250A_N2P1_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/8		1	F	SP_Server_P02	16 Gb	0	0	0	0	☐
fc1/9		1	F	GTSFS9150A_N1_ADP1_P3_...	16 Gb	0	0	0	44	☐
fc1/10		1	F	GTSFS9150A_N2_ADP1_P3_...	16 Gb	0	0	0	2	☐
fc1/11		1	F	GTSFS9150A_N1_ADP1_P1_...	16 Gb	0	0	0	91	☐
fc1/12		1	F	GTSFS9150A_N2_ADP1_P1_...	16 Gb	0	0	0	3	☐
fc1/13		1	F	RHBARE04_P1	16 Gb	0	0	0	17	☐
fc1/14		1	F	RHBARE03_P1	16 Gb	0	0	0	8	☐
fc1/15		1	F	RHBARE02_P1	16 Gb	0	0	0	20	☐
fc1/17		1	F	GTSFS9150B_N4_ADP2_P1_...	16 Gb	0	0	0	15	☐
fc1/18		1	F	GTSFS9150B_N3_ADP2_P1_...	16 Gb	0	0	0	171	☐
fc1/21		1	F	FS9250A_N1P3_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/22		1	F	FS9250A_N2P3_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/23		1	F	aristaesx01-vmhba2	16 Gb	0	0	0	5	☐
fc1/24		1	F	aristaesx02-vmhba2	16 Gb	0	0	0	3	☐
fc1/25		1	F	aristaesx04-vmhba2	16 Gb	0	0	0	2	☐
fc1/26		1	F	aristaesx03-vmhba2	16 Gb	0	0	0	3	☐
fc1/27		1	F	NewFlash5200_N2_ADP2_P1...	32 Gb	0	0	0	24	☐
fc1/28		1	F	NewFlash5200_N1_ADP2_P1...	32 Gb	0	0	0	105	☐
fc1/29		1	F	NewFlash5200_N2_ADP1_P1...	32 Gb	0	0	0	0	☐
fc1/30		1	F	NewFlash5200_N1_ADP1_P1...	32 Gb	0	0	0	1	☐

Figure 8-4 Summary view configured settings

You can select the **Poll Interval** options from the drop-down list, as shown Figure 8-5.

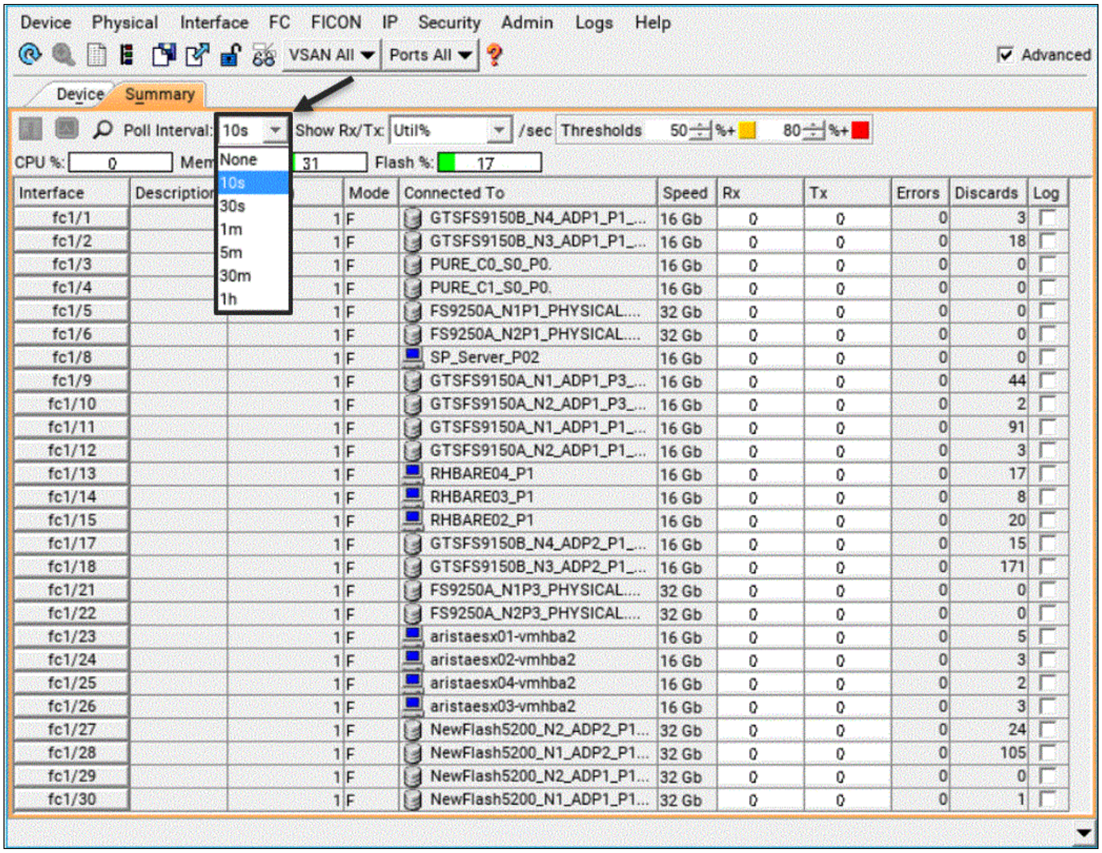

Poll Interval: 10s Show Rx/Tx: Util% /sec Thresholds 50 %+ 80 %+

CPU %: 0 Mem 31 Flash %: 17

Interface	Description		Mode	Connected To	Speed	Rx	Tx	Errors	Discards	Log
fc1/1		1	F	GTSFS9150B_N4_ADP1_P1_...	16 Gb	0	0	0	3	☐
fc1/2		1	F	GTSFS9150B_N3_ADP1_P1_...	16 Gb	0	0	0	18	☐
fc1/3		1	F	PURE_C0_S0_P0.	16 Gb	0	0	0	0	☐
fc1/4		1	F	PURE_C1_S0_P0.	16 Gb	0	0	0	0	☐
fc1/5		1	F	FS9250A_N1P1_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/6		1	F	FS9250A_N2P1_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/8		1	F	SP_Server_P02	16 Gb	0	0	0	0	☐
fc1/9		1	F	GTSFS9150A_N1_ADP1_P3_...	16 Gb	0	0	0	44	☐
fc1/10		1	F	GTSFS9150A_N2_ADP1_P3_...	16 Gb	0	0	0	2	☐
fc1/11		1	F	GTSFS9150A_N1_ADP1_P1_...	16 Gb	0	0	0	91	☐
fc1/12		1	F	GTSFS9150A_N2_ADP1_P1_...	16 Gb	0	0	0	3	☐
fc1/13		1	F	RHBARE04_P1	16 Gb	0	0	0	17	☐
fc1/14		1	F	RHBARE03_P1	16 Gb	0	0	0	8	☐
fc1/15		1	F	RHBARE02_P1	16 Gb	0	0	0	20	☐
fc1/17		1	F	GTSFS9150B_N4_ADP2_P1_...	16 Gb	0	0	0	15	☐
fc1/18		1	F	GTSFS9150B_N3_ADP2_P1_...	16 Gb	0	0	0	171	☐
fc1/21		1	F	FS9250A_N1P3_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/22		1	F	FS9250A_N2P3_PHYSICAL....	32 Gb	0	0	0	0	☐
fc1/23		1	F	aristaesx01-vmhba2	16 Gb	0	0	0	5	☐
fc1/24		1	F	aristaesx02-vmhba2	16 Gb	0	0	0	3	☐
fc1/25		1	F	aristaesx04-vmhba2	16 Gb	0	0	0	2	☐
fc1/26		1	F	aristaesx03-vmhba2	16 Gb	0	0	0	3	☐
fc1/27		1	F	NewFlash5200_N2_ADP2_P1...	32 Gb	0	0	0	24	☐
fc1/28		1	F	NewFlash5200_N1_ADP2_P1...	32 Gb	0	0	0	105	☐
fc1/29		1	F	NewFlash5200_N2_ADP1_P1...	32 Gb	0	0	0	0	☐
fc1/30		1	F	NewFlash5200_N1_ADP1_P1...	32 Gb	0	0	0	1	☐

Figure 8-5 Poll Interval

When using DM to set the error thresholds, select **Threshold Manager**, as shown in Figure 8-6 on page 403.

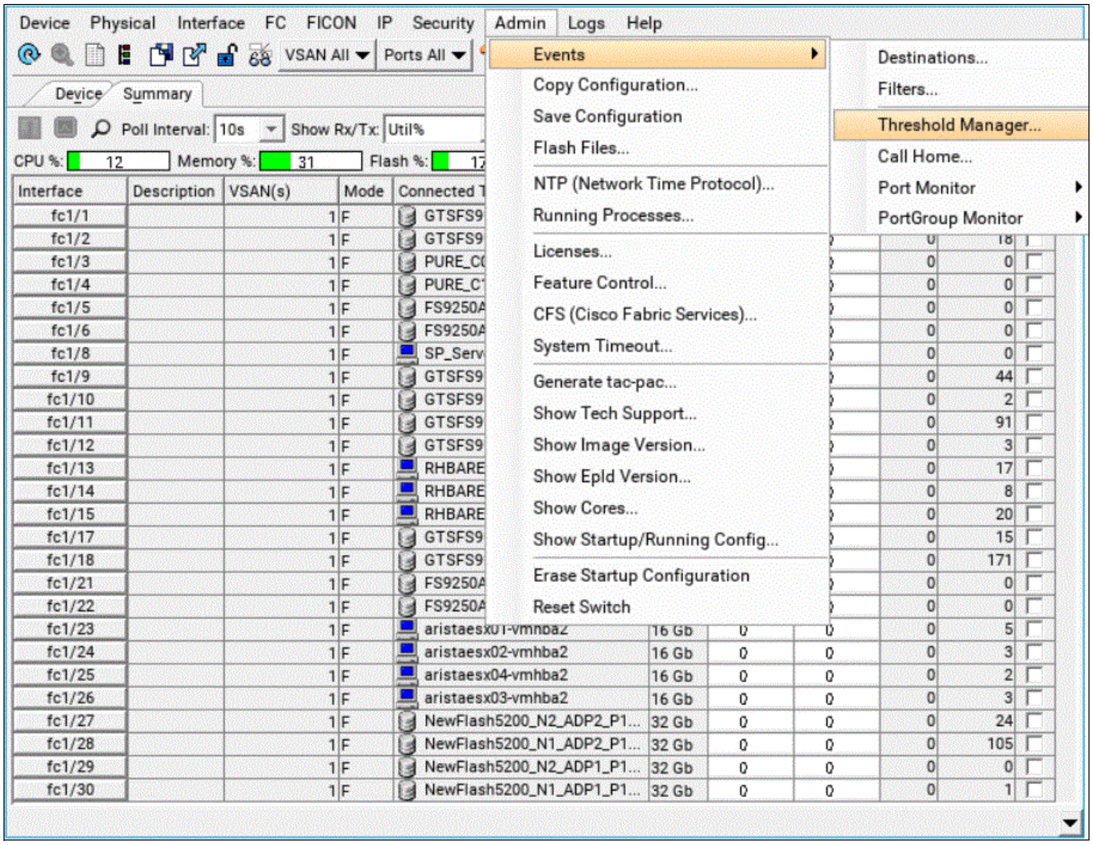

Interface	Description	VSAN(s)	Mode	Connected T						
fc1/1		1	F	GTSFS9						
fc1/2		1	F	GTSFS9				0	18	☐
fc1/3		1	F	PURE_CC				0	0	☐
fc1/4		1	F	PURE_CT				0	0	☐
fc1/5		1	F	FS9250A				0	0	☐
fc1/6		1	F	FS9250A				0	0	☐
fc1/8		1	F	SP_Serv				0	0	☐
fc1/9		1	F	GTSFS9				0	44	☐
fc1/10		1	F	GTSFS9				0	2	☐
fc1/11		1	F	GTSFS9				0	91	☐
fc1/12		1	F	GTSFS9				0	3	☐
fc1/13		1	F	RHBARE				0	17	☐
fc1/14		1	F	RHBARE				0	8	☐
fc1/15		1	F	RHBARE				0	20	☐
fc1/17		1	F	GTSFS9				0	15	☐
fc1/18		1	F	GTSFS9				0	171	☐
fc1/21		1	F	FS9250A				0	0	☐
fc1/22		1	F	FS9250A				0	0	☐
fc1/23		1	F	aristaesx01-vmhba2	16 Gb	0	0	0	5	☐
fc1/24		1	F	aristaesx02-vmhba2	16 Gb	0	0	0	3	☐
fc1/25		1	F	aristaesx04-vmhba2	16 Gb	0	0	0	2	☐
fc1/26		1	F	aristaesx03-vmhba2	16 Gb	0	0	0	3	☐
fc1/27		1	F	NewFlash5200_N2_ADP2_P1...	32 Gb	0	0	0	24	☐
fc1/28		1	F	NewFlash5200_N1_ADP2_P1...	32 Gb	0	0	0	105	☐
fc1/29		1	F	NewFlash5200_N2_ADP1_P1...	32 Gb	0	0	0	0	☐
fc1/30		1	F	NewFlash5200_N1_ADP1_P1...	32 Gb	0	0	0	1	☐

Figure 8-6 Selecting Threshold Manager

The Threshold Monitor can trigger an SNMP alert and log messages when a selected statistic reaches its configured threshold value.

> **Best practice:** Configure the DM thresholds on your IBM c-type Family switches so that you can monitor the performance of your SAN environment and troubleshoot any potential problems that exceed the specified thresholds.

The following values are considered industry best practices:

- Link Failures: Value =1 and Sample = 60
- Sync Losses: Value =1 and Sample = 60
- InvalidTxWords: Value =1 and Sample = 60
- InvalidCrcs: Value =1 and Sample = 60

There are more variables and thresholds that you can select and apply to a single port, multiple ports, or all ports, as shown in Figure 8-7.

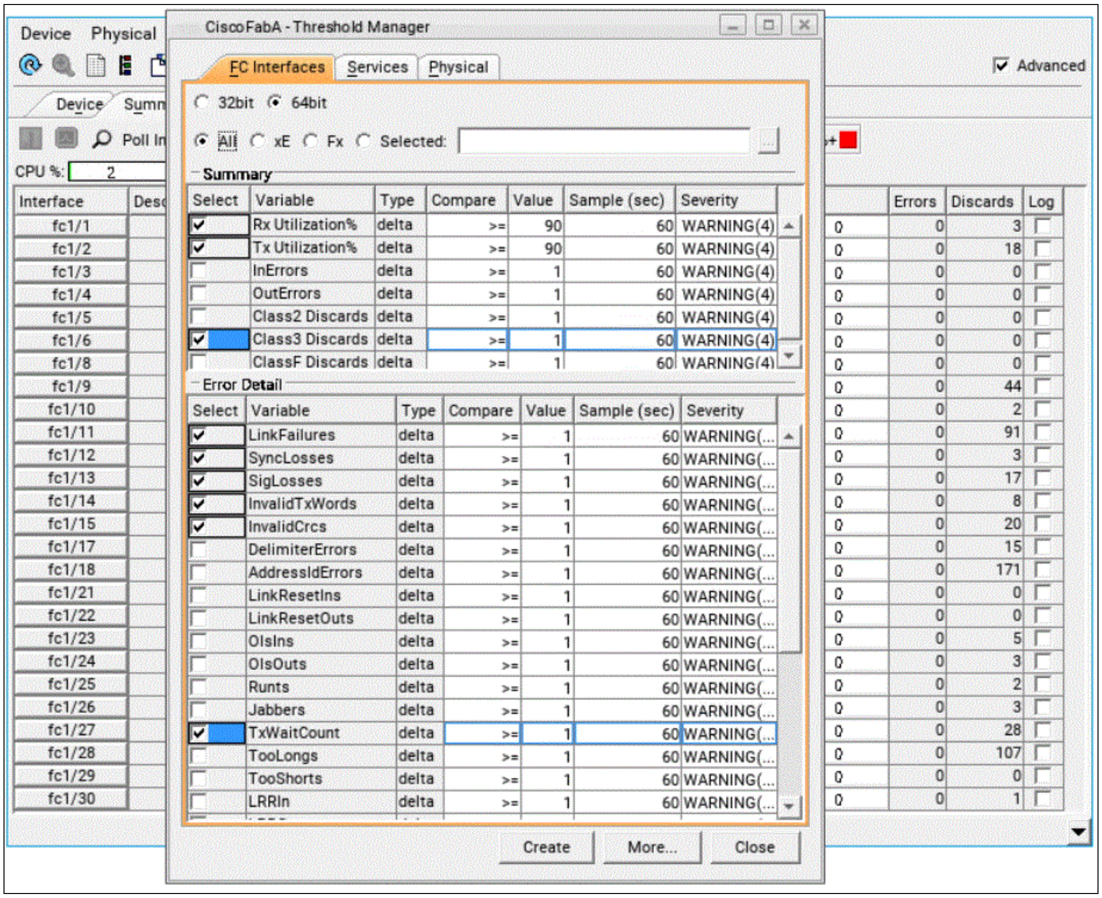

CiscoFabA - Threshold Manager

○ 32bit ● 64bit

● All ○ xE ○ Fx ○ Selected:

Summary

Select	Variable	Type	Compare	Value	Sample (sec)	Severity
☑	Rx Utilization%	delta	>=	90	60	WARNING(4)
☑	Tx Utilization%	delta	>=	90	60	WARNING(4)
□	InErrors	delta	>=	1	60	WARNING(4)
□	OutErrors	delta	>=	1	60	WARNING(4)
□	Class2 Discards	delta	>=	1	60	WARNING(4)
☑	Class3 Discards	delta	>=	1	60	WARNING(4)
□	ClassF Discards	delta	>=	1	60	WARNING(4)

Error Detail

Select	Variable	Type	Compare	Value	Sample (sec)	Severity
☑	LinkFailures	delta	>=	1	60	WARNING(...
☑	SyncLosses	delta	>=	1	60	WARNING(...
☑	SigLosses	delta	>=	1	60	WARNING(...
☑	InvalidTxWords	delta	>=	1	60	WARNING(...
☑	InvalidCrcs	delta	>=	1	60	WARNING(...
□	DelimiterErrors	delta	>=	1	60	WARNING(...
□	AddressIdErrors	delta	>=	1	60	WARNING(...
□	LinkResetIns	delta	>=	1	60	WARNING(...
□	LinkResetOuts	delta	>=	1	60	WARNING(...
□	OlsIns	delta	>=	1	60	WARNING(...
□	OlsOuts	delta	>=	1	60	WARNING(...
□	Runts	delta	>=	1	60	WARNING(...
□	Jabbers	delta	>=	1	60	WARNING(...
☑	TxWaitCount	delta	>=	1	60	WARNING(...
□	TooLongs	delta	>=	1	60	WARNING(...
□	TooShorts	delta	>=	1	60	WARNING(...
□	LRRIn	delta	>=	1	60	WARNING(...

Interface	Desc		Errors	Discards	Log
fc1/1		0	0	3	□
fc1/2		0	0	18	□
fc1/3		0	0	0	□
fc1/4		0	0	0	□
fc1/5		0	0	0	□
fc1/6		0	0	0	□
fc1/8		0	0	0	□
fc1/9		0	0	44	□
fc1/10		0	0	2	□
fc1/11		0	0	91	□
fc1/12		0	0	3	□
fc1/13		0	0	17	□
fc1/14		0	0	8	□
fc1/15		0	0	20	□
fc1/17		0	0	15	□
fc1/18		0	0	171	□
fc1/21		0	0	0	□
fc1/22		0	0	0	□
fc1/23		0	0	5	□
fc1/24		0	0	3	□
fc1/25		0	0	2	□
fc1/26		0	0	3	□
fc1/27		0	0	28	□
fc1/28		0	0	107	□
fc1/29		0	0	0	□
fc1/30		0	0	1	□

Figure 8-7 Threshold Manager

8.1.1 Data Center Network Manager

The IBM c-type Family provides several advanced licensed features that you can use for analytics and telemetry streaming to help you sustain resiliency in your environment:

- DCNM Advanced
- SAN analytics and telemetry data streaming

DCNM is a management tool that is used for provisioning, monitoring, and troubleshooting IBM c-type Family SAN environments. It provides a command and control style structured regime that gives you complete visibility into your entire IBM c-type Family fabric infrastructures. DCNM provides a centralized, high-level, and web-based view that includes a complete feature set that meets administrative requirements in data centers by streamlining IBM c-type management, provisioning, monitoring, and troubleshooting SAN devices.

Best practice: IBM c-type DCNM Advanced includes SAN Insights, which is the recommended web UI when using the SAN Analytics feature.

Figure 8-8 on page 405 shows the DCNM Advanced login window.

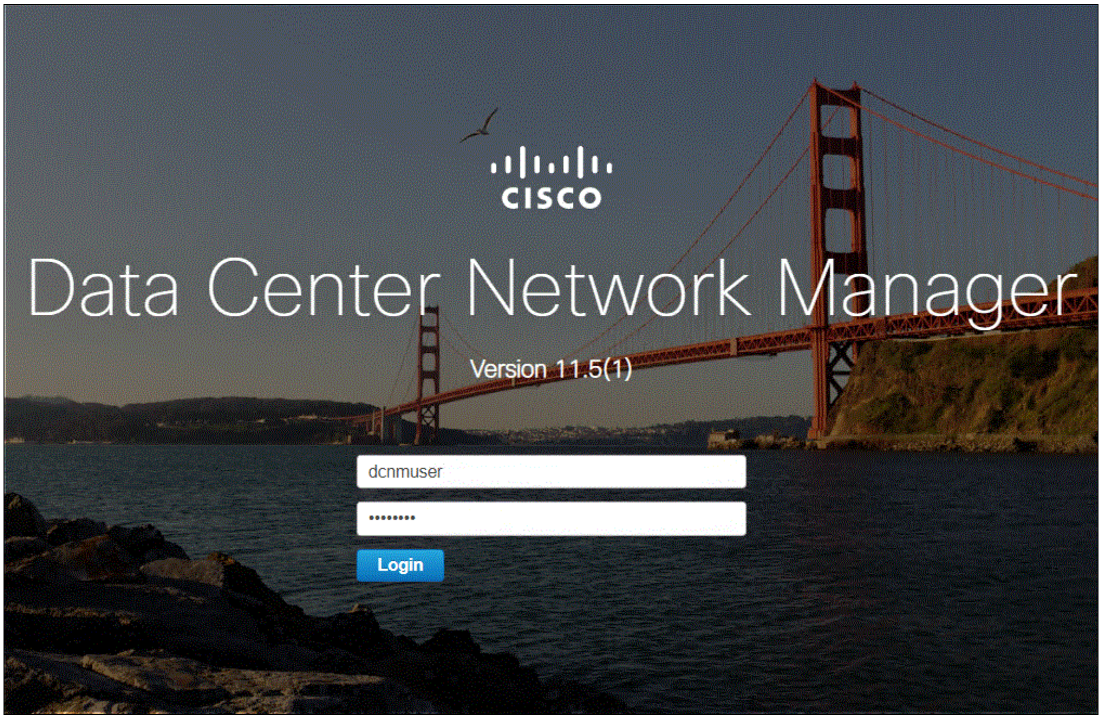

Figure 8-8 DCNM login window

After you log in to DCNM, the dashboard summary opens, which provides storage administrators with a 24-hour snapshot of their SAN fabric and the ability to focus on key health and performance metrics on your IBM c-type SAN fabric.

There are many default dashlets that can be customized to provide a visual look into your SAN environment. These dashlets range from an inventory of switches and modules to ones like Top CPU, Top ISLs, Link traffic, and Alerts, as shown in Figure 8-9.

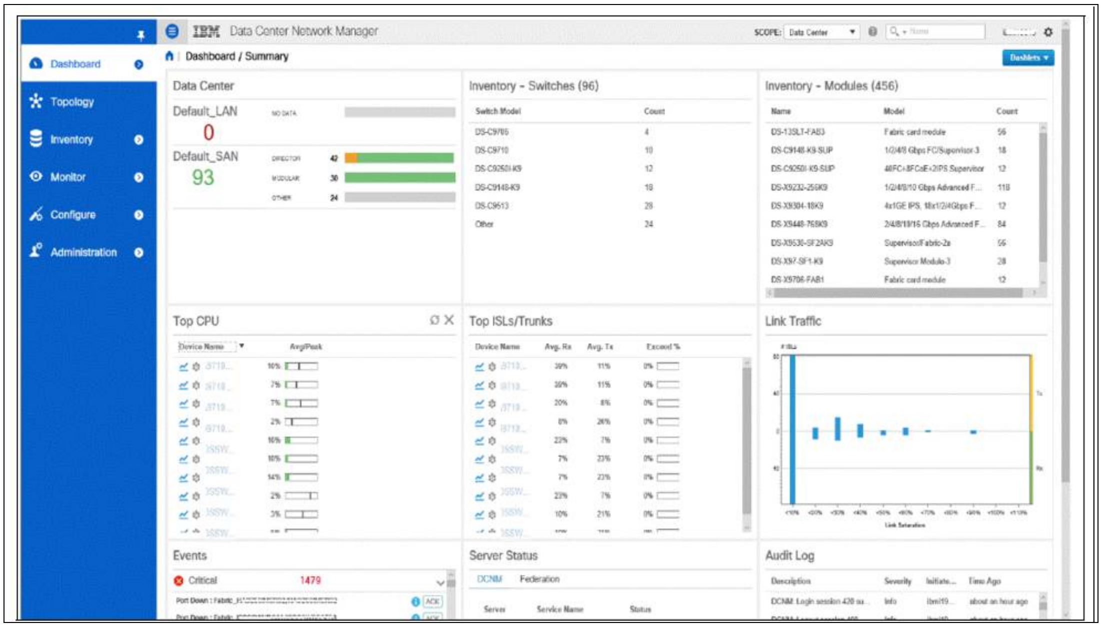

Inventory - Switches (96)

Switch Model	Count
DS-C9706	4
DS-C9710	10
DS-C9250I-K9	12
DS-C9148-K9	18
DS-C9513	28
Other	24

Inventory - Modules (456)

Name	Model	Count
DS-13SLT-FAB3	Fabric card module	56
DS-C9148-K9-SUP	1/2/4/8 Gbps FC/Supervisor-3	18
DS-C9250I-K9-SUP	40FC+8FCoE+2IPS Supervisor	12
DS-X9232-256K9	1/2/4/8/10 Gbps Advanced F...	118
DS-X9304-18K9	4x1GE IPS, 18x1/2/4Gbps F...	12
DS-X9448-768K9	2/4/8/10/16 Gbps Advanced F...	84
DS-X9530-SF2AK9	Supervisor/Fabric-2a	56
DS-X97-SF1-K9	Supervisor Module-3	28
DS-X9706-FAB1	Fabric card module	12

Figure 8-9 DCNM summary dashboard

There are various scopes that are available in the DCNM web interface. In this example, we focus on the Default_SAN scope and select **Topology** in the left pane, which shows our fabric topology view, as shown in Figure 8-10.

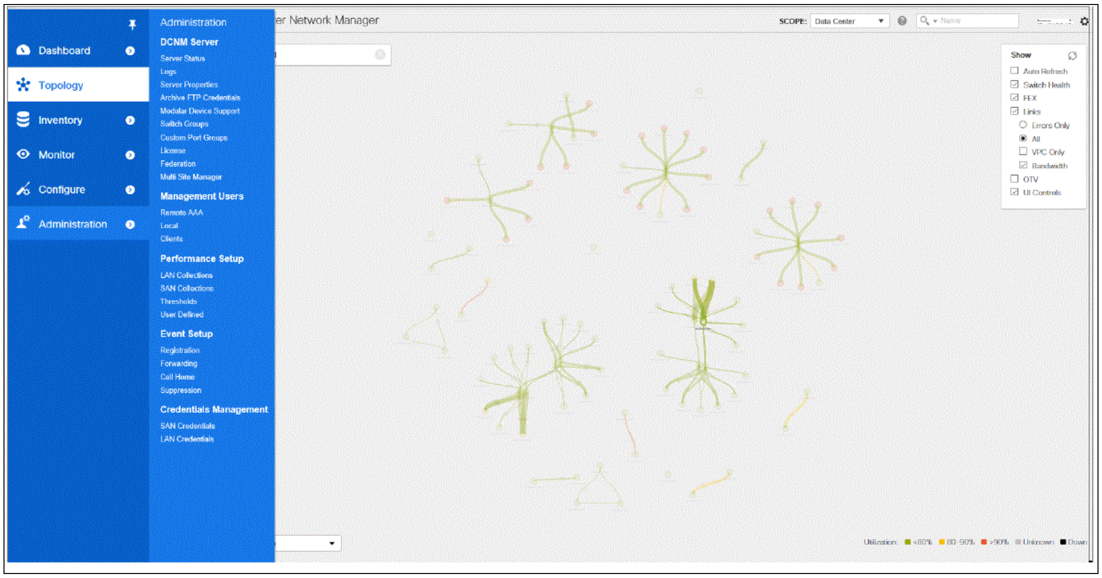

SCOPE: Data Center

Show
- ☐ Auto Refresh
- ☑ Switch Health
- ☑ FEX
- ☑ Links
  - ○ Errors Only
  - ● All
  - ☐ VPC Only
  - ☑ Bandwidth
- ☐ OTV
- ☑ UI Controls

Figure 8-10 Topology view

Suggested reading:

- DCNM SAN Management Configuration Guide 11.5(1)
- Cisco MDS SAN Analytics and Telemetry Streaming Configuration Guide

8.1.2 Port Monitor

Port Monitor (PMON) is used to monitor the status of individual ports or groups of ports. Various parameters can be configured and thresholds set. You can set rising or falling thresholds. When a threshold is breached, an alert or syslog message is generated. PMON can be configured by using DCNM, DM, or the command-line interface (CLI). In this section, we show a configuration that is set by using DCNM. One of the major advantages of using DCNM is that you can configure a consistent policy that can be applied across multiple switches.

To configure PMON by using DCNM, launch DCNM and select **Configure** → **SAN** → **Port Monitoring**, as shown in Figure 8-11 on page 407.

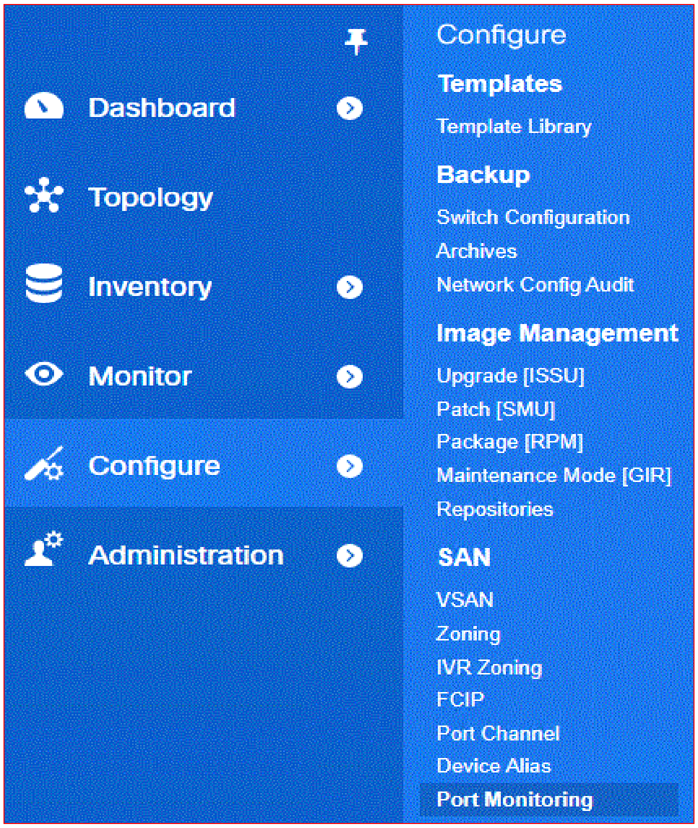

Figure 8-11 PMON menu option

The window that is shown in Figure 8-12 opens and shows one of the default policies, which is named Normal_accessPort.

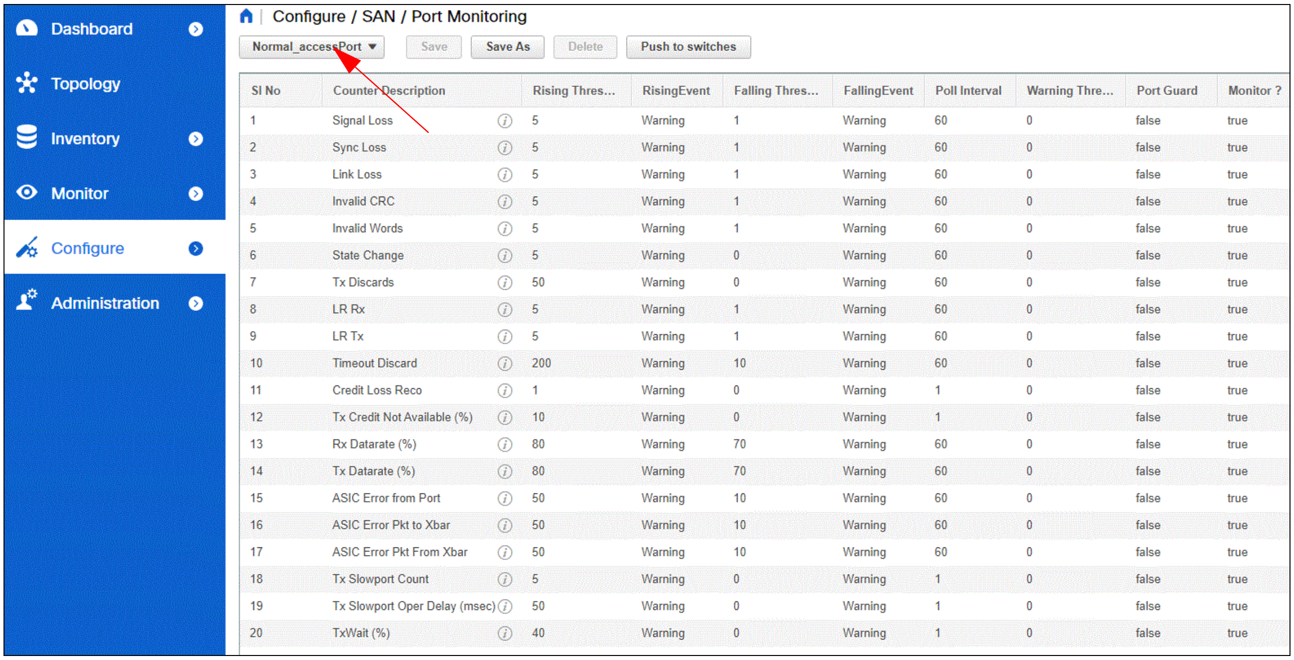

Sl No	Counter Description	Rising Thres...	RisingEvent	Falling Thres...	FallingEvent	Poll Interval	Warning Thre...	Port Guard	Monitor ?
1	Signal Loss	5	Warning	1	Warning	60	0	false	true
2	Sync Loss	5	Warning	1	Warning	60	0	false	true
3	Link Loss	5	Warning	1	Warning	60	0	false	true
4	Invalid CRC	5	Warning	1	Warning	60	0	false	true
5	Invalid Words	5	Warning	1	Warning	60	0	false	true
6	State Change	5	Warning	0	Warning	60	0	false	true
7	Tx Discards	50	Warning	0	Warning	60	0	false	true
8	LR Rx	5	Warning	1	Warning	60	0	false	true
9	LR Tx	5	Warning	1	Warning	60	0	false	true
10	Timeout Discard	200	Warning	10	Warning	60	0	false	true
11	Credit Loss Reco	1	Warning	0	Warning	1	0	false	true
12	Tx Credit Not Available (%)	10	Warning	0	Warning	1	0	false	true
13	Rx Datarate (%)	80	Warning	70	Warning	60	0	false	true
14	Tx Datarate (%)	80	Warning	70	Warning	60	0	false	true
15	ASIC Error from Port	50	Warning	10	Warning	60	0	false	true
16	ASIC Error Pkt to Xbar	50	Warning	10	Warning	60	0	false	true
17	ASIC Error Pkt From Xbar	50	Warning	10	Warning	60	0	false	true
18	Tx Slowport Count	5	Warning	0	Warning	1	0	false	true
19	Tx Slowport Oper Delay (msec)	50	Warning	0	Warning	1	0	false	true
20	TxWait (%)	40	Warning	0	Warning	1	0	false	true

Figure 8-12 Port Monitor policies

There are various counters that have rising for falling thresholds and an event type set. Event types can be seen by selecting one of the row entries and then selecting the event type, as shown in Figure 8-13.

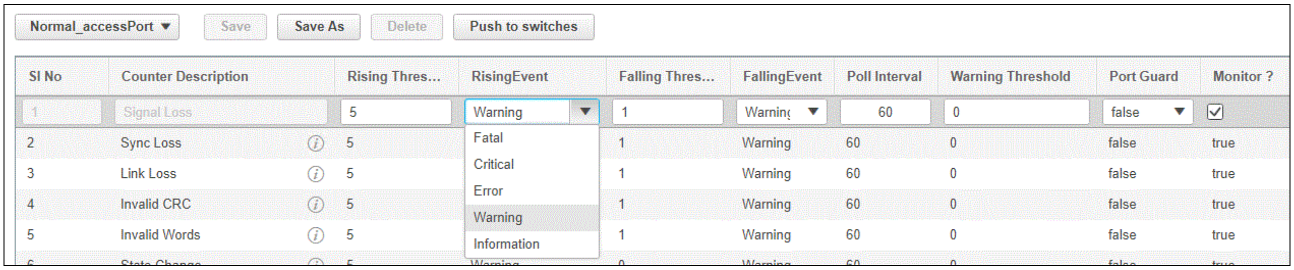

Sl No	Counter Description	Rising Thres...	RisingEvent	Falling Thres...	FallingEvent	Poll Interval	Warning Threshold	Port Guard	Monitor ?
1	Signal Loss	5	Warning ▼	1	Warning ▼	60	0	false ▼	☑
2	Sync Loss	5		1	Warning	60	0	false	true
3	Link Loss	5		1	Warning	60	0	false	true
4	Invalid CRC	5		1	Warning	60	0	false	true
5	Invalid Words	5		1	Warning	60	0	false	true
6	State Change	5	Warning	0	Warning	60	0	false	true

Figure 8-13 Port Monitor event types

We create a policy by selecting an existing policy from the drop-down menu that is shown in Figure 8-12 on page 407 that is closest to the policy that you require, which you then modify as required. Figure 8-14 shows that we selected the Most-Aggressive_allPort policy and made several changes.

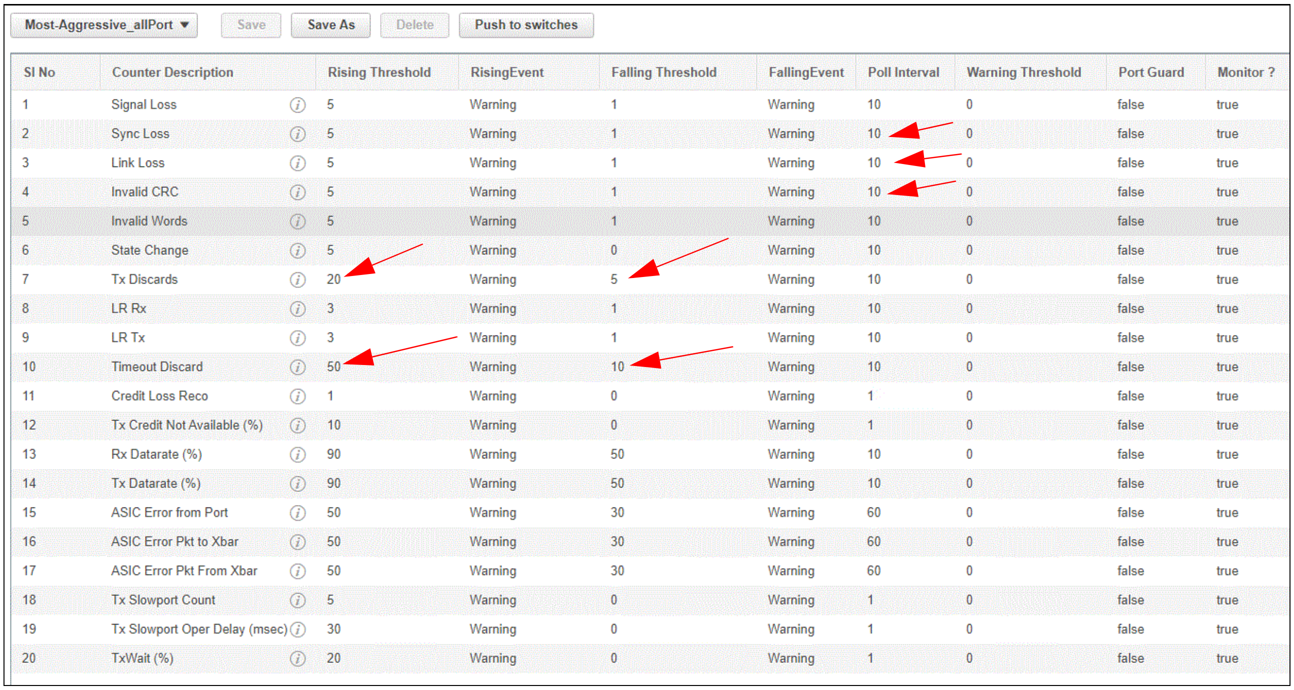

Sl No	Counter Description	Rising Threshold	RisingEvent	Falling Threshold	FallingEvent	Poll Interval	Warning Threshold	Port Guard	Monitor ?
1	Signal Loss	5	Warning	1	Warning	10	0	false	true
2	Sync Loss	5	Warning	1	Warning	10	0	false	true
3	Link Loss	5	Warning	1	Warning	10	0	false	true
4	Invalid CRC	5	Warning	1	Warning	10	0	false	true
5	Invalid Words	5	Warning	1	Warning	10	0	false	true
6	State Change	5	Warning	0	Warning	10	0	false	true
7	Tx Discards	20	Warning	5	Warning	10	0	false	true
8	LR Rx	3	Warning	1	Warning	10	0	false	true
9	LR Tx	3	Warning	1	Warning	10	0	false	true
10	Timeout Discard	50	Warning	10	Warning	10	0	false	true
11	Credit Loss Reco	1	Warning	0	Warning	1	0	false	true
12	Tx Credit Not Available (%)	10	Warning	0	Warning	1	0	false	true
13	Rx Datarate (%)	90	Warning	50	Warning	10	0	false	true
14	Tx Datarate (%)	90	Warning	50	Warning	10	0	false	true
15	ASIC Error from Port	50	Warning	30	Warning	60	0	false	true
16	ASIC Error Pkt to Xbar	50	Warning	30	Warning	60	0	false	true
17	ASIC Error Pkt From Xbar	50	Warning	30	Warning	60	0	false	true
18	Tx Slowport Count	5	Warning	0	Warning	1	0	false	true
19	Tx Slowport Oper Delay (msec)	30	Warning	0	Warning	1	0	false	true
20	TxWait (%)	20	Warning	0	Warning	1	0	false	true

Figure 8-14 Changing the Port Monitor policy

Now, save the policy under a new name, as shown in Figure 8-15 on page 409.

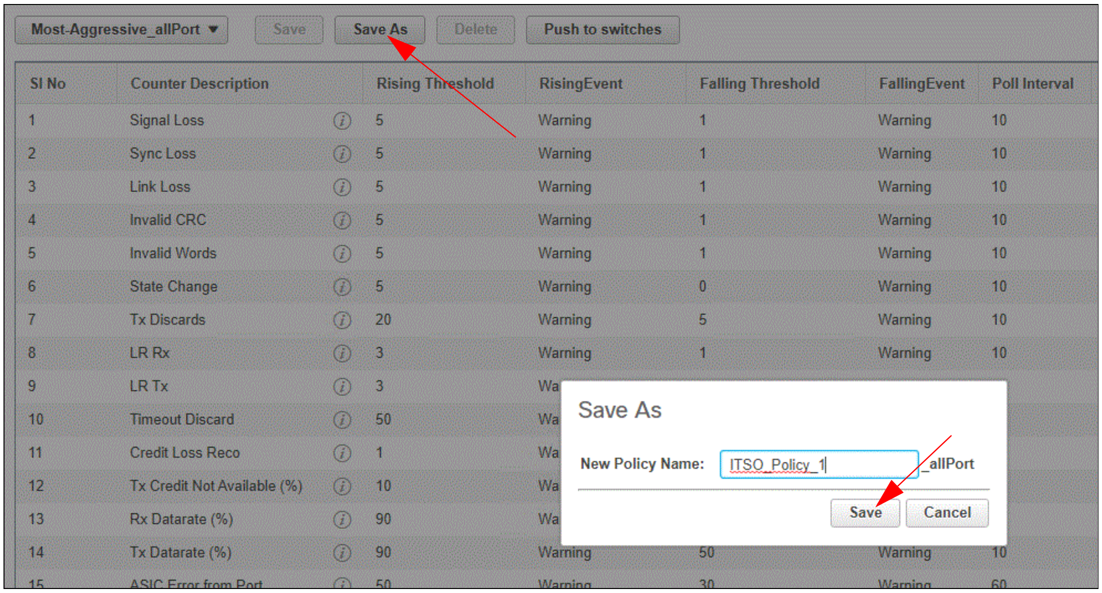

Sl No	Counter Description		Rising Threshold	RisingEvent	Falling Threshold	FallingEvent	Poll Interval
1	Signal Loss	ⓘ	5	Warning	1	Warning	10
2	Sync Loss	ⓘ	5	Warning	1	Warning	10
3	Link Loss	ⓘ	5	Warning	1	Warning	10
4	Invalid CRC	ⓘ	5	Warning	1	Warning	10
5	Invalid Words	ⓘ	5	Warning	1	Warning	10
6	State Change	ⓘ	5	Warning	0	Warning	10
7	Tx Discards	ⓘ	20	Warning	5	Warning	10
8	LR Rx	ⓘ	3	Warning	1	Warning	10
9	LR Tx	ⓘ	3	Wa			
10	Timeout Discard	ⓘ	50	Wa			
11	Credit Loss Reco	ⓘ	1	Wa			
12	Tx Credit Not Available (%)	ⓘ	10	Wa			
13	Rx Datarate (%)	ⓘ	90	Wa			
14	Tx Datarate (%)	ⓘ	90	Warning	50	Warning	10
15	ASIC Error from Port	ⓘ	50	Warning	30	Warning	60

Save As

New Policy Name: ITSO_Policy_1 _allPort

Figure 8-15 Saving a new policy

After the policy is saved, it is available under the CustomPolicy list of policies, as shown in Figure 8-16.

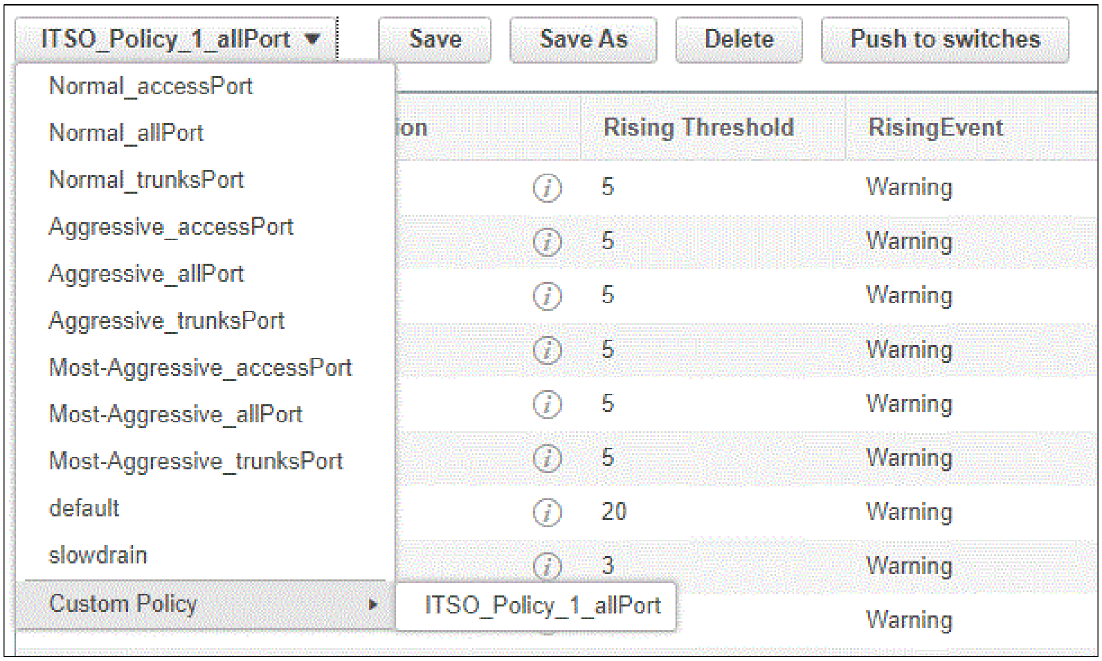

ion		Rising Threshold	RisingEvent
	ⓘ	5	Warning
	ⓘ	5	Warning
	ⓘ	5	Warning
	ⓘ	5	Warning
	ⓘ	5	Warning
	ⓘ	5	Warning
	ⓘ	20	Warning
	ⓘ	3	Warning
			Warning

Figure 8-16 Custom Policy list

To apply a policy to the switches, click **Push to switches**, as shown in Figure 8-17.

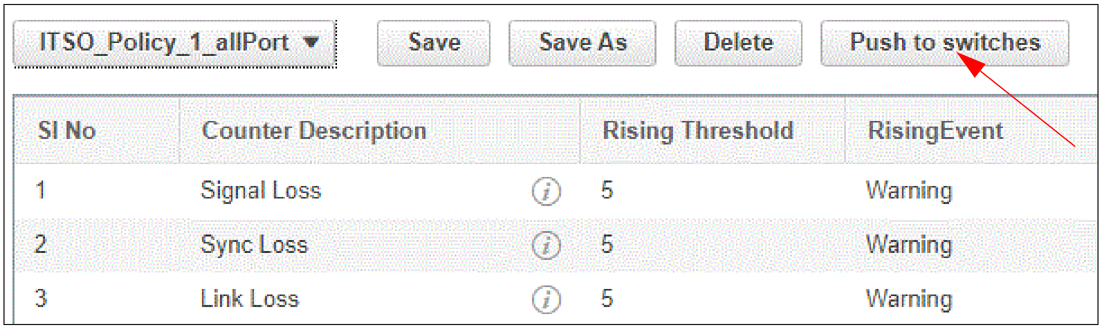

Sl No	Counter Description		Rising Threshold	RisingEvent
1	Signal Loss	ⓘ	5	Warning
2	Sync Loss	ⓘ	5	Warning
3	Link Loss	ⓘ	5	Warning

Figure 8-17 Push to switches menu

A window of your environment opens. Our environment is shown in Figure 8-18, where we selected both fabrics. When all switches are chosen, click **Push**.

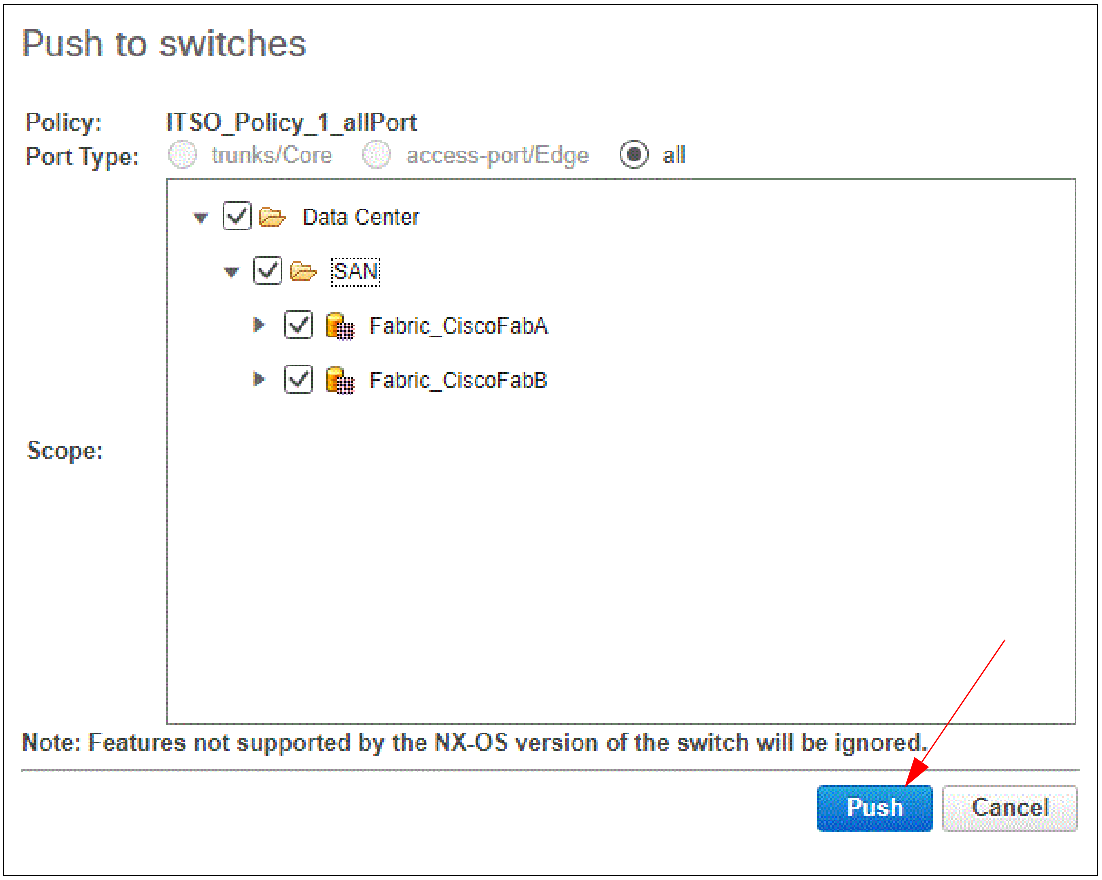
Push to switches

Policy: ITSO_Policy_1_allPort

Port Type: ○ trunks/Core ○ access-port/Edge ● all

Scope:

- ☑ Data Center
  - ☑ SAN
    - ☑ Fabric_CiscoFabA
    - ☑ Fabric_CiscoFabB

Note: Features not supported by the NX-OS version of the switch will be ignored.

Figure 8-18 Pushing a policy to selected switches

A warning dialog box opens if a policy exists on a switch and must be overwritten, as shown in Figure 8-19. In our case, we click **Yes** to overwrite.

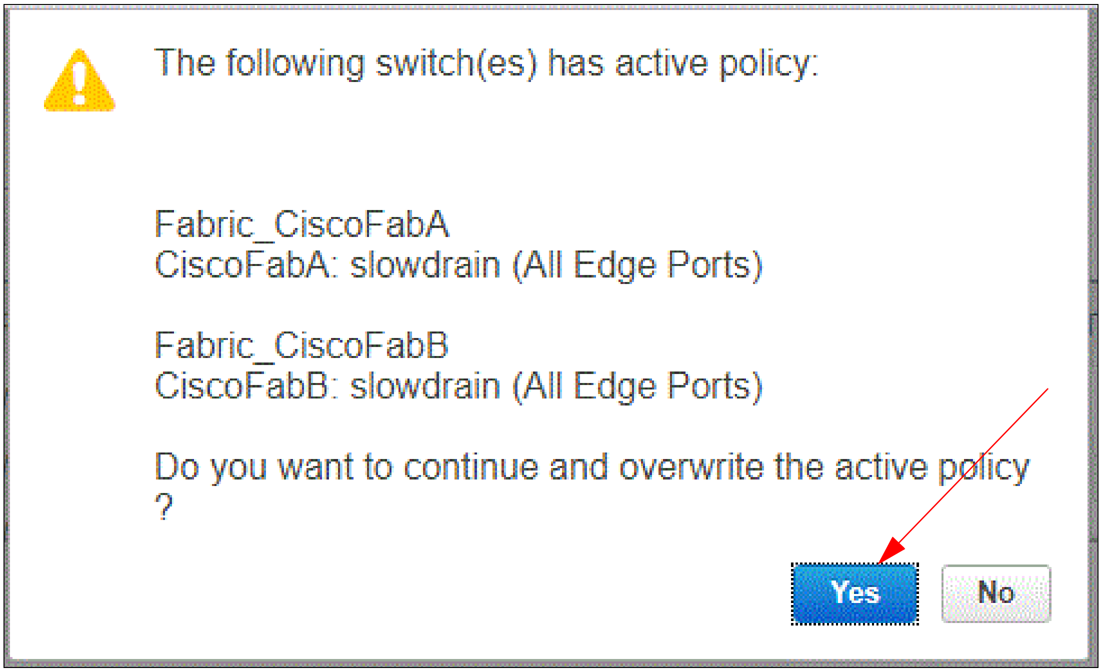
The following switch(es) has active policy:

Fabric_CiscoFabA
CiscoFabA: slowdrain (All Edge Ports)

Fabric_CiscoFabB
CiscoFabB: slowdrain (All Edge Ports)

Do you want to continue and overwrite the active policy
?

Figure 8-19 Overwriting a warning message

A results window opens. If everything worked, the window shows a status of Success, as shown in Figure 8-20.

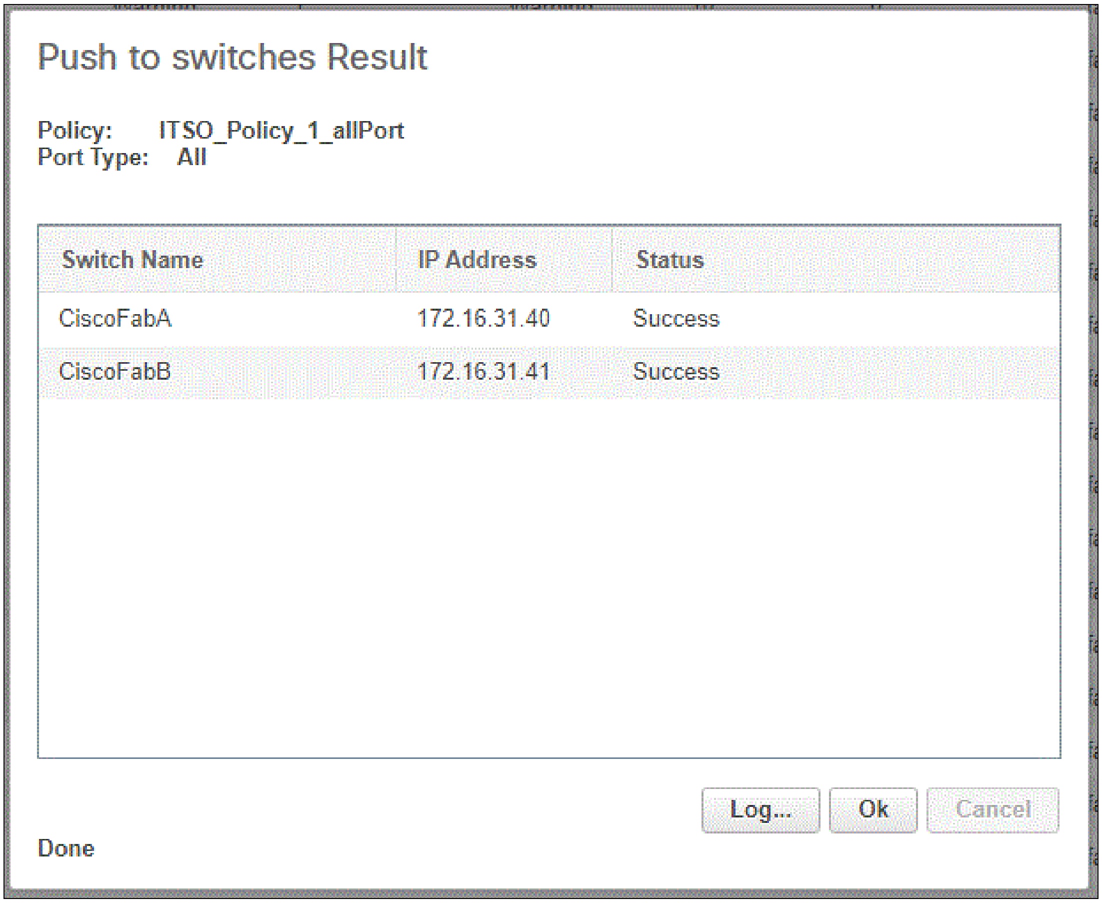

Push to switches Result

Policy: ITSO_Policy_1_allPort
Port Type: All

Switch Name	IP Address	Status
CiscoFabA	172.16.31.40	Success
CiscoFabB	172.16.31.41	Success

Done

Figure 8-20 Push to switches results

For more information about the results, click **Log**. If the push to the switches was not successful, the log provides more information to help with troubleshooting. Figure 8-21 shows an example of log details. Use the scroll bar on the right to view the full log.

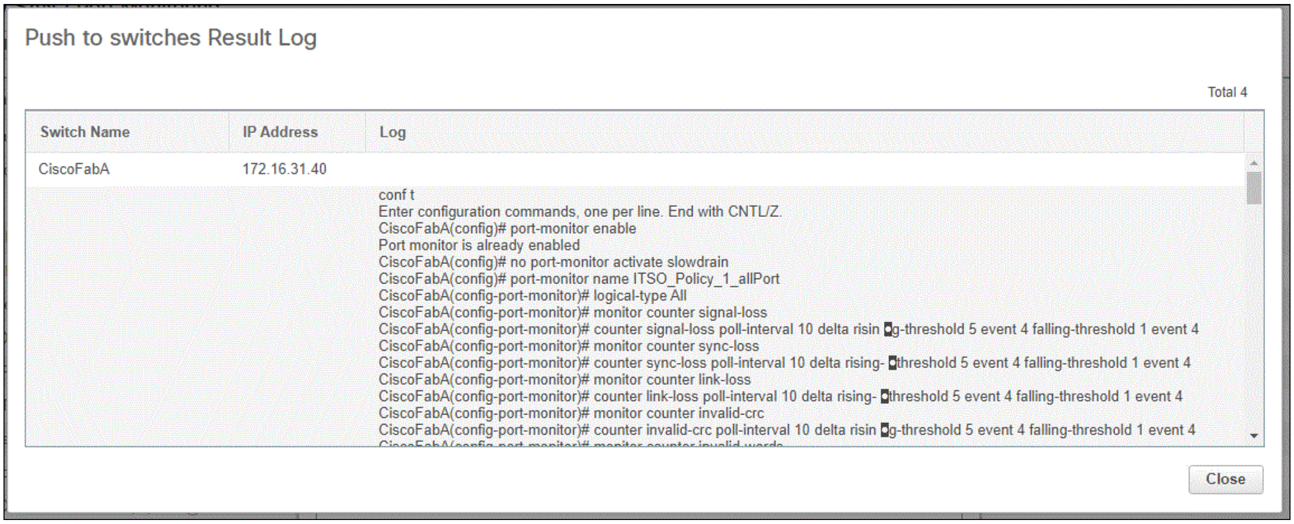

Push to switches Result Log

Total 4

Switch Name	IP Address	Log
CiscoFabA	172.16.31.40	

```
conf t
Enter configuration commands, one per line. End with CNTL/Z.
CiscoFabA(config)# port-monitor enable
Port monitor is already enabled
CiscoFabA(config)# no port-monitor activate slowdrain
CiscoFabA(config)# port-monitor name ITSO_Policy_1_allPort
CiscoFabA(config-port-monitor)# logical-type All
CiscoFabA(config-port-monitor)# monitor counter signal-loss
CiscoFabA(config-port-monitor)# counter signal-loss poll-interval 10 delta risin ▯g-threshold 5 event 4 falling-threshold 1 event 4
CiscoFabA(config-port-monitor)# monitor counter sync-loss
CiscoFabA(config-port-monitor)# counter sync-loss poll-interval 10 delta rising- ▯threshold 5 event 4 falling-threshold 1 event 4
CiscoFabA(config-port-monitor)# monitor counter link-loss
CiscoFabA(config-port-monitor)# counter link-loss poll-interval 10 delta rising- ▯threshold 5 event 4 falling-threshold 1 event 4
CiscoFabA(config-port-monitor)# monitor counter invalid-crc
CiscoFabA(config-port-monitor)# counter invalid-crc poll-interval 10 delta risin ▯g-threshold 5 event 4 falling-threshold 1 event 4
CiscoFabA(config-port-monitor)# monitor counter invalid-words
```

Figure 8-21 Log details

To view PMON events, select **Monitor** → **Switch** → **Events**, as shown in Figure 8-22.

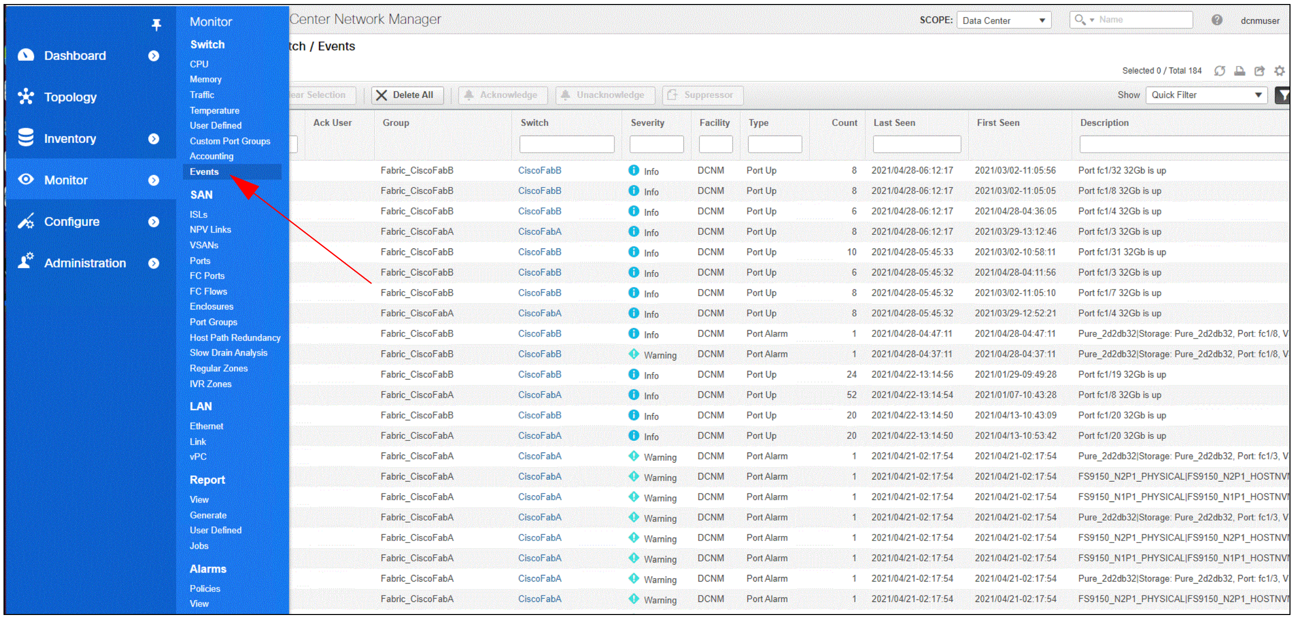

Ack User	Group	Switch	Severity	Facility	Type	Count	Last Seen	First Seen	Description
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	8	2021/04/28-06:12:17	2021/03/02-11:05:56	Port fc1/32 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	8	2021/04/28-06:12:17	2021/03/02-11:05:05	Port fc1/8 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	6	2021/04/28-06:12:17	2021/04/28-04:36:05	Port fc1/4 32Gb is up
	Fabric_CiscoFabA	CiscoFabA	Info	DCNM	Port Up	8	2021/04/28-06:12:17	2021/03/29-13:12:46	Port fc1/3 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	10	2021/04/28-05:45:33	2021/03/02-10:58:11	Port fc1/31 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	6	2021/04/28-05:45:32	2021/04/28-04:11:56	Port fc1/3 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	8	2021/04/28-05:45:32	2021/03/02-11:05:10	Port fc1/7 32Gb is up
	Fabric_CiscoFabA	CiscoFabA	Info	DCNM	Port Up	8	2021/04/28-05:45:32	2021/03/29-12:52:21	Port fc1/4 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Alarm	1	2021/04/28-04:47:11	2021/04/28-04:47:11	Pure_2d2db32|Storage: Pure_2d2db32, Port: fc1/8, V
	Fabric_CiscoFabB	CiscoFabB	Warning	DCNM	Port Alarm	1	2021/04/28-04:37:11	2021/04/28-04:37:11	Pure_2d2db32|Storage: Pure_2d2db32, Port: fc1/8, V
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	24	2021/04/22-13:14:56	2021/01/29-09:49:28	Port fc1/19 32Gb is up
	Fabric_CiscoFabA	CiscoFabA	Info	DCNM	Port Up	52	2021/04/22-13:14:54	2021/01/07-10:43:28	Port fc1/8 32Gb is up
	Fabric_CiscoFabB	CiscoFabB	Info	DCNM	Port Up	20	2021/04/22-13:14:50	2021/04/13-10:43:09	Port fc1/20 32Gb is up
	Fabric_CiscoFabA	CiscoFabA	Info	DCNM	Port Up	20	2021/04/22-13:14:50	2021/04/13-10:53:42	Port fc1/20 32Gb is up
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	Pure_2d2db32|Storage: Pure_2d2db32, Port: fc1/3, V
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N2P1_PHYSICAL|FS9150_N2P1_HOSTNVI
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N1P1_PHYSICAL|FS9150_N1P1_HOSTNVI
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	Pure_2d2db32|Storage: Pure_2d2db32, Port: fc1/3, V
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N2P1_PHYSICAL|FS9150_N2P1_HOSTNVI
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N1P1_PHYSICAL|FS9150_N1P1_HOSTNVI
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	Pure_2d2db32|Storage: Pure_2d2db32, Port: fc1/3, V
	Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N2P1_PHYSICAL|FS9150_N2P1_HOSTNVI

Figure 8-22 PMON events

Events can be filtered by selecting the **Quick Filter** option, as shown in Figure 8-23, and then selecting a filter. Here, we filter on **Warning**.

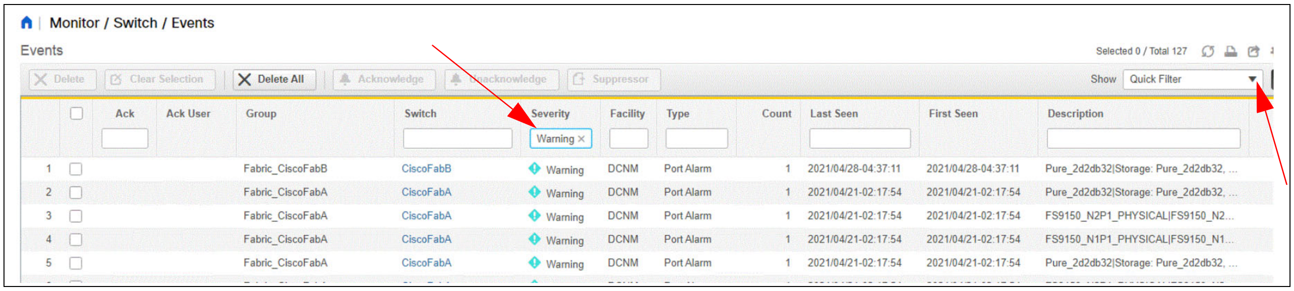

	□	Ack	Ack User	Group	Switch	Severity	Facility	Type	Count	Last Seen	First Seen	Description
						Warning ×						
1	□			Fabric_CiscoFabB	CiscoFabB	Warning	DCNM	Port Alarm	1	2021/04/28-04:37:11	2021/04/28-04:37:11	Pure_2d2db32|Storage: Pure_2d2db32, ...
2	□			Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	Pure_2d2db32|Storage: Pure_2d2db32, ...
3	□			Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N2P1_PHYSICAL|FS9150_N2...
4	□			Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	FS9150_N1P1_PHYSICAL|FS9150_N1...
5	□			Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1	2021/04/21-02:17:54	2021/04/21-02:17:54	Pure_2d2db32|Storage: Pure_2d2db32, ...

Figure 8-23 Quick Filter

By selecting none, one, or multiple rows in the left column, different actions appear to show what filters can be applied to the row, such as:

- Delete
- Clear Selection
- Delete All
- Acknowledge
- Unacknowledge
- Suppressor

These actions are highlighted in Figure 8-24.

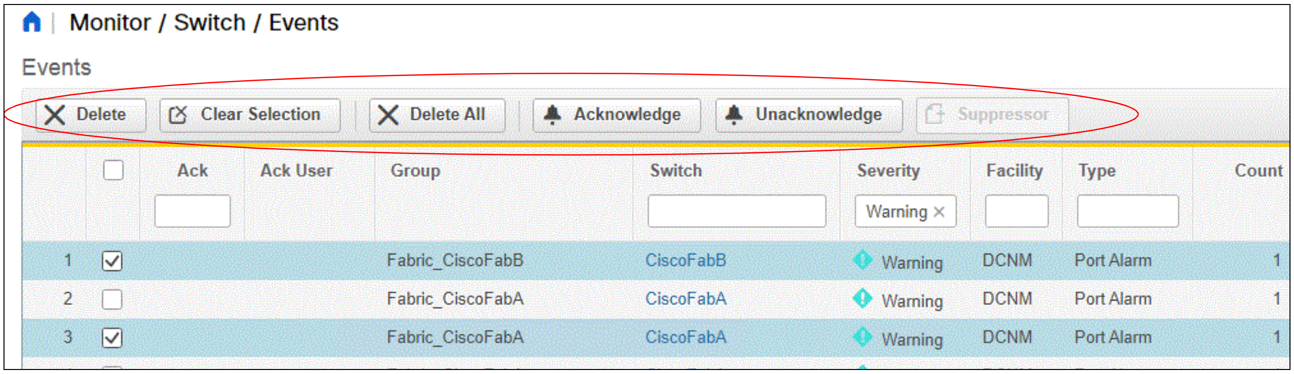

Monitor / Switch / Events

Events

	□	Ack	Ack User	Group	Switch	Severity	Facility	Type	Count
						Warning ×			
1	☑			Fabric_CiscoFabB	CiscoFabB	Warning	DCNM	Port Alarm	1
2	□			Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1
3	☑			Fabric_CiscoFabA	CiscoFabA	Warning	DCNM	Port Alarm	1

Figure 8-24 Event actions

You can view the active PMON policy on a switch by using DM or the CLI. The following examples show how to do this task by using DM. Launch DM from DCNM, as shown in Figure 8-25.

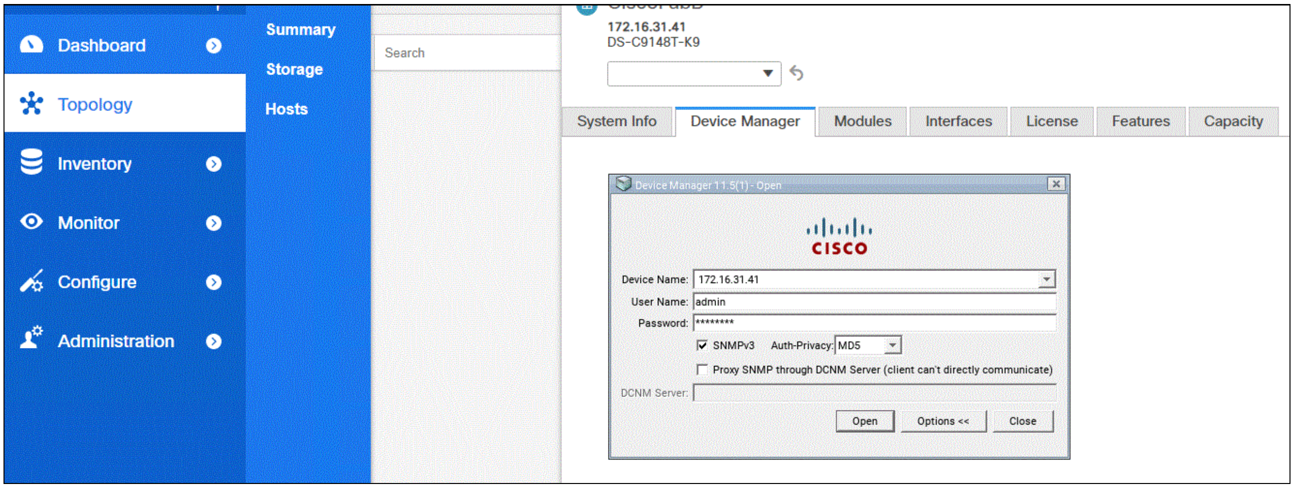

172.16.31.41
DS-C9148T-K9

Device Manager 11.5(1) - Open

Device Name: 172.16.31.41
User Name: admin
Password: ********
☑ SNMPv3 Auth-Privacy: MD5
□ Proxy SNMP through DCNM Server (client can't directly communicate)
DCNM Server:

Figure 8-25 Launch Device Manager

Select **Admin** → **Events** → **Port Monitor** → **Show**, as shown in Figure 8-26 on page 415.

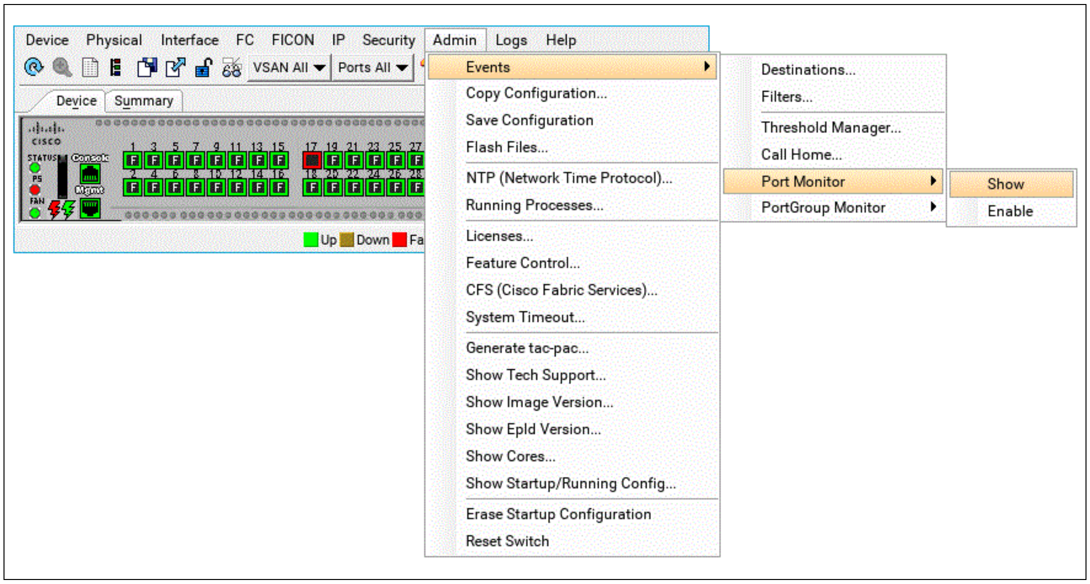

Figure 8-26 Port Monitor Show menu

A window opens and runs the **show port-monitor** command. Example 8-1 is truncated to show only the active policy, but inactive policies also can be shown by using this command.

Example 8-1 The show port-monitor command

```
Opening CLI connection to 172.16.31.41 and running commands ...

Waiting for output, CLI Command:show port-monitor
-------------------------------------------------------------------------------------------
Port Monitor : enabled
-------------------------------------------------------------------------------------------
Congestion-Isolation : disabled
-------------------------------------------------------------------------------------------

Policy Name  : ITSO_Policy_1_allPort
Admin status : Active
Oper status  : Active
Logical type : All Ports
--------------------------------------------------------      ----------------------------------------------------------------------
Counter                  Threshold  Interval Rising Threshold event Falling Threshold  event Warning Threshold     PMON Portguard
-------                  ---------  -------- ---------------- ----- ------------------ ----- --------------------- --------------
Link Loss                Delta      10       5                4     1                  4     Not enabled           Not enabled
Sync Loss                Delta      10       5                4     1                  4     Not enabled           Not enabled
Signal Loss              Delta      10       5                4     1                  4     Not enabled           Not enabled
Invalid Words            Delta      10       5                4     1                  4     Not enabled           Not enabled
Invalid CRC's            Delta      10       5                4     1                  4     Not enabled           Not enabled
State Change             Delta      10       5                4     0                  4     Not enabled           Not enabled
TX Discards              Delta      10       20               4     5                  4     Not enabled           Not enabled
LR RX                    Delta      10       3                4     1                  4     Not enabled           Not enabled
LR TX                    Delta      10       3                4     1                  4     Not enabled           Not enabled
Timeout Discards         Delta      10       50               4     10                 4     Not enabled           Not enabled
Credit Loss Reco         Delta      1        1                4     0                  4     Not enabled           Not enabled
TX Credit Not Available  Delta      1        10%              4     0%                 4     Not enabled           Not enabled
RX Datarate              Delta      10       90%              4     50%                4     Not enabled           Not enabled
TX Datarate              Delta      10       90%              4     50%                4     Not enabled           Not enabled
ASIC Error Pkt from Port Delta      60       50               4     30                 4     Not enabled           Not enabled
ASIC Error Pkt to xbar   Delta      60       50               4     30                 4     Not enabled           Not enabled
ASIC Error Pkt from xbar Delta      60       50               4     30                 4     Not enabled           Not enabled
TX-Slowport-Oper-Delay   Absolute   1        30ms             4     0ms                4     Not enabled           Not enabled
TXWait                   Delta      1        20%              4     0%                 4     Not enabled           Not enabled
-----------------------------------------------------------------------------------------------------------
Command completed
```

Click **Close** to close the window, as shown in Figure 8-27.

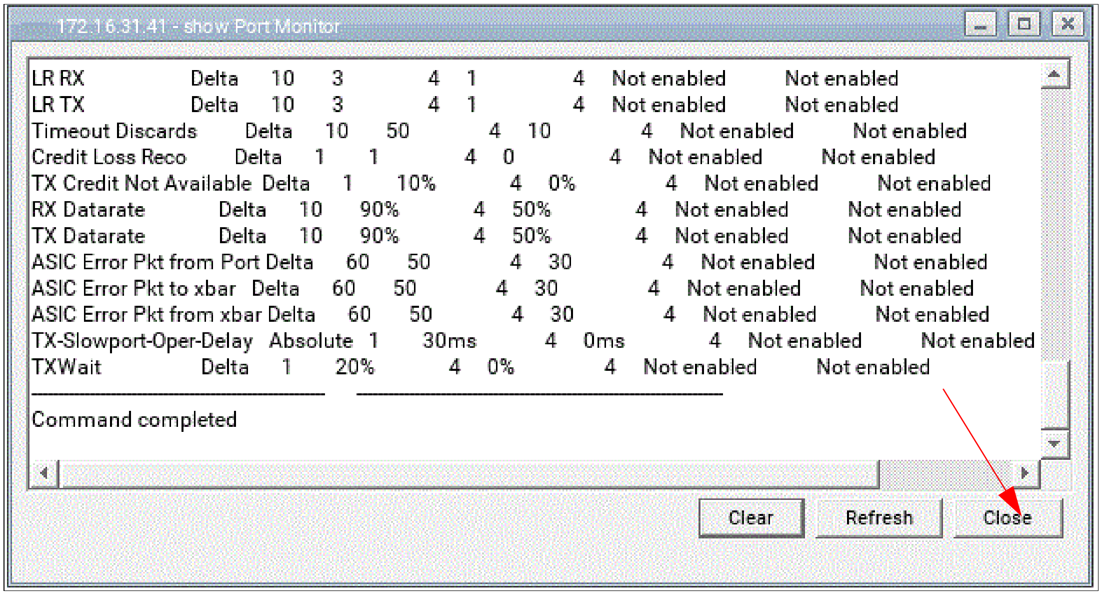

```
LR RX                 Delta    10     3          4    1           4   Not enabled        Not enabled
LR TX                 Delta    10     3          4    1           4   Not enabled        Not enabled
Timeout Discards        Delta    10     50          4    10          4   Not enabled        Not enabled
Credit Loss Reco        Delta    1      1           4    0           4   Not enabled        Not enabled
TX Credit Not Available  Delta    1      10%         4    0%          4   Not enabled        Not enabled
RX Datarate             Delta    10     90%         4    50%         4   Not enabled        Not enabled
TX Datarate             Delta    10     90%         4    50%         4   Not enabled        Not enabled
ASIC Error Pkt from Port Delta    60     50          4    30          4   Not enabled        Not enabled
ASIC Error Pkt to xbar   Delta    60     50          4    30          4   Not enabled        Not enabled
ASIC Error Pkt from xbar Delta    60     50          4    30          4   Not enabled        Not enabled
TX-Slowport-Oper-Delay  Absolute  1     30ms         4    0ms         4   Not enabled        Not enabled
TXWait                  Delta    1      20%         4    0%          4   Not enabled        Not enabled
-----------------------------------   ---------------------------------------------
Command completed
```

Figure 8-27 Close window

8.2 Backing up a switch configuration

When performing configuration changes, upgrades, or hardware replacements on IBM c-type switches, it is a best practice to perform a backup of the switch configuration. The IBM c-type Family can leverage Cisco NX-OS software that is on the switches to create backups of a switch configuration. The backup copy of a configuration file is stored in the internal memory to a remote server as a backup or to use for configuring other IBM c-type devices in your fabric. The commands are run by the software when the device is started by the `startup-config` file or when commands are entered at the CLI in configuration mode.

Cisco NX-OS software has two types of configuration files:

- Startup configuration: Used during a device start to configure the software features.
- Running configuration: Contains the configuration and eventual changes that you want to make and then save to the `startup-configuration` file.

These two configuration files can be different in instances where you want to change the device configuration temporarily without saving the running configuration changes to `startup-configuration`.

Before changing the startup configuration file, save the `running-configuration` file to the startup configuration by using the **`copy running-config startup-config`** command or copy a configuration file from a backup copy that is a file server to the startup configuration.

To change the running configuration, use the **`configure terminal`** command to enter configuration mode. After you enter global configuration mode, commands generally run immediately and then are saved to the running configuration file immediately after the command runs or when you exit configuration mode.

> **Best practice:** Back up your switch configuration and save a copy to an external location before making changes.

In Example 8-2, we save `running-config` by running the **`copy running-config startup-config`** command to store the running configuration. Then, we copy the configuration that we saved to a TFTP server, which accomplishes three things:

- Verifies that you have an operational TFTP server in your environment.
- Verifies that you can communicate to the server over the IP network.
- Allows you to store a copy of the configuration in a location that is external to the switch so that you have a backup in the event of a switch failure.

Example 8-2 The copy running-config startup-config command

```
switch# copy running-config startup-config
switch# copy startup-config tftp://0.0.0.0/switchname_config_date
```

Alternatively, you can perform a configuration backup of your switch by using DCNM. With this feature, you can back up device configurations from a running configuration. The backup files can be stored on the DCNM server or an external location, which is recommended.

Important: For more information about how to back up a device configuration by using DCNM, see Backup.

8.3 Call Home

IBM highly recommends that you enable Call Home for all IBM c-type switches. Call Home is a support function that is embedded in all IBM c-type products. By enabling Call Home, the health and stability of your system is monitored every hour of every day throughout the year by the industry’s top troubleshooting specialists at IBM Support. The Call Home service provides reduced risk by alerting you about a system defect through My Notifications, which automatically open a Problem Management Record (PMR). This service can decrease system downtime through faster problem determination and resolution.

Call Home provides email-based notification of critical system events, which can go to your SAN administrators. The Call Home function is available directly through the IBM c-type Family switches. Call Home provides multiple Call Home messages, separate potential destinations, and you can define your own destination profiles, use predefined profiles, and configure up to 50 email addresses per destination profile.

Best practice: Configure Call Home on all c-type switches in your environment as a preventive maintenance feature. For more information, see "Configuring Call Home" in the *Cisco MDS 9000 Family Configuration Guide* and IBM c-type Family and Cisco MDS 9000 Series Remote Support Overview.

8.4 Read Diagnostic Parameters

The IBM c-type Series is equipped with a large set of diagnostic and troubleshooting tools. In fact, FC technology often is chosen for its reliability and performance for mission-critical applications. Concurrently, an FC SAN must be manageable. Standard bodies have continually improved the protocol definition and product vendors have implemented multiple solutions to accommodate those needs and help manage storage networks.

FC fabric connectivity requires multiple electrical and optical components to function correctly, including cables, transceivers, port ASICs, switching ASICs, and communication buses internal to the switches. If any of these components are faulty, they affect I/O operations over the fabric. Today, FC is deployed in mission-critical networks where resiliency and throughput are high-priority requirements. In such networks, early identification of any faults is critical to gaining customer trust. For this reason, the IBM c-type Series provides a comprehensive set of system-and link-level diagnostic capabilities.

This software-based suite of tools, hardware-enabled for some tests, can dynamically verify whether everything is working as expected. The Generic Online Diagnostics (GOLD) capability offers a complete suite of tests to verify that supervisor engines, switching modules, ASICs, communication buses, optics, and interconnections are functioning properly. GOLD tests can be run at initial system start, periodically at run time, and on demand when invoked by the administrator.

The start diagnostics run during the start procedure and detect faulty hardware when a new networking device is brought online. These tests represent an evolution of the power-on self-test (POST) capabilities once present on similar switches. They verify the checksum on the boot and switch firmware images, perform internal data loopback testing on all FC ports, and perform access and integrity checks on management ports and nonvolatile memory components. During the diagnostics phase, the switch logs any errors that are encountered.

Runtime and on-demand tests are even more specific and implement health-monitoring diagnostics. Enabled by default, they verify the health of a live system at user-configurable periodic intervals. The health-monitoring diagnostic tests detect possible hardware errors and data-path problems without disrupting the data or control traffic. ISL diagnostics are available to help check the health and performance of ISLs (E and TE ports) before the links are activated for production traffic, measuring frame round-trip latencies and cable lengths.

Figure 8-28 shows the ISL diagnostics capability.

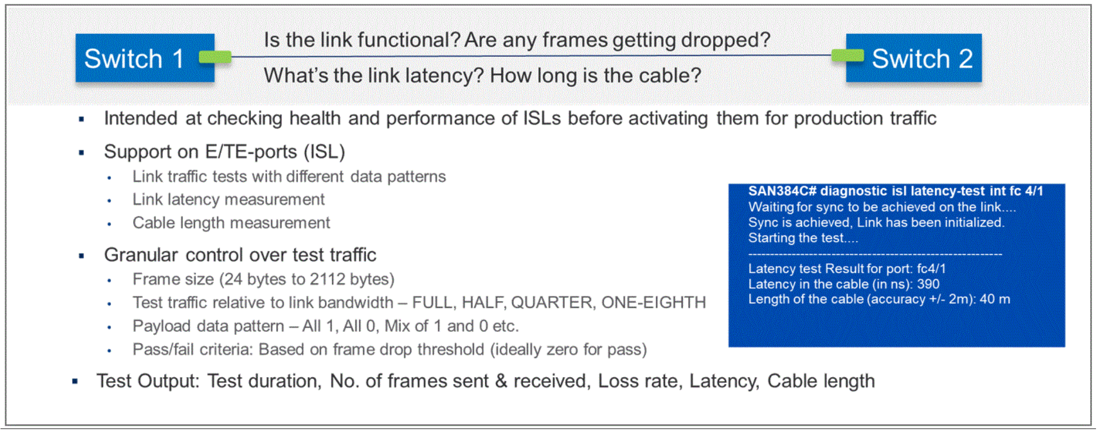
Switch 1
Is the link functional? Are any frames getting dropped?
What's the link latency? How long is the cable?
Switch 2

- Intended at checking health and performance of ISLs before activating them for production traffic
- Support on E/TE-ports (ISL)
  - Link traffic tests with different data patterns
  - Link latency measurement
  - Cable length measurement
- Granular control over test traffic
  - Frame size (24 bytes to 2112 bytes)
  - Test traffic relative to link bandwidth – FULL, HALF, QUARTER, ONE-EIGHTH
  - Payload data pattern – All 1, All 0, Mix of 1 and 0 etc.
  - Pass/fail criteria: Based on frame drop threshold (ideally zero for pass)
- Test Output: Test duration, No. of frames sent & received, Loss rate, Latency, Cable length

```
SAN384C# diagnostic isl latency-test int fc 4/1
Waiting for sync to be achieved on the link....
Sync is achieved, Link has been initialized.
Starting the test....
-------------------------------------------------------------
Latency test Result for port: fc4/1
Latency in the cable (in ns): 390
Length of the cable (accuracy +/- 2m): 40 m
```

Figure 8-28 ISL diagnostics capability

Single-hop and multihop tests are also available for further analysis along the path, including link-saturation stress tests.

Host bus adapter (HBA) diagnostic capability is also available. It is like ISL diagnostics but supported on F ports. These capabilities can be configured from a CLI or DCNM. Figure 8-29 shows how to configure HBA diagnostics from DCNM.

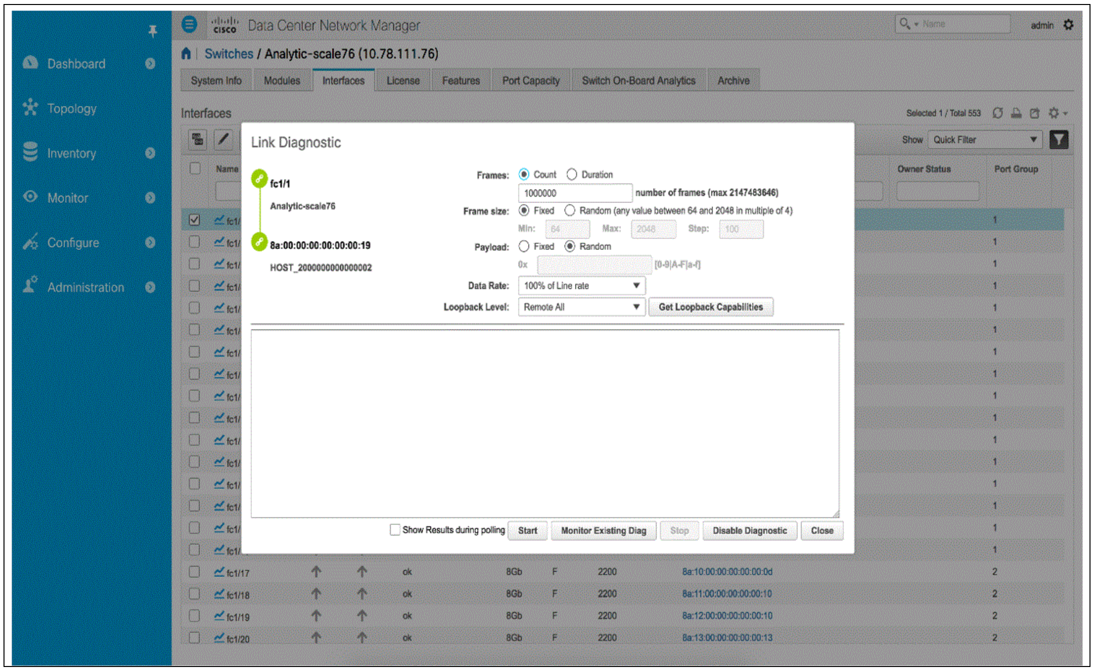

Link Diagnostic

fc1/1
Analytic-scale76
8a:00:00:00:00:00:00:19
HOST_2000000000000002

Frames: ● Count ○ Duration
1000000 number of frames (max 2147483646)
Frame size: ● Fixed ○ Random (any value between 64 and 2048 in multiple of 4)
Min: 64 Max: 2048 Step: 100
Payload: ○ Fixed ● Random
0x [0-9|A-F|a-f]
Data Rate: 100% of Line rate
Loopback Level: Remote All
□ Show Results during polling

fc1/17	↑	↑	ok	8Gb	F	2200	8a:10:00:00:00:00:00:0d	2
fc1/18	↑	↑	ok	8Gb	F	2200	8a:11:00:00:00:00:00:10	2
fc1/19	↑	↑	ok	8Gb	F	2200	8a:12:00:00:00:00:00:10	2
fc1/20	↑	↑	ok	8Gb	F	2200	8a:13:00:00:00:00:00:13	2

Figure 8-29 HBA diagnostic configuration

Host-to-switch connectivity (N and F ports) tests are also available to IBM c-type Family devices as an extension to the diagnostics suite. For host connectivity probing, the International Committee for Information Technology Standards (INCITS) T11 FC-LS-4 standard refers to a specific implementation for beaconing the peer port for ease of identification and a capability to gather detailed information from end nodes. This solution is based on Link Cable Beacon Extended Link Service (LCB-ELS) and the Read Diagnostic Parameters (RDP) Link Service command, which is used to query N_port-related link-and port-level diagnostic parameters. In addition to these intelligent diagnostics features, the IBM c-type Family offers hardware-based slow-drain port detection and remediation capabilities that go beyond the capabilities that are offered by competing products.

There are two versions of the RPD feature. The RDP query can be host-originated or switch-originated, and provides visibility into the operational port and link characteristics of any other port in the SAN. In both cases, the feature must be supported by the HBA and the switch and is included at no cost on IBM c-type switches.

The host-originated feature is intended for periodic housekeeping for health and performance of switch ports locally connecting to HBAs or remote switch ports. Typically, hosts initiate the RDP request to query the diagnostic parameters of the N_port of the target device.

Switch-originated RDP works in the opposite way. An RDP request can be sent from an IBM c-type switch to any end device and request the diagnostic parameters of the N_port. The queried device can be locally or remotely connected to the switch from where the RDP request is sent.

In our example, we focus on switch-originated RDP, as shown in Figure 8-30.

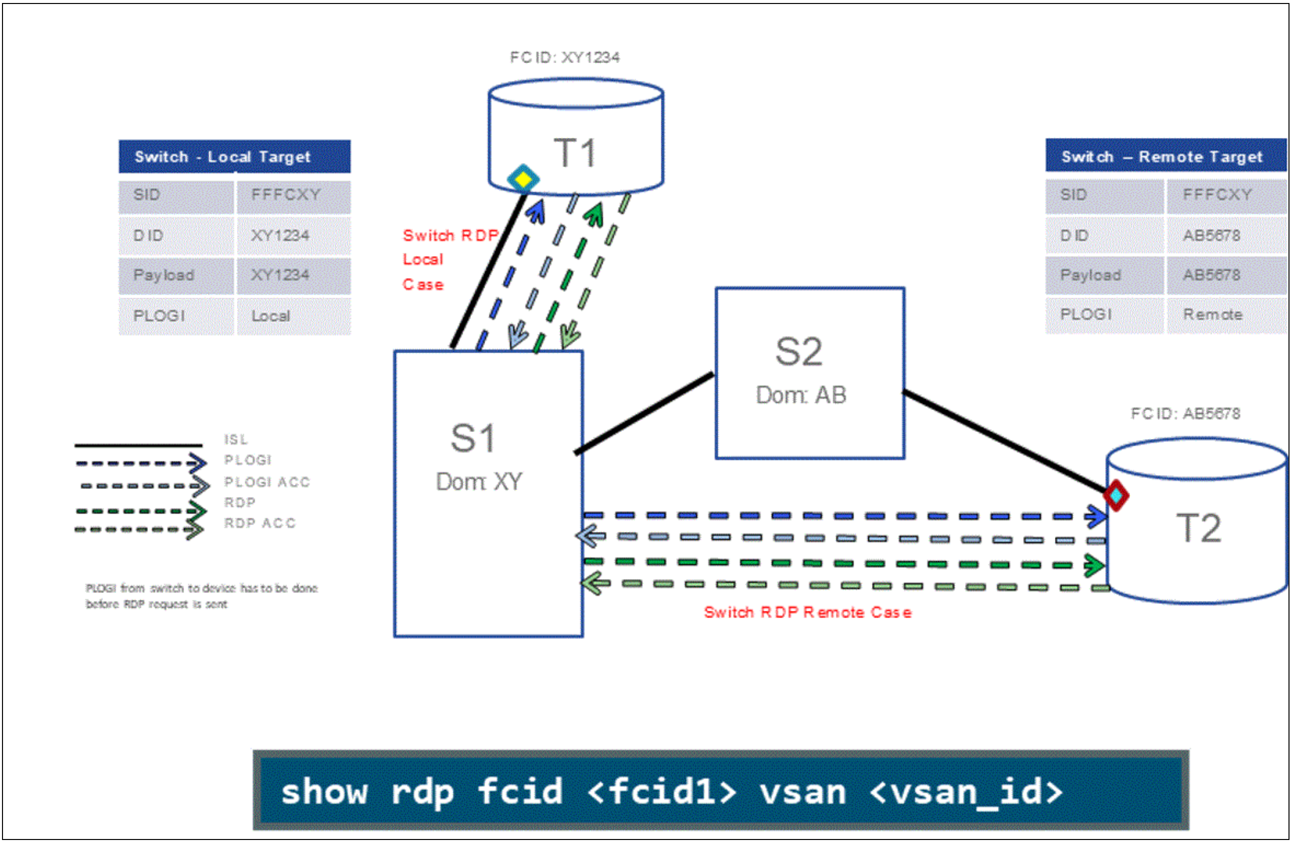

Switch - Local Target	
SID	FFFCXY
DID	XY1234
Payload	XY1234
PLOGI	Local

Switch – Remote Target	
SID	FFFCXY
DID	AB5678
Payload	AB5678
PLOGI	Remote

```
show rdp fcid <fcid1> vsan <vsan_id>
```

Figure 8-30 RDP workflow

The FC RDP feature can read port and link diagnostic parameters like link errors, congestion counters, port names, port speeds, Small Form-factor Pluggable (SFP) diagnostics, temperatures, Rx power, Tx power, electrical current, Forward Error Correction (FEC) status, buffer credits, serial number, vendor details, model number, and manufacture date.

The benefit of RDP feature is that link issues can be diagnosed centrally without sending someone inside data center rooms with optical power meters or other measurements tools, which potentially might cause further disruption on the links. Congestion situations can be identified, and appropriate values of buffer credits can be determined for the distances that are involved. It is also possible to identify links where auto-negotiation did not operate properly, and the operating speed is lower than expected. This information, when available for each end of a link or each port in the path from server to disk, gives visibility into the health of the transport infrastructure. The SAN can be monitored for current failing links and transceivers, and investigated for proactive maintenance with predictive analyses.

Example 8-3 shows where the switch gets information from the connected host HBA.

Example 8-3 The show rdp fcid command

```
switch# show rdp fcid 0xaa0260 vsan 1
-------------------------------------------------------------
                   RDP frame details
-------------------------------------------------------------
Link Service Request Info:
------------------------------
```

```
Port Speed Descriptor Info:
------------------------------
Port speed capabilities : 16GbE 8GbE 4GbE
Port Oper speed         : Unknown Oper speed

Link Error Status:
------------------------------
VN PHY port type               : FC
Link failure count             : 0
Loss of sync count             : 0
Loss of signal count           : 0
Primitive sequence proto error : 0
Invalid Transmission word      : 0
Invalid CRC count              : 0

Port Name Descriptor:
------------------------------
Node WWN            : 20:00:8c:60:4f:54:54:00
Port WWN            : 21:01:8c:60:4f:54:54:00
Attached Node WWN   : 50:08:01:60:00:89:07:51
Attached Port WWN   : 50:08:01:60:00:89:08:51

SFP Diag params:
------------------------------
SFP flags       : SFP+  Optical
SFP Tx Type     : Short Wave

FEC Status:
------------------------------
Corrected blocks   : 0
Uncorrected blocks : 0

Buffer Credit Descriptor:
------------------------------
Rx B2B credit   : 1
Tx B2B credit   : 16
Port RTT        : 0 ns

Optical Product Data:
------------------------------
Vendor Name    : CISCO-AVAGO
Model No.      : AFBR-57F5PZ-CS1
Serial No.     : AVA1602J0FY
Revision       : B2
Date           : 120112

Port Congestion:
------------------------------
Tx Zero Credit Count      : 3
Rx Zero Credit Count      : 0
Tx Delay Count            : 0
Delay Interval            : 2500
Tx Discard Count          : 0
Tx Discard Interval       : 500
Active State Tx LR Count  : 0
Active State Rx LR Count  : 0

  ----------------------------------------------------------------------------
                Current              Alarms                  Warnings
                Measurement     High        Low         High          Low
  ----------------------------------------------------------------------------
  Temperature   26.89 C         75.00 C     -5.00 C     70.00 C        0.00 C
  Voltage        3.28 V          3.63 V      2.97 V      3.46 V        3.13 V
  Current        7.37 mA        10.50 mA     2.50 mA    10.50 mA       2.50 mA
  Tx Power      -2.49 dBm        1.70 dBm  -13.01 dBm   -1.30 dBm     -9.00 dBm
  Rx Power     -23.87 dBm        3.00 dBm  -15.92 dBm    0.00 dBm    -11.90 dBm
  ----------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning

ssss
```

The RDP feature can be applied to FICON environments, and the same set of parameters can be collected.

The capabilities in the diagnostics suite of IBM c-type switches might reduce operational costs. Some of the capabilities are as follows:

- Verify infrastructure readiness before going live into production (pre-production).
- Provide key insights to troubleshoot production connections (production).
- Find issues before they become critical situations.
- Reduce operational costs by pinpointing and resolving issues fast.

Figure 8-31 summarizes the benefits of the IBM c-type diagnostics suite.

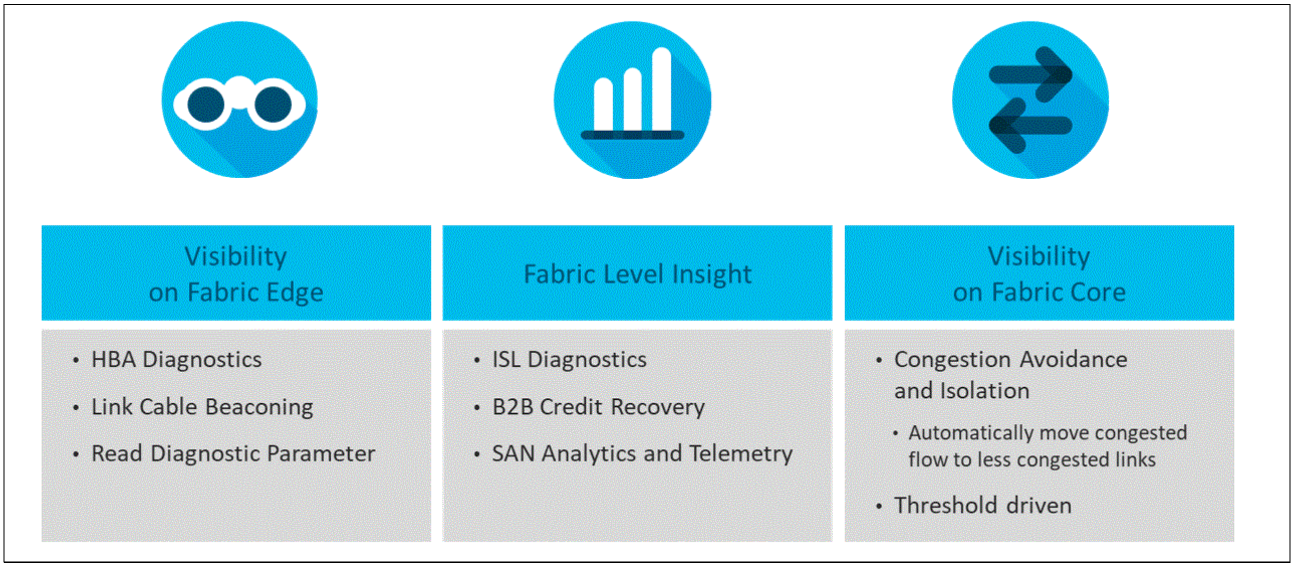

Visibility on Fabric Edge

- HBA Diagnostics
- Link Cable Beaconing
- Read Diagnostic Parameter

Fabric Level Insight

- ISL Diagnostics
- B2B Credit Recovery
- SAN Analytics and Telemetry

Visibility on Fabric Core

- Congestion Avoidance and Isolation
  - Automatically move congested flow to less congested links
- Threshold driven

Figure 8-31   Benefits of the diagnostics suite

8.5 Port beacon and location LED

Administrators of storage networks strive to achieve the higher possible uptime and network stability. Data must be always accessible. As part of this effort, a couple of hidden gems deserve to be mentioned because they are important differentiators of the IBM c-type Family:

- Ease of module replacement
- Locator ID on all modules and beacon LED

Ease of module replacement

Ease of replacement of failed components was considered in the design of the IBM c-type Series of mission-critical directors. Because availability improves when the time to restore service after a fault is short, all modules are hot-swappable, and all include fast ejectors or handles and can be replaced in less than 1 minute when spare parts are available onsite. Modules are easily accommodated in available slots in the chassis and receive electricity from dual connectors that are hosted in the passive backplane. In this way, single points of failure (SPOFs) are eliminated. IBM c-type switches have fewer parts than must be replaced in a failure, like fans and power supplies.

Optical transceivers can be easily replaced by extracting them from hosting modules with their mylar tab latch or bale-clasp latch.

Figure 8-32 shows SFPs and their different latch types.

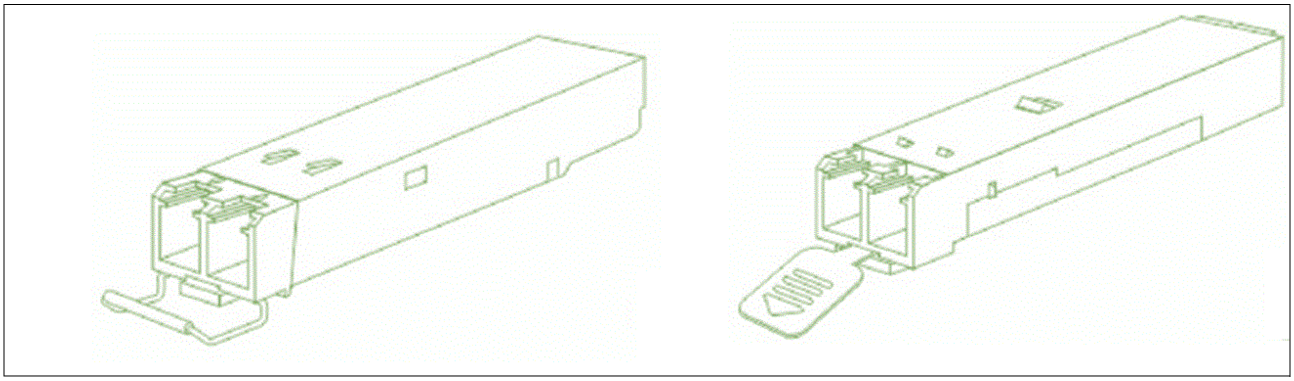

Figure 8-32 Latch types

Locator ID LED on all modules and beacon LEDs

Before replacing a failed module, you must identify it among its peers in the chassis and across different chassis in the data center. To help you, the IBM c-type Family of FC switches provides LEDs for easy identification of ports and modules. The beacon mode displays a flashing green light that helps you identify the port that you are seeking.

The locator ID LED helps you identify line cards, supervisors, power supplies, fans, or crossbar fabric units. The IBM c-type Family is the only one that offers locator IDs for all system modules. The administrator can turn on the beacon mode or locator ID LED from the remote central management station so that the support engineer can quickly identify the component that requires attention. Enabling the beacon mode or locator ID LED has no effect on the operation of the interface or module.

The locator ID LED for the power supply in slot 3 can be turned on from the NX-OS CLI by sending the command that is shown in Example 8-4.

Example 8-4 The locator-led powersupply 3 command

```
switch(config)# locator-led powersupply 3
```

Figure 8-33 shows a locator ID LED on the 3-kW AC Power Supply Module.

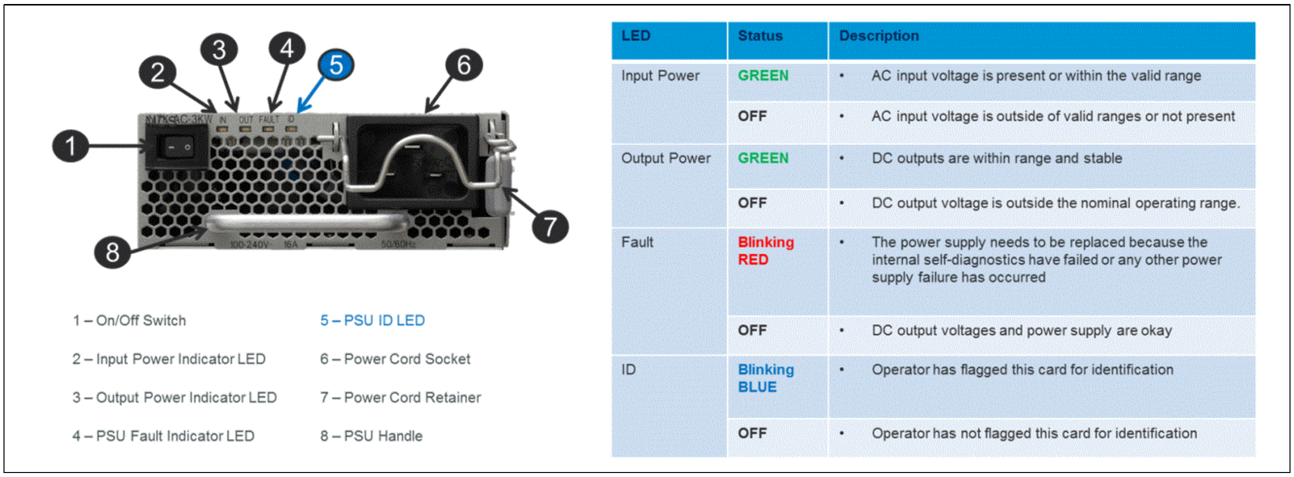

LED	Status	Description
Input Power	GREEN	• AC input voltage is present or within the valid range
	OFF	• AC input voltage is outside of valid ranges or not present
Output Power	GREEN	• DC outputs are within range and stable
	OFF	• DC output voltage is outside the nominal operating range.
Fault	Blinking RED	• The power supply needs to be replaced because the internal self-diagnostics have failed or any other power supply failure has occurred
	OFF	• DC output voltages and power supply are okay
ID	Blinking BLUE	• Operator has flagged this card for identification
	OFF	• Operator has not flagged this card for identification

Figure 8-33 Locator ID LED

8.6 On-board Failure Logging

IBM c-type Family switch and director line cards can log failure data to persistent storage, which can be later retrieved and displayed for analysis. The On-board Failure Logging (OBFL) feature stores failure and environmental information in non-volatile memory on the module or switch. The information helps in post-mortem analysis of failed cards or failed switches. Consider it the black box of your switch, like what was introduced years ago on aircraft.

Specific critical events, error conditions, and important statistics are automatically recorded with their timestamps in non-volatile random access memory (NVRAM) onboard the IBM c-type Family switch and director line cards. This OBFL capability provides an event data recorder for networking devices and is useful for performing root-cause analyses of slow-drain situations even after they are cleared. Post-mortem analysis of failed cards or failed switches is possible by retrieving the stored information. OBFL is enabled by default on all IBM c-type Family switches and director line cards.

The OBFL process on each line card runs separately at (typically) 20-second intervals and records any counter that changed value in the last interval. When it detects a counter that changed value, it records the following information:

- Interface or interface range
- Counter name
- Current counter value
- Date and time of when OBFL detected the counter's changed values

To determine the amount that the counter incremented in the OBFL interval, the previous counter value for the same counter name for the same interface must be subtracted from the current counter value. There are various sections in OBFL, and they have different purposes. The main sections are as follows

- cpuhog: Information about processes that use excessive CPU.
- environmental-history: Information about temperature sensors.
- error-stats: Information about errors that are related to performance, congestion, and slow drain.
- interrupt-stats: Information about various module interrupts.
- slowport-monitor-events: Information about slow PMON.
- txwait: Information about the amount of time that interfaces spend at zero Tx credits.
- stack-trace: Information about process crashes.

Each of these recorded events can be displayed starting at a specific date and time and ending at a specific date and time. This capability allows problems that occurred even months ago to be investigated. These events are often the first place to look after a problem occurs. OBFL is a unique feature of IBM c-type storage networking devices and is considered valuable by support specialists. It is one of those features that under normal conditions is often ignored but becomes critical when you need it.

Example 8-5 shows how frame drops and the TxWait counter would be timestamped so that it is easier to correlate frame drops within the switch with external notification of drops and application issues. A simple counter of drops with no timestamp would not serve this purpose. Example 8-5 shows that two counters of the F32 ASIC keep incrementing over time.

Example 8-5 The show logging onboard error-stats command

```
F241-SAN384C# show logging onboard module 9 error-stats
----------------------------
 Show Clock
----------------------------
2021-04-13 16:08:23
---------------------------------
 Module: 9 error-stats
---------------------------------

------------------------------------------------------------------------------------------------
 ERROR STATISTICS INFORMATION FOR DEVICE DEVICE: FCMAC
------------------------------------------------------------------------------------------------
    Interface      |                                                |         |    Time Stamp
      Range        |            Error Stat Counter Name             |  Count  |MM/DD/YY HH:MM:SS
                   |                                                |         |
------------------------------------------------------------------------------------------------
fc9/17             |F32_TMM_PORT_TIMEOUT_DROP                       |18032    |04/13/21 16:08:18
fc9/17             |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO             |4357     |04/13/21 16:08:18
fc9/17             |F32_TMM_PORT_TIMEOUT_DROP                       |11817    |04/13/21 16:07:58
fc9/17             |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO             |4206     |04/13/21 16:07:58
fc9/17             |F32_TMM_PORT_TIMEOUT_DROP                       |6161     |04/13/21 16:07:38
fc9/17             |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO             |4055     |04/13/21 16:07:38
fc9/17             |F32_TMM_PORT_TIMEOUT_DROP                       |223      |04/13/21 16:07:18
fc9/17             |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO             |3933     |04/13/21 16:07:18
fc9/17             |F32_TMM_PORT_TIMEOUT_DROP                       |195      |04/13/21 16:06:58
fc9/17             |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO             |3808     |04/13/21 16:06:58
```

```
fc9/17              |F32_TMM_PORT_TIMEOUT_DROP                   |177       |04/13/21 16:06:38
fc9/17              |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO         |3656      |04/13/21 16:06:38
fc9/17              |F32_TMM_PORT_TIMEOUT_DROP                   |155       |04/13/21 16:06:18
fc9/17              |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO         |3505      |04/13/21 16:06:18
fc9/17              |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO         |3352      |04/13/21 16:05:57
fc9/17              |F32_TMM_PORT_TIMEOUT_DROP                   |139       |04/13/21 16:05:37
fc9/17              |F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO         |3199      |04/13/21 16:05:37
fc9/17              |F32_TMM_PORT_TIMEOUT_DROP                   |123       |04/13/21 16:05:17
```

8.7 FICON Director Activity Report

The System Management Facility (SMF) collects and records system and job-related information that can be used for reporting, analyzing, and evaluating activity on an IBM z/OS system, including the processor, control units (CUs), and FICON Directors. Users (through programming exits) and purchased software products can create their own records.

Resource Measurement Facility (RMF) is the IBM strategic product to present the system activity, and it uses SMF records and z/OS monitoring services for its functions.

Comprehensive Management Facility (CMF) is a performance monitoring product that competes with IBM RMF. CMF was created in the late 1970s by Boole and Babbage, but it is now owned by BMC. CMF can use SMF records and creates its own SMF records, and other sources of system information.

SMF record type 74 has several subtypes. Subtype 7 is used to collect FICON Director data, which it gets by communicating with the FICON IBM Control Unit Port (CUP). Both RMF and CMF can produce a FICON Director Activity Report from that data.

By default, SMF 74.7 records are not saved and FICON Director Activity reports are not produced.

Figure 8-34 on page 427 shows the SMF Record Types for I/O.

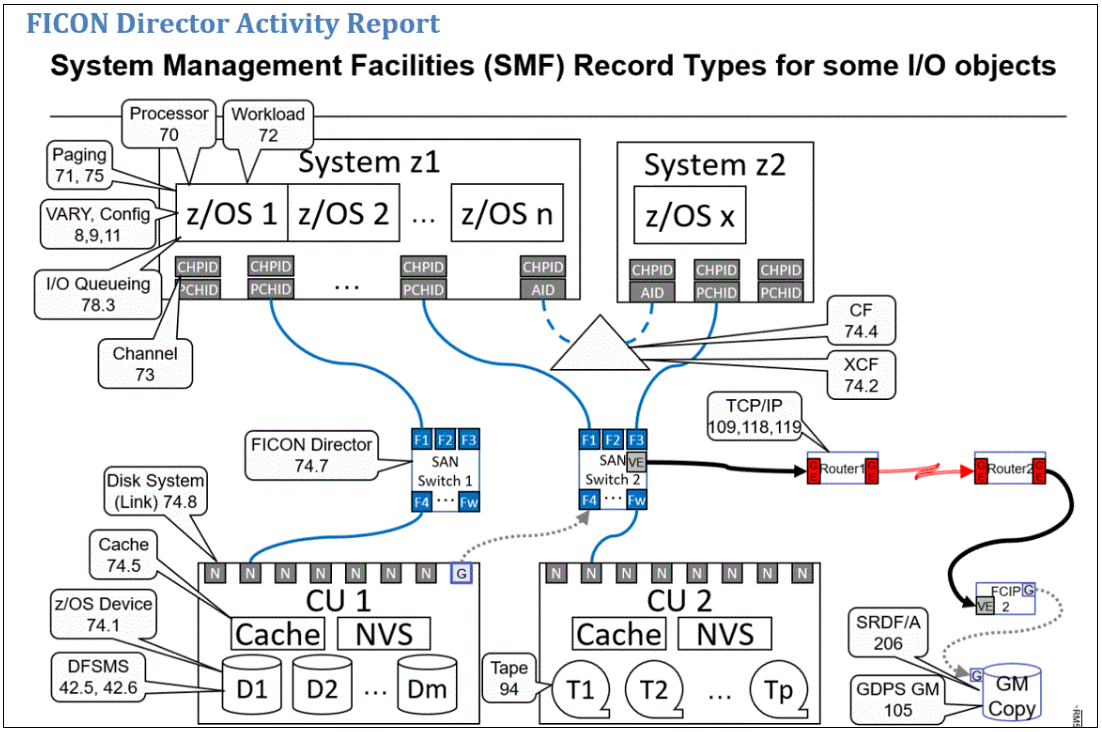

Figure 8-34 SMF Record types for I/O

To capture SMF 74.7 records and create FICON Director Activity reports, you must complete the following tasks:

1. Capture SMF 74.7 records:

   a. Parmlib entries:

      Parmlib member SMFPRMxx:

      Add (or change) a parameter to a record:

      ```
      LSNAME(SYS1.SMF.PERF,TYPE(30,89,74.7))RECORDING(LOGSTREAM)
      ```

   b. Operator commands:

      i. Use the **SETSMF** command to dynamically change (add) the TYPEs that are collected.

      ii. Use the **D SMF** command to verify what you specified.

2. Enable FICON Director Activity Report by using RMF:

   a. Parmlib entries:

      i. Parmlib member ERBRMFxx:

         Add (or change) a parameter to **FCD** (FICON Director Analysis). It is not on by default (ERBfRMF00) when the member contains **NOFCD /* NO FICON DIRECTOR MEASURED */**.

      ii. Optionally, parmlib member IECIOSxx:

         Add the parameter **FICON STATS=NO** on any system where you do *not* want these records collected. You can put FCD in all systems, and identify which focal point system collects them.

   b. Operator commands:

      Run the **D IOS,FICON** command to see what you have.

c. Job (three steps):

   i. Copy interesting records from the RMF data set:

      `OUTDD(B,TYPE(74(7)))`

   ii. Sort (required).

   iii. Report:

      ```
      a. SYSOUT(X)
      b. REPORTS(FCD)
      ```

3. Enable FICON Director Activity Report by using CMF:

   a. Parmlib member CMFCPMxx:

      Add (or change) a parameter to **`FICONSW`**.

   b. Report:

      Select FICON Director Activity Report.

Figure 8-35 shows an example of a FICON Director Activity Report from RMF. The `UNIT` column identifies the following items:

- `SWITCH` for an ISL
- `CHP` for a FICON Channel Path ID (CHPID)
- `CHP-H` for an IBM High-Performance FICON for System z (IBM zHPF) FICON CHPID
- `CU` for a CU interface.

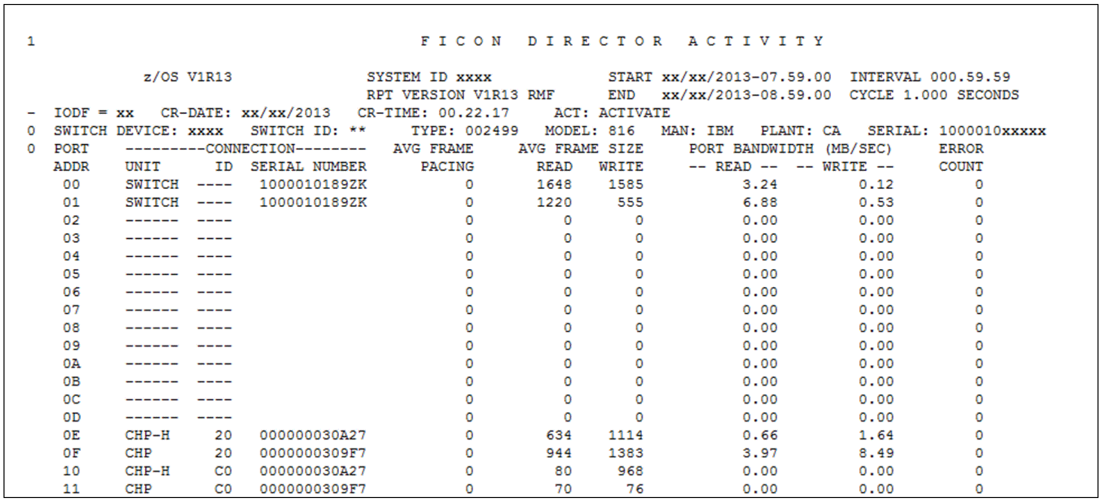

```
1                                            F I C O N   D I R E C T O R   A C T I V I T Y

            z/OS V1R13               SYSTEM ID xxxx            START xx/xx/2013-07.59.00  INTERVAL 000.59.59
                                     RPT VERSION V1R13 RMF     END   xx/xx/2013-08.59.00  CYCLE 1.000 SECONDS
-  IODF = xx   CR-DATE: xx/xx/2013   CR-TIME: 00.22.17     ACT: ACTIVATE
0  SWITCH DEVICE: xxxx   SWITCH ID: **     TYPE: 002499   MODEL: 816   MAN: IBM   PLANT: CA   SERIAL: 1000010xxxxx
0  PORT    ---------CONNECTION--------   AVG FRAME     AVG FRAME SIZE     PORT BANDWIDTH (MB/SEC)     ERROR
   ADDR    UNIT      ID  SERIAL NUMBER     PACING       READ   WRITE      -- READ --   -- WRITE --     COUNT
    00     SWITCH  ----   1000010189ZK          0       1648    1585            3.24          0.12         0
    01     SWITCH  ----   1000010189ZK          0       1220     555            6.88          0.53         0
    02     ------  ----                         0          0       0            0.00          0.00         0
    03     ------  ----                         0          0       0            0.00          0.00         0
    04     ------  ----                         0          0       0            0.00          0.00         0
    05     ------  ----                         0          0       0            0.00          0.00         0
    06     ------  ----                         0          0       0            0.00          0.00         0
    07     ------  ----                         0          0       0            0.00          0.00         0
    08     ------  ----                         0          0       0            0.00          0.00         0
    09     ------  ----                         0          0       0            0.00          0.00         0
    0A     ------  ----                         0          0       0            0.00          0.00         0
    0B     ------  ----                         0          0       0            0.00          0.00         0
    0C     ------  ----                         0          0       0            0.00          0.00         0
    0D     ------  ----                         0          0       0            0.00          0.00         0
    0E     CHP-H     20   000000030A27          0        634    1114            0.66          1.64         0
    0F     CHP       20   0000000309F7          0        944    1383            3.97          8.49         0
    10     CHP-H     C0   000000030A27          0         80     968            0.00          0.00         0
    11     CHP       C0   0000000309F7          0         70      76            0.00          0.00         0
```

Figure 8-35 Sample FICON Director Activity Report

Figure 8-36 on page 429 shows an example of a FICON Director Activity Report from RMF where there are two switches with ISLs.

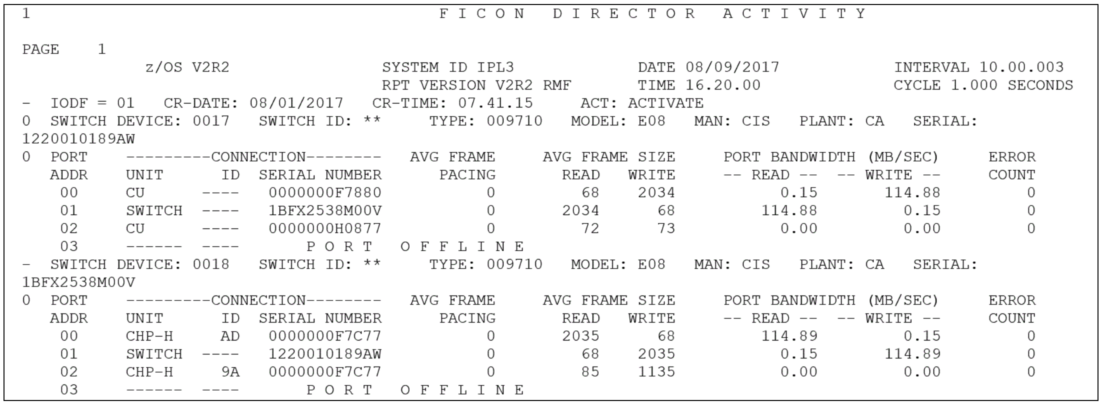

```
1                                                      F I C O N   D I R E C T O R   A C T I V I T Y

PAGE    1
                z/OS V2R2                  SYSTEM ID IPL3            DATE 08/09/2017             INTERVAL 10.00.003
                                           RPT VERSION V2R2 RMF      TIME 16.20.00               CYCLE 1.000 SECONDS
-  IODF = 01   CR-DATE: 08/01/2017   CR-TIME: 07.41.15     ACT: ACTIVATE
0  SWITCH DEVICE: 0017   SWITCH ID: **     TYPE: 009710   MODEL: E08   MAN: CIS   PLANT: CA   SERIAL:
1220010189AW
0  PORT    ---------CONNECTION--------   AVG FRAME     AVG FRAME SIZE     PORT BANDWIDTH (MB/SEC)     ERROR
   ADDR    UNIT       ID  SERIAL NUMBER     PACING      READ   WRITE     -- READ --  -- WRITE --     COUNT
    00     CU       ----   0000000F7880          0        68    2034          0.15       114.88          0
    01     SWITCH   ----   1BFX2538M00V          0      2034      68        114.88         0.15          0
    02     CU       ----   0000000H0877          0        72      73          0.00         0.00          0
    03     ------   ----        P O R T   O F F L I N E
-  SWITCH DEVICE: 0018   SWITCH ID: **     TYPE: 009710   MODEL: E08   MAN: CIS   PLANT: CA   SERIAL:
1BFX2538M00V
0  PORT    ---------CONNECTION--------   AVG FRAME     AVG FRAME SIZE     PORT BANDWIDTH (MB/SEC)     ERROR
   ADDR    UNIT       ID  SERIAL NUMBER     PACING      READ   WRITE     -- READ --  -- WRITE --     COUNT
    00     CHP-H      AD   0000000F7C77          0      2035      68        114.89         0.15          0
    01     SWITCH   ----   1220010189AW          0        68    2035          0.15       114.89          0
    02     CHP-H      9A   0000000F7C77          0        85    1135          0.00         0.00          0
    03     ------   ----        P O R T   O F F L I N E
```

Figure 8-36 Sample FICON Director Activity Report with two ISL connected switches

The report can be used to develop a diagram of connectivity and data flow, as shown in Figure 8-37.

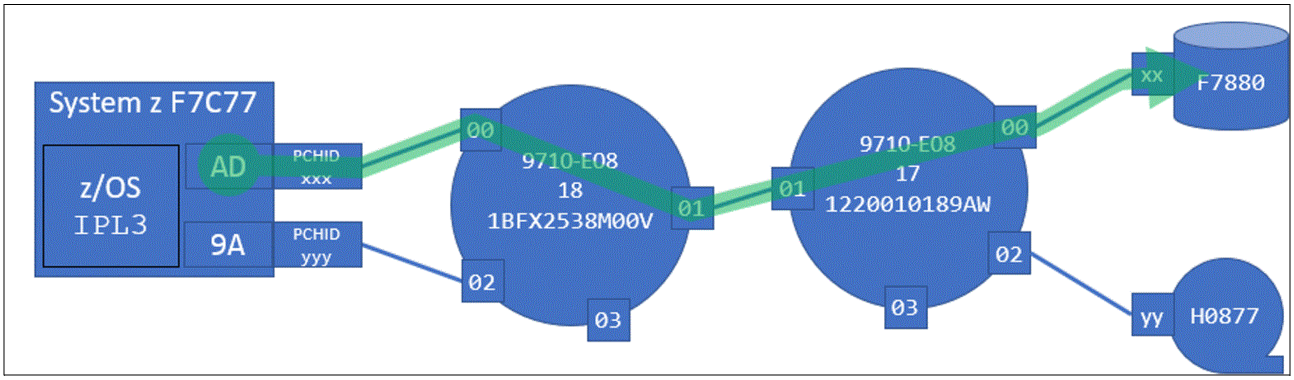

Figure 8-37 Diagram of I/O flow based on FICON Director Activity Report

The read/write information is from the FICON Director perspective. For example, Switch 18 port 00 is reading at 114.89 MBps, which means that Channel AD is writing. Switch 18 port 01 is an ISL (the unit is `SWITCH`) and is writing 114.89 MBps to switch 17 port 01, which is reading 114.88 MBps. Switch 17 port 00 is writing 114.88 MBps to the device, which is reading.

You can get a good idea about how your I/O flows are performing by examining several RMF or CMF reports:

- FICON Directory Activity Report.
- Device Activity Report
- I/O Queueing Report

Look for port utilization (port bandwidth divided by link speed) and I/O frame pacing, which can indicate that the port is overutilized. Adding paths to the CU, if possible, can help. Errors might indicate a physical problem. Long I/O queues indicate a heavily loaded device that might need more paths. For more information, see the following resources:

- *z/OS Version 2 Release 3MVS System Management Facility (SMF)*, SA38-0667-30
- *z/OS Version 2 Release 3Resource Measurement Facility User's Guide*, SC34-2664-30
- *z/OS Version 2 Release 3Resource Measurement Facility Report Analysis*, SC34-2665-30
- *z/OS Version 2 Release 3 MVS Initialization and Tuning Guide*, SA23-1379-30

- *z/OS Version 2 Release 3 MVS Initialization and Tuning Reference*, SA23-1380-30
- Using CMF Monitor 6.2
- *ABCs of IBM z/OS System Programming Volume 1*, SG24-6981

Abbreviations and acronyms

AAA	authentication, authorization, and accounting
ACL	access control list
AES	Advanced Encryption Standard
AES-GCM	AES-Galois Counter Mode
AES-GMAC	AES-Galois Message Authentication Code
AL_PA	arbitrated loop physical address
ALA	Arbitrated Loop Address or Arbitrated Loop (Station) Address
B2B	buffer-to-buffer
CCW	Channel Command Word
CFM	cubic feet per minute
CFS	Cisco Fabric Services
CHPID	Channel Path ID
CHS	Cylinder, Head, Sector
CKD	Count Key Data
CLI	command-line interface
CMF	Comprehensive Management Facility
CMR	Command Response
CMT	CHPID Mapping Tool
CP	central processor
CPC	Central Processor Complex
CRC	cyclic redundancy check
CSS	channel subsystem
CSW	Channel Status Word
CTC	channel-to-channel
CU	control unit
CUADDR	CU address
CUP	IBM Control Unit Port
CWDM	Coarse Wavelength-Division Multiplexing
DASD	Direct Access Storage Device
DCM	Dynamic Channel Path Management
DCNM	Data Center Network Manager
DH-CHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
DID	destination ID
DM	Device Manager
DoD	Department of Defense
DR	disaster recovery
DWDM	Dense Wavelength-Division Multiplexing
E_Port	expansion port
EBCDIC	Extended Binary-Coded Decimal Interchange Code
ECKD	Extended Count Key Data
ECMP	Equal Cost Multipath
EDIF	Encryption for Data in Flight
EFMD	Exchange Fabric Membership Data
EIA	Electronic Industries Alliance
ELW SFP	Extended Long Wave Small Form-factor Pluggable
EMIF	ESCON Multiple Image Facility
ESCON	Enterprise I/O System Connection
ESP	Encapsulating Security Payload
F_Port	fabric port
FBA	Fixed-Block Architecture
FC	Fibre Channel
FC-FS	Fibre Channel Framing and Signaling
FC-SB	Fibre Channel Single Byte
FC-SP	Fibre Channel Security Protocol
FCES	Fibre Channel Endpoint Security
FCIP	Fibre Channel over IP
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FCS	Fabric Configuration Server
FCTC	FICON Channel-to-Channel
FEC	Forward Error Correction
FICON	Fibre Connection
FIDR	FICON Dynamic Routing
FSPF	Fabric Shortest Path First
GbE	gigabit Ethernet
GC	Global Copy
GDPS	Geographically Distributed Parallel Sysplex
GM	Global Mirror
GMT	Greenwich Mean Time
GOLD	Generic Online Diagnostics
GPL	GNU General Public License

HA	high availability or highly available
HBA	host bus adapter
HCD	Hardware Configuration Definition
HSA	Hardware Storage Area
IBM	International Business Machines Corporation
IC	Internal Coupling
IDID	insistent domain ID
IETF	Internet Engineering Task Force
IFCC	interface control check
IKE	Internet Key Exchange
IML	Initial Machine Load
INCITS	International Committee for Information Technology Standards
IOA	I/O Acceleration
IOCDS	Input/Output Configuration Data Set
IOCP	Input/Output Configuration Program
IOD	In Order Delivery
IODF	input/output definition file
IOS	I/O subsystem
IPFM	IP Fabric for Media
IPL	Initial Program Load
IPS	IP Storage
IPsec	IP Security
IPv6	IP Version 6
iSCSI	internet Small Computer Systems Interface
ISL	Inter-Switch Link
iSNS	Internet Storage Name Server
ISSD	in-service software downgrade
ISSU	in-service software upgrade
IU	Information Unit
IVR	Inter-VSAN Routing
JRE	Java Runtime Engine
LAN	local area network
LCP	Logical Corruption Protection
LDAP	Lightweight Directory Access Protocol
LFM	linear feet per minute
LGPL	Lesser General Public License
LIOD	Lossless In Order Delivery
LPAR	logical partition
LSR	Link State Record
LUN	logical unit number
LWL	long wavelength
MAN	metropolitan area network
MDS	Multilayer Director Switch
MES	Miscellaneous Equipment Specification
MGM	Metro/Global Mirror
MIF ID	Multiple Image Facility Identifier
MM	Metro Mirror
MOTD	message of the day
MPIO	multipath I/O
MTP	Multi-fiber Termination Push-on
MTU	maximum transmission unit
MVS	Multiple Virtual Storage
MVSCP	MVS Configuration Program
NEBS	Network Equipment Building Standards
NPIV	N_Port ID Virtualization
NPU	Network Processing Unit
NTP	Network Time Protocol
NVMe	Non-Volatile Memory Express
NVMe-FC	Non-Volatile Memory Express over Fibre Channel
NVRAM	non-volatile random-access memory
nWWN	node worldwide name
OBFL	On-board Failure Logging
OS	operating system
OSA	Open Systems Adapter
OVA	Open Virtual Appliance
OXID	originator exchange ID
PAK	Product Activation Key
PCHID	Physical Channel ID
PCIe	PCI Express
PFS	perfect forward secrecy
PMON	Port Monitor
PMR	Problem Management Record
POAP	Power On Auto Provisioning
POR	power-on reset
POST	power-on self-test
PPRC	Peer-to-Peer Remote Copy
PSU	power supply unit
PU	processor unit
pWWN	port worldwide name
QoS	quality of service
QSA	Query Security Attributes

RAC	Real Application Cluster
RADIUS	Remote Authentication Dial-In User Service
RBAC	role-based access control
RDBMS	Relational Database Management System
RDP	Read Diagnostic Parameters
RLIR	Registered Link Incident Report
RMF	Resource Measurement Facility
RNID	Request Node Identification
RPQ	Request for Price Quotation
RSCN	Registered State Change Notification
RSPAN	Remote Switched Port Analyzer
RTP	Real-Time Protocol
RTT	round-trip time
Rx	receive
SA	security association
SAD	security association database
SAN	storage area network
SBCCS	Single-Byte Command Code Sets
SDM	Sysplex Data Mover
SE	Service Element
SFP	Small Form-factor Pluggable
SFTP	Secure File Transfer Protocol
SID	source ID
SIO	Start I/O
SIST	single initiator single target
SLA	service-level agreement
SMC-R	Shared Memory Communications over RMDA
SME	subject matter expert
SMF	System Management Facility
SMI-S	Storage Management Initiative Specification
SNMPv3	Simple Network Management Protocol Version 3
SPAN	Switched Port Analyzer
SPD	Security Policy Database
SPI	security parameter index
SPOF	single point of failure
SRAM	static random-access memory
SRDF	Symmetrix Remote Data Facility
SSCH	Start Subchannel
SSD	solid-state drive
SSIC	IBM System Storage Interoperation Center
STP	Server Time Protocol
SUDI	secure unique device identification
SVC	SAN Volume Controller
SWL	short wavelength
sWWN	switch worldwide name
TACACS	Terminal Access Controller Access Control System
TAM	Trust Anchor Module
TCAM	Ternary Content Addressable Memory
TCO	total cost of ownership
TIC	Transfer In Channel
Tx	transmit
ULP	upper layer protocol
USB	Universal Serial Bus
UTC	Universal Coordinated Time
VM	virtual machine
VMID	virtual machine ID
VOQ	virtual output queue
VPN	virtual private network
VRRP	Virtual Routing Redundancy Protocol
VSAN	virtual storage area network
VTL	virtual tape library
VTS	virtual tape server
WAN	wide area network
WWN	worldwide name
XRC	eXtended Remote Copy
zDAC	System z Dynamic Auto Discovery
zHPF	High-Performance FICON for System z

Redbooks

IBM Storage Networking c-type FICON Implementation Guide

SG24-8468-00

ISBN 0738460214

IBM

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

SG24-8468-00

ISBN 0738460214

Printed in U.S.A.