Redbooks
ibm.com/redbooks

Configuring Fibre Channel Protocol on IBM LinuxONE and IBM Z to IBM Storage FlashSystem

Jay Brenneman

IBM LinuxONE

Storage

IBM

Redpaper

Introduction

IBM® LinuxONE and Z platforms offer support for open systems Fibre Channel Protocol (FCP) networks and subsystems, enabling their use for block storage provisioning for Linux and select z/VM® workloads. This IBM Redpaper provides a comprehensive guide to the hardware configuration and enablement of FCP for Linux environments.

This document is intended for IBM Z® and LinuxONE system administrators, storage administrators, infrastructure architects, and implementers.

FCP overview

Fibre Channel Protocol (FCP) is a modern storage fabric interconnect that provides transport from Server host systems to their Storage devices. IBM FICON® Is an extension of IBM's previous ESCON interconnect protocol implemented as an additional layer on top of FCP infrastructure.

IBM's implementation of the Fibre Channel Protocol (FCP) on IBM LinuxONE and IBM Z platforms provides Linux and z/VM environments with access to a comprehensive portfolio of storage technologies from diverse vendors, capitalizing on the inherent flexibility and robust connectivity options characteristic of the FCP ecosystem.

FCP allows you to build a network of FCP switches, hosts, and storage controllers that connect to each other and create a Storage Area Network (SAN). See Figure 1 on page 2.

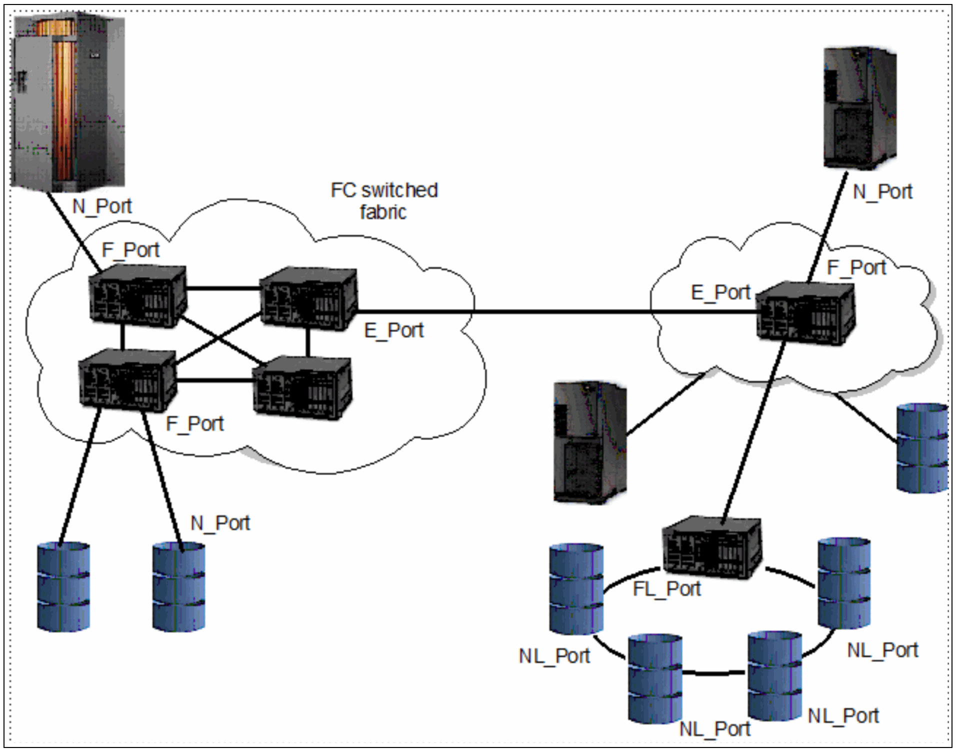

Figure 1 A typical SAN

- N Ports connect host and storage devices to the SAN.
- F Ports connect the SAN to host and storage devices.
- E Ports connect SAN switches to other SAN switches.
- L_Ports represent Arbitrated Loop ports and encompass both N_Port and F_Port functionalities. These port types are relatively uncommon in modern Fibre Channel deployments.

Devices in the network are identified by World Wide Node Names (WWNN), and individual ports within devices are identified by World Wide Port Names (WWPN). See Figure 2 on page 3.

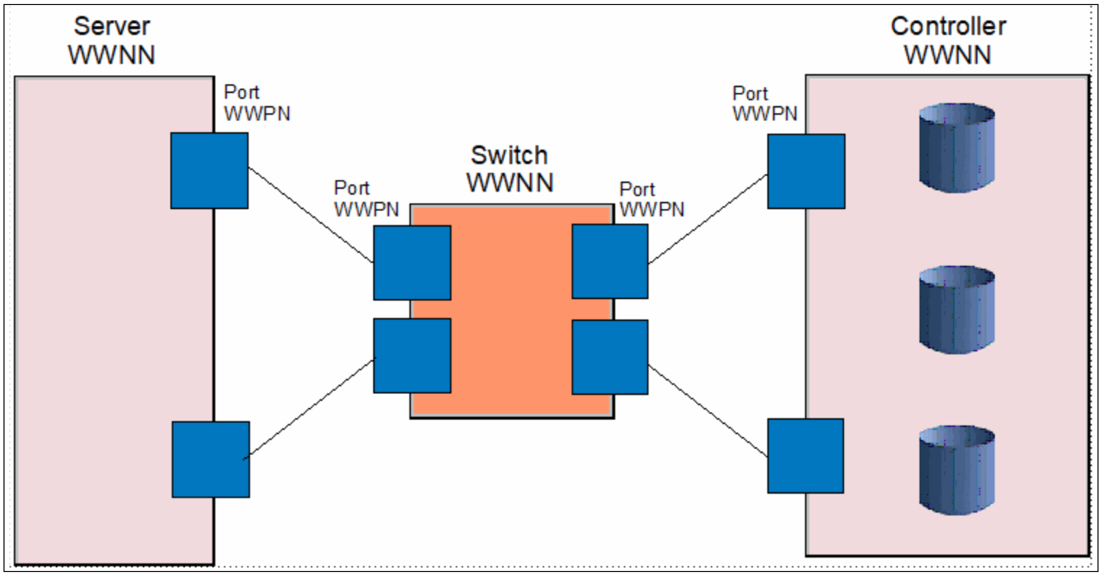

Figure 2 Example of worldwide names

All ports - host ports, storage ports, and switch ports - have a globally unique WWPN. There are rules about how WWPNs are generated similar to Ethernet MAC addresses to prevent name collisions across device vendors.

Access control is managed by lists of WWPNs that are permitted to communicate both in the Switch and in the Storage controller.

In the switch the access list is called a zone. A zone is a list of the WWPNs of ports which are allowed to talk to each other.

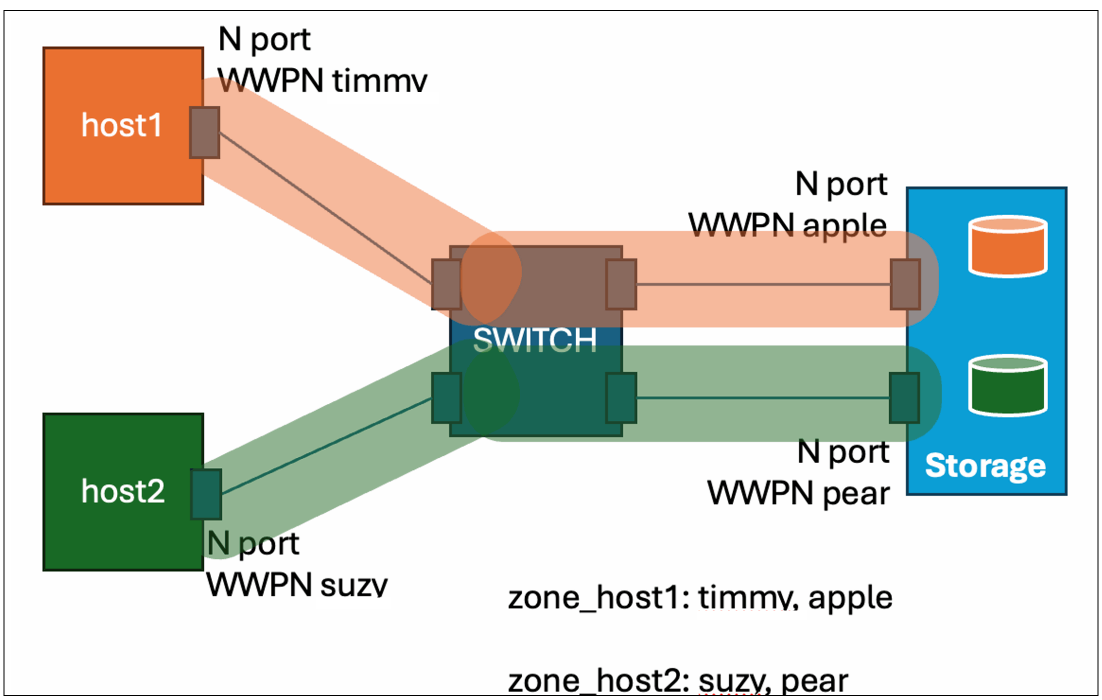

Figure 3 Zones

In Figure 3 on page 3 there are 2 zones. One zone is named zone_host1, and it permits WWPN timmy to talk to WWPN apple and vice versa. The other zone is named zone_host2, and it permits WWPN suzy to talk to WWPN pear, and vice versa. The switch will not permit any WWPN to converse with or even know the existence of a WWPN that it is not zoned with. The physical ports on the switch where host and storage fibers are connected do not affect zoning. However, these connections significantly impact availability and backplane bandwidth.

Zoning is a fabric level construct and extends across all switches in a multi switch fabric, so if SWITCH is actually a network like SWITCH-to-SWITCH-to-SWITCH the links in the middle don't really matter as far as the zones are concerned. Zones are only concerned with the endpoints.

Storage controllers present their virtual storage volumes to the network as logical units, and are accessed from the host via their logical unit number (LUN). See Figure 4.

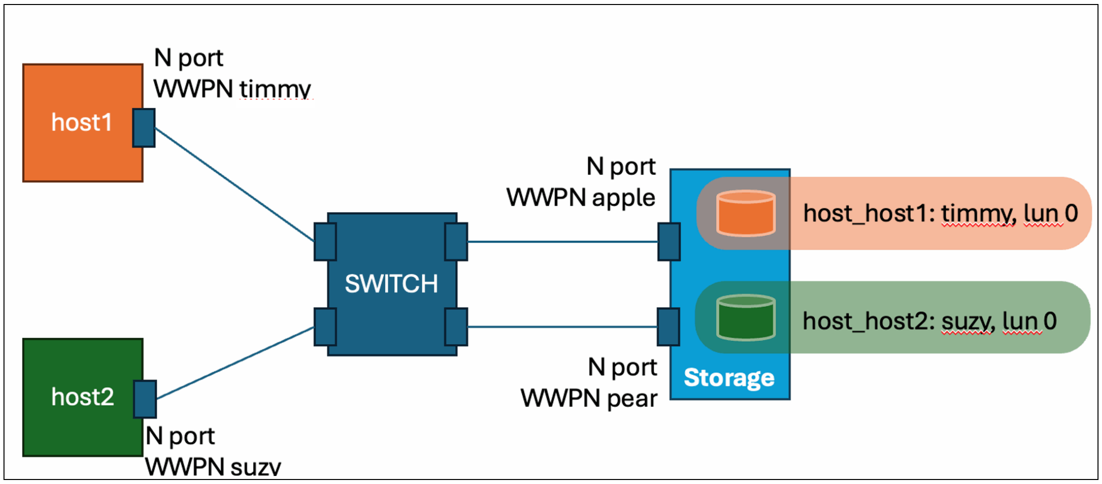

Figure 4 LUNs

In the storage controller the access list is called a host definition with an associated LUN mapping. It is a list of WWPNs associated with a specific server or virtual machine and the list of virtual volumes in the storage controller that the system is permitted to access. In the above diagram there is a host definition named host_host1 with the WWPN timmy listed, and a LUN map that associates a volume with that host as LUN 0. There is another host definition named host_host2 which the WWPN suzy listed, and it maps a different volume to that host, but also as LUN 0.

Both hosts have a LUN 0 in the storage controller, but the volumes accessed at LUN 0 are different.

A host definition can also be a cluster of systems with multiple WWPNs from all cluster members. This allows you to make sure that the LUN numbers are consistent across all the cluster members if that's important to you.

Some storage controllers enumerate the virtual volumes as they map them to a defined host starting at LUN number 0 and counting up, such as the IBM FlashSystem® devices.

Other storage controllers enumerate the virtual volumes using a fixed virtual controller and virtual device number that is intrinsic to the virtual volume itself as they map the volumes to a defined host. The IBM DS8000® class of systems do this.

Tip: This difference has some implications for configuration. IBM FlashSystem allows administrators to easily assign LUN 0 to specific devices, like boot volumes. IBM DS8000 systems, however, do not use LUN 0; they assign identifiers, such as `40104022` (indicating the 22nd virtual volume on logical controller 10).

FCP versus FICON

For those familiar with FICON, there are key differences compared to FCP storage configurations. In FICON, the host's I/O configuration describes the entire storage network, from the host through the switches to the storage systems. The network topology is defined within this I/O configuration, and defined control units (CUs) and devices directly correspond to logical control units and devices on the storage controller. For example, a CU 7000 definition on CHPIDs 40 and 41 with devices 7000-700F maps directly to a specific logical control unit on an ECKD storage server with 16 corresponding devices. Device address 7000 typically represents the same 3390 storage volume across all LPARs with access.

Additionally, the path from the Host port to the storage device is defined in the IO configuration using the physical port numbers of the switch. The FICON IO config reads like a roadmap describing how to get to the storage: Take CHPID 20 to Switch 33, Leave Switch 33 via port 09 to get to the DS8000 and then take the 9th right from the roundabout to get to LCU 9 in the DS8000. This means that if you move a link between the storage controller and the switch to a different port on the switch, that whole defined link from the Z system to the storage controller goes offline because the exit port from the switch is now unplugged.

FCP functions differently. FCP control units and devices on the host behave more like OSA devices than physical disks. From the host's perspective, the FCP device is simply an interface used by the zfcp driver for network communication. The SAN switches and storage controllers manage all network routing and storage interactions. Thus, even if FCP devices in different LPARs (for example, device 7000 in LPAR 1 and device 7000 in LPAR 2) share a physical port in the I/O configuration, they are treated as distinct entities.

N Port ID virtualization

In a SAN fabric, N_Ports designate the host or storage endpoints. If a single physical port is used by multiple logical entities with different needs, the SAN requires a mechanism to distinguish between these entities to ensure proper zoning and LUN mapping.

N_Port ID Virtualization (NPIV) implements a virtualization technique that generates distinct virtual World Wide Port Names (WWPNs) for each logical entity utilizing a shared physical SAN port. This virtualization is essential for ensuring the proper operation of zoning and storage masking mechanisms.

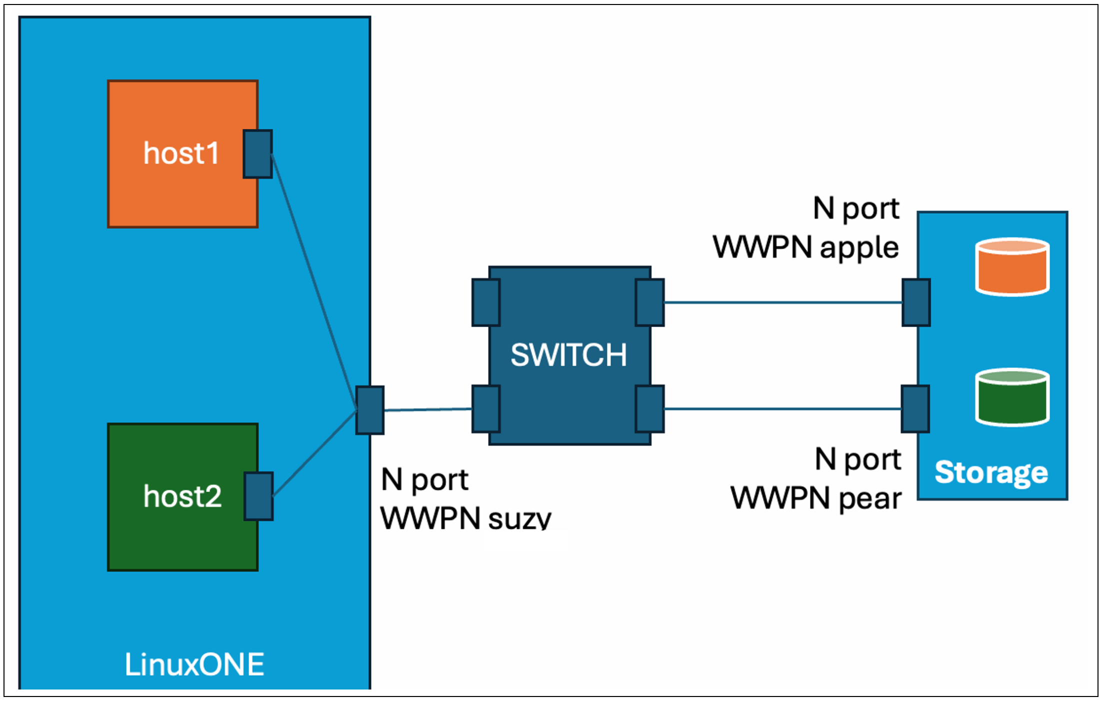

Figure 5 LPARs on a LinuxONE machine, sharing access to a FCP Channel without NPIV active

In Figure 5 host1 and host2 are LPARs on a LinuxONE machine, sharing access to a FCP Channel without NPIV active. Without NPIV, there is just the physical WWPN of the port, which host1 and host2 are sharing. The switch cannot tell the difference between traffic from host1 and host2 since it all comes from WWPN suzy, so all users of that port will be able to access everything that WWPN suzy is zoned to.

The same problem happens in the Storage controller. The Storage controller only sees IOs coming from WWPN suzy, and so they can get to whatever resources are mapped to suzy.

This is not unique to IBM Z and LinuxONE systems, this problem applies to basically any hypervisor platform.

Once you activate NPIV on FCP ports, you have the ability to generate unique WWPNs for each LPAR's devices sharing the physical port.

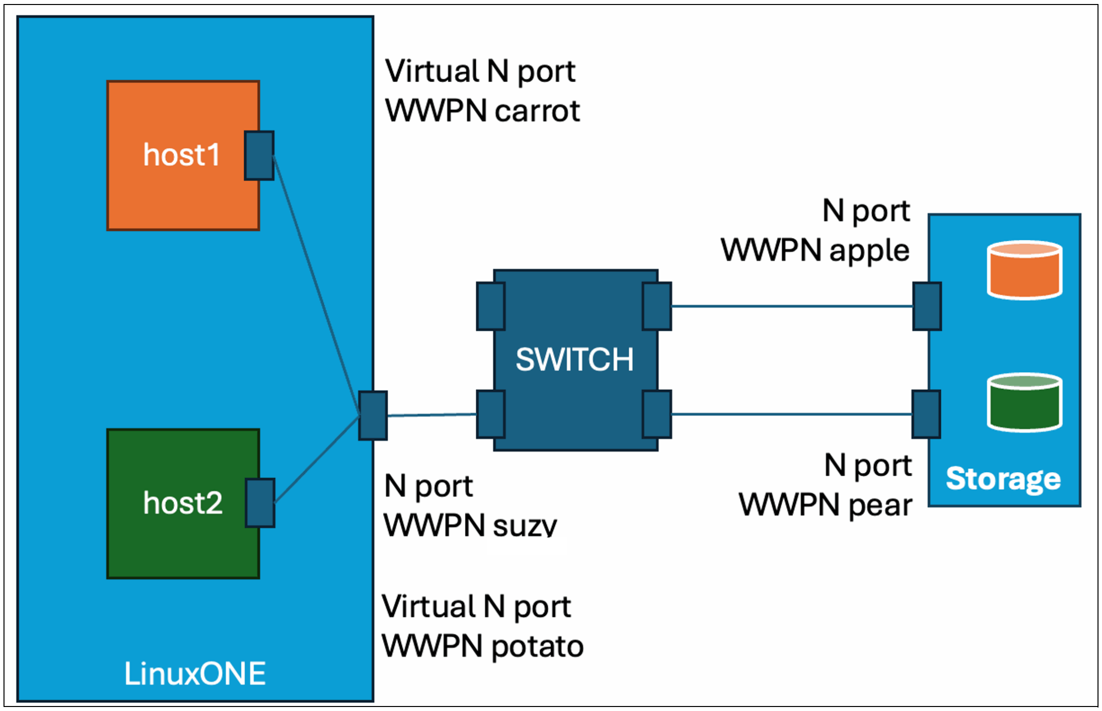

Figure 6 NPIV is activated for the host1 and host2 LPARs

In Figure 6 we have activated NPIV on that port for the host1 and host2 LPARs. Now IOs coming from host1 will be seen by the SAN as coming from WWPN carrot, and IOs coming from host2 will be seen as having come from WWPN potato. The physical port WWPN is still suzy, but it is no longer used for zoning or LUN mapping.

With N_Port ID Virtualization (NPIV) implemented, the FCP logical devices presented to LPAR host1 and LPAR host2 are treated as distinct entities by the SAN fabric. This necessitates discrete zoning configurations within the SAN switches and the creation of individual host definitions on the storage controller to facilitate access to their designated logical unit numbers (LUNs).

Supported storage services

The IBM Storage dev/test organization is responsible for testing and determining what IBM Storage products are supported for attachment to LinuxONE and Z systems.

The IBM System Storage Interoperation Center (SSIC), provided by IBM Storage, helps determine which storage devices are supported for connection to different server platforms, such as IBM LinuxONE and Z. The SSIC also provides information on compatible switch interconnections and tested firmware bundles.

Storage vendors are responsible for testing and confirming compatibility between their products and IBM LinuxONE or Z systems. To determine if a specific storage product is supported, consult the storage vendor.

Defining FCP devices on IBM Z and LinuxONE

Example 1 shows a simple example of defining FCP devices in IOCP that are shared across all LPARs on a LinuxONE Emperor system with 6 Logical Channel Subsystems running in classic (not DPM) mode.

Example 1   Defining FCP devices in IOCP

```
CHPID PATH=(CSS(0,1,2,3,4,5),40),SHARED,PCHID=180,TYPE=FCP
CHPID PATH=(CSS(0,1,2,3,4,5),45),SHARED,PCHID=181,TYPE=FCP
CHPID PATH=(CSS(0,1,2,3,4,5),46),SHARED,PCHID=100,TYPE=FCP
CHPID PATH=(CSS(0,1,2,3,4,5),4C),SHARED,PCHID=101,TYPE=FCP

CNTLUNIT CUNUMBR=A002,                                          *
      PATH=((CSS(0),40),(CSS(1),40),(CSS(2),40),(CSS(3),40),(C*
      SS(4),40),(CSS(5),40)),UNIT=FCP
IODEVICE ADDRESS=(A000,003),CUNUMBR=(A002),UNIT=FCP

CNTLUNIT CUNUMBR=A102,                                          *
      PATH=((CSS(0),46),(CSS(1),46),(CSS(2),46),(CSS(3),46),(C*
      SS(4),46),(CSS(5),46)),UNIT=FCP
IODEVICE ADDRESS=(A100,003),CUNUMBR=(A102),UNIT=FCP

CNTLUNIT CUNUMBR=A202,                                          *
      PATH=((CSS(0),4C),(CSS(1),4C),(CSS(2),4C),(CSS(3),4C),(C*
      SS(4),4C),(CSS(5),4C)),UNIT=FCP
IODEVICE ADDRESS=(A200,003),CUNUMBR=(A202),UNIT=FCP

CNTLUNIT CUNUMBR=A302,                                          *
      PATH=((CSS(0),45),(CSS(1),45),(CSS(2),45),(CSS(3),45),(C*
      SS(4),45),(CSS(5),45)),UNIT=FCP
IODEVICE ADDRESS=(A300,003),CUNUMBR=(A302),UNIT=FCP
```

The following provides a more detailed explanation:

- PCHID 180 is shared by all LPARs as CHPID 40.
- CHPID 40 connects to CU A002, which provides devices A000-A002.
- PCHID 181 is shared by all LPARs as CHPID 45.
- CHPID 45 connects to CU A302, which provides devices A300-A302.
- PCHID 100 is shared by all LPARs as CHPID 46.
- CHPID 46 connects to CU A102, which provides devices A100-A102.
- PCHID 101 is shared by all LPARs as CHPID 4C.
- CHPID 4C connects to CU A202, which provides devices A200-A202.

For example, while device A000 is visible as an FCP device on all LPARs, A000 on LPAR LP01 is distinct from A000 on LPAR LP17. FCP devices are not DASD volumes; they simply provide a network interface for SAN communication.

Consult IBM Publication Input/Output Configuration Program User's Guide for ICP IOCP, SB10-7177-02 and scroll to Appendix F "*Machine Limits and Rules*" and look for the table named "*Constant Machine Limits*". In that table is listed a value for "`Maximum devices or subchannels per FCP channel path`", which for our test system is 480. That means we can configure up to 480 FCP devices on a single FCP Channel.

An example configuration

Consider the I/O configuration and the "Maximum devices or subchannels per FCP channel path" limit of 480 for channel 180. A common misconception is that three devices are configured. However, the actual number depends on the number of LPARs utilizing the channel. Given a maximum of 85 LPARs for this machine class, the upper limit is 255 (3 devices/LPAR * 85 LPARs).

However, a note in the table states: "If an FCP channel path is dedicated to a logical partition, the maximum number of devices is 254." As our channel is not dedicated, this limitation does not apply.

Machine-specific limits, detailed in the "Machine Limits (by Machine Type)" table, must also be considered. See Figure 7.

Machine Limit	Machine type (Note "1" on page 315)											
	2817,1	2818,1	2827,1	2828,1	2964,1	2965,1	3906,1	3907,1	8561,1	8562,1	3931,1	3932,1
Maximum FCP channel paths	256	160	256	160	256	160	256	128	384	128	384	96

Figure 7 Machine limits (by Machine type)

As shown in Figure 7, the maximum number of configurable FCP ports varies by machine type. Our example configuration uses only four ports, which is within the limits for our machine type.

Keep scrolling till you find the table shown in Figure 8.

Table 21. FCP Channel Path Limits

Limit or Recommendation	FICON Express 8/8S	FICON Express 16S/16S+/16SA/32SA/32S
Maximum configurable devices or subchannels per FCP channel path	480	480
Recommended maximum number of active NPIV-enabled subchannels per FCP channel path	32	64 or 128 (Note "11" on page 316)
Maximum number of remote N_Ports per FCP channel path	512	1024
Maximum concurrent I/O operations per FCP channel path	960	1528
Maximum number of LUNs that can be open at the same time	4096	8192

Figure 8 FCP Channel Path Limits table

Note 11 states: "A maximum of 128 NPIV-enabled subchannels are supported on 3932 and 3931 models (with an LMC (Licensed Machine Code) patch installed) for FICON Express32S. Otherwise, a maximum of 64 NPIV-enabled subchannels are supported."

Therefore, whether our example configuration exceeds the supported limits depends on the number of LPARs sharing the devices and the specific technologies in use on both the processor and switch sides.

Regarding the processor side, the machine types are z16/LinuxONE v.4, and the adapters are those shipped with that system. Assuming a z16™ configuration with 16S+ cards and device sharing across 85 LPARs, the limit is exceeded. However, with only 21 LPARs (`21 LPARs * 3 devices/LPAR = 63 devices`), the configuration remains within the most restrictive NPIV limit for 16S+ cards.

Regarding switch support for host adapter NPIV, each vendor publishes documentation detailing the specific capabilities and limits of their devices. Some models support a maximum of 64 NPIV addresses per port, while others support up to 256. Some models may not support NPIV-enabled host ports at all. Consult the relevant switch documentation to determine the NPIV capabilities and limits of your specific switches. If a switch supports only 32 NPIV addresses per port, this becomes the effective limit for that link.

It is important to address the ambiguity surrounding the "Recommended Maximum" of 64 or 128 NPIV-enabled devices. IBM product testing has generally shown that 64 NPIV-enabled devices per channel are compatible with most director-class (high-end) switch chassis that explicitly support NPIV host ports.

While configuring more than 64 NPIV devices per FCP channel is technically possible, its suitability depends heavily on the specific workload characteristics. Defining a definitive hard limit is challenging due to the recovery processes triggered by channel failures (for example, cable or optic failures). During channel recovery, each NPIV address undergoes a re-login process to both the switch and the storage. While 64 concurrent recoveries are typically handled effectively by supported switches, 128 concurrent recoveries, although achievable, require meticulous planning of target port distribution within zones and careful consideration of the "FCP Channel Path Limits" table. Configurations exceeding 128 concurrent recoveries, even with switches claiming support for 256 NPIV addresses per port, have presented operational challenges. These challenges can manifest as incomplete fabric logins for some NPIV links, switch blade firmware crashes, or other unpredictable behavior detrimental to enterprise storage network stability.

Returning to our configuration, if we reduced it to one FCP device per channel, it could potentially be shared across all 85 potential LPARs. Further limiting sharing would definitively allow channel sharing across 64 LPARs, provided the switch ports support 64 NPIV addresses per port.

Should a configuration exceeding 64 NPIV addresses per port be considered, rigorous testing is mandatory. This testing must replicate the scale and concurrent IOPS of the production environment and include cable pull tests performed under peak workload conditions to validate the clean recovery of the switch and FCP channel NPIV configuration. Due to the numerous variables involved, IBM cannot provide a specific support statement for configurations exceeding 64 NPIV addresses per port beyond recommending thorough and comprehensive testing, including diverse failure recovery scenarios under at least 200% load, prior to production deployment.

When designing the SAN configuration, consider all relevant FCP channel limits. Exceeding these limits can result in unpredictable connectivity issues that may be difficult to diagnose.

For the purpose of this discussion, we assume that all concerns have been addressed and the FCP connectivity has been designed to provide each LPAR with the required number of FCP devices (whether zero or many) while adhering to all applicable limits.

Enabling NPIV on FCP channels

Once a valid I/O configuration is active on the system, NPIV must be enabled on the FCP channels. Unless there are specific, atypical requirements, enabling NPIV on all FCP channels is strongly recommended. For example, in IBM Internal comprehensive test configurations encompassing multiple CPCs, switches, and storage subsystems, NPIV is enabled on all but one FCP channel, which is intentionally configured without NPIV for demonstration purposes.

Should NPIV not yet be enabled for an LPAR's devices, the following procedure will activate it. The overall process involves these steps:

1. Configure the channel offline from the specific LPAR we want NPIV enabled in.
2. Configure NPIV active on that channel for that LPAR.
3. Configure the channel online to the LPAR we want NPIV enabled in.

Configure the FCP channel offline from one specific LPAR

Perform the following steps:

1. From the top HMC menu select your Central Processor Complex (CPC).

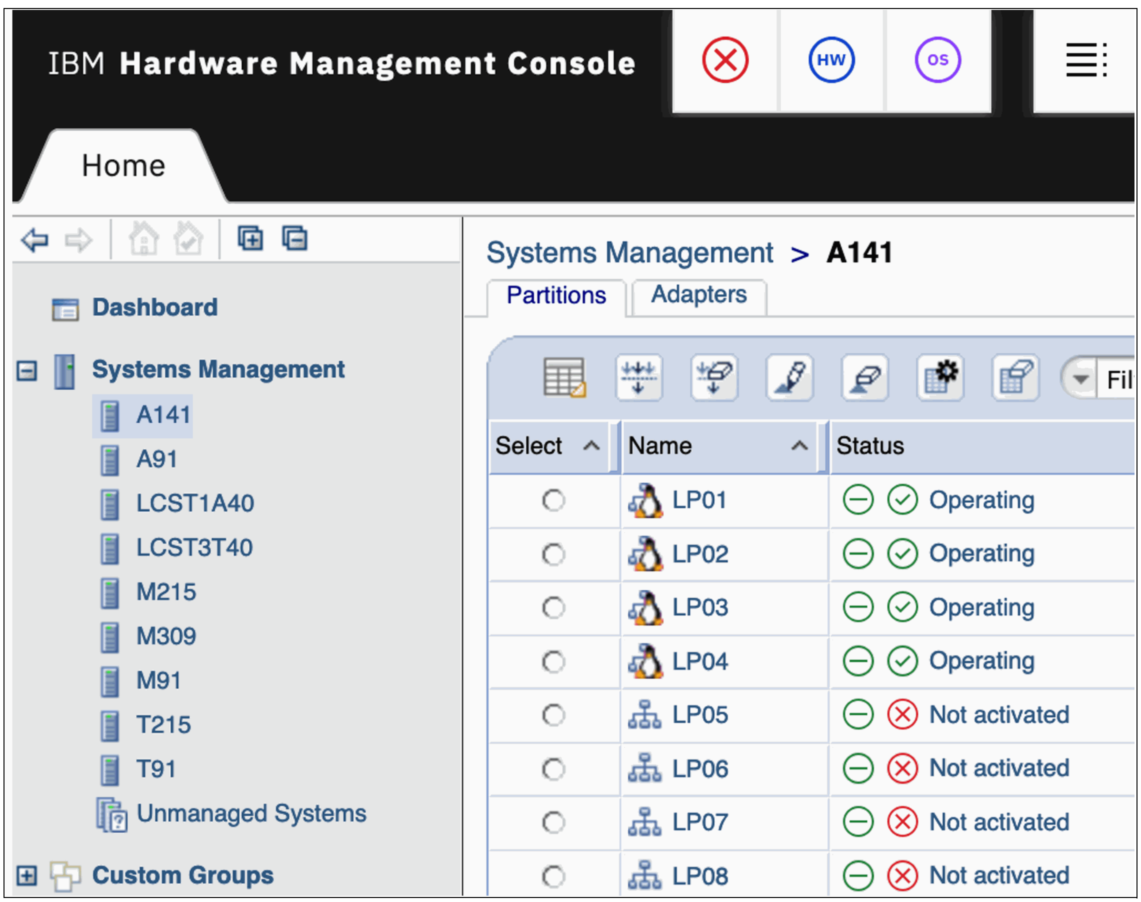

Figure 9 Select your CPC

2. Open the **Recovery** menu on the bottom center pane and select **Single Object Operations** to open the Service Element interface for the selected CPC. See Figure 10 on page 12.

Figure 10   Single Object Operations

Note: Using "Single Object Operations" is only required for z14 / LinuxONE II systems. For z15® and z16 systems you can select the **Adapters** tab on the HMC view with the CPC selected and jump to the **Select the Channel ID** step (step 3).

The Service Element window will open in a new browser window. See Figure 11.

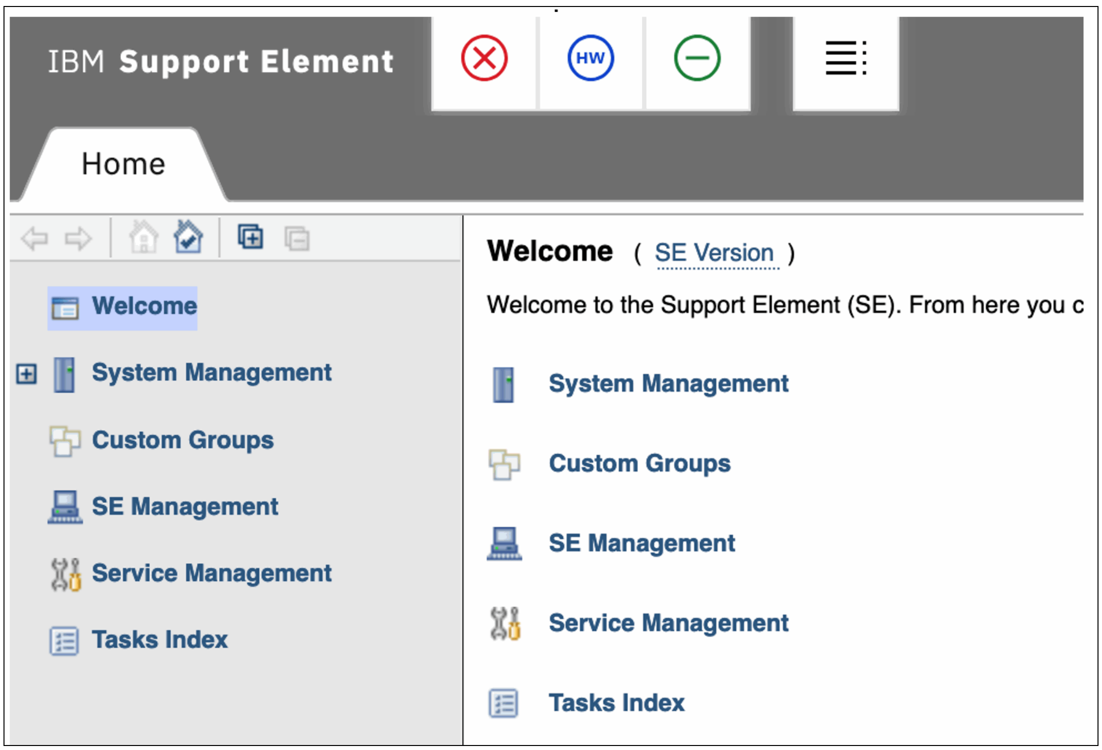

Figure 11   Service Element window

3. In the System Management menu, navigate to the **CPC** menu and select the **Channels** view. Locate the channel ID for the FCP channel requiring NPIV activation for a specific LPAR and select it. This view is similar to the "Adapters" tab on the Hardware Management Console (HMC) for z15 and z16 systems. See Figure 12 on page 13.

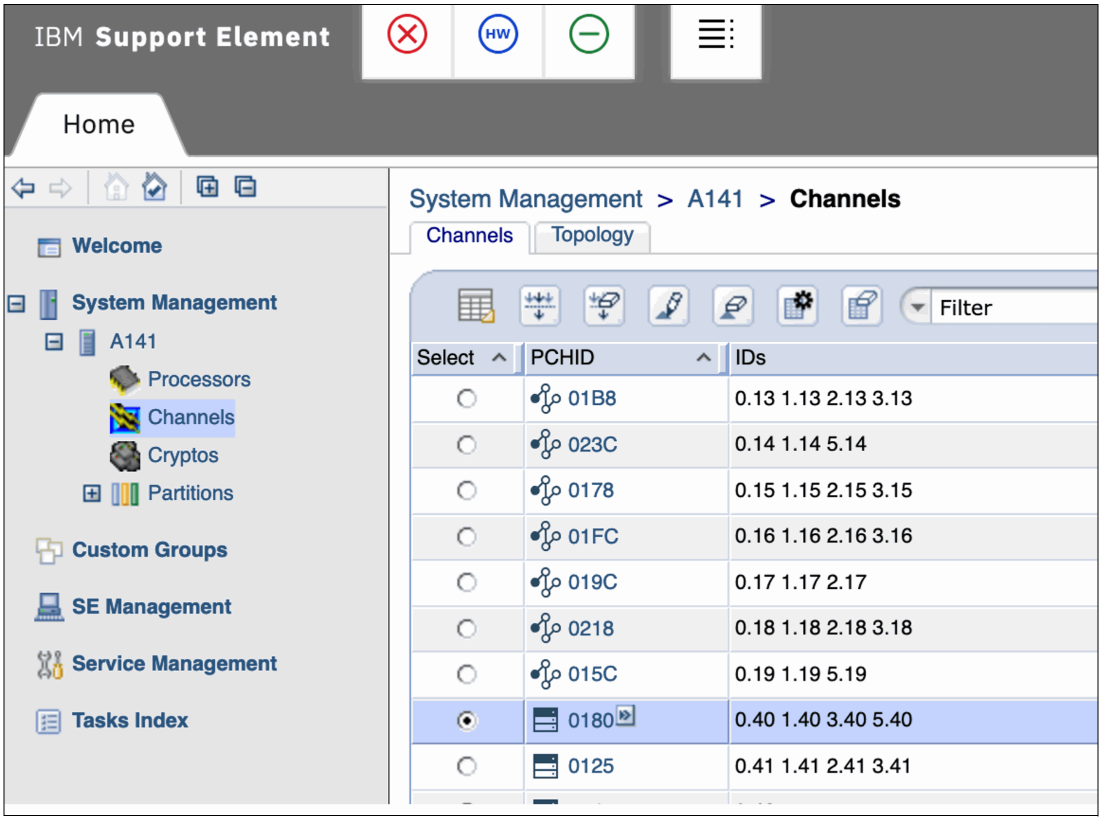

Figure 12 Channels view

4. With the Channel selected, open the **CHPID Operations** menu in the bottom center pane and select **Configure On/Off**.

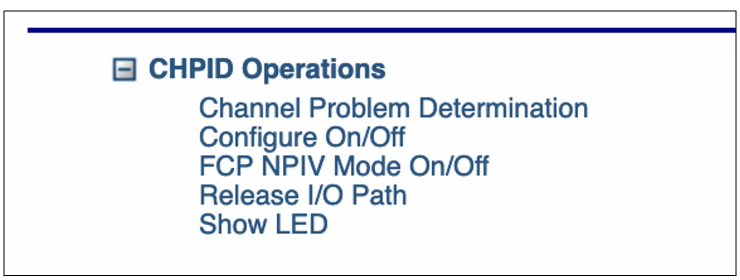

Figure 13 Select Configure On/Off

5. A listing of every LPAR that shares this channel opens. Scroll to and select the entry for the LPAR you wish to configure the channel Offline from. See Figure 14 on page 14.

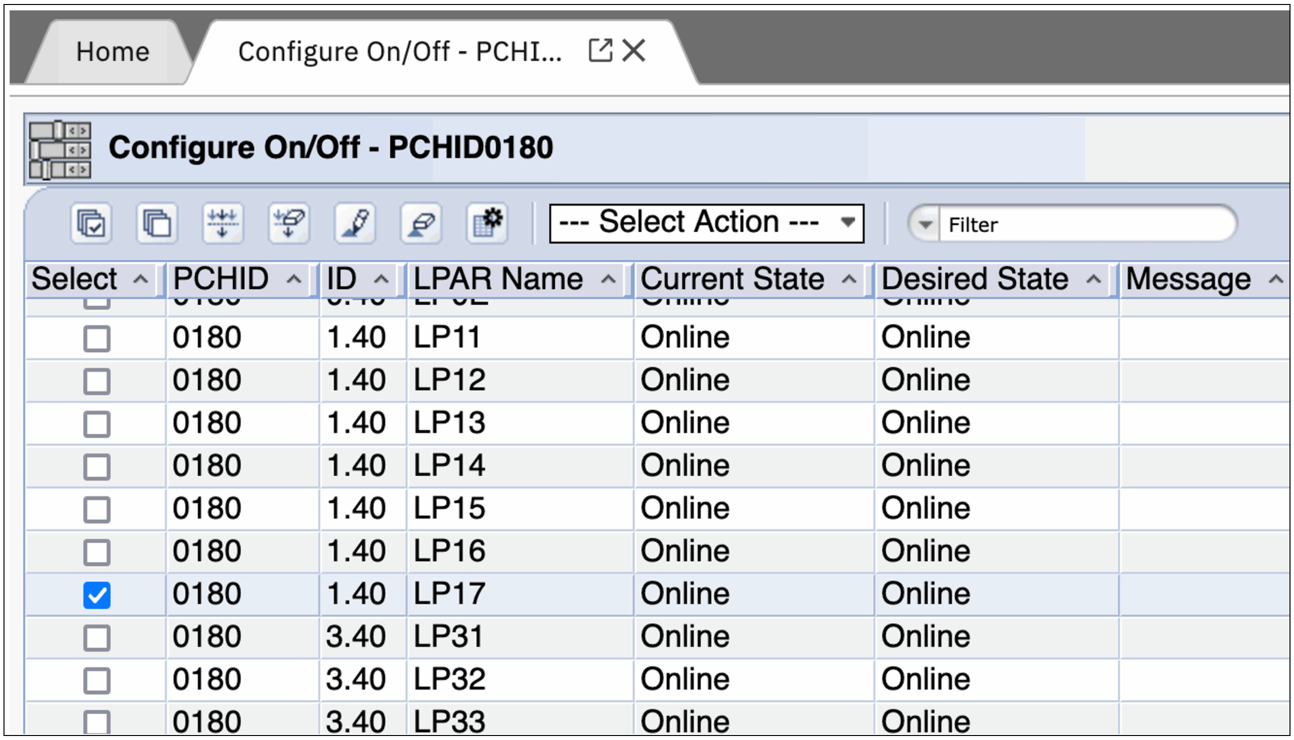

Figure 14 select the entry for the LPAR you wish to configure the channel Offline from

6. n the **Select Action** drop-down menu, select the **Toggle** option. Exercise caution and avoid the **Toggle All…** options, as these will deactivate the channel for all LPARs, which is disruptive if any LPARs are actively using it. Toggle only the CHPID associated with the selected LPAR unless it has been definitively confirmed that no LPARs are currently utilizing the channel. See Figure 15 on page 15.

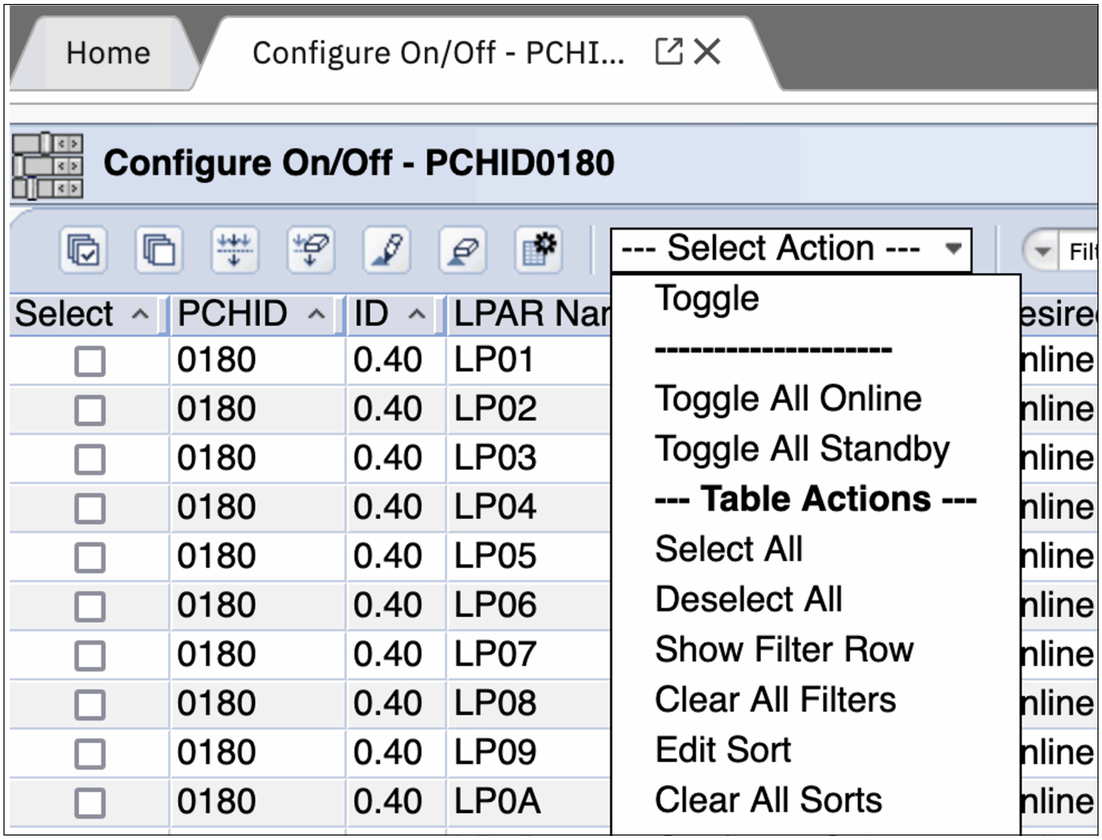

Figure 15 Select the Toggle option

7. Verify that our desired LPAR CHPID is the only one selected with a changed Desired State. Click **OK** at the bottom of the panel. See Figure 16 on page 16.

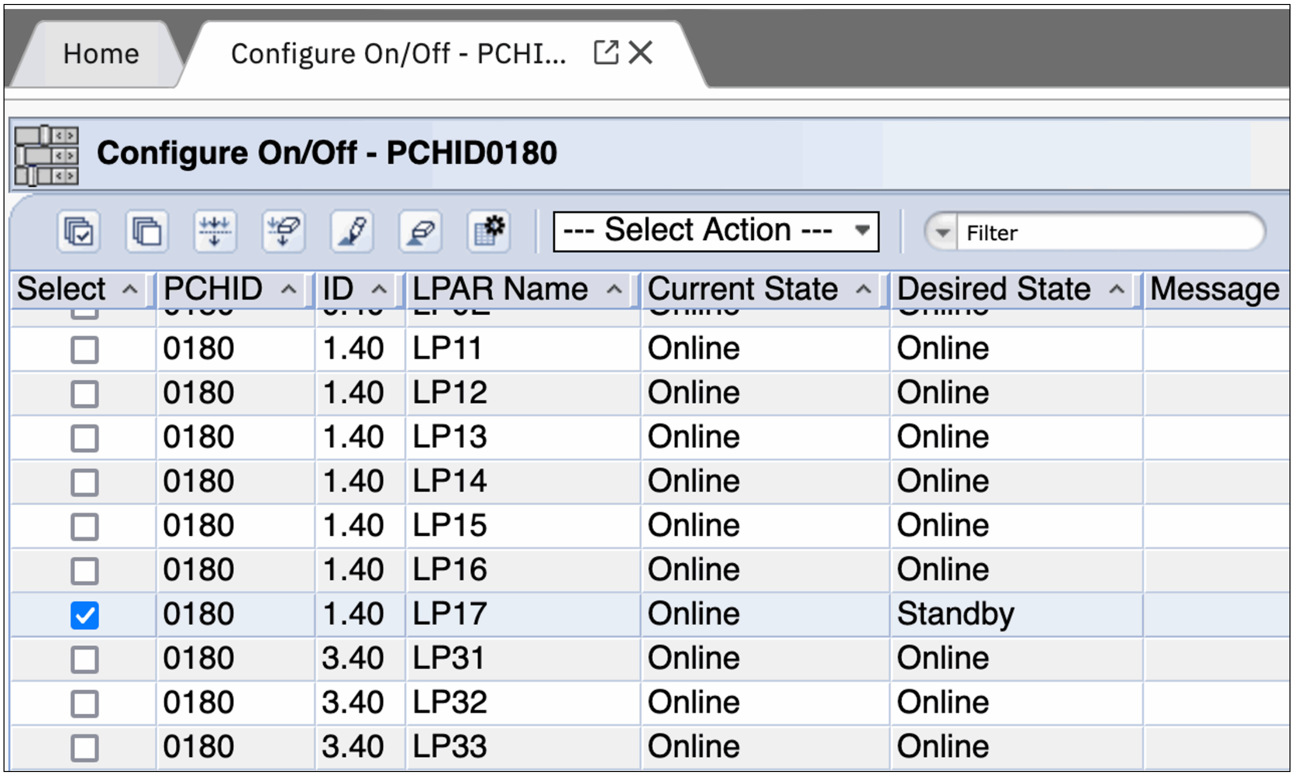

Figure 16 Verify that our desired LPAR CHPID is the only one selected with a changed Desired State

8. A confirmation pane appears. Verify that you are configuring only the intended LPARs CHPID. Then. click **Yes**. See Figure 17.

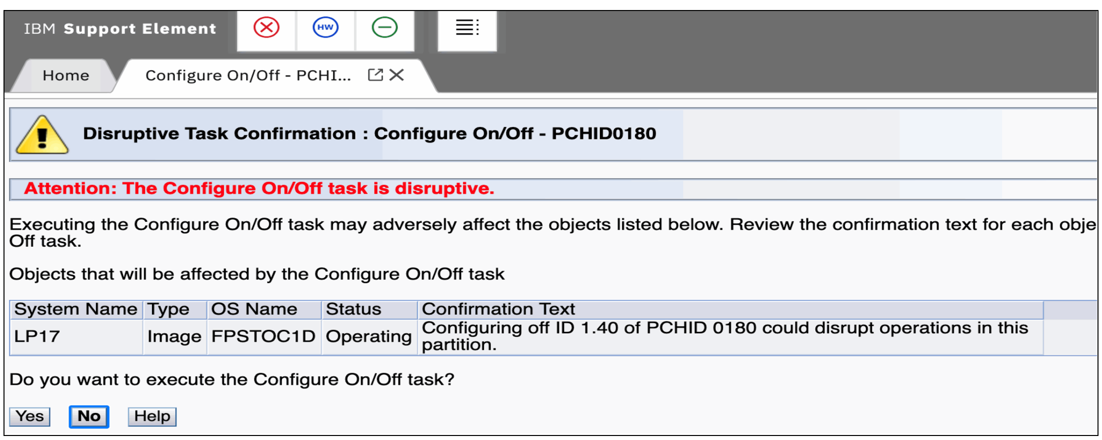

Figure 17 Confirmation pane

Configure NPIV active on that channel for that LPAR

The CHPID is now Offline for our LPAR, and now we can enable NPIV.

1. From the channels view, make sure our Channel is still selected.

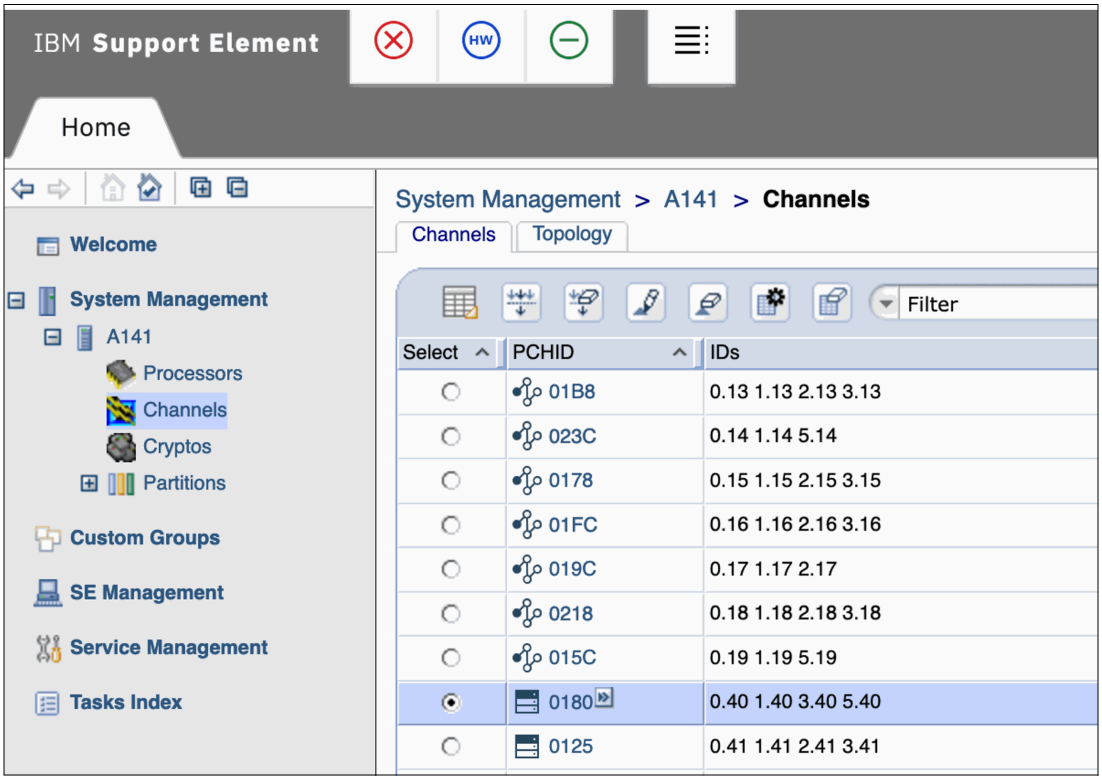

Figure 18 IBM Support Element menu

2. Select the **FCP NPIV Mode On/Off** entry from the CHPID Operations menu in the bottom center pane. See Figure 19.

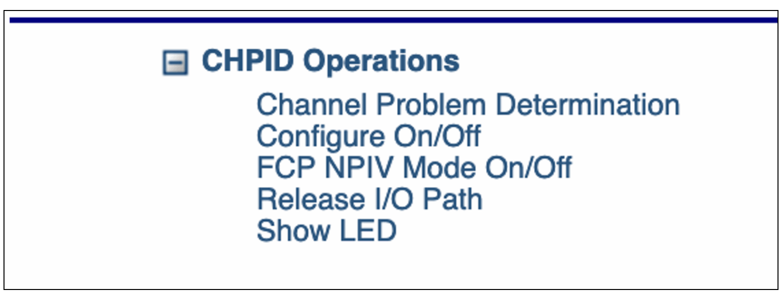

Figure 19 Select the FCP NPIV Mode On/Off entry

3. The FCP NPIV Mode On/Off tab opens. Scroll to and select the **NPIV Mode Enabled** checkbox on our LPAR's row for this CHPID.

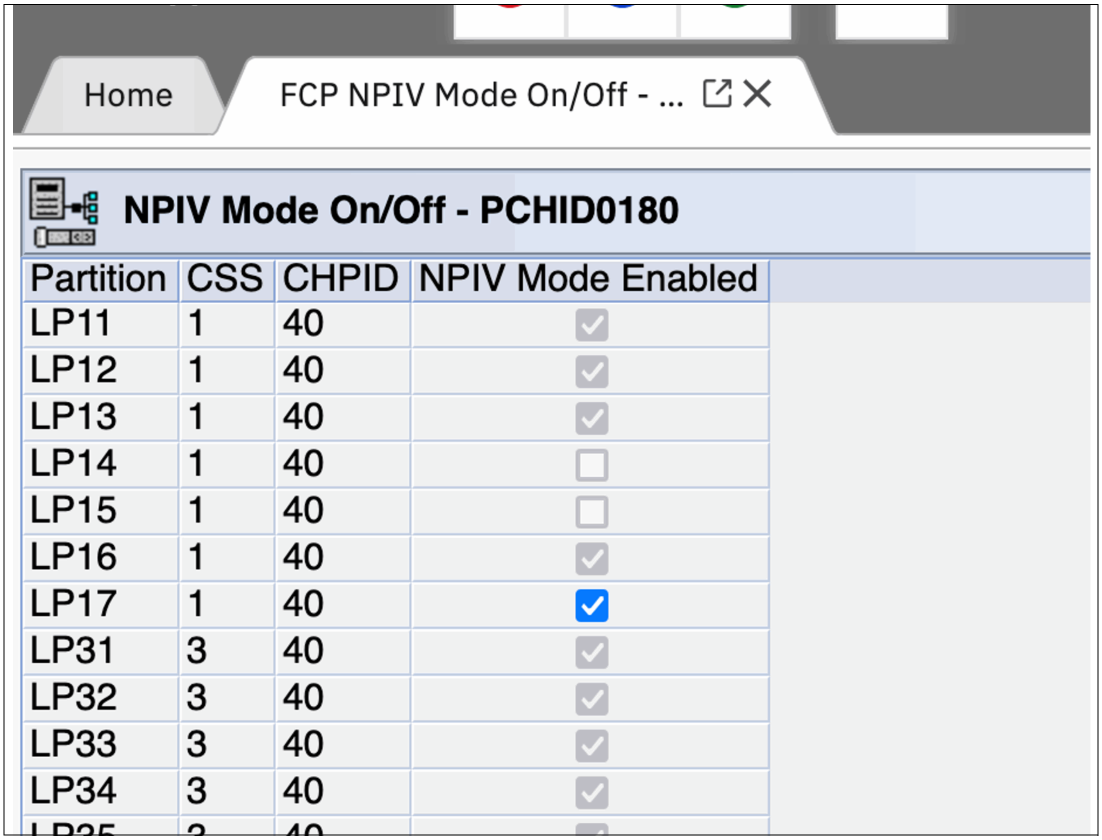

Figure 20 Select the NPIV Mode Enabled checkbox

4. If we had not first configured LP17's CHPID 40 offline the box would be greyed out and we would not be able to change it. Click the **Apply** button at the bottom of the pane to make the change, then click the **Cancel** button to close out of the panel.

Configure the channel online to the LPAR

NPIV is now active on our CHPID for our LPAR.

1. Now go back and Configure your LPAR's CHPID back online. From the Channels listing. See Figure 21 on page 19.

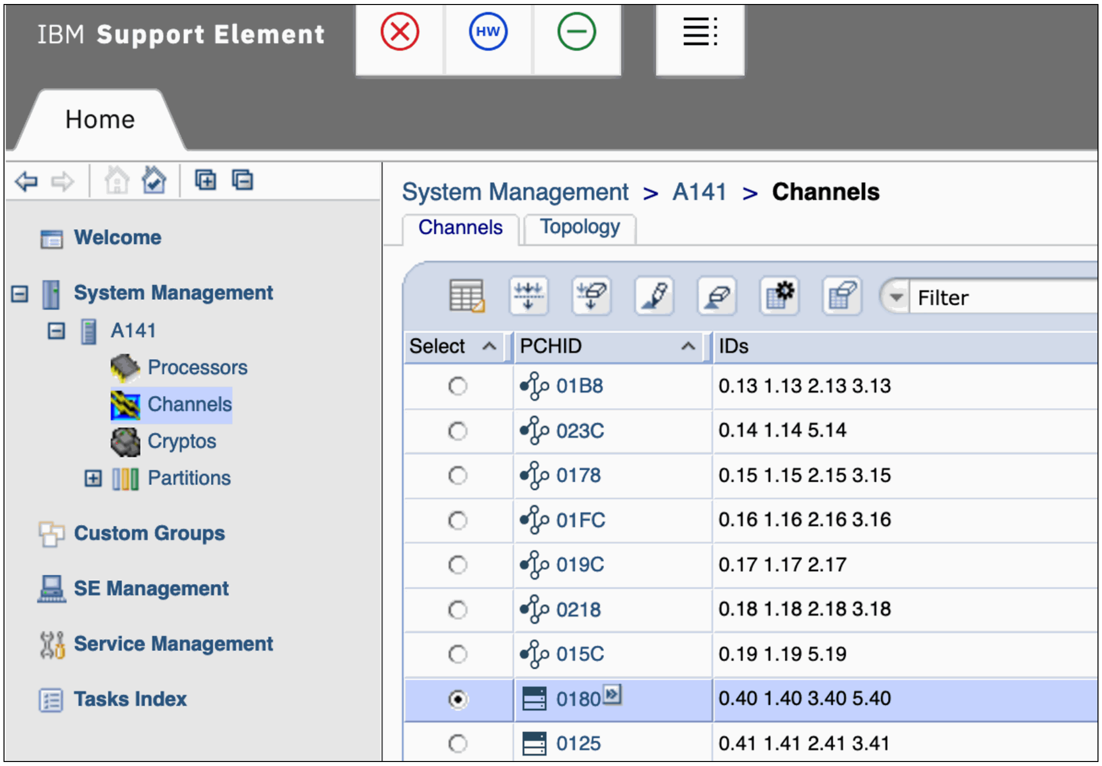

Figure 21 IBM Support Element view

2. Select the **CHPID Operations** menu in the bottom center pane and select **Configure On/Off**. See Figure 22.

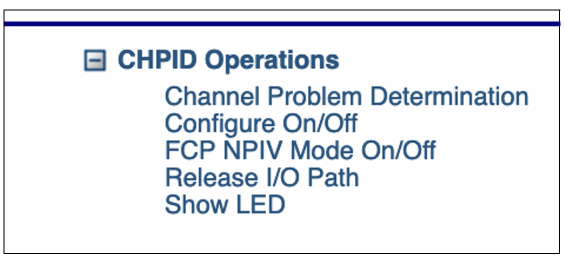

Figure 22 Select Configure On/Off

3. Locate and select the entry corresponding to the LPAR you wish to bring back online. See Figure 23 on page 20.

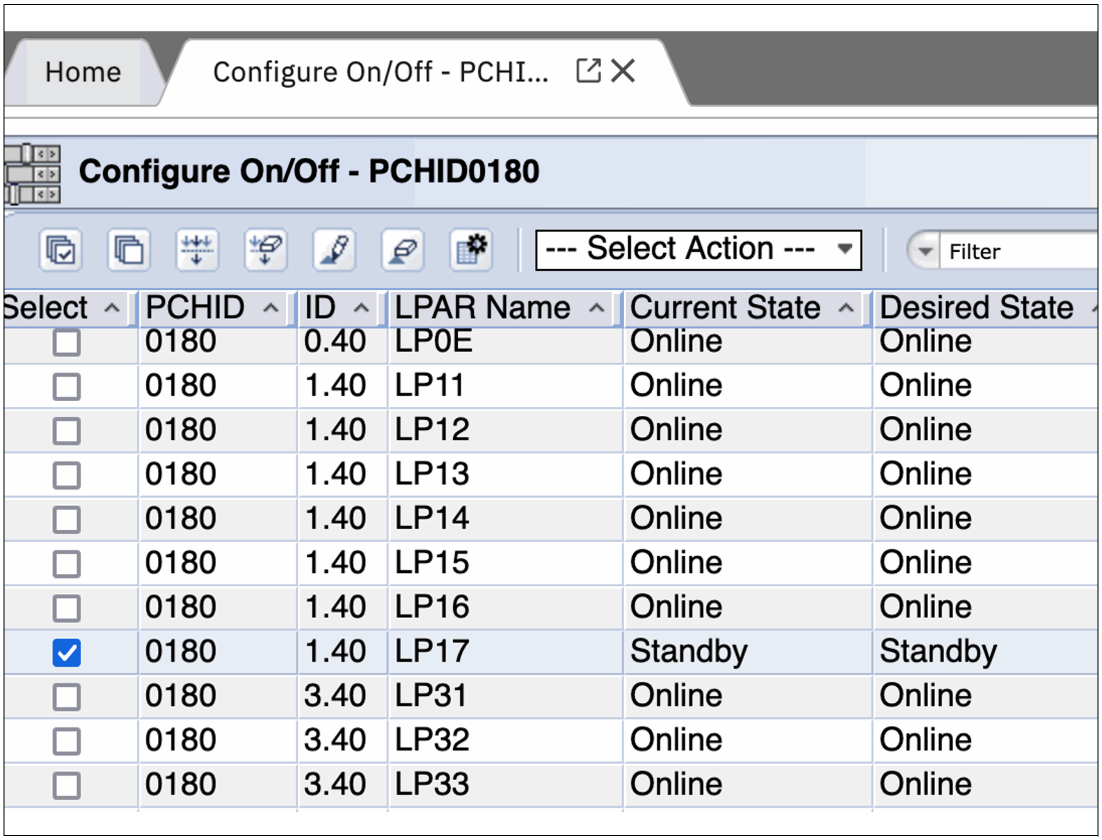

Select	PCHID	ID	LPAR Name	Current State	Desired State
☐	0180	0.40	LP0E	Online	Online
☐	0180	1.40	LP11	Online	Online
☐	0180	1.40	LP12	Online	Online
☐	0180	1.40	LP13	Online	Online
☐	0180	1.40	LP14	Online	Online
☐	0180	1.40	LP15	Online	Online
☐	0180	1.40	LP16	Online	Online
☑	0180	1.40	LP17	Standby	Standby
☐	0180	3.40	LP31	Online	Online
☐	0180	3.40	LP32	Online	Online
☐	0180	3.40	LP33	Online	Online

Figure 23 Locate and select the entry corresponding to the LPAR you wish to bring back online

4. Select **Toggle** from the drop down menu. Only Toggle the one CHPID for our specific LPAR we have selected. See Figure 24 on page 21.

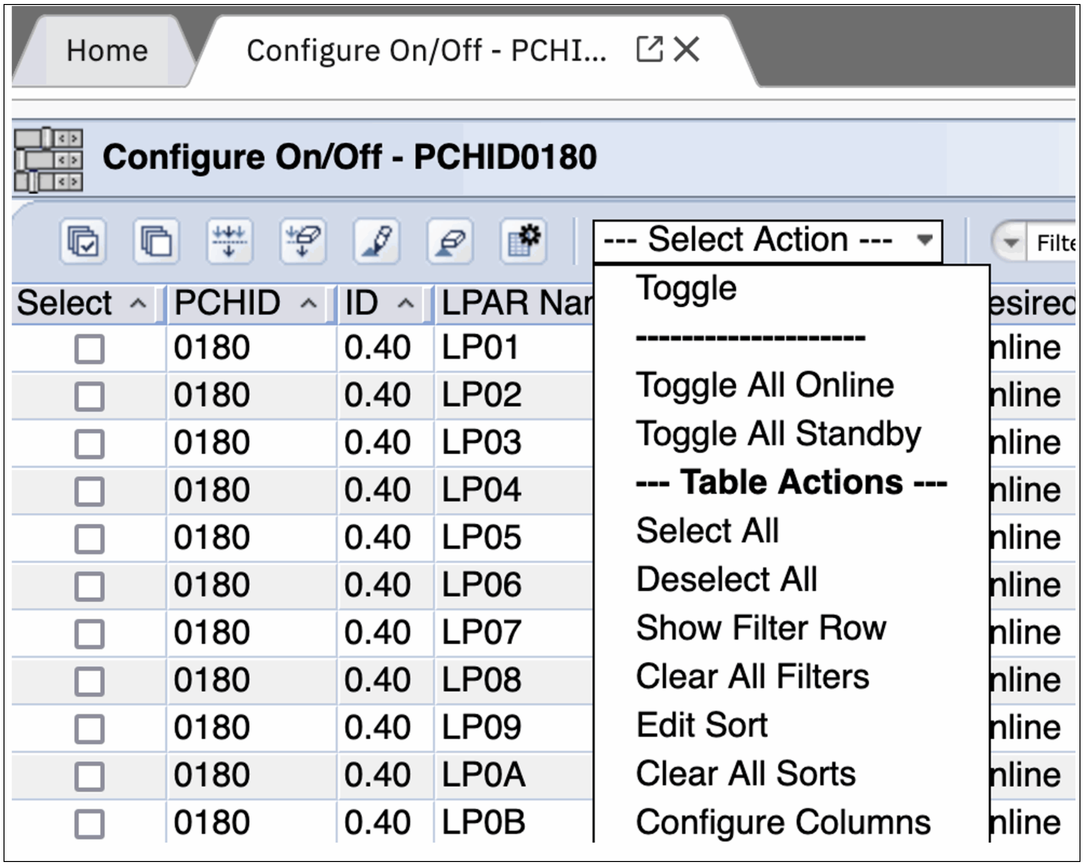

Figure 24   Select Toggle from the drop down menu

5. The panel updates to show our desired change, Click **OK**. See Figure 25.

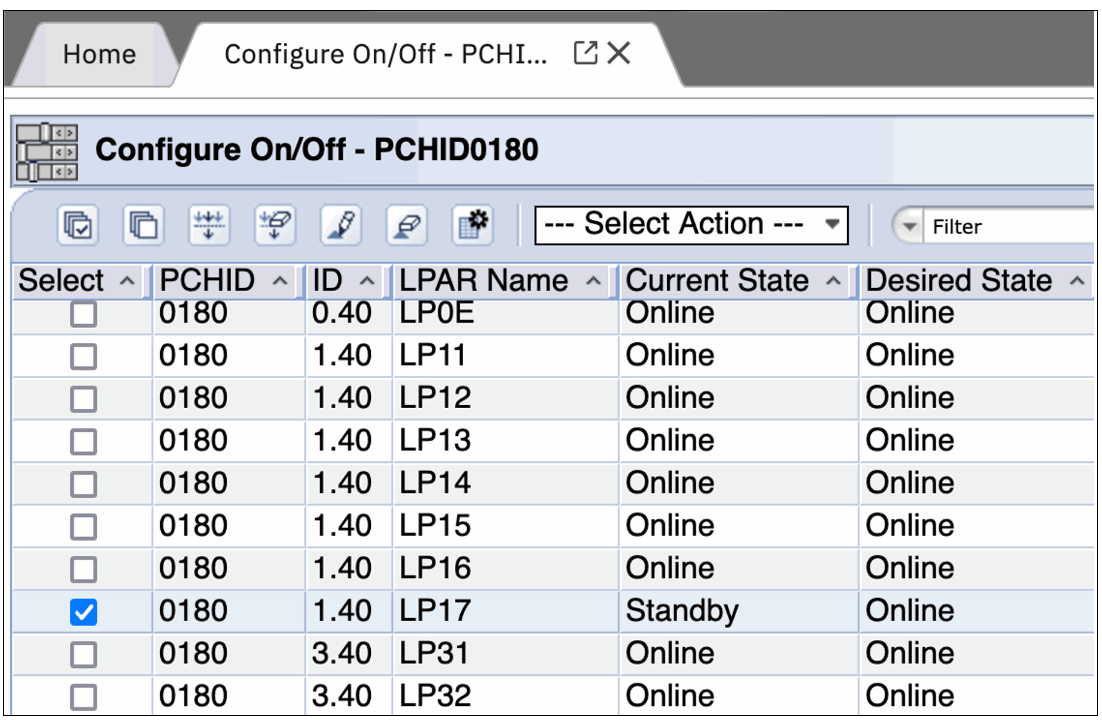

Figure 25   Updated panel

6. Click **OK** to dismiss the completion notification. Note that a confirmation prompt may or may not appear, as enabling this setting is not always disruptive.
7. With the CHPID now online for the LPAR and NPIV enabled, log out of the Service Element web interface (if applicable) using the logout function in the top right corner.
8. Repeat the above process to enable NPIV on all CHPIDs and LPARs where it is required.

> **Tip:** You can configure a CHPID offline from multiple LPARs at one time, enable NPIV on all of them, and then configure the CHPIDs back online all at once if you have many changes to make.

Finding an LPAR's NPIV WWPNs

Having enabled NPIV where necessary, we will now identify the corresponding NPIV WWPNs using the Hardware Management Console (HMC).

1. Log on to the HMC. See Figure 26.

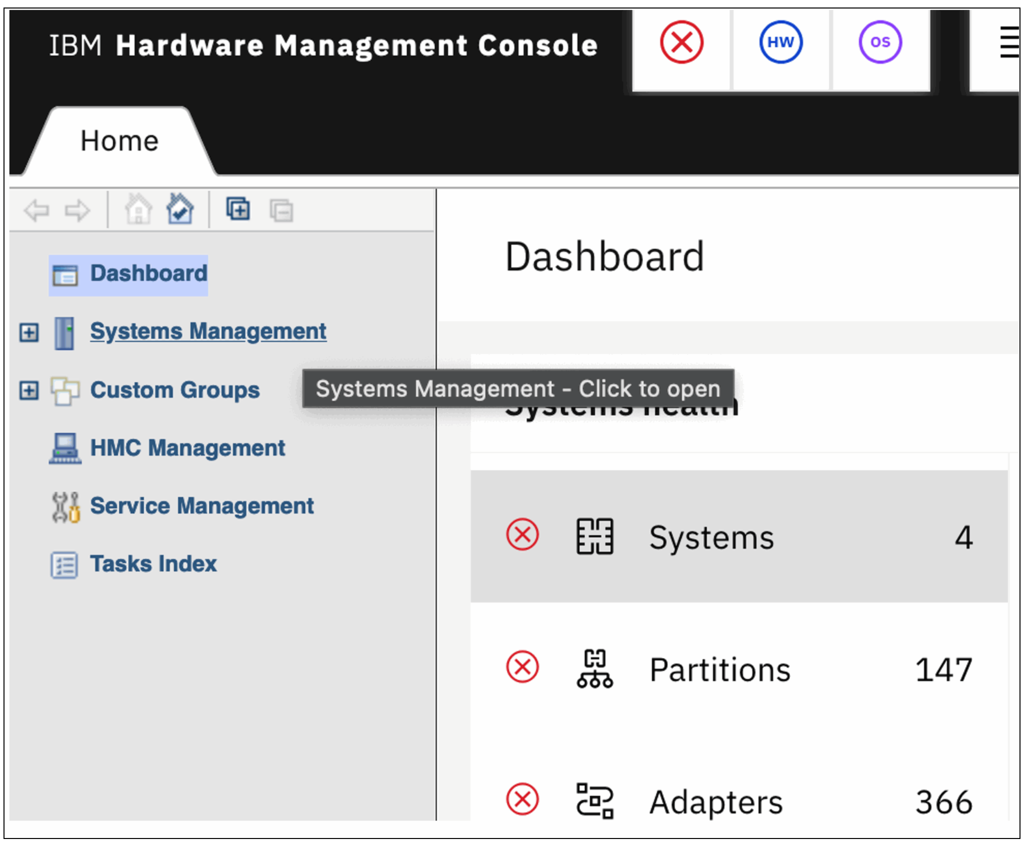

Figure 26 Hardware Management Console

2. Open the **Systems Management** tree and Select the **CPC** you want to examine. See Figure 27 on page 23.

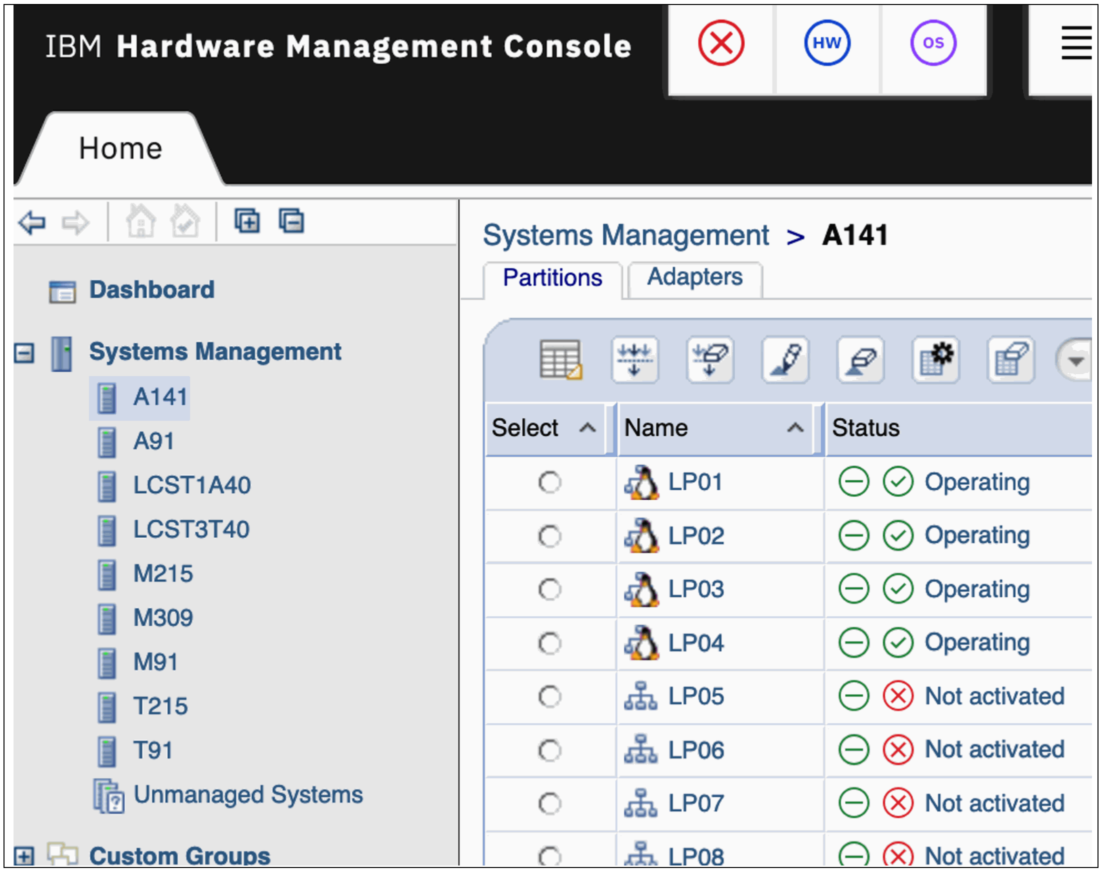

Figure 27 Select the CPC you want to examine

3. Open the **Configuration** menu in the bottom center section, and select **FCP Configuration**. See Figure 28.

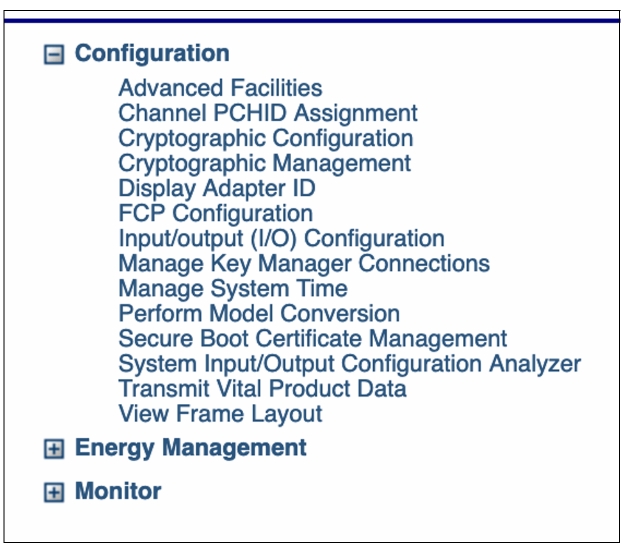

Figure 28 Select FCP Configuration

4. The FCP Configuration tab will open. Leave the option set to **Display all NPIV port names that are currently assigned to FCP subchannels…** and click **OK**. See Figure 29 on page 24.

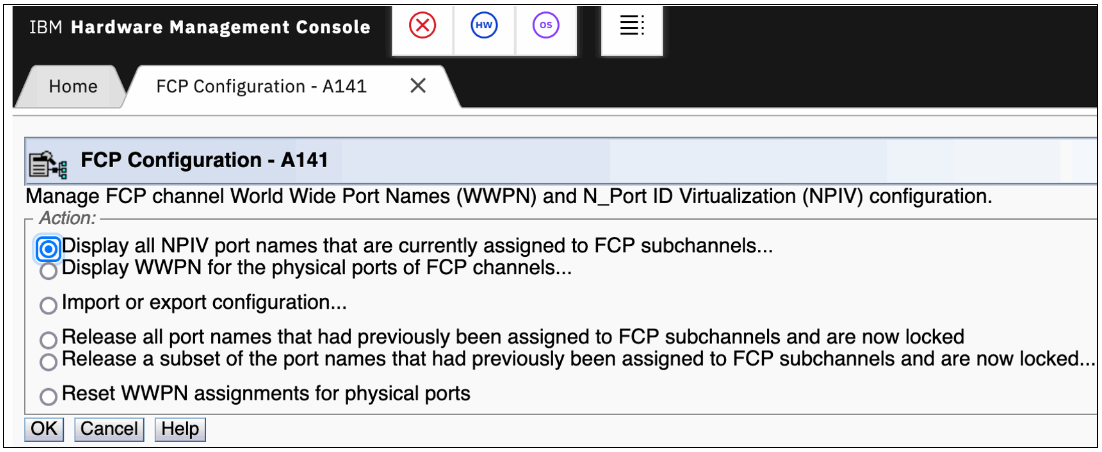

Figure 29 FCP Configuration tab

5. The next menu opens. Set the option to **Display all assigned ports for an LPAR** and click **OK**. See Figure 30.

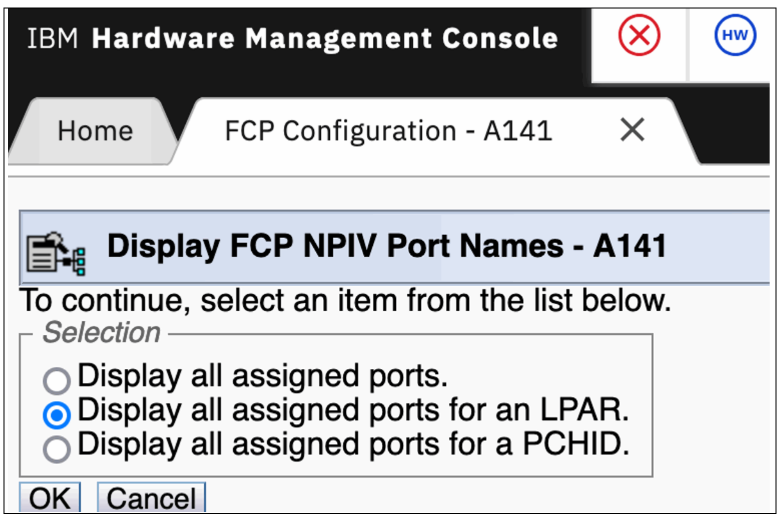

Figure 30 Set the option to Display all assigned ports for an LPAR

6. The next menu opens. Select the LPAR you wish to view the NPIV addresses for. From the list of LPARs (scroll as needed), select the desired entry and click **OK**.

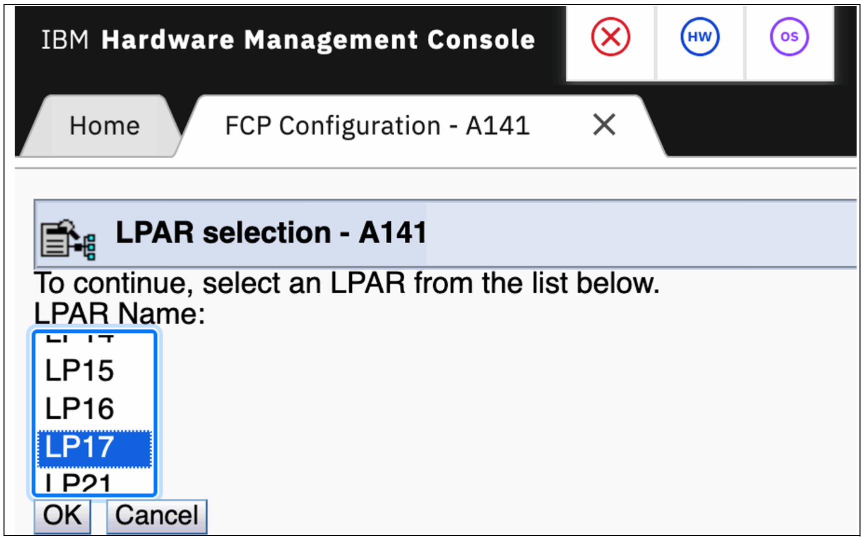

Figure 31 Select the LPAR you wish to view the NPIV addresses for

The device listing for your selected LPAR opens. See Figure 32.

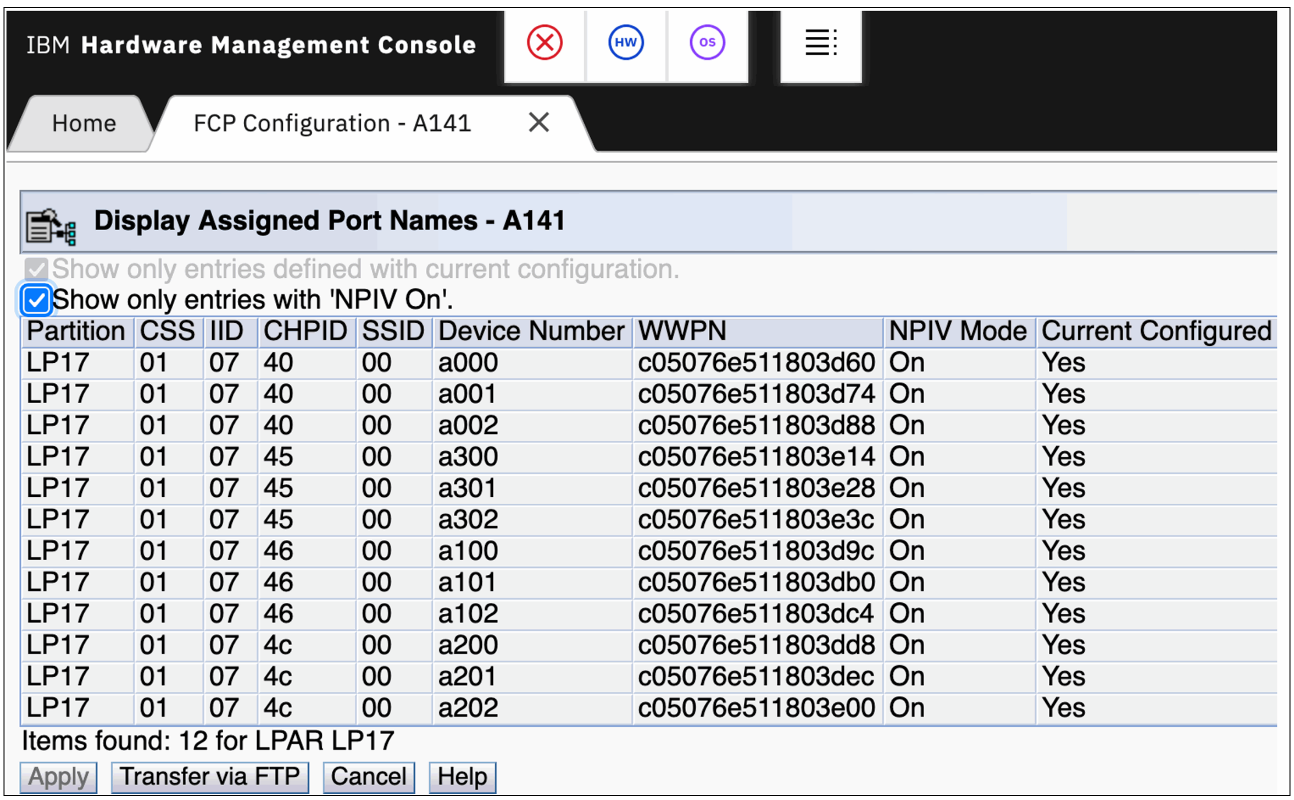

Partition	CSS	IID	CHPID	SSID	Device Number	WWPN	NPIV Mode	Current Configured
LP17	01	07	40	00	a000	c05076e511803d60	On	Yes
LP17	01	07	40	00	a001	c05076e511803d74	On	Yes
LP17	01	07	40	00	a002	c05076e511803d88	On	Yes
LP17	01	07	45	00	a300	c05076e511803e14	On	Yes
LP17	01	07	45	00	a301	c05076e511803e28	On	Yes
LP17	01	07	45	00	a302	c05076e511803e3c	On	Yes
LP17	01	07	46	00	a100	c05076e511803d9c	On	Yes
LP17	01	07	46	00	a101	c05076e511803db0	On	Yes
LP17	01	07	46	00	a102	c05076e511803dc4	On	Yes
LP17	01	07	4c	00	a200	c05076e511803dd8	On	Yes
LP17	01	07	4c	00	a201	c05076e511803dec	On	Yes
LP17	01	07	4c	00	a202	c05076e511803e00	On	Yes

Items found: 12 for LPAR LP17

Apply | Transfer via FTP | Cancel | Help

Figure 32 The device listing for your selected LPAR

Each device number on each CHPID has a unique WWPN, enabling the SAN and storage devices to differentiate between virtualized systems sharing the same physical host adapter. Without NPIV, devices a000, a001, and a002 within LPAR LP17 would be indistinguishable to the SAN, granting them all access to the same SAN resources. Furthermore, without NPIV, device a000 on LP17 would be identical to device a000 on any other LPAR.

> **Tip:** For later use, save this data to a text file. While screenshots are helpful, they do not facilitate easy text extraction.

For our zoning and storage example we will be using the following devices:

– LP17 a000 c05076e511803d60

– LP17 a100 c05076e511803d9c

– LP17 a200 c05076e511803dd8

– LP17 a300 c05076e511803e14

7. Click the **X** on the FCP Configuration tab to close this pane. Continue to click **X** on the tab until it closes to back all the way out of FCP Configuration.

   If your CHPIDs are not enabled for NPIV in this LPAR, you will see an empty screen similar to Figure 33.

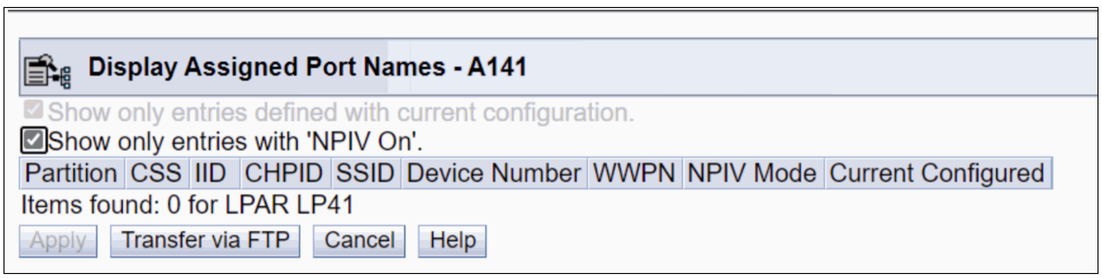

Figure 33 CHPIDs are not enabled for NPIV in this LPAR

8. Click **Cancel** to close the window.

Finding the correct WWPNs in a FlashSystem storage device

As storage systems have evolved to offer increased functionality, their internal complexity has also grown to ensure backward compatibility. IBM FlashSystem devices have gained numerous capabilities over time, with NPIV implementation introducing significant complexity. It is important to distinguish this storage subsystem NPIV implementation from the server-side NPIV discussed in “N Port ID virtualization” on page 5.

IBM FlashSystem utilizes NPIV-enabled SAN ports for management isolation across various protocols, including SCSI, NVMe over Fibre Channel, and Metro/Global Mirror replication, as well as for hardware failover high availability (HA).

An IBM FS9500 FlashSystem has a port configuration that looks like Figure 34 on page 27, if it is fully populated with FCP adapters.

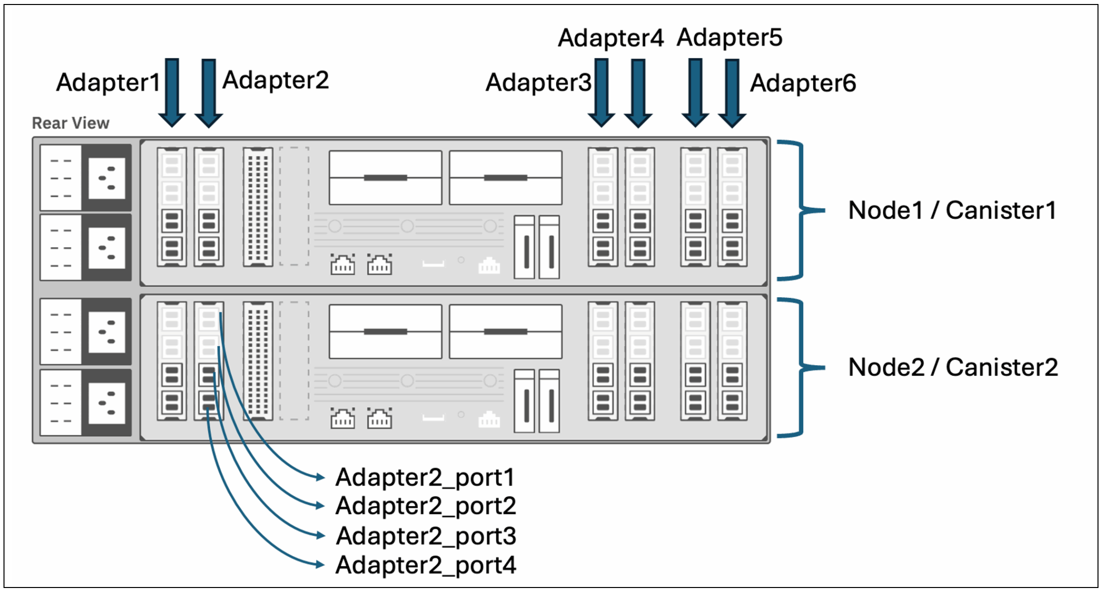

Figure 34 IBM FS9500 FlashSystem

Each IBM FlashSystem 9500 node consists of an upper and a lower node. Each node contains six I/O adapters, numbered 1 through 6 from left to right. Each adapter has multiple ports, numbered sequentially from top to bottom, starting with port 1.

It is important to note that other IBM FlashSystem models may have different physical orientations. For example, the lower node may be inverted, resulting in reversed port numbering relative to the upper node. Consult the specific product documentation for each model to determine the correct port enumeration.

Identifying a specific port on the FS9500 requires all three components: Node, Adapter, and Port. Within the FS9500 management interface, FCP ports are assigned sequential port IDs within each node. This enumeration begins with adapter 1, port 1 (top left), proceeds down that adapter, then moves to adapter 2, port 1, and continues down that adapter, and so forth, until the bottom port of adapter 6 is reached. Each node on our systems has 24 ports (6 adapters * 4 ports/adapter).

Figure 34 of the web management interface highlights active, connected ports in a darker shade. The following output from the FlashSystem command-line interface (CLI), generated using `lsportfc -filtervalue "type=fc:status=active"`, identifies the active FCP ports on our FS9500. Correlate the `Node:Adapter:Port` designations with the sequential port IDs to understand their mapping. See Figure 35 on page 28.

port_id	node	WWPN	nportid	status	adapter_location	adapter_port_id
3	node1	50050768131345AD	01C000	active	1	3
4	node1	50050768131445AD	E3D040	active	1	4
7	node1	50050768132345AD	9D0000	active	2	3
8	node1	50050768132445AD	E3D840	active	2	4
11	node1	50050768135345AD	01C100	active	5	3
12	node1	50050768135445AD	E5A3C0	active	5	4
15	node1	50050768136345AD	9D0020	active	6	3
16	node1	50050768136445AD	E3D800	active	6	4
19	node1	50050768137345AD	01C200	active	7	3
20	node1	50050768137445AD	E5A380	active	7	4
23	node1	50050768138345AD	9D0040	active	8	3
24	node1	50050768138445AD	E5A0C0	active	8	4
3	node2	50050768131345AE	12040	active	1	3
4	node2	50050768131445AE	E3D000	active	1	4
7	node2	50050768132345AE	9D0060	active	2	3
8	node2	50050768132445AE	E3BBC0	active	2	4
11	node2	50050768135345AE	12140	active	5	3
12	node2	50050768135445AE	E5A340	active	5	4
15	node2	50050768136345AE	9D0080	active	6	3
16	node2	50050768136445AE	E3BB80	active	6	4
19	node2	50050768137345AE	12240	active	7	3
20	node2	50050768137445AE	E5A300	active	7	4
23	node2	50050768138345AE	9D00A0	active	8	3
24	node2	50050768138445AE	E5A080	active	8	4

Figure 35 Port mappings

The WWPNs listed in Figure 35 are the Physical WWPNs on those FCP ports, and they are *Not For Host IO*. Do not zone any of the above WWPNs to a host for running any IO. *In current FlashSystems firmware, the physical WWPNs are only for the node failover communications zones and for the Global and Metro Mirror zones*. We will zone some of them to each other in the zoning section, but not to anything else.

With NPIV enabled, each FCP port on an IBM FlashSystem storage unit presents three WWPNs to the SAN, as illustrated in Figure 36. The trimmed output of `lstargetportfc -filtervalue "port_id=3:owning_node_id=1"` shows all logical target ports on node 1's port 3, which is the third port down on adapter 1 on node 1.

id	WWPN	WWNN	port_id	node_id	nportid	host_io_perm	virt	prot
7	50050768131345AD	50050768130045AD	3	1	01C000	no	no	scsi
8	50050768131745AD	50050768130045AD	3	1	01C001	yes	yes	scsi
9	50050768131B45AD	50050768130045AD	3	1	01C002	yes	yes	nvme

Figure 36 Logical target ports

The table in Figure 36 on page 28 displays three WWPNs associated with a single physical port logged into the SAN switch. The physical WWPN, `50050768131345AD` (previously shown in Figure 35 on page 28), has a "Host_IO_Permitted" status of "No." Additionally, two

NPIV-enabled ports are logged into the fabric: one for FCP SCSI protocol I/O (`50050768131745AD`) and the other for NVMe over Fibre Channel I/O (`50050768131B45AD`).

The physical port WWPN should not be zoned to a host, as it is not authorized for host I/O. Similarly, neither virtual WWPN should be zoned to a host unless that host will be using the corresponding protocol. Specifically, do not zone an FCP SCSI host to the NVMe port, nor an NVMe over Fibre Channel host to an FCP SCSI port.

Tip: Although zoning a host to inappropriate ports on the FlashSystem may not immediately generate errors, it will significantly complicate subsequent troubleshooting. To facilitate future diagnostics, ensure that hosts are zoned only to appropriate ports on the storage controller.

Figure 37 on page 30 shows the output of `lstargetportfc -filtervalue "virtualized=yes:protocol=scsi"`, which lists all the valid NPIV target ports for SCSI protocol hosts.

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
2	50050768131545AD	50050768130045AD	1	1	0	yes	scsi
5	50050768131645AD	50050768130045AD	2	1	0	yes	scsi
8	50050768131745AD	50050768130045AD	3	1	01C001	yes	scsi
11	50050768131845AD	50050768130045AD	4	1	E3D041	yes	scsi
14	50050768132545AD	50050768130045AD	5	1	0	yes	scsi
17	50050768132645AD	50050768130045AD	6	1	0	yes	scsi
20	50050768132745AD	50050768130045AD	7	1	9D0001	yes	scsi
23	50050768132845AD	50050768130045AD	8	1	E3D841	yes	scsi
26	50050768135545AD	50050768130045AD	9	1	0	yes	scsi
29	50050768135645AD	50050768130045AD	10	1	0	yes	scsi
32	50050768135745AD	50050768130045AD	11	1	01C101	yes	scsi
35	50050768135845AD	50050768130045AD	12	1	E5A3C1	yes	scsi
38	50050768136545AD	50050768130045AD	13	1	0	yes	scsi
41	50050768136645AD	50050768130045AD	14	1	0	yes	scsi
44	50050768136745AD	50050768130045AD	15	1	9D0021	yes	scsi
47	50050768136845AD	50050768130045AD	16	1	E3D801	yes	scsi
50	50050768137545AD	50050768130045AD	17	1	0	yes	scsi
53	50050768137645AD	50050768130045AD	18	1	0	yes	scsi
56	50050768137745AD	50050768130045AD	19	1	01C201	yes	scsi
59	50050768137845AD	50050768130045AD	20	1	E5A381	yes	scsi
62	50050768138545AD	50050768130045AD	21	1	0	yes	scsi
65	50050768138645AD	50050768130045AD	22	1	0	yes	scsi
68	50050768138745AD	50050768130045AD	23	1	9D0041	yes	scsi
71	50050768138845AD	50050768130045AD	24	1	E5A0C1	yes	scsi
74	50050768131545AE	50050768130045AE	1	2	0	yes	scsi
77	50050768131645AE	50050768130045AE	2	2	0	yes	scsi
80	50050768131745AE	50050768130045AE	3	2	012041	yes	scsi
83	50050768131845AE	50050768130045AE	4	2	E3D001	yes	scsi
86	50050768132545AE	50050768130045AE	5	2	0	yes	scsi
89	50050768132645AE	50050768130045AE	6	2	0	yes	scsi
92	50050768132745AE	50050768130045AE	7	2	9D0061	yes	scsi
95	50050768132845AE	50050768130045AE	8	2	E3BBC1	yes	scsi
98	50050768135545AE	50050768130045AE	9	2	0	yes	scsi
101	50050768135645AE	50050768130045AE	10	2	0	yes	scsi
104	50050768135745AE	50050768130045AE	11	2	012141	yes	scsi
107	50050768135845AE	50050768130045AE	12	2	E5A341	yes	scsi
110	50050768136545AE	50050768130045AE	13	2	0	yes	scsi
113	50050768136645AE	50050768130045AE	14	2	0	yes	scsi
116	50050768136745AE	50050768130045AE	15	2	9D0081	yes	scsi
119	50050768136845AE	50050768130045AE	16	2	E3BB81	yes	scsi
122	50050768137545AE	50050768130045AE	17	2	0	yes	scsi
125	50050768137645AE	50050768130045AE	18	2	0	yes	scsi
126	50050768137745AE	50050768130045AE	19	2	012241	yes	scsi
131	50050768137845AE	50050768130045AE	20	2	E5A301	yes	scsi
134	50050768138545AE	50050768130045AE	21	2	0	yes	scsi
137	50050768138645AE	50050768130045AE	22	2	0	yes	scsi
140	50050768138745AE	50050768130045AE	23	2	9D00A1	yes	scsi
143	50050768138845AE	50050768130045AE	24	2	E5A081	yes	scsi

Figure 37 Valid NPIV target ports for SCSI protocol hosts

For our configuration, only the WWPNs listed above are valid for zoning with our LinuxONE hosts. Examining the `nportid` column reveals that the device has connections to three separate switches. Specifically, Port 4 on all adapters connects to one switch, while Port 3 on adapters 1, 3, and 5 connects to a second switch, and Port 3 on adapters 2, 4, and 6 connects to a third.

The following examples will focus solely on the two switches connected to Port 3. Consequently, the table in Figure 37 will be simplified and divided into two tables, one for each of these switches.

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
8	50050768131745AD	50050768130045AD	3	1	01C001	yes	scsi
32	50050768135745AD	50050768130045AD	11	1	01C101	yes	scsi
56	50050768137745AD	50050768130045AD	19	1	01C201	yes	scsi
80	50050768131745AE	50050768130045AE	3	2	012041	yes	scsi
104	50050768135745AE	50050768130045AE	11	2	012141	yes	scsi
126	50050768137745AE	50050768130045AE	19	2	012241	yes	scsi

Figure 38 WWPNs we are going to use for FCP target ports - Switch one

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
20	50050768132745AD	50050768130045AD	7	1	9D0001	yes	scsi
44	50050768136745AD	50050768130045AD	15	1	9D0021	yes	scsi
68	50050768138745AD	50050768130045AD	23	1	9D0041	yes	scsi
92	50050768132745AE	50050768130045AE	7	2	9D0061	yes	scsi
116	50050768136745AE	50050768130045AE	15	2	9D0081	yes	scsi
140	50050768138745AE	50050768130045AE	23	2	9D00A1	yes	scsi

Figure 39 WWPNs we are going to use for FCP target ports - Switch two

The WWPNs presented in Figure 38 on page 31 and Figure 39 constitute the exclusive set of ports that will be employed for FCP target port zoning within the subsequent section. The construction of these tables was based on the previously generated output from the `lstargetportfc -filtervalue "virtualized=yes:protocol=scsi"` command (show me all NPIV SCSI ports), corroborated with documentation detailing the physical fiber link connections, and subsequently verified against the nportid values, which serve as indicators of connections to distinct SAN fabrics.

For the configuration of FlashSystem chassis internal connectivity zones, it is also necessary to obtain the physical WWPNs of all connected ports on the two target switches. The output in Figure 40 and Figure 41 on page 32 was generated using the `lstargetportfc -filtervalue "virtualized=no"` command and has been further processed to categorize the results by switch connection.

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
7	50050768131345AD	50050768130045AD	3	1	01C000	no	scsi
31	50050768135345AD	50050768130045AD	11	1	01C100	no	scsi
55	50050768137345AD	50050768130045AD	19	1	01C200	no	scsi
79	50050768131345AE	50050768130045AE	3	2	012040	no	scsi
103	50050768135345AE	50050768130045AE	11	2	012140	no	scsi
127	50050768137345AE	50050768130045AE	19	2	012240	no	scsi

Figure 40 Physical WWPNs of all connected ports - Switch one

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
19	50050768132345AD	50050768130045AD	7	1	9D0000	no	scsi
43	50050768136345AD	50050768130045AD	15	1	9D0020	no	scsi
67	50050768138345AD	50050768130045AD	23	1	9D0040	no	scsi
91	50050768132345AE	50050768130045AE	7	2	9D0060	no	scsi
115	50050768136345AE	50050768130045AE	15	2	9D0080	no	scsi
139	50050768138345AE	50050768130045AE	23	2	9D00A0	no	scsi

Figure 41 Physical WWPNs of all connected ports - Switch two

Important: As stated earlier, the physical WWPNs listed in these tables should not be used for connections from remote Fibre Channel hosts. Their purpose is limited to traffic associated with port-to-port failover within the chassis or inter-chassis communication for Peer-to-Peer Remote Copy, Metro Mirror, and Global Mirror replication.

Creating SAN zones to permit host ports to communicate with storage ports

The host-side WWPNs for LP17, with the corresponding switch connections added based on infrastructure documentation are shown in Example 2.

Example 2 Host-side WWPNs for LP17

```
lpardevicewwpn            switch
LP17 a000 c05076e511803d60 sw one
LP17 a100 c05076e511803d9c sw two
LP17 a200 c05076e511803dd8 sw one
LP17 a300 c05076e511803e14 sw two
```

Let us also include these WWPNs for LP16. See Example 3.

Example 3 Host-side WWPNs for LP16

```
lpardevicewwpn            switch
LP16 a000 c05076e511803d5c sw one
LP16 a100 c05076e511803d98 sw two
LP16 a200 c05076e511803dd4 sw one
LP16 a300 c05076e511803e10 sw two
```

Also, we have these ports on our FS9500 plugged and eligible for our use (See Figure 42 and Figure 43 on page 33).

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
8	50050768131745AD	50050768130045AD	3	1	01C001	yes	scsi
32	50050768135745AD	50050768130045AD	11	1	01C101	yes	scsi
56	50050768137745AD	50050768130045AD	19	1	01C201	yes	scsi
80	50050768131745AE	50050768130045AE	3	2	012041	yes	scsi
104	50050768135745AE	50050768130045AE	11	2	012141	yes	scsi
126	50050768137745AE	50050768130045AE	19	2	012241	yes	scsi

Figure 42 Ports on our FS9500 - Switch one

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
20	50050768132745AD	50050768130045AD	7	1	9D0001	yes	scsi
44	50050768136745AD	50050768130045AD	15	1	9D0021	yes	scsi
68	50050768138745AD	50050768130045AD	23	1	9D0041	yes	scsi
92	50050768132745AE	50050768130045AE	7	2	9D0061	yes	scsi
116	50050768136745AE	50050768130045AE	15	2	9D0081	yes	scsi
140	50050768138745AE	50050768130045AE	23	2	9D00A1	yes	scsi

Figure 43 Ports on our FS9500 - Switch two

The overall connectivity is illustrated in Figure 44.

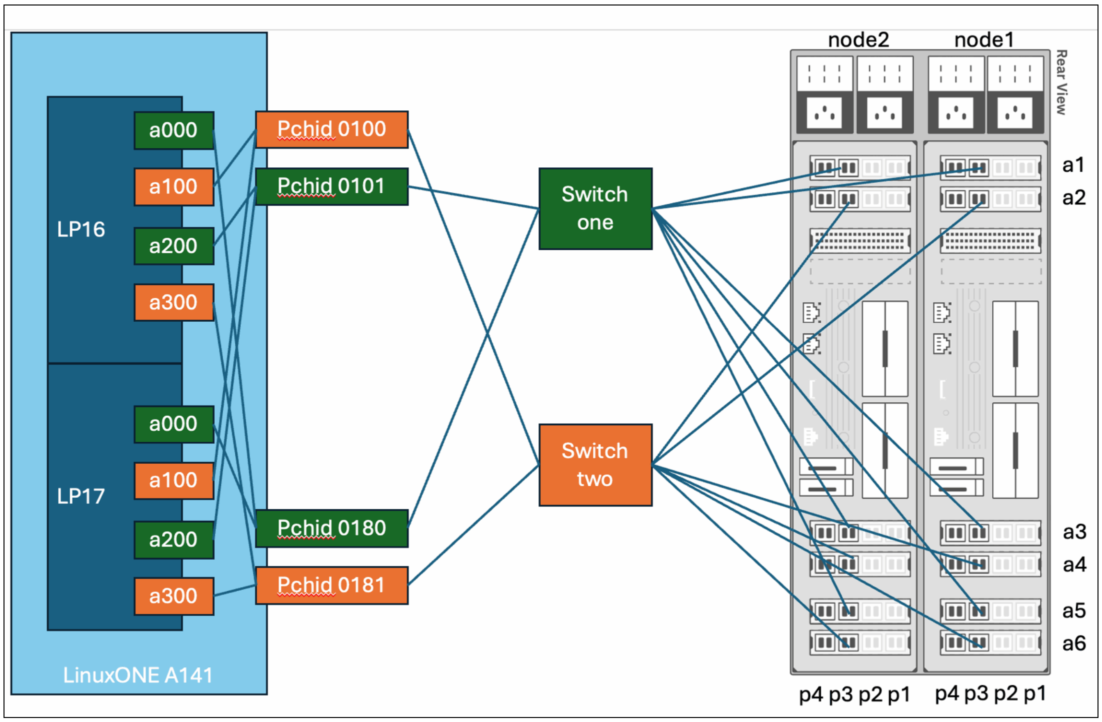

Figure 44 Overall connectivity diagram

Based on these connections, we can create zones that look like Example 4 and Example 5 on page 33 on switch one.

Example 4 Creating zones for FS9500 - switch one

```
fs9500_a141_path0_switch1
  host LP16 a000
  host LP17 a000
  target node1:adapter3:port3 - id 32
  target node2:adapter3:port3 - id 104
```

Example 5 Creating zones for FS9500 - switch one

```
fs9500_a141_path2_switch1
  host LP16 a200
  host LP17 a200
  target node1:adapter5:port3 - id 56
  target node2:adapter5:port3 - id 126
```

Figure 45 illustrates the connectivity diagram resulting from the configuration in Example 4 on page 33.

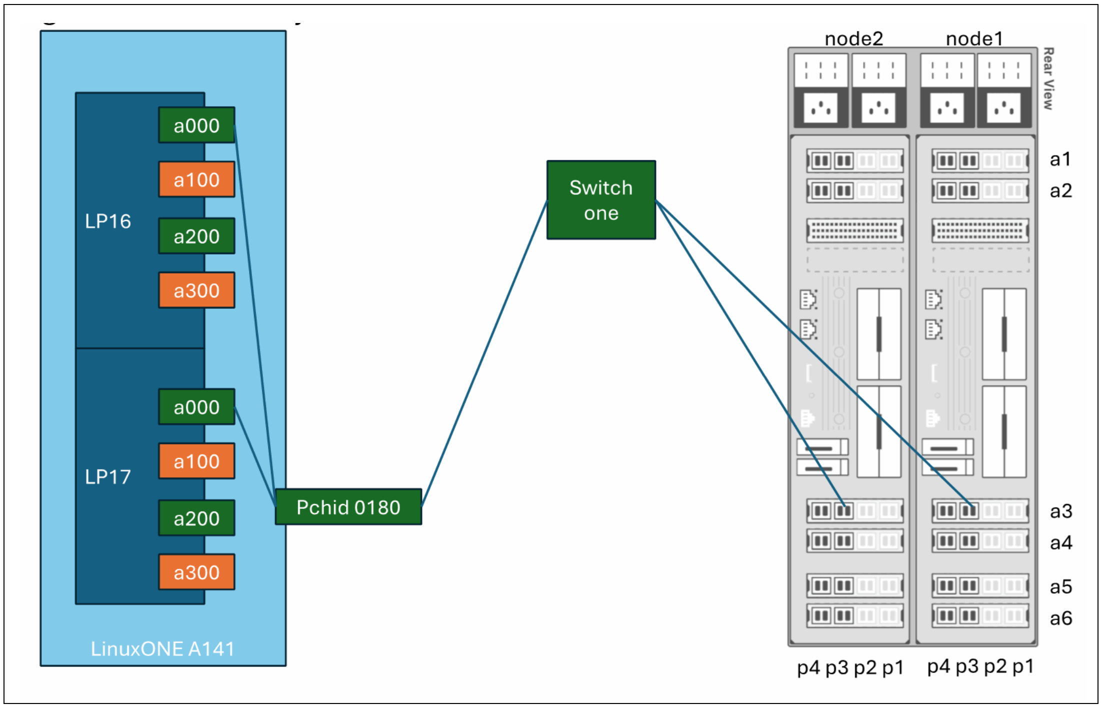

Figure 45 Connectivity diagram resulting from the configuration in Example 4

Figure 46 on page 34 illustrates the connectivity diagram resulting from the configuration in Example 5 on page 33.

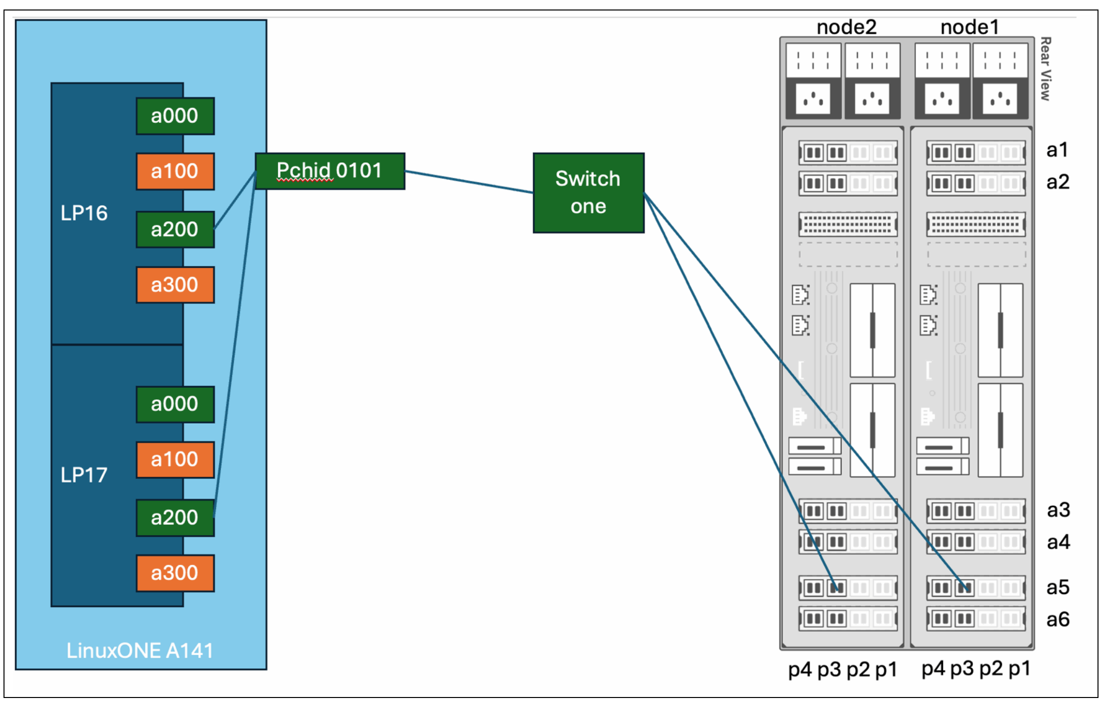

Figure 46 Connectivity diagram resulting from the configuration in Example 5

We also create zones that look like Example 6 and Example 7 on page 35 on switch two.

Example 6 Creating zones for FS9500 - switch two

```
fs9500_a141_path1_switch2
  host LP16 a100
  host LP17 a100
  target node1:adapter2:port3 - id 20
  target node2:adapter2:port3 - id 92
```

Example 7 Creating zones for FS9500 - switch two

```
fs9500_a141_path3_switch2
  host LP16 a300
  host LP17 a300
  target node1:adapter4:port3 - id 44
  target node2:adapter4:port3 - id 116
```

Figure 47 on page 35 shows the connectivity diagram resulting from the configuration in Example 6 on page 35.

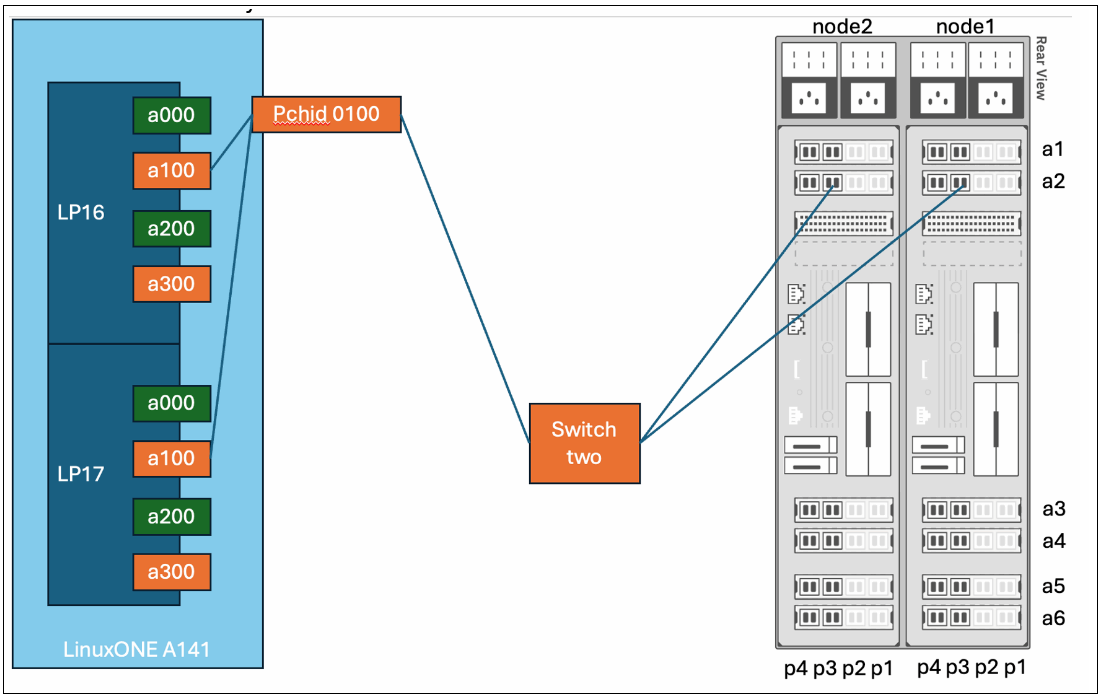

Figure 47 Connectivity diagram resulting from the configuration in Example 6

Figure 48 shows the connectivity diagram resulting from the configuration in Example 7 on page 35.

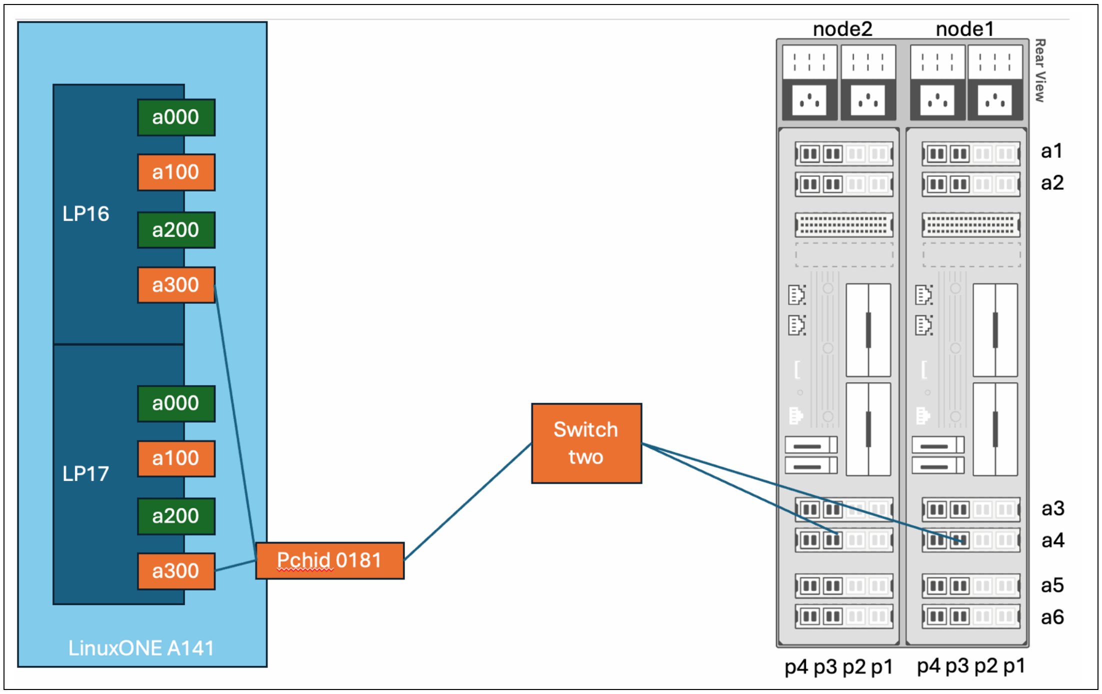

Figure 48 Connectivity diagram resulting from the configuration in Example 7

In our configuration, all ports on adapters 1 and 6 in both FlashSystem nodes are dedicated to node-to-node failover zones and Metro/Global Mirror communication. Host connections are exclusively configured to use ports on adapters 2, 3, 4, and 5. Odd-numbered adapters connect to switch one, while even-numbered adapters connect to switch two.

Each zone includes two target ports on the storage unit, located on the same adapter and port number on both nodes. For example, node 1, adapter 4, port 3, and node 2, adapter 4, port 3, are paired in a zone. This configuration is intentional. The FlashSystem firmware may relocate NPIV WWPNs between nodes during certain failure or service events. This relocation occurs on a corresponding adapter-port basis; therefore, the NPIV SCSI WWPN for adapter 2, port 3 on node 1 will only be relocated to adapter 2, port 3 on node 2. It will not be relocated to a different port on adapter 2 or to a different adapter on node 2.

Our switches support an advanced zoning feature: Peer Zones on Brocade devices and Smart Zones on Cisco devices. This feature enables administrators to classify zone members into two distinct categories. While the specific terminology varies between vendors, the fundamental concept involves two member types, designated here as Type 1 and Type 2. When adding a member to a zone, it is assigned as either Type 1 or Type 2. Members of Type 1 are aware of all Type 2 members, and vice versa. Critically, Type 1 members will not be aware of other Type 1 members.

If your switches support Peer Zones or Smart Zones, their utilization is strongly recommended. These advanced zoning capabilities significantly simplify the management of complex storage networks and reduce the overall number of zones required.

In the absence of these advanced zoning features, Single Initiator Zoning represents the next most effective approach. In our previous switch configuration, this method resulted in the management of thousands of zones. Single Initiator Zoning involves creating a zone containing one host initiator and a set of target adapters. This zone is then replicated for each host requiring access to those same target adapters. Consequently, in a SAN with 1,000 hosts, each with four initiators communicating with the same set of eight target ports, 4,000 zones would be necessary. In contrast, using Peer/Smart Zones, the same scenario could be achieved with as few as four zones, each containing 1,000 initiators and two targets.

Important: On IBM LinuxONE and Z systems, the host port NPIV WWPN will not register with the SAN fabric until the corresponding LPAR is activated, an operating system is loaded, and the operating system subsequently configures the device online. The storage administrator requires the necessary host-side WWPNs to configure zones. These WWPNs must be communicated to the storage administrator.

NPIV storage target WWPNs on an IBM FlashSystem register with the network shortly after the storage system is powered on and are therefore visible within the fabric. However, the storage administrator still requires explicit direction regarding which of the available NPIV WWPNs should be utilized.

The actual zone configuration implemented on Brocade switch one for these hosts is as shown in Example 8.

Example 8 Zone configuration implemented on Brocade switch one

```
lpardevicewwpn           switch
LP17 a000 c05076e511803d60 sw one
LP17 a200 c05076e511803dd8 sw one
LP16 a000 c05076e511803d5c sw one
LP16 a200 c05076e511803dd4 sw one
```

We will connect the hosts in Example 8 to the storage ports in Figure 49 on page 37.

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
32	50050768135745AD	50050768130045AD	11	1	01C101	yes	scsi
56	50050768137745AD	50050768130045AD	19	1	01C201	yes	scsi
104	50050768135745AE	50050768130045AE	11	2	012141	yes	scsi
126	50050768137745AE	50050768130045AE	19	2	012241	yes	scsi

Figure 49 Storage ports

Prior to configuring zones, we create aliases to enhance manageability and clarity. See Example 9.

Example 9 Creating aliases

```
alicreate "a141_lp17_a000_hostname0","c0:50:76:e5:11:80:3d:60"
alicreate "a141_lp17_a200_hostname0","c0:50:76:e5:11:80:3d:d8"
alicreate "a141_lp16_a000_hostname1","c0:50:76:e5:11:80:3d:5c"
alicreate "a141_lp16_a200_hostname1","c0:50:76:e5:11:80:3d:d4"
alicreate "fs95fp1_node1_a3p3_npiv","50:05:07:68:13:57:45:ad"
alicreate "fs95fp1_node2_a3p3_npiv","50:05:07:68:13:57:45:ae"
alicreate "fs95fp1_node1_a5p3_npiv","50:05:07:68:13:77:45:ad"
alicreate "fs95fp1_node2_a5p3_npiv","50:05:07:68:13:77:45:ae"
```

We now define the peer zones. Storage target WWPNs are designated as `-principal` members, and host initiator WWPNs are designated as `-member` members.

It is important to note the selective approach to port connectivity. While the Linux multipath driver effectively handles up to eight paths to a LUN, adding more paths beyond this threshold primarily increases SAN overhead without proportionally improving performance or resilience. A full mapping of all target ports to all initiator ports using Smart/Peer zones would create 16 paths, which is suboptimal. Our configuration establishes a mapping of only two target ports to each LPAR's initiator on a shared adapter. See Example 10.

Example 10 Defining the peer zones

```
zonecreate --peerzone "a141hosts_fs95fp1_port_33" \
-principal "fs95fp1_node1_a3p3_npiv;fs95fp1_node2_a3p3_npiv" \
-members "a141_lp17_a000_hostname0;a141_lp16_a000_hostname1"

zonecreate --peerzone "a141hosts_fs95fp1_port_53" \
-principal "fs95fp1_node1_a5p3_npiv;fs95fp1_node2_a5p3_npiv" \
-members"a141_lp17_a200_hostname0;a141_lp16_a200_hostname1"
```

It is also necessary to configure connectivity within the FlashSystem storage controller itself to support various service and failover recovery scenarios.

Since these ports are dedicated for internal controller communication and are not used by external systems, aliases are not required for these zones.

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
7	50050768131345AD	50050768130045AD	3	1	01C000	no	scsi
31	50050768135345AD	50050768130045AD	11	1	01C100	no	scsi
55	50050768137345AD	50050768130045AD	19	1	01C200	no	scsi
79	50050768131345AE	50050768130045AE	3	2	012040	no	scsi
103	50050768135345AE	50050768130045AE	11	2	012140	no	scsi
127	50050768137345AE	50050768130045AE	19	2	012240	no	scsi

Figure 50 Configuring and the zone is:

The zone configuration is shown in Example 11.

Example 11 Zone configuration

```
zonecreate "fs95fp1_local_only" \
"50050768131345AD; 50050768135345AD;50050768137345AD;50050768131345AE; \
  50050768135345AE; 50050768137345AE"
```

To reiterate, these physical WWPNs for the FlashSystem storage ports are exclusively for interstorage communication (either within the chassis or between FlashSystem systems). They must never be zoned to host systems for host I/O.

The newly created zones should now be added to the zoneset. Save the changes and enable the new zoneset. See Example 12.

Example 12 Add zones to zoneset

```
cfgadd "ZONE14_0","a141hosts_fs95fp1_port_33;a141hosts_fs95fp1_port_53, \
  fs95fp1_local_only"
cfgsave
cfgenable "ZONE14_0"
```

Similarly, on our other switch, which is a Cisco, we are going to zone the hosts in Example 13 to the storage ports shown in Figure 51 on page 39.

Example 13 Create zones on the Cisco switch

```
LP17 a100 c05076e511803d9c sw two
LP17 a300 c05076e511803e14 sw two
LP16 a100 c05076e511803d98 sw two
LP16 a300 c05076e511803e10 sw two
```

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
20	50050768132745AD	50050768130045AD	7	1	9D0001	yes	scsi
44	50050768136745AD	50050768130045AD	15	1	9D0021	yes	scsi
92	50050768132745AE	50050768130045AE	7	2	9D0061	yes	scsi
116	50050768136745AE	50050768130045AE	15	2	9D0081	yes	scsi

Figure 51 Storage ports on the Cisco switch

We will begin by creating aliases to improve readability. See Example 14 on page 39.

Example 14 Creating aliases

```
config t
device-alias database
device-alias name a141_lp17_a100_hostname0 pwwn C0:50:76:E5:11:80:3D:9C
device-alias name a141_lp17_a300_hostname0 pwwn C0:50:76:E5:11:80:3E:14
device-alias name a141_lp16_a100_hostname1 pwwn C0:50:76:E5:11:80:3D:98
device-alias name a141_lp16_a300_hostname1 pwwn C0:50:76:E5:11:80:3E:10
device-alias name fs95fp1_node1_a2p3_npiv pwwn 50:05:07:68:13:27:45:AD
device-alias name fs95fp1_node2_a2p3_npiv pwwn 50:05:07:68:13:27:45:AE
device-alias name fs95fp1_node1_a4p3_npiv pwwn 50:05:07:68:13:67:45:AD
device-alias name fs95fp1_node2_a4p3_npiv pwwn 50:05:07:68:13:67:45:AE
device-alias commit
```

Now we create our new zones. The `target` argument indicates these are storage target WWPNs, and the `init` argument indicates these are host initiator WWPNs. See Example 15.

Example 15 Creating new zones

```
config t
zone name a141hosts_fs95sp1_port_23 vsan 15
member device-alias fs95fp1_node1_a2p3_npiv target
member device-alias fs95fp1_node2_a2p3_npiv target
member device-alias a141_lp17_a100_hostname0 init
member device-alias a141_lp16_a100_hostname1 init

zone name a141hosts_fs95sp1_port_43 vsan 15
member device-alias fs95fp1_node1_a4p3_npiv target
member device-alias fs95fp1_node2_a4p3_npiv target
member device-alias a141_lp17_a300_hostname0 init
member device-alias a141_lp16_a300_hostname1 init
```

As we did on the other switch will also create the local only zone for the chassis to talk to itself. Figure 52 shows the physical ports plugged to this switch.

id	WWPN	WWNN	port_id	node	nportid	virtualized	protocol
19	50050768132345AD	50050768130045AD	7	1	9D0000	no	scsi
43	50050768136345AD	50050768130045AD	15	1	9D0020	no	scsi
67	50050768138345AD	50050768130045AD	23	1	9D0040	no	scsi
91	50050768132345AE	50050768130045AE	7	2	9D0060	no	scsi
115	50050768136345AE	50050768130045AE	15	2	9D0080	no	scsi
139	50050768138345AE	50050768130045AE	23	2	9D00A0	no	scsi

Figure 52 Physical ports plugged to switch two

We create a standard zone for these ports as shown in Example 16.

Example 16 Create a standard zone

```
zone name fs95fp1_local_only vsan 15
member pwwn 50050768132345AD
member pwwn 50050768136345AD
member pwwn 50050768138345AD
member pwwn 50050768132345AE
member pwwn 50050768136345AE
member pwwn 50050768138345AE
```

To reiterate again, those are the physical WWPNs of the FlashSystem, and are only ever zoned to each other or other FlashSystem physical WWPNs. They are not for host IO.

Finally, we add our new zones to a zoneset, activate the changes, and save them. See Example 17.

Example 17 Add our new zones to a zoneset, activate the changes, and save them

```
zoneset name ZONE15_0 vsan 15
member a141hosts_fs95sp1_port_23
member a141hosts_fs95sp1_port_43
member fs95fp1_local_only
zoneset activate name ZONE15_0 vsan 15
end
copy running-config startup-config
```

As of this point, we have completed the zoning of our IBM LinuxONE or Z system's NPIV host ports to our FS9500 storage ports on our Brocade and Cisco switches.

While the initial configuration includes only two hosts in these zones, Peer/Smart Zones offer the significant benefit of simplified scaling. To provide additional hosts with access to the storage controller, it is sufficient to add their respective NPIV WWPNs to the pre-existing zones. This method also streamlines the identification of initiator ports communicating with a particular target port. Rather than searching across numerous zones to locate a specific WWPN and count the single initiator zones where it appears, one can simply locate the relevant zone containing the storage WWPN and count the number of initiator WWPNs within that zone.

Creating host definitions and mapping volumes in the storage controller

The FS9500 offers both a web UI and a CLI. These examples will utilize the CLI, as it provides a more robust foundation for automation.

Recall that our two LPARs NPIV WWPNs are listed in Example 18 and Example 19 on page 41.

Example 18 WWPNs for LP17

```
lpardevicewwpn
LP17 a000 c05076e511803d60
LP17 a200 c05076e511803dd8
LP16 a000 c05076e511803d5c
LP16 a200 c05076e511803dd4
```

Example 19 WWPNs for LP16

```
lpardevicewwpn
LP16 a000 c05076e511803d5c
LP16 a100 c05076e511803d98
LP16 a200 c05076e511803dd4
LP16 a300 c05076e511803e10
```

Host definitions must be created on the FS9500 storage controller before volumes can be mapped to host systems. This presents a potential bootstrapping challenge when the target systems boot from these volumes. The web UI offers a simplified method for creating host definitions by presenting a list of previously unseen WWPNs that have connected to the storage system. Selecting these "unknown" WWPNs creates a new host definition, enabling volume mapping. However, if the target systems have not yet been zoned to the storage or lack a running operating system, these WWPNs will not be available for selection.

Given that our zone creation process involves offline WWPNs, we will maintain consistency by using the FS9500 CLI. We use the `mkhost` command with the `-force` parameter because the FlashSystem code will attempt to validate all the supplied WWPNs which will fail if the LPAR is not Activated with an OS Loaded and running and that OS has configured the devices online. See Example 20.

Example 20 mkhost command

```
mkhost -name A141_LP17_hostname0 -protocol fcscsi -fcwwpn \
c05076e511803d60:c05076e511803d9c:c05076e511803dd8:c05076e511803e14 -force
mkhost -name A141_LP16_hostname1 -protocol fcscsi -fcwwpn \
c05076e511803d5c:c05076e511803d98:c05076e511803dd4:c05076e511803e10 -force
```

Then, we map an existing volume to each host definition:

```
mkvdiskhostmap -host A141_LP17_hostname0 <some_volume_id>
mkvdiskhostmap -host A141_LP16_hostname1 <some_other_volume_id>
```

This completes the basic storage configuration on the SAN storage controller for simple volume mappings (one volume per host). For clustered file systems such as IBM Storage Scale (formerly GPFS), mapping a volume or set of volumes to multiple hosts requires the

`-force` option or the `mkhostcluster` command. These methods enable the FlashSystem to gracefully manage shared volume access across multiple hosts.

Booting new systems from SAN volumes while simultaneously sharing a set of volumes across a host cluster necessitates careful planning and design. Allowing multiple systems to access each other's boot volumes using non-clustered file systems, such as ext4, xfs, or btrfs, must be strictly avoided. Inadvertent concurrent access will invariably corrupt the data on these volumes, potentially rendering the operating system unusable if the affected volume is a boot volume.

One method to implement a "shared volumes, but not all" configuration is to utilize the `-ignoreseedvolume` parameter when creating the host cluster, thereby preventing shared access to boot volumes. However, for maximum assurance, it is recommended to employ multiple host adapters per channel, dedicating one set for operating system volume access and another set for cluster volume access.

For example, on LP17 we would continue to use the ports in Example 21 for access to the boot volume:

Example 21   Ports for access to the boot volume

```
lpardevicewwpn
LP17 a000 c05076e511803d60
LP17 a100 c05076e511803d9c
LP17 a200 c05076e511803dd8
LP17 a300 c05076e511803e14
```

And we could use the devices in Example 22 for access to a set of shared cluster volumes.

Example 22   Devices

```
lpardevicewwpn
LP17 a001 c05076e511803d74
LP17 a101 c05076e511803db0
LP17 a201 c05076e511803dec
LP17 a301 c05076e511803e28
```

This approach requires additional zoning, host definitions, and LUN mappings for the a*01 devices, but ensures that accidental access to host boot volumes is prevented by the explicit configuration of the host definitions on the storage controller.

Verifying SAN configuration using HMC functions

This concludes the end-to-end configuration of the storage system:

- NPIV is enabled on both the host and storage ports
- The host and storage NPIV ports are zoned to each other
- The host is defined in the storage system and it has a LUN mapped to it

Now let us use the HMC web interface to examine the SAN configuration to see if we can determine if our changes are correct.

Access the Channels view using the HMC's **Single Object Operations** function or, for z15 or later systems, by selecting the **Adapters** tab while the CPC is selected. Refer to the beginning of the **Enabling NPIV on FCP Channels** section for navigation instructions. See Figure 53 on page 43.

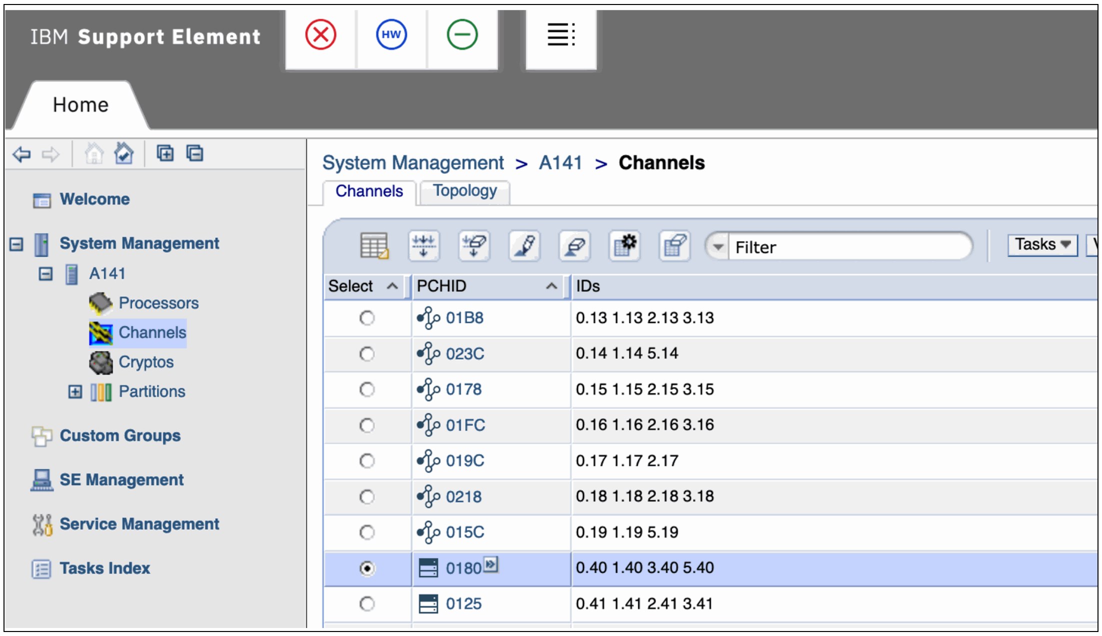

Figure 53 Channels view

With the channel selected, run the **Channel Problem Determination** action in the **Service** menu. See Figure 54.

Figure 54 Channel Problem Determination action

From the displayed menu, select the LPAR for which you want to verify connectivity, and then click **OK**. See Figure 55 on page 44.

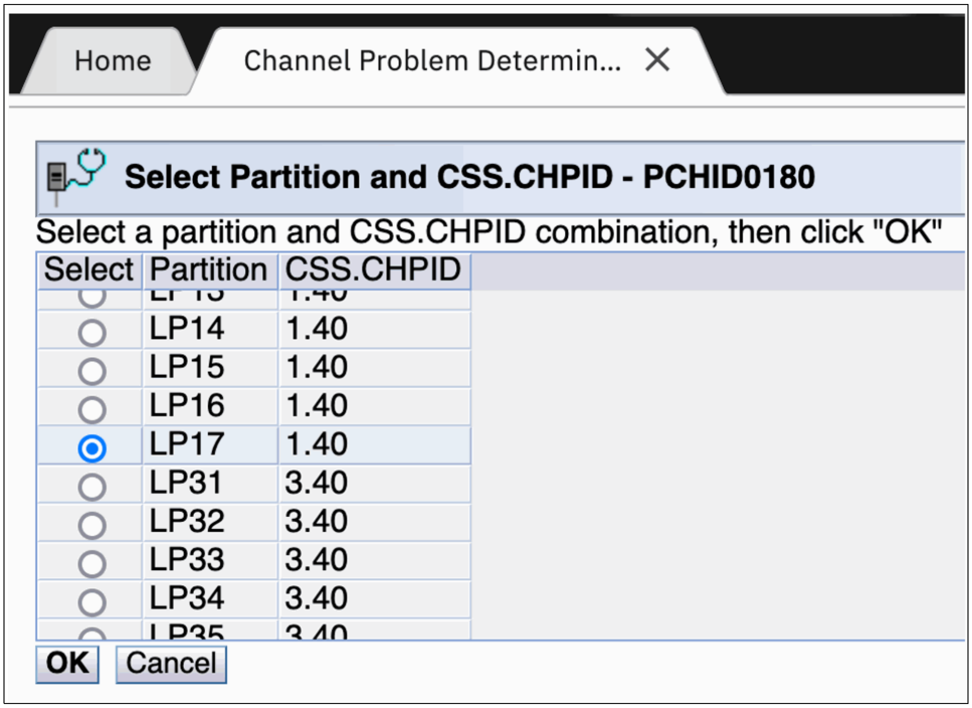

Figure 55 IBM Support Element menu

In the following menu select **SAN explorer** and click **OK**. See Figure 56.

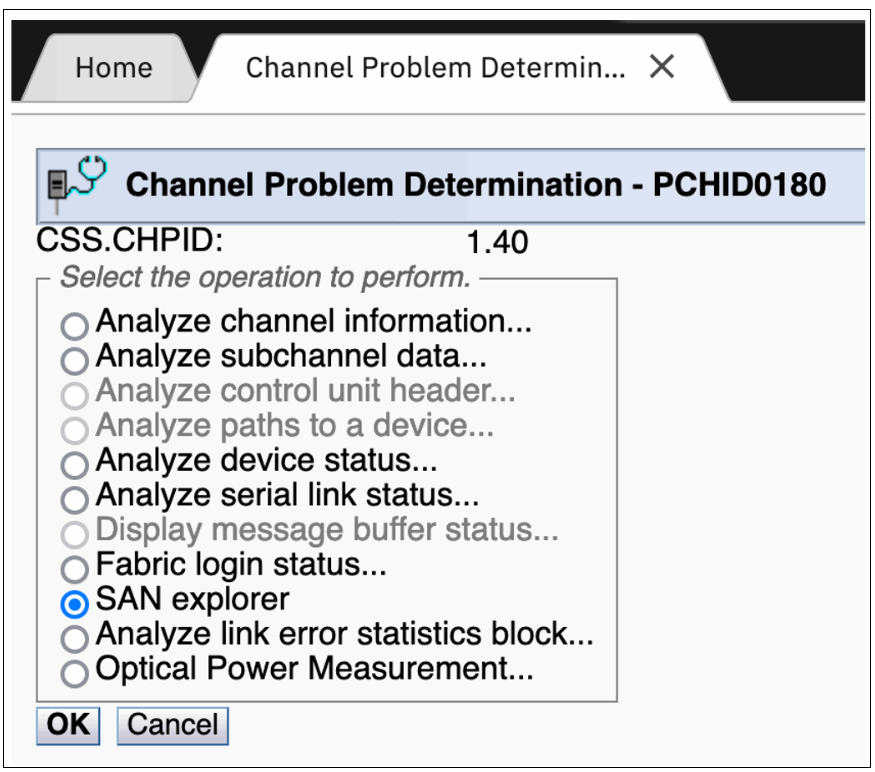

Figure 56 Select SAN explorer

At this juncture, the examples diverge from a "from-scratch" configuration. These examples were generated using an existing, operational configuration rather than a blank, unconfigured system. This difference is first apparent in Figure 57 on page 45.

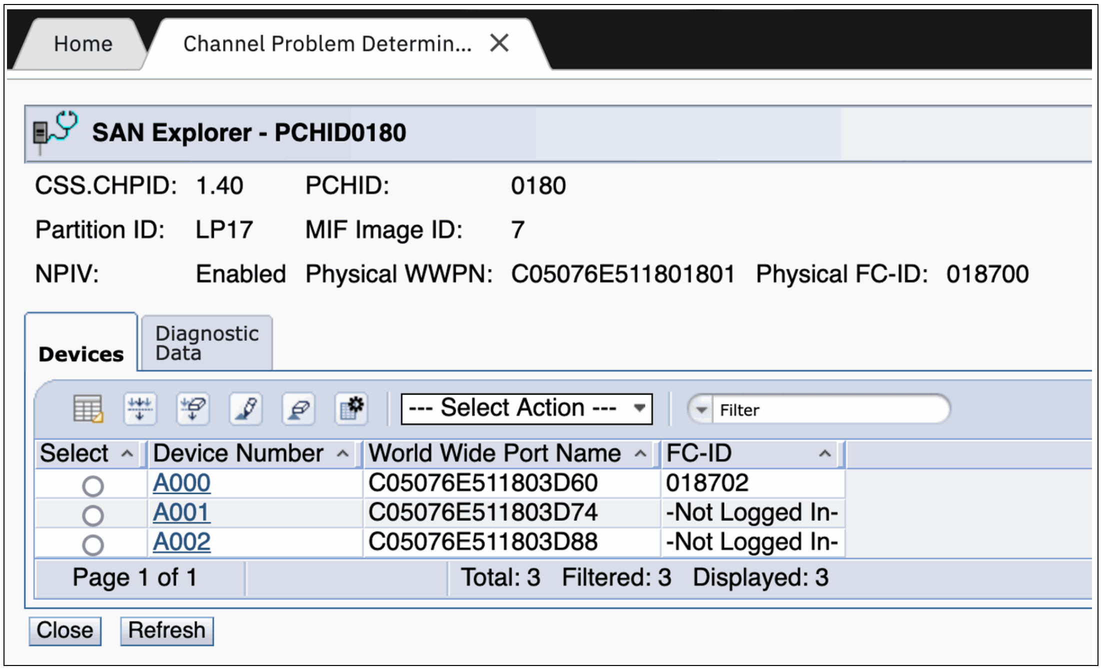

Figure 57 SAN Explorer menu

The presence of an FC-ID for device A000 indicates that it was actively registered with the SAN fabric at the time this panel was accessed. This implies that LP17 had a running operating system with device A000 varied online.

Prior to this point, the displayed panels utilized the current I/O configuration to populate menus and guide subsequent actions.

In Figure 57, selecting an LPAR's channel in SAN Explorer causes the system to examine the state of each FCP device on that channel within the LPAR. The NPIV World Wide Port Name (WWPN) of each device is then displayed, along with its current registration status with the SAN fabric. If the device is active and registered, its corresponding FC-ID is shown. Otherwise, the status "-Not Logged In-" is displayed.

If your LPAR is not running you will see a panel that looks more like Figure 58 on page 46.

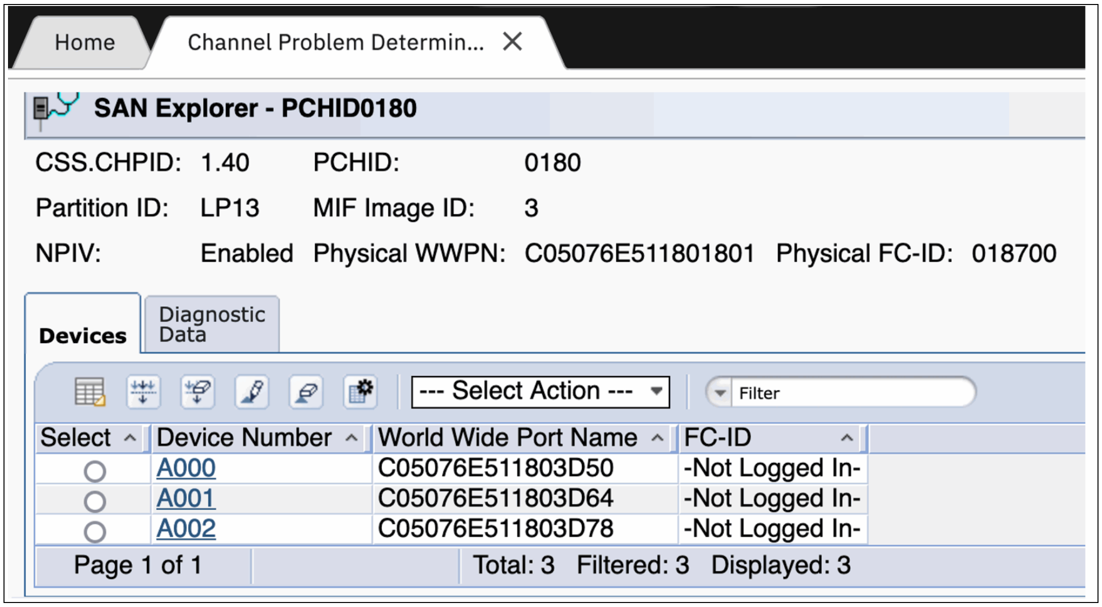

Figure 58   SAN Explorer menu

Going back to our LP17 panel with our running Linux OS in it, click on the **A000** device number. See Figure 59.

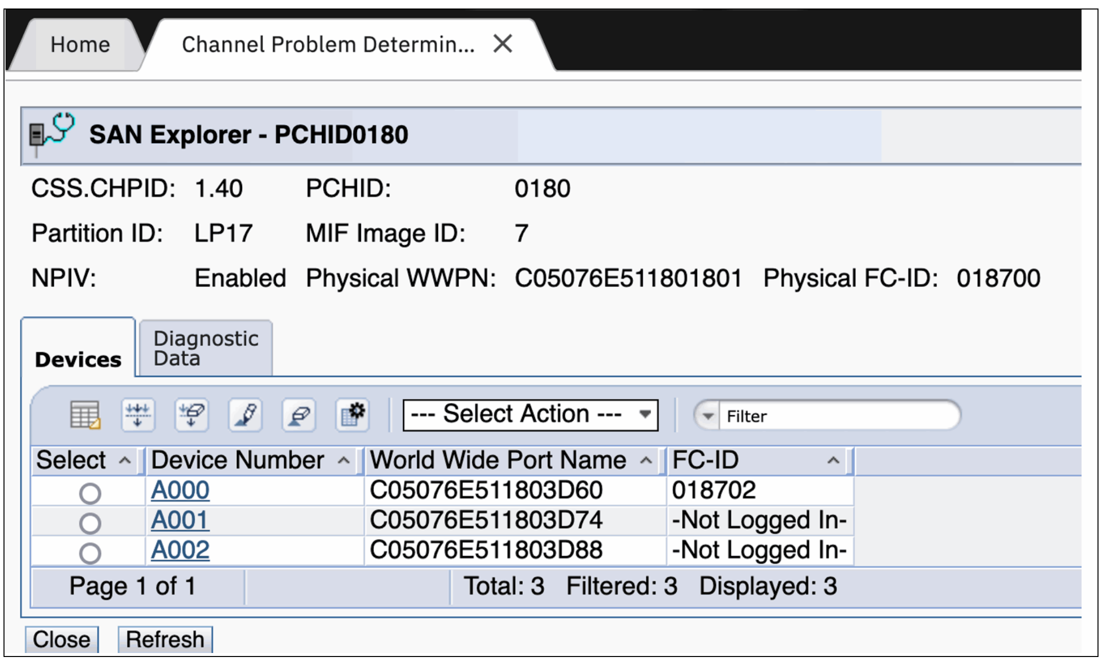

Figure 59   Click on the A000 device number

Selecting the LPAR's channel triggers the SAN Explorer function in the Service Element (SE) to initiate a Fabric Login from the LPAR's A000 device. The SE then queries the SAN Name Server for all accessible zoned ports before performing a Fabric Logout. A correctly zoned configuration will result in a display similar to Figure 60 on page 47.

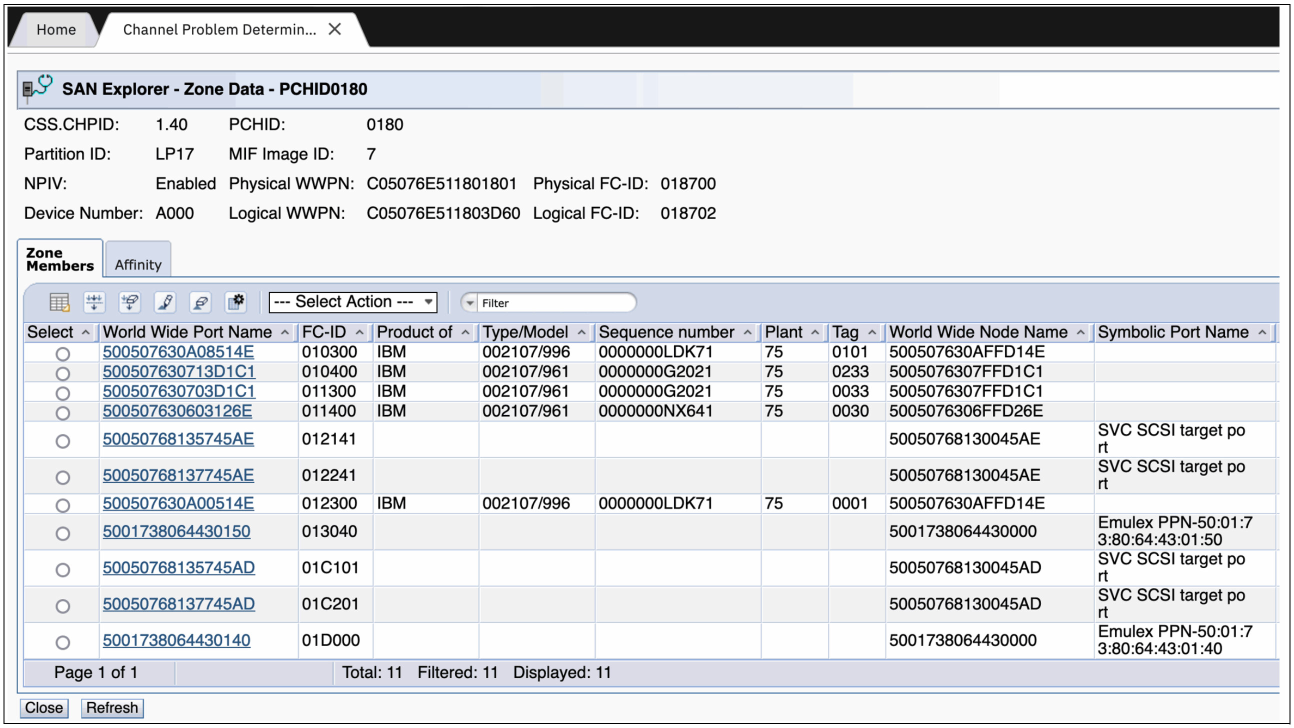

Select	World Wide Port Name	FC-ID	Product of	Type/Model	Sequence number	Plant	Tag	World Wide Node Name	Symbolic Port Name
○	500507630A08514E	010300	IBM	002107/996	0000000LDK71	75	0101	500507630AFFD14E	
○	500507630713D1C1	010400	IBM	002107/961	0000000G2021	75	0233	5005076307FFD1C1	
○	500507630703D1C1	011300	IBM	002107/961	0000000G2021	75	0033	5005076307FFD1C1	
○	500507630603126E	011400	IBM	002107/961	0000000NX641	75	0030	5005076306FFD26E	
○	50050768135745AE	012141						50050768130045AE	SVC SCSI target po rt
○	50050768137745AE	012241						50050768130045AE	SVC SCSI target po rt
○	500507630A00514E	012300	IBM	002107/996	0000000LDK71	75	0001	500507630AFFD14E	
○	5001738064430150	013040						5001738064430000	Emulex PPN-50:01:7 3:80:64:43:01:50
○	50050768135745AD	01C101						50050768130045AD	SVC SCSI target po rt
○	50050768137745AD	01C201						50050768130045AD	SVC SCSI target po rt
○	5001738064430140	01D000						5001738064430000	Emulex PPN-50:01:7 3:80:64:43:01:40

Figure 60 SAN Explorer - Zone data

The following example deviates from the ideal configuration presented in "An example configuration" on page 9. In contrast to the two zoned ports to the FS9500 illustrated in that example, the actual system configuration includes four zoned ports to the FS9500, along with two ports to each of two DS8000 systems, one port to a third DS8000, and two ports to an XIV®.

This discrepancy serves as an example of practical implementation differing from ideal recommendations.

The recommended configuration for FS9500 connectivity is four paths to the first node and four paths to the second node. The zoning diagrams and procedures outlined in the "Creating SAN zones to permit host ports to communicate with storage ports" on page 32" section accurately depict this recommended configuration. The actual configuration shown Figure 60 provides an excessive number of paths, which is suboptimal for both the Linux multipath driver and the FS9500 itself.

If an NPIV WWPN is not zoned to any target ports, an empty panel, as shown Figure 61 on page 48, will be displayed.

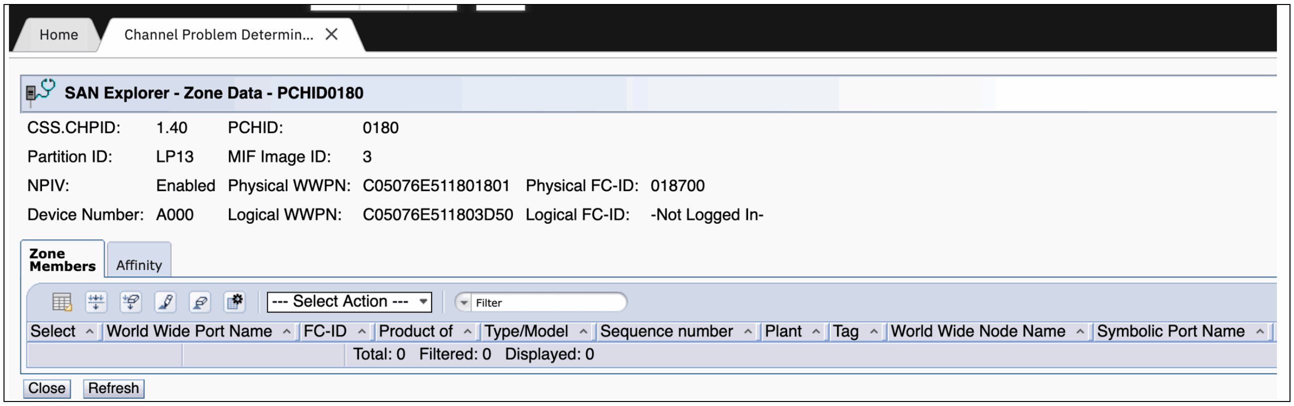

Figure 61 Empty panel

If the selected FCP device's NPIV WWPN within the specified LPAR and channel is not zoned to any target ports, the panel will display no entries. If the panel displays unrecognized WWPNs, the NPIV Logical WWPN for the selected device is incorrectly zoned and requires correction by the SAN administrator before proceeding.

Returning to the primary example, select one of the displayed target WWPNs corresponding to your storage controller to view the presented LUNs. See Figure 62.

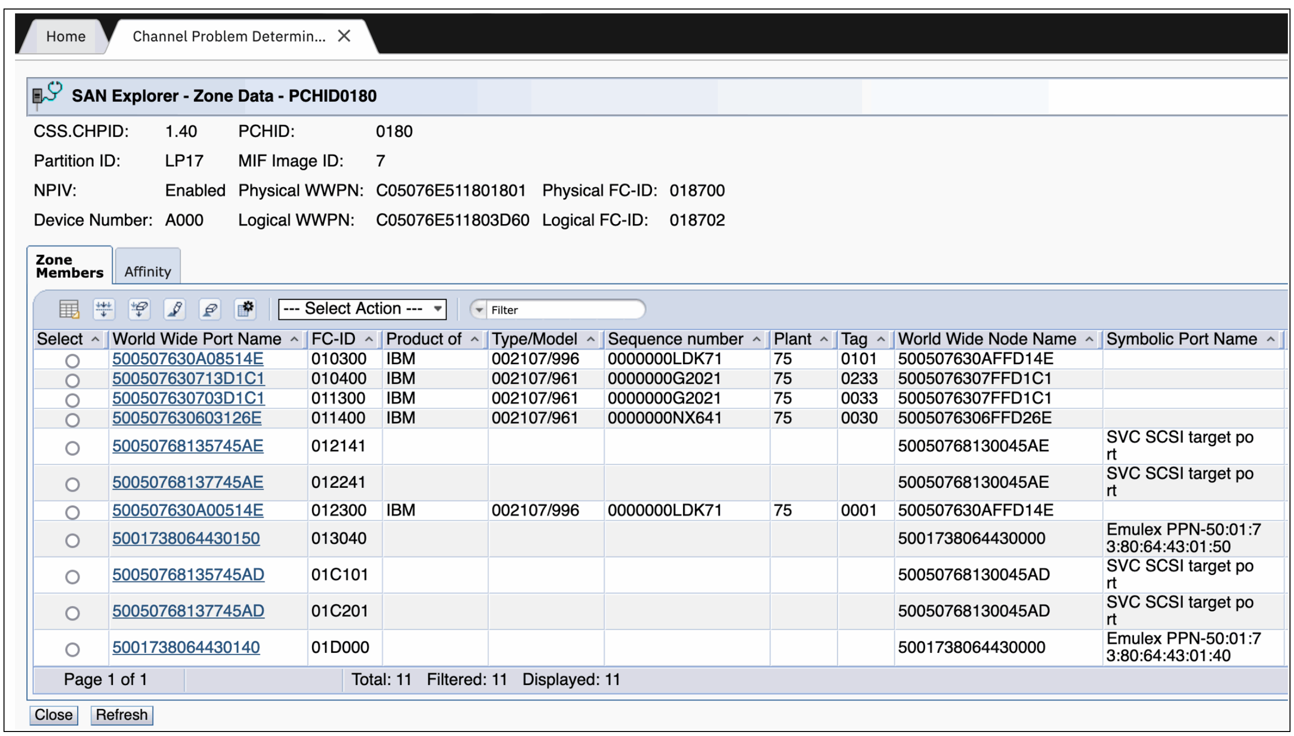

Figure 62 SAN Explorer - Zone Data

This action displays the LUNs, if any, that the storage controller has assigned to the host. The displayed LUNs are determined by the host definition and the corresponding LUN map configured within the storage controller. See Figure 63 on page 49.

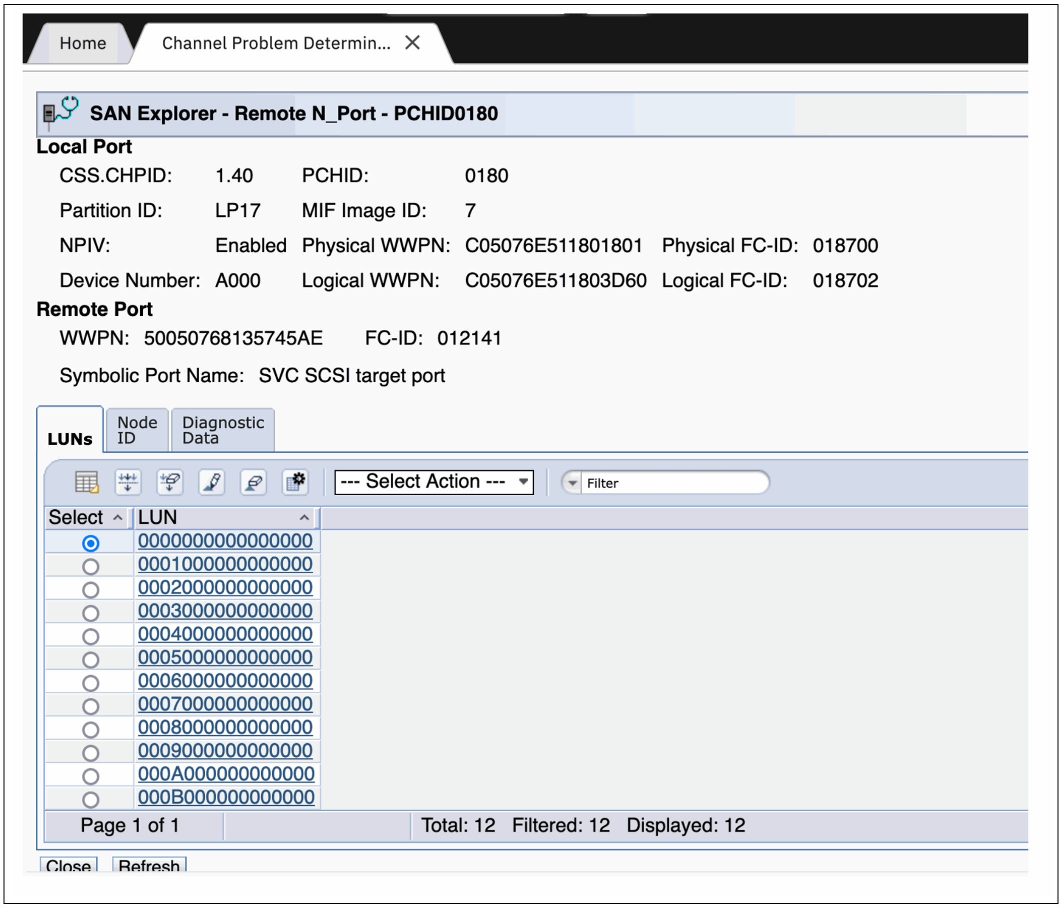

Figure 63 LUNs

We can see that our FlashSystem is presenting 12 LUNs: 0-B.

Note: IBM FlashSystem and SVC devices, by default, present Logical Unit Numbers (LUNs) sequentially starting at LUN 0. On IBM LinuxONE and Z systems, the LUN is represented by a 16-digit hexadecimal field, with the rightmost 12 digits set to 0. Storage administrators retain the option to explicitly assign LUNs during volume mapping on the storage controller, thereby overriding the default sequential assignment.

In contrast, IBM DS8000 systems derive LUNs based on the Logical Control Unit (LCU) and unit address, padded with the hexadecimal value '40' and with only the rightmost 8 digits set to 0. For example, a volume at unit address 55 within LCU 33 on a DS8000 would be assigned the LUN 4033405500000000. On DS8000 systems, the LUN presented to the host is determined by the system and cannot be modified by the storage administrator.

Click on one of the 16 digit LUN numbers. See Figure 64 on page 50.

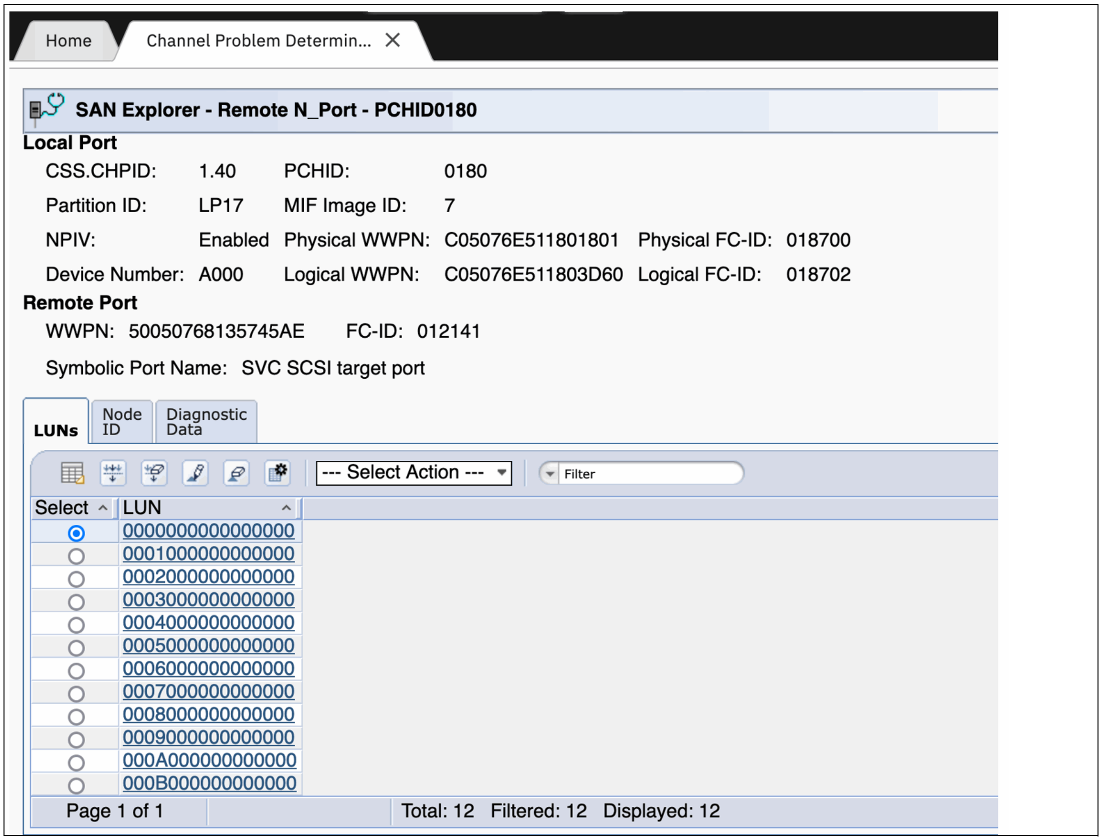

Figure 64

This opens the next screen with details about that LUN. See Figure 65 on page 51.

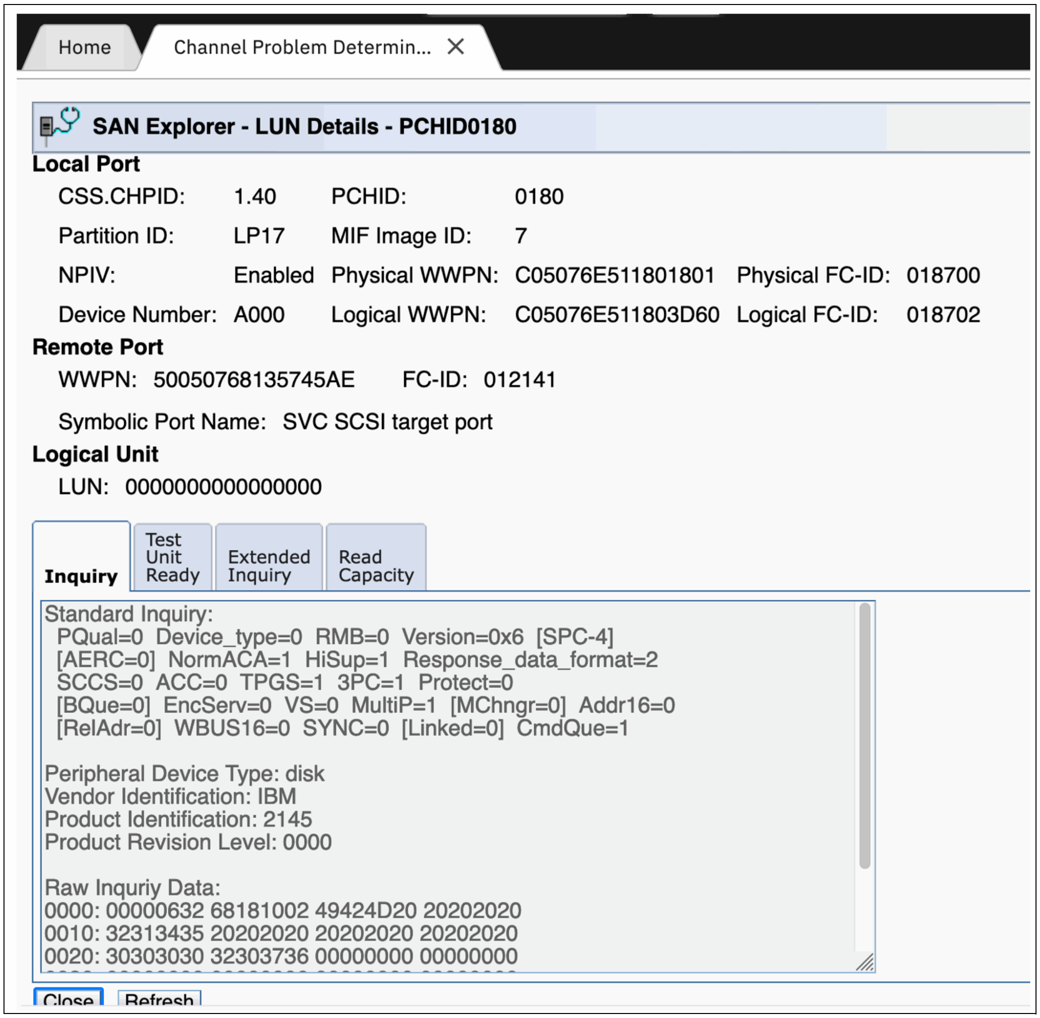

Figure 65 LUN details

In this screen we can see the LUN details that SAN explorer was able to retrieve from the FlashSystem when we clicked the LUN number on the previous panel. We can verify the Vendor and Product version strings in the Inquiry tab as shown in Figure 65.

If you click the **Test Unit Ready** tab you should see “Ready” in the big data field if the volume is ready for IO.

If you click the **Read Capacity** tab you will see capacity details about the device. See Figure 66 on page 52.

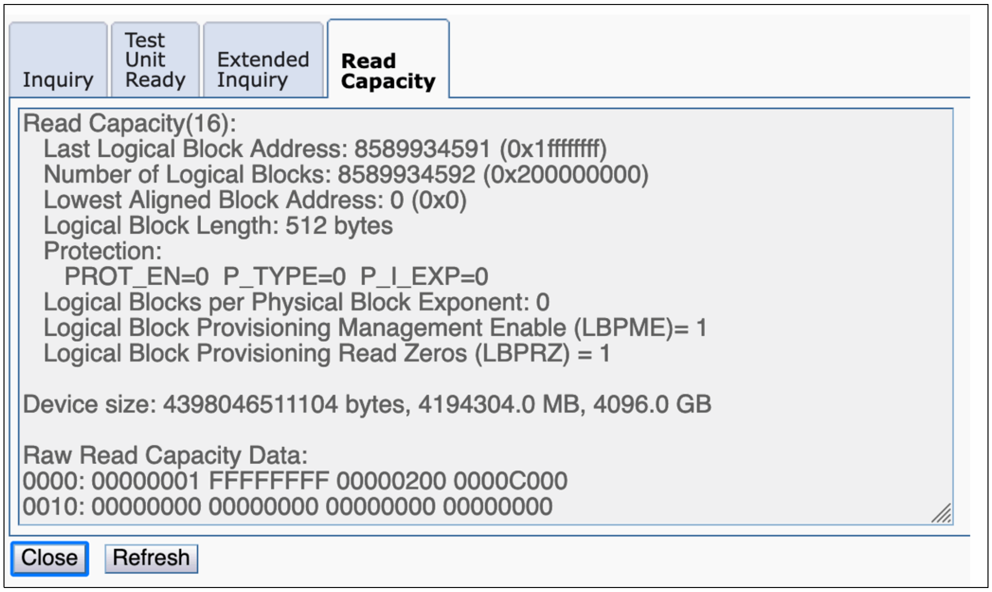

Figure 66 Read Capacity

Click **Close** until you get back to the Zone Data panel, and then click the other **WWPN** for your FlashSystem controller. See Figure 67.

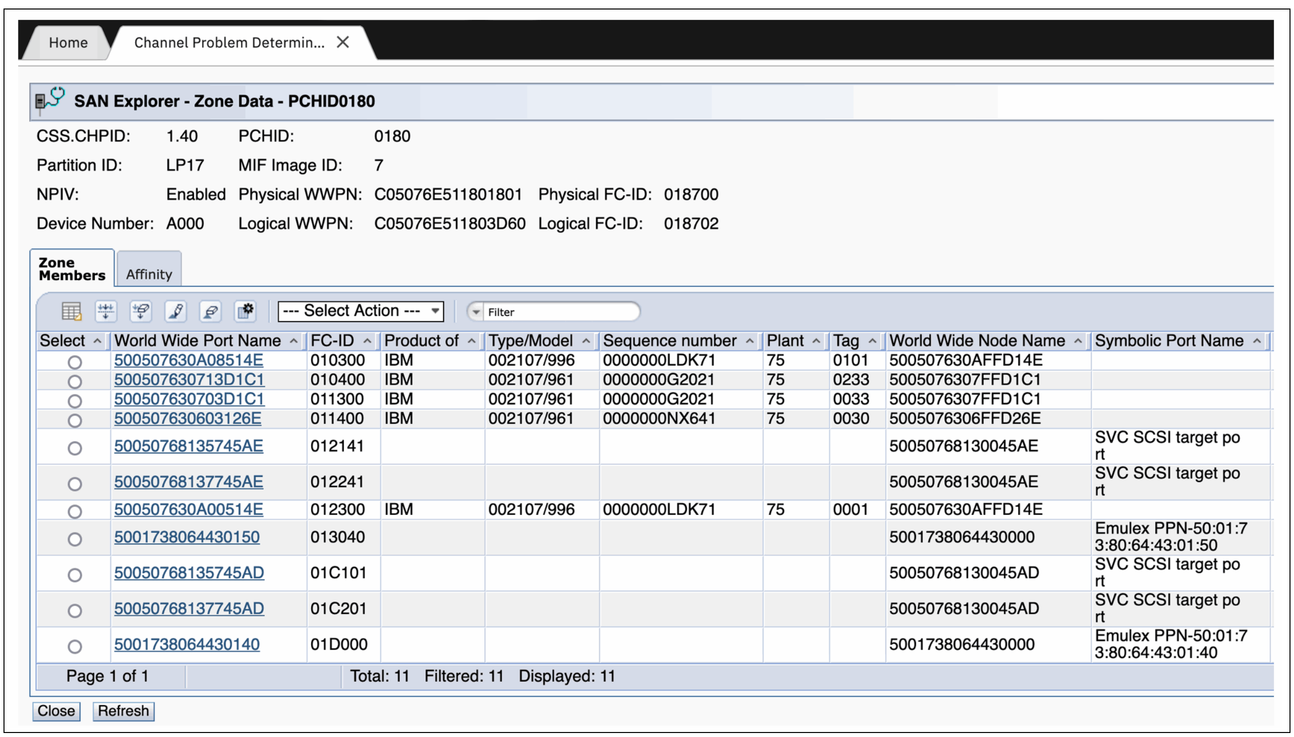

Figure 67 Zone Data panel

That will open up the LUN list view. See Figure 68 on page 53.

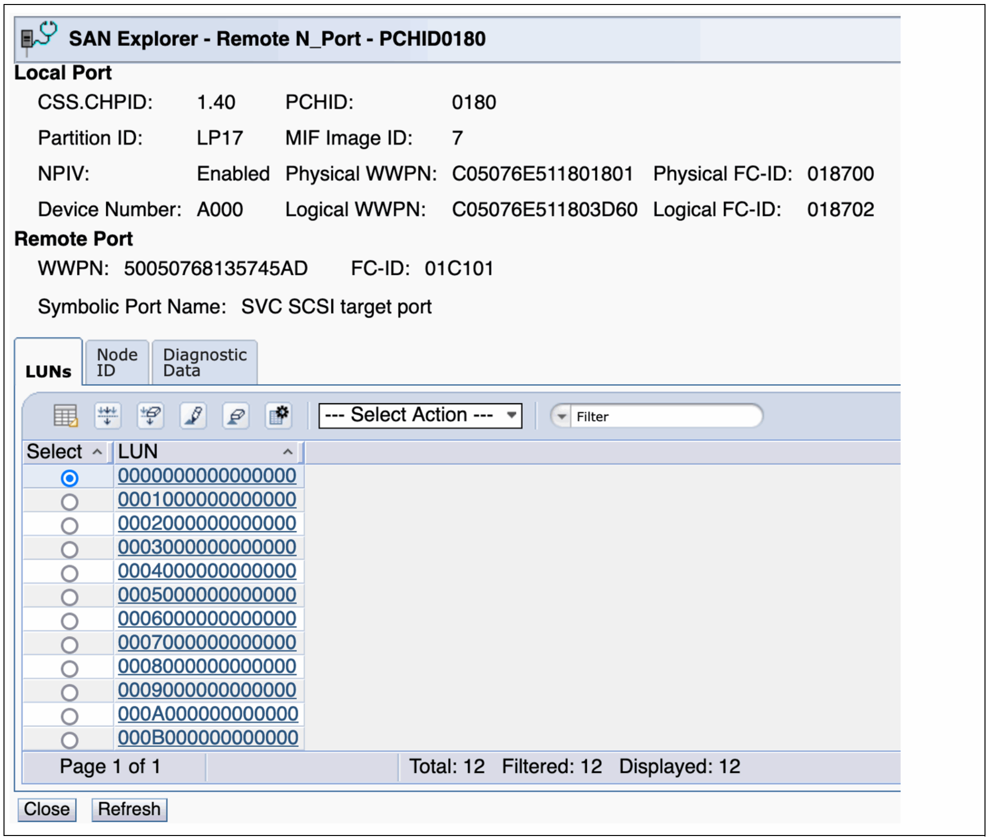

Figure 68 LUN list view

Verify that you see all your expected LUNs on this path to the other node as well.

Click **Close** or **Cancel** until you are back at the top level "Channel" or "Adapters" listing, and select another FCP Channel and repeat the process to make sure its ports are all zoned as required. For each storage WWPN also compare the LUNs provided by the storage controller to make sure they are consistent across all configured paths.

At this point we have used the HMC San explorer task to verify all our configured zones in our SAN switches as well as our host definitions and LUN maps in the storage controller.

Summary

In this Redpaper we have:

- Provided an overview of FCP, and a little bit of how it contrasts versus FICON.
- Provided an example IO configuration of FCP channels shared by multiple LPARs.
- Provided an example of how to enable NPIV on FCP channels, and discover the generated WWPNs.
- Provided an example of how to discover the correct WWPNs to use for zoning on a FS9500.
- Provided an example for both Cisco and Brocade zoning using the new zoning features.
- Provided an example for creating Host definitions and LUN maps on a FS9500.

- Provided an example of using the SAN explorer HMC function to ensure that all of the above was done as expected.

At this point the LPARs SAN configurations should be usable for Linux. Refer to your Linux flavour's documentation for details on how to configure that.

Appendix - Handy switch command cheat sheets

Brocade command cheat sheet:

- Check if FCP dev is logged in (This will also show if the FCP device has an alias):

  ```
  nodefind <fcp_dev_wwpn> or <alias>
  ```

- Get WWPN for a given alias:

  ```
  alishow <alias> <---Can be wildcarded, for example, alishow *fs95fp*
  ```

- Find zone config:

  ```
  zoneshow <zone_name> <---Can be wildcarded for example zoneshow *fs95fp*
  ```

- Find if fcp_dev is in a zone:

  ```
  zoneshow | grep <wwpn>
  ```

- Find if alias is in a zone:

  ```
  zoneshow --alias <alias_name>
  ```

- Alias management:
  - Create alias:

    ```
    alicreate "<alias-name>", "<fcp_dev_wwpn>"
    ```

  - Change alias:

    ```
    zoneobjectrename "<old_alias_name>", "<new_alias_name>"
    ```

  - Remove alias:

    ```
    alidelete "<alias-name>"
    ```

- Create zone:

  ```
  zonecreate --peerzone <zone_name> -principal
  "<storage_controller_1;storage_controller_2> -members
  "<alias_name_1;alias_name2>"
  ```

- Add zone to zoneset:

  ```
  cfgadd <zoneset>,<newly_added_zone_name>
  ```

- Save changes:

  ```
  cfgsave -f
  cfgenable <zoneset>
  ```

Check if zone is in active config:

```
cfgactvshow | grep <zone_name>
```

Cisco command cheat sheet:

- Check if FCP dev is logged in:

```
show fcs port pwwn <wwpn> vsan <vsan_number>
```

- Check if alias exists for a FCP dev:

```
show device-alias database | grep <your:wwpn:number>
```

- Get WWPN for a given alias:

```
show device-alias name <alias_name>
```

- Find if alias or fcp_dev is in a zone:

show zone | grep <alias_name or wwpn>

- Get configuration for a given zone:

```
show zone name <zone_name>
```

- Alias management:

```
config -t
device-alias database
```

  - Create alias:

```
device-alias name <alias-name> <your:wwpn:number>
```

  - Change alias:

```
device-alias rename <old-alias> <new-alias>
```

  - Remove alias:

```
no device-alias name <alias-name>
```

  - Save alias changes:

```
device-alias commit
```

- Create zone:

```
config -t
zone name <zone-name> vsan <vsan_number>
```

- Add members to zone:

```
member device-alias <alias-name> init <--(if the alias is an FCP device) target
(if the alias is a npiv storage device)
```

- Remove members to zone:

```
no member device-alias <alias-name>
```

- Adding zones to a zoneset and activating changes:

```
zoneset <zoneset> vsan <vsan_number>
member <new_zone_name>
zoneset activate name <zoneset> vsan <vsan_number>
```

Author

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

Jay Brenneman is an IBM Senior Technical Staff Member who leads the IBM LinuxONE Hybrid Cloud Integration Test team, based in Poughkeepsie, NY. He has extensive experience testing Linux on IBM z Systems® (previously known as s390x), starting with the first release of SuSE Linux 7 that included s390 support in the year 2000.

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

Stay connected to IBM Redbooks

- Find 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

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.

DS8000®
FICON®
IBM®
IBM FlashSystem®
IBM Z®
IBM z Systems®
Redbooks®
Redbooks (logo) ®
XIV®
z Systems®
z/VM®
z15®
z16™

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.

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

REDP-5751-00

ISBN

Printed in U.S.A.