Redbooks Paper

Frederick Wong

Introducing N_Port Identifier Virtualization for IBM System z9

Introduction

This IBM® Redpaper describes the N_Port ID Virtualization (NPIV) feature in a Fibre Channel Protocol (FCP) fabric. NPIV is a new feature available with IBM System z9™ 109 servers. NPIV allows a single FCP port to register multiple Worldwide Port Names (WWPN) with a fabric name server. Each registered WWPN is assigned a unique N_Port ID. With NPIV, a single FCP port can appear as multiple WWPNs in the FCP fabric. This Redpaper focuses on NPIV concepts and implementation; for details about general FCP implementation, see *Linux on zSeries: Fibre Channel Protocol Implementation Guide*, SG24-6344.

Overview of the NPIV feature

System z™ FCP channels require a FICON® Express adapter. FCP channels can be shared by multiple LPARs. Each port on the adapter is assigned a permanent 64-bit WWPN by the manufacturer; this is used at Fabric Login (FLOGI).

Without the NPIV feature, each operating system image that has an FCP port is identified to the fabric by the permanent WWPN of the port. In this case, all operating system images have the same access rights in the fabric. The permanent WWPN of the port determines:

- Zone membership for all images sharing the port
- Logical Unit Number (LUN) access rights for all images sharing the port

With the NPIV feature, the Service Element (SE) creates new WWPNs for the FCP port at FLOGI. A unique WWPN is then assigned to each operating system image sharing the port. The generated NPIV WWPN is registered with the fabric switch and uniquely identifies each image for fabric zoning and LUN masking. Figure 1 on page 2 illustrates the NPIV feature.

Note: A, FCP port can be shared by LPARs running in either NPIV mode or non-NPIV mode. The IOCDS is same for both modes.

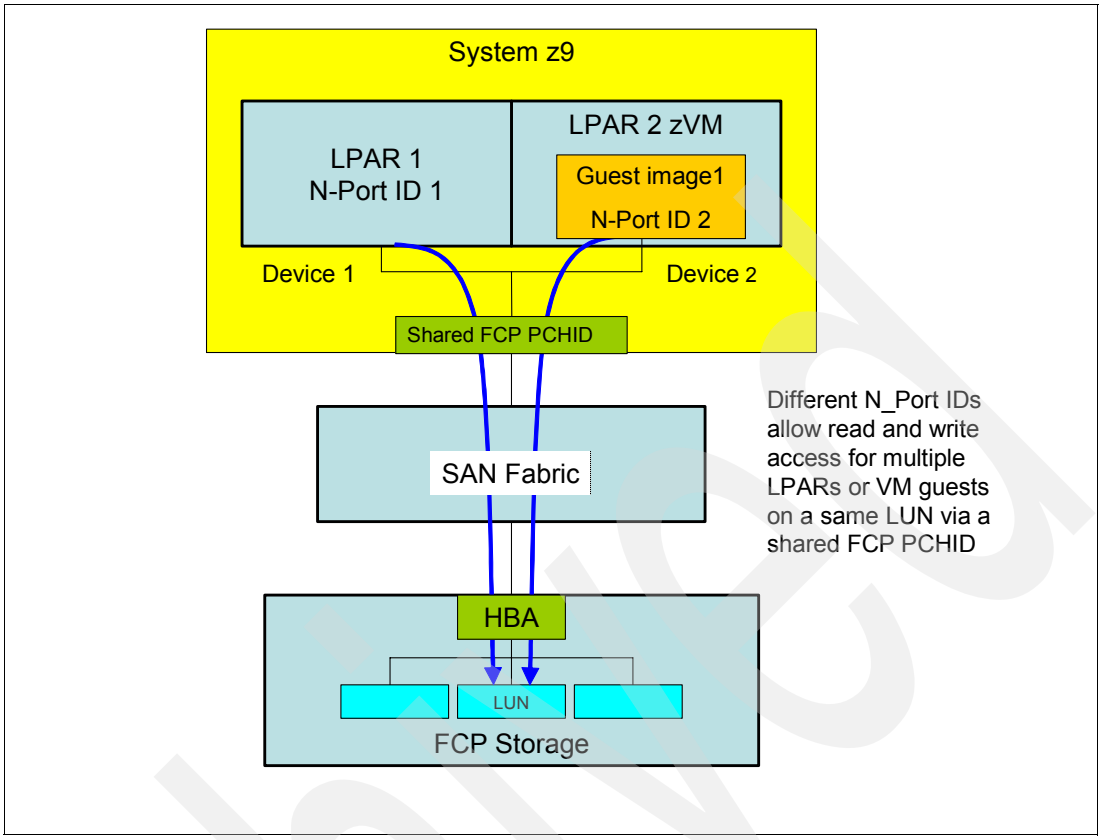

Figure 1 NPIV provides unique WWPNs to servers sharing an FCP port

In the figure, two LPARs share a single physical FCP port. Each instance registers with the name server. The NPIV WWPN is supported in the Fabric Discovery (FDISC) process.

During Power On Reset (POR) or dynamic I/O activation, each FCP subchannel is assigned a WWPN by the Support Element (SE) regardless of whether the LPAR is NPIV-enabled. If the LPAR is not enabled for NPIV, the microcode does not use the NPIV WWPNs. The SE retains, on its hard drive, the information about the assigned WWPN (to prevent the data from being lost if the system is shut down or the FCP adapter is replaced).

Each LPAR receives a different N_Port ID. This allows multiple LPARs or VM guests to read and write to the same Logical Unit Number (LUN) using the same physical port. Without NPIV, writing to the same LUN over a shared port is not allowed.

Format of the NPIV WWPN

An NPIV WWPN is 64 bits in length. Its format is shown in Figure 2 on page 3.

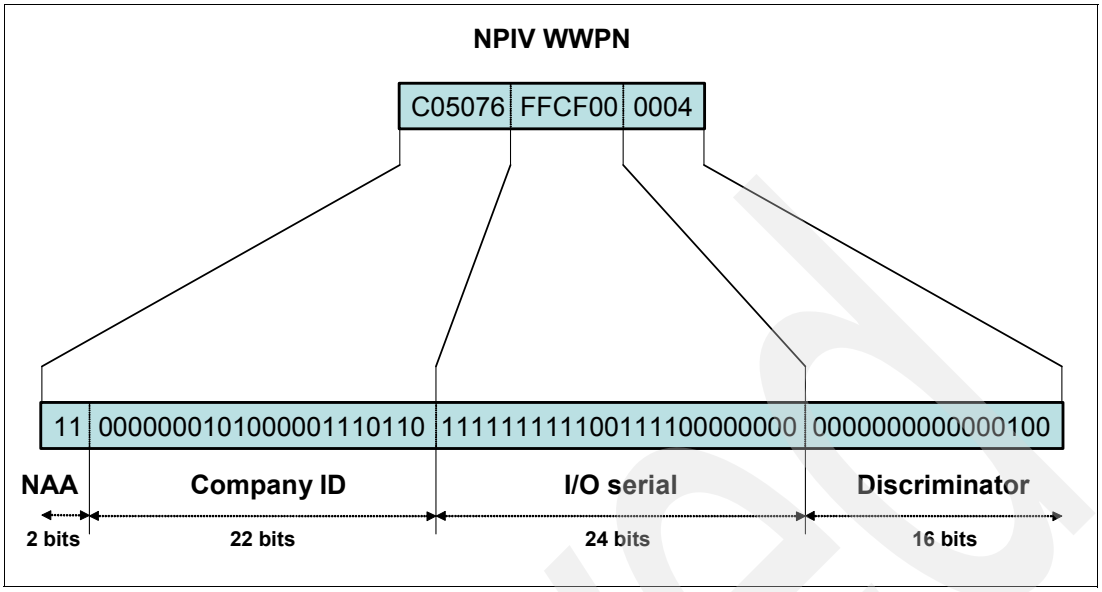

Figure 2 Format of an NPIV WWPN

The fields for an NPIV WWPN are:

- Network Address Authority (NAA)
  This two-bit field is always set to binary “11”.
- Company ID
  The 22-bit field identifies the server manufacturer and is assigned by IEEE. For NPIV, this value differs from the company ID field in the permanent WWPN of the adapter. The value of the company ID is the same for IBM System z, System p™, and System i™ servers.
- I/O serial
  The 24-bit field uniquely identifies a specific server. A single I/O serial number is assigned to each System p and System i server. System z servers are assigned a range of 128 consecutive I/O serial numbers.
- Discriminator
  The 16-bit NPIV discriminator field is used to generate a unique NPIV WWPN when sharing an FCP port.

A 16-bit discriminator provides up to over 64 000 NPIV WWPNs for a shared FCP port. Although this is sufficient for both System p and System i servers, System z systems require a larger address space. To increase the address space, the low order 7 bits of the I/O serial field are combined with the 16-bit discriminator field. The resulting 23-bit effective discriminator field increases the address space to over 8 000 000 unique WWPNs for a shared FCP port.

Prerequisites

NPIV is available on System z9 servers. A FICON Express 2 adapter running with MCL003 on EC J99658 is required. Currently, two FICON adapters are available:

- FC 3319: QD8L Ficon Express2 LW 3.0 2G (4 Ports), Single mode fiber
- FC 3320: QD8A Ficon Express2 SW 3.0 2G (4 Ports), Multi-mode fiber

NPIV is supported for Linux® systems running SUSE Linux Enterprise Server 9 (SLES9) with SP3. The Linux system can run in either LPAR mode or as a z/VM® guest. When the system is running under z/VM, the following z/VM versions are supported:

- z/VM 5.2
- z/VM 5.1 with the PTF for APAR VM63744

z/VM 5.2 supports Linux installation to FCP- attached SCSI disks that are accessed by NPIV-managed WWPNs.

Configuration considerations

Some general recommendations for using NPIV include:

- Do not use more than 32 subchannels per physical channel in NPIV mode. Also, do not perform more than 128 total target logins (for example, in a configuration with 32 subchannels, limit the number of target logins to no more than an average of 4). Using more subchannels, target logins, or both can create timeouts.
- Zone each NPIV WWPN individually. This can reduce fabric traffic because all participants in the zone are notified when another N_Port joins or leaves the zone.
- Consider using multipathing performance and availability concern (multipathing is discussed in *Linux on zSeries: Fibre Channel Protocol Implementation Guide*, SG24-6344).
- Enable NPIV on the SAN switch before enabling it on the System z9 server. If NPIV is not enabled on the switch, the attempt to establish a connection to the fabric will fail for all subchannels that are operated in NPIV mode.
- Be aware that each login from a NPIV-mode subchannel into a storage subsystem counts as a separate host login. The IBM TotalStorage® Enterprise Storage Server® (ESS) model 800 supports up to 16 Host Bus Adapters (HBAs). Each HBA supports up to 124 host logins. The ESS itself supports up to 512 logins. Consult the vendor's documentation for limits in your configuration.
- Switches typically limit the number of supported N_Port IDs. Because each NPIV WWPN is assigned an N_Port ID at login, this limit can be exceeded.
- Some switches limit the number of N_Port IDs that can be assigned to a physical port (this limit can be configurable on some switches).
- FCP microcode MCL003 on EC J99658 requires a special activation procedure. It provides the minimum microcode level (6.0) needed to support the NPIV feature on the z9-109. If your system is running FCP on a FICON Express 2 adapter, all FCP PCHIDs should be configured off *before* activating the MCL (otherwise, activating the MCL may cause an operating system abend).

Configuring the NPIV feature

To demonstrate how to configure NPIV, we used the sample configuration shown in Figure 3 on page 5.

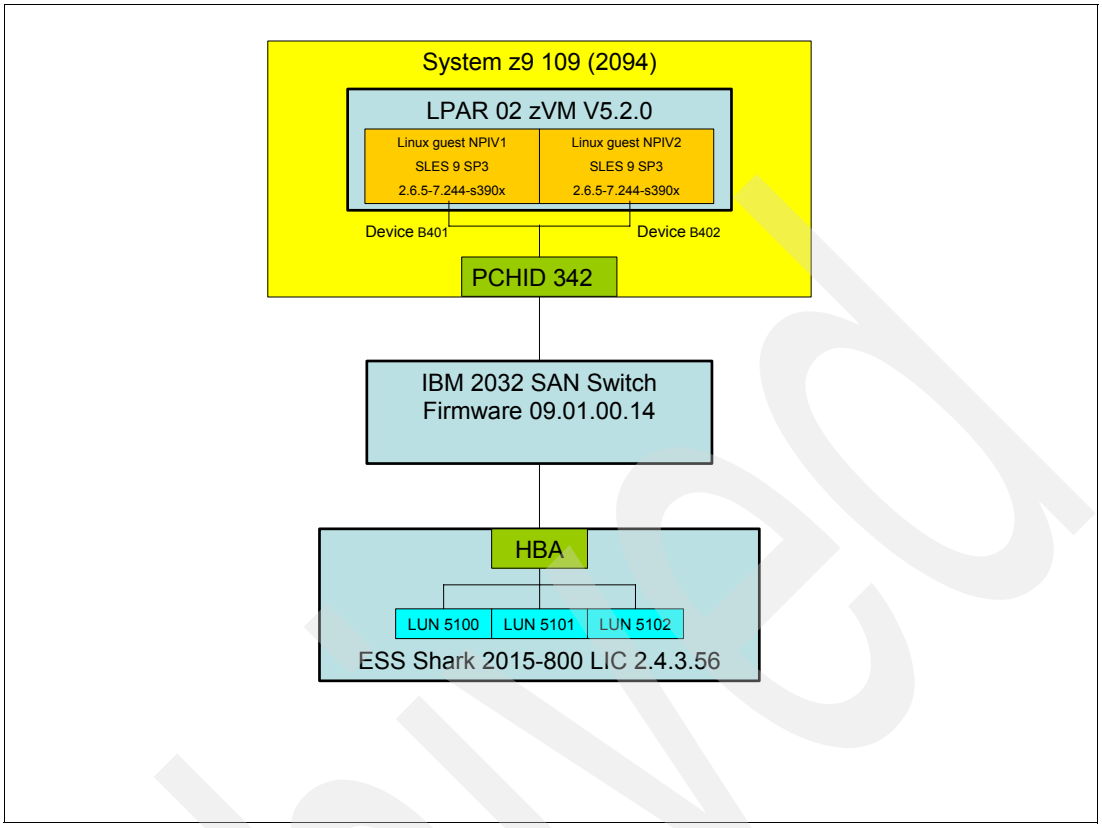

Figure 3 Configuration to demonstration NPIV

In this example, two Linux guests access SCSI disks over an NPIV-enabled FCP adapter. The steps we demonstrate include:

- Configuring NPIV on the SAN switch
- Configuring NPIV on the System z9 server
- Configuring fabric security
- Configuring the Linux server

Configuring NPIV on the SAN switch

Before enabling NPIV on the System z9 server, you must enable NPIV on the switch.

Note: The following examples are specific to our IBM 2032 switch. Other switches that support NPIV work similarly. Follow your switch vendor documentation to enable NPIV on your specific switch.

To enable NPIV on the switch:

1. Adding the NPIV feature to the switch.
2. Activating the NPIV feature on switch.
3. Configuring NPIV on an individual switch port.

Adding the NPIV feature to the switch

To add the NPIV feature to the switch:

1. Start the EFCM switch management application and log on to the switch. Select **Configure** →**Features** as shown in Figure 4 on page 6.

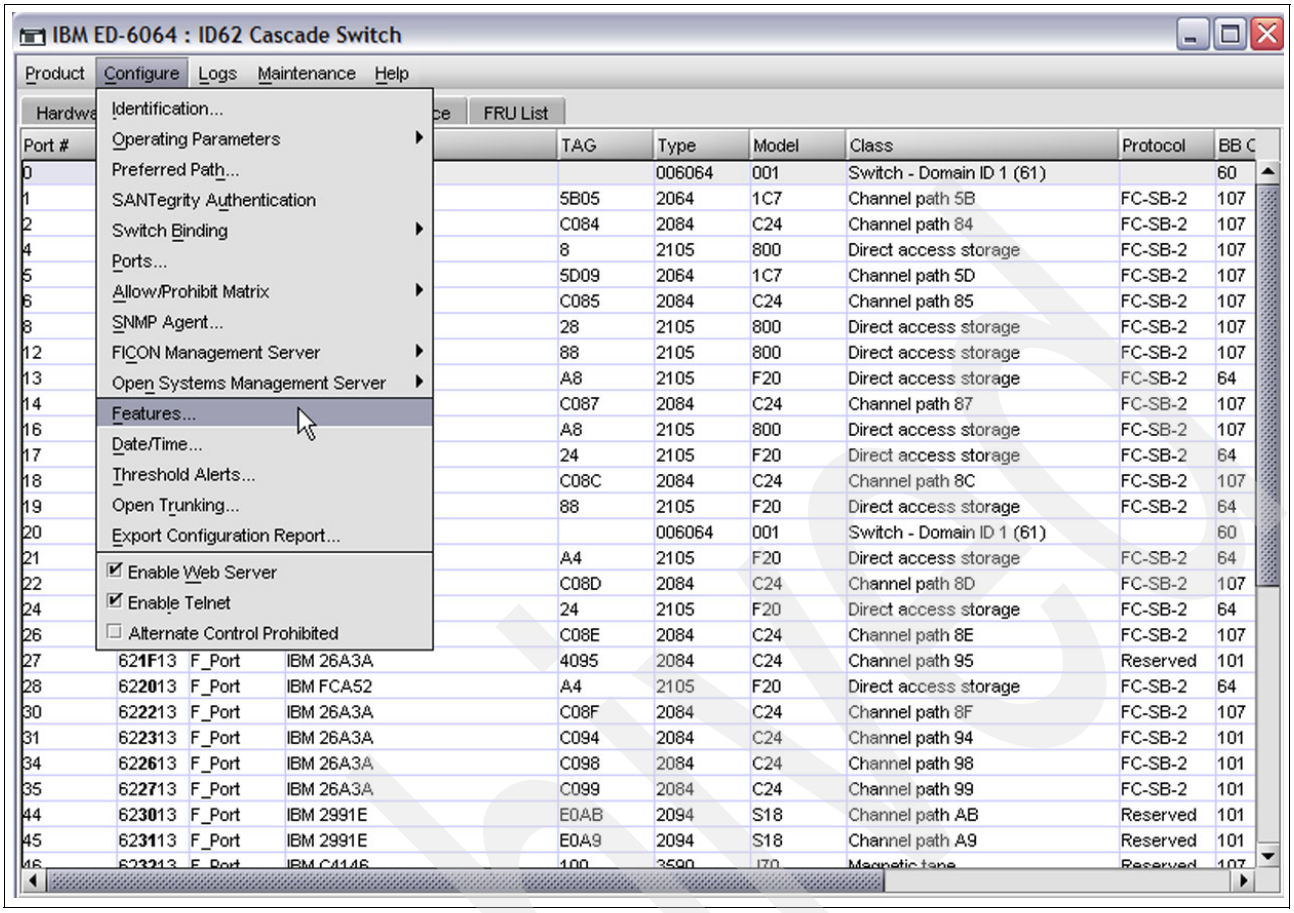

Figure 4 Navigation to the Features configuration menu

2. In the Features configuration menu, click **New** and enter the NPIV feature key provided by the switch manufacturer. This menu is shown in Figure 5.

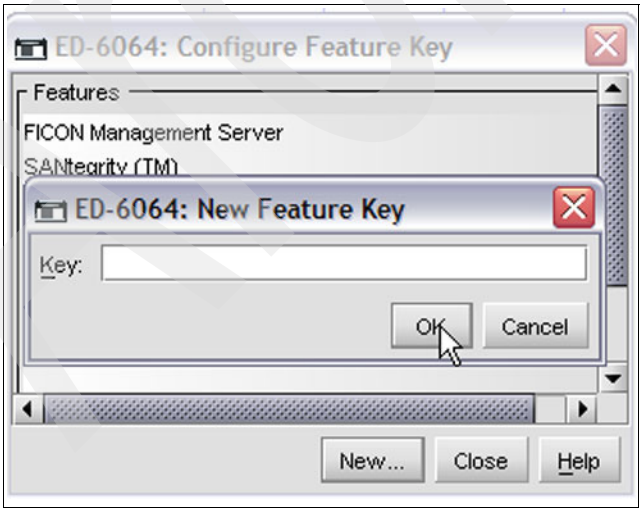

Figure 5 Provide the NPIV feature key

Activating the NPIV feature on switch

To activate the NPIV feature on the switch:

1. Select **Configure →Operating Parameters →Switch Parameters** to open the Configure Switch Parameters menu shown in Figure 6.

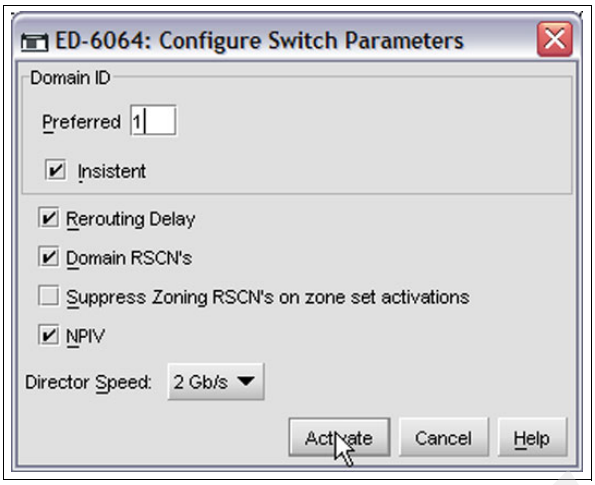

Figure 6 Activate the NPIV feature on the switch

2. Select the NPIV option.
3. Click **Activate**.

Configuring NPIV on an individual switch port

To configure NPIV for an individual port:

1. Select **Configure →Ports** to open the port configuration menu shown in Figure 7.

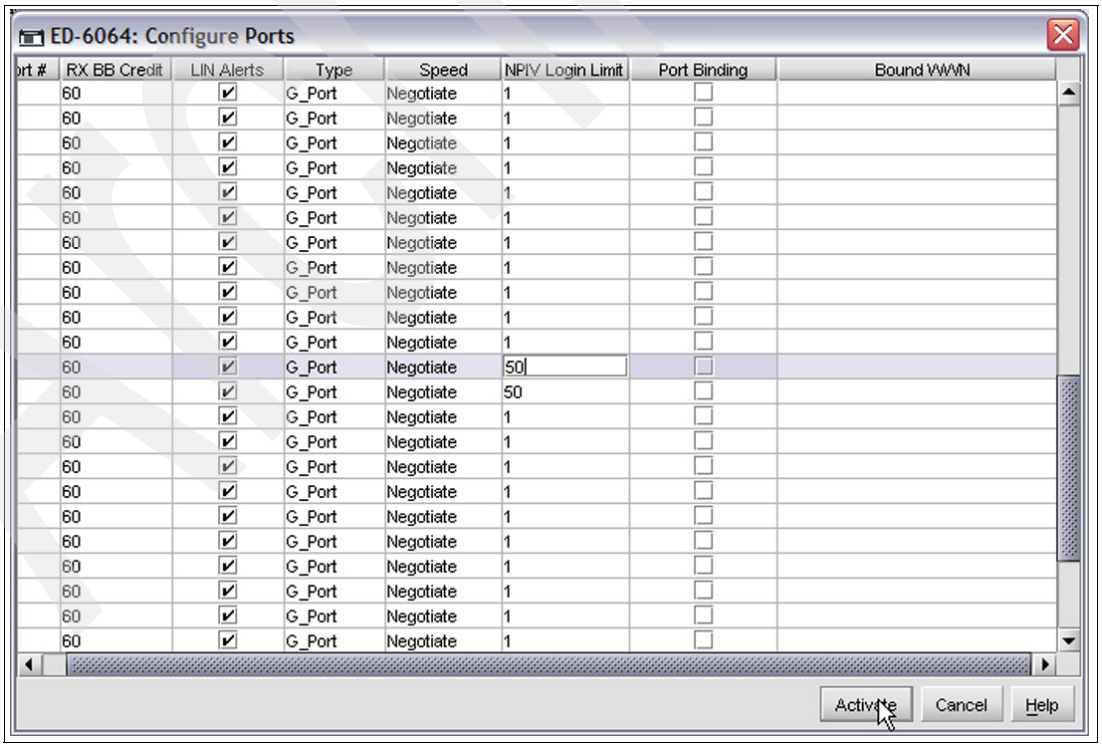

Figure 7 Configure NPIV for an individual switch port

2. Double-click **NPIV Login Limit** for the port and enter the desired login limit. This limits the total number of WWPN logins (both NPIV logins and the default WWPN login).
3. Click **Activate** to complete configuration.

Note: The switch that we used accepts values between 1 and 256. In the examples that follow, three logins are required (two NPIV logins and the default login). When configuring the port, we arbitrarily assign a login limit of 50. If this value is set too small, an out of resource in fabric error message is reported on the Linux host.

Configuring NPIV on the System z9 server

Once NPIV is configured on the switch, NPIV can be enabled on the System z9 server. The FCP CHPID must be taken offline to enable the NPIV feature.

Important: NPIV should be enabled on the fabric switch *before* configuring NPIV on the System z9 server. If NPIV is enabled on the System z9 server, but not on the switch, the FCP CHPID reverts to non-NPIV mode on fabric login. This could be resolved by turning the corresponding switch port off and on (block and unblock the switch port on our Mcdata switch). Alternatively, but more disruptively, a CHPID off/on is required to enable NPIV once the switch is enabled for NPIV.

To enable the NPIV, use the HMC to:

- Enable the NPIV feature on the System z9 server
- Find the NPIV WWPNs for the FCP CHPID
- Find the permanent WWPN for the FCP CHPID

Enabling the NPIV feature on the System z9 server

The NPIV feature can be enabled from the CHPID Operations menu (Figure 8 on page 9) in the SE. From the Hardware Management Console (HMC), select the Single Object to navigate to the SE:

1. Select your CPC, right-click it, and select the Channels option.
2. Scroll to the CHPID Operations task on the right.

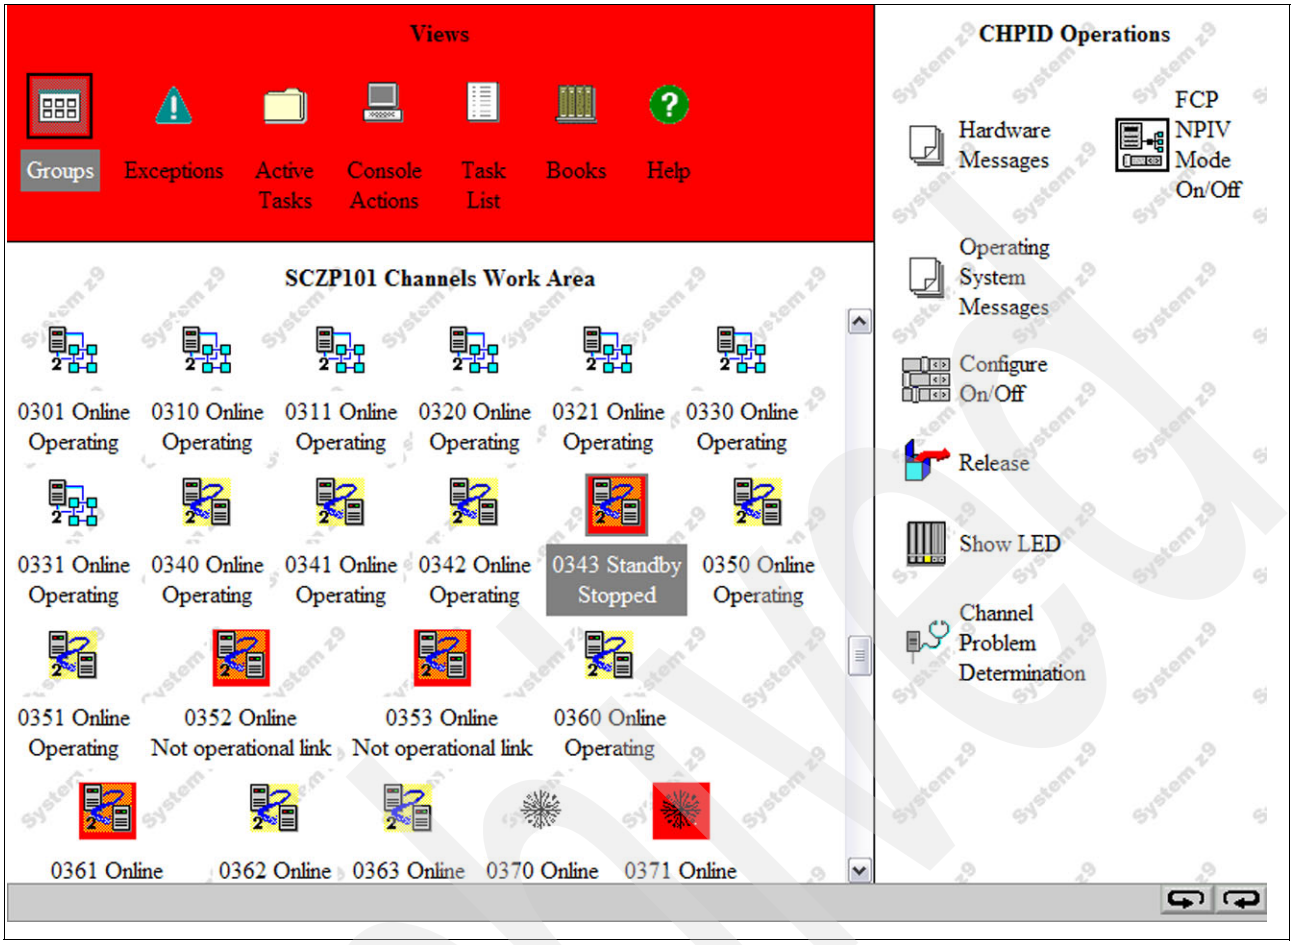

Figure 8 The CHPID Operations menu in the SE

The NPIV feature can be selectively enabled for individual LPARs. From the CHPID Operations menu:

1. Set the FCP CHPID to **standby** as follows:
   a. Double-click **Configure On/Off** to open the menu shown in Figure 9.

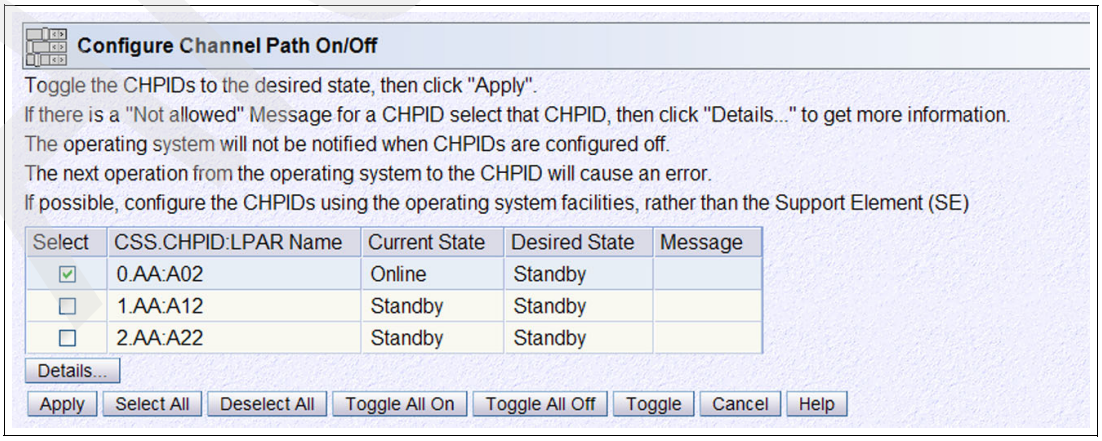

Figure 9 Configure Channel Path On/Off menu

b. Select the appropriate LPARs and click **Toggle** to change the Desired State option to Standby.

c. Click **Apply** to commit the changes.

2. Enable the NPIV feature as follows:

a. Select your PCHID number in **Channel Work Area**.

- From the CHPID Operations menu, double-click **FCP NPIV Mode On/Off** to open the NPIV Mode On/Off menu in Figure 10.

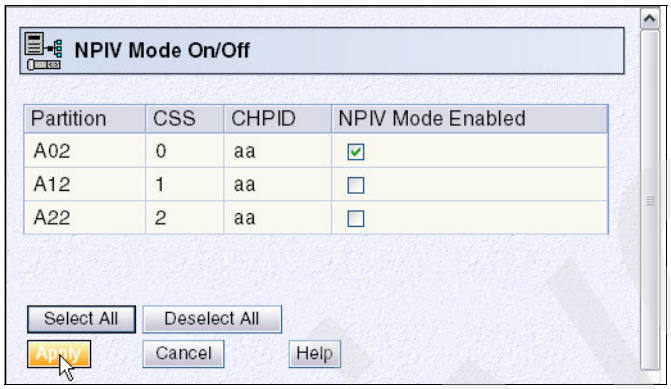

Figure 10 The NPIV Mode On/Off menu of a PCHID under CPC

- Alternatively, you could select the CHPID number in **Chpids Work Area** under a LPAR image. From the CHPID Operations menu, double-click **FCP NPIV Mode On/Off** to navigate to the NPIV Mode On/Off menu in Figure 11.

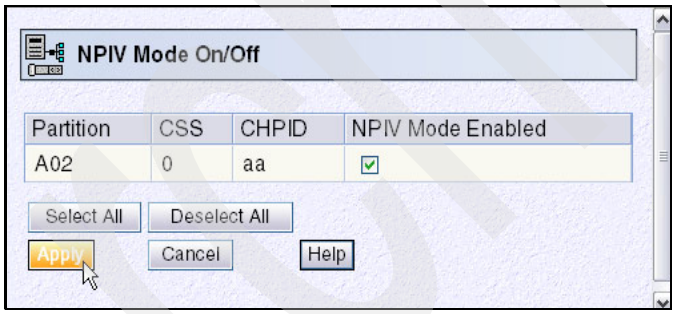

Figure 11 The NPIV Mode On/Off Menu of a CHPID under a LPAR image

b. Select the NPIV Mode Enable option for each LPAR.

c. Click **Apply** to commit the changes.

Note: To enable NPIV, the CHPID must be in standby state for the LPAR. If not, the NPIV Mode On/Off option is disabled to prevent any mode changes.

3. Set the FCP CHPID online as follows:

a. From the CHPID Operations menu, double-click **Configure Channel Path On/Off**.

b. Select the appropriate LPARs and click **Toggle** to change the Desired State option to **Online**.

c. Click **Apply** to commit the changes.

Finding the NPIV WWPNs for the FCP CHPID

Once enabled, the NPIV WWPNs that are assigned to an FCP CHPID can be found in the SE. These WWPNs are needed to configure LUN masking on the storage server and zoning in the fabric.

The NPIV WWPNs are accessible from the CPC Configuration menu in the SE (shown in Figure 12).

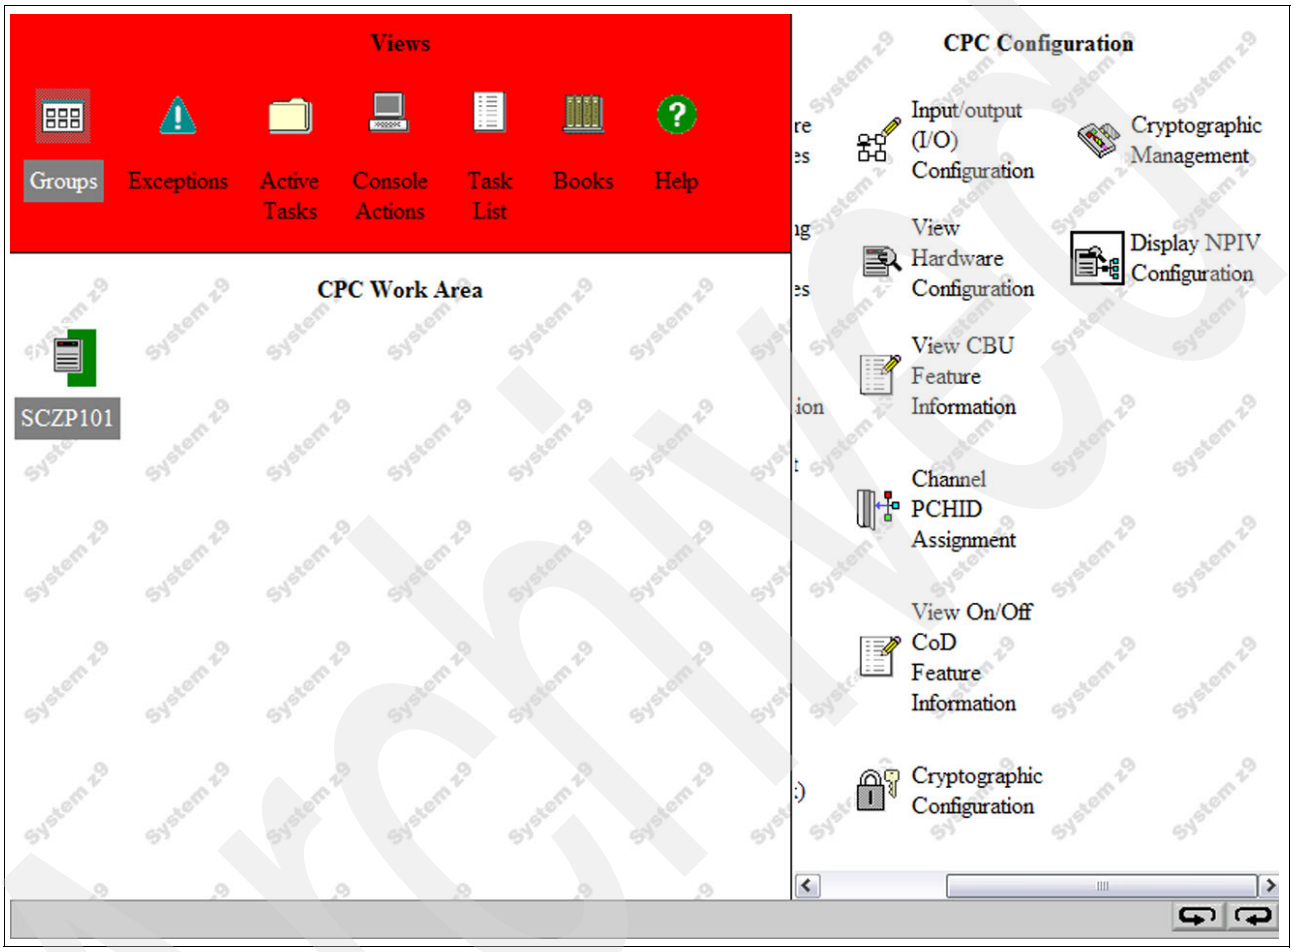

Figure 12 The CPC configuration menu in the SE

To find the WWPNs:

1. Click **Display NPIV Configuration** to navigate to the FCP Channel - FCP NPIV Port Names menu shown in Figure 13 on page 12.

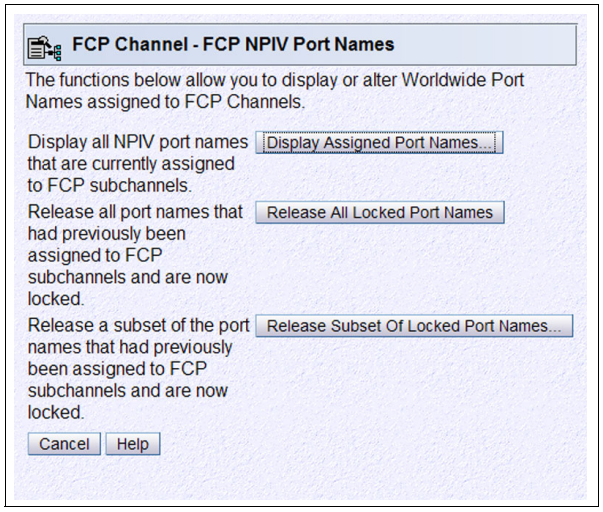

Figure 13 The FCP Channel - FCP NPIV Port Names menu

2. Click **Display Assigned Port Name** to open the Display Assigned Port Names menu shown in Figure 14. It will be extremely helpful to restrict the number of WWPNs display by selecting the (Show NPIV=On) option.

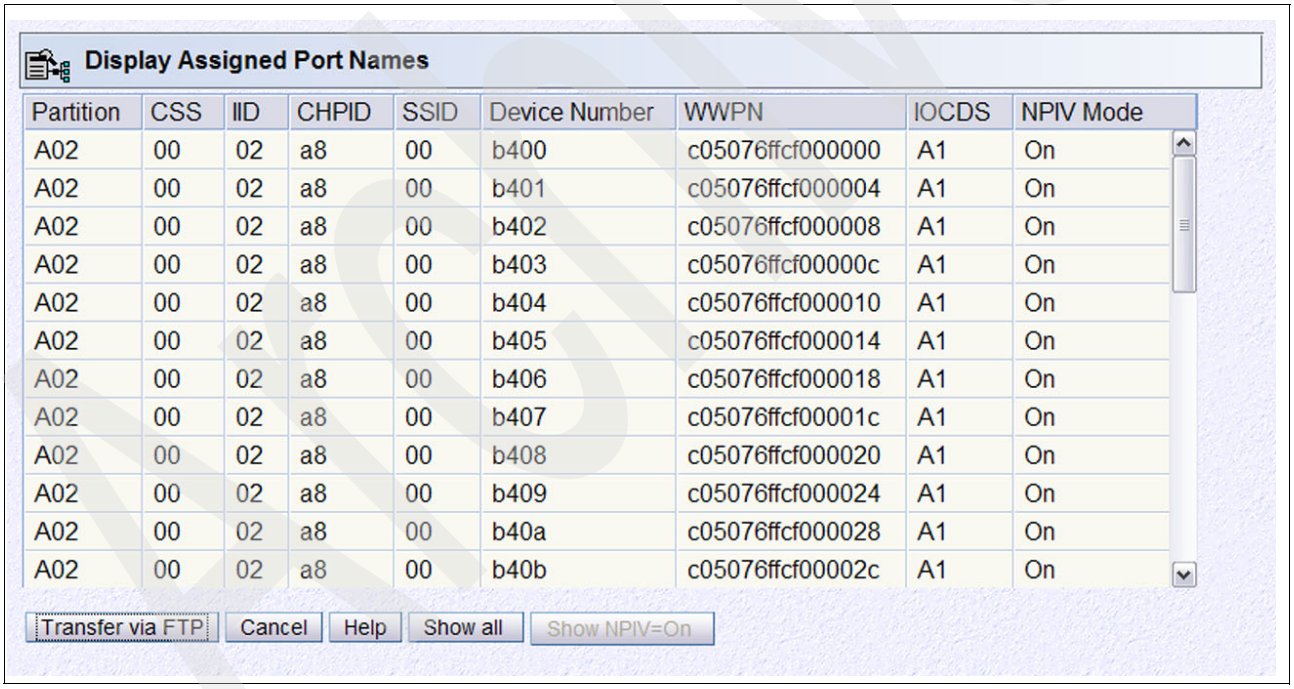

Figure 14 The Display Assigned Port Names menu

3. Each device number in an LPAR is assigned a unique NPIV WWPN. Click **Transfer via FTP** to copy a text version of this menu to an FTP server.

Finding the permanent WWPN for the FCP CHPID

The permanent WWPN is also needed for LUN masking and zoning. To find it:

1. From the Channel Operations menu in the SE, click **Channel Problem Determination** as shown in Figure 15 on page 13.

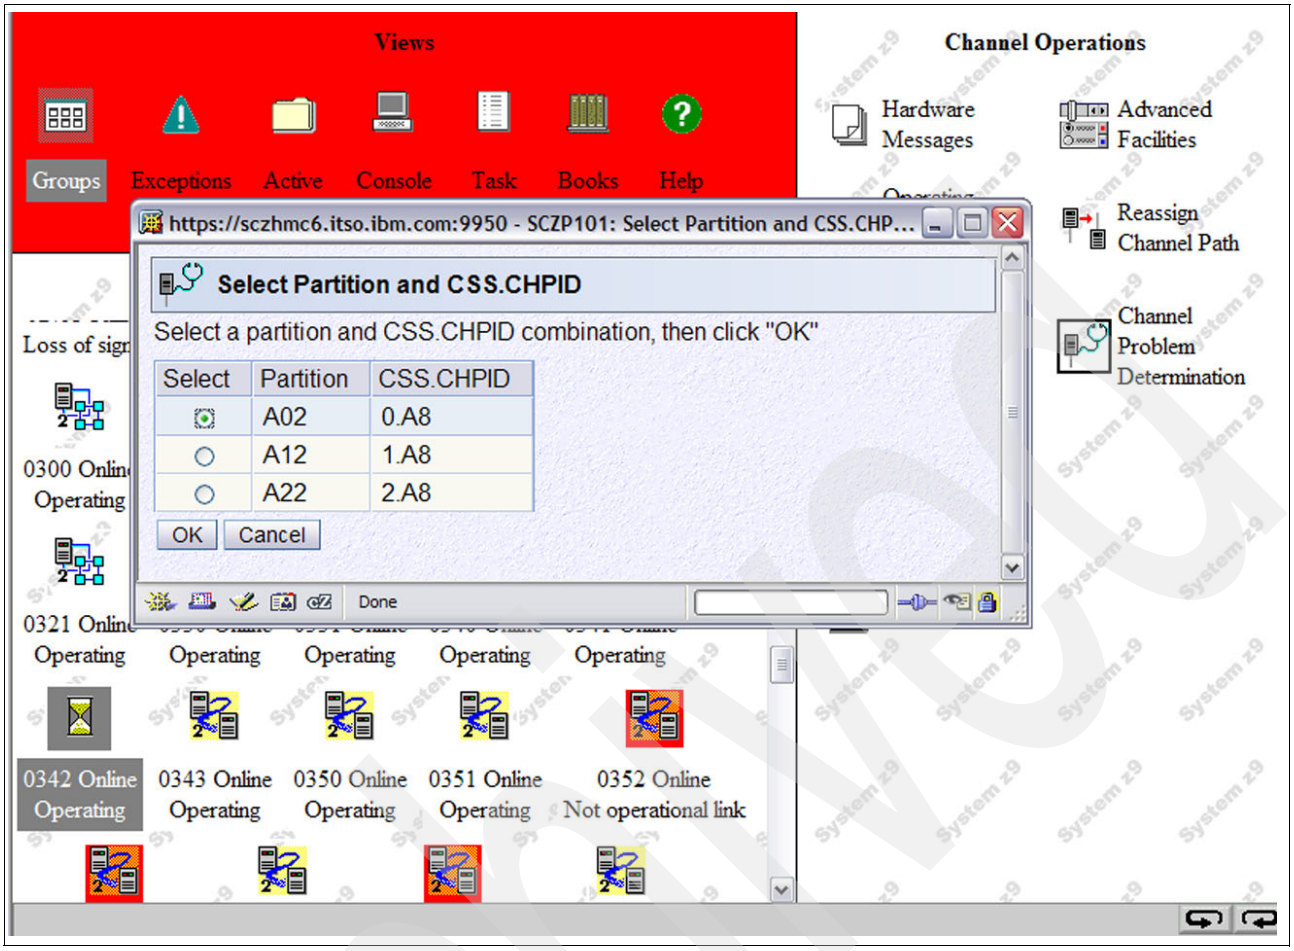

Figure 15 Select the LPAR in the Channel Problem Determination menu

2. Select the desired LPAR and click **OK** to open the Channel Problem Determination menu shown in Figure 16.

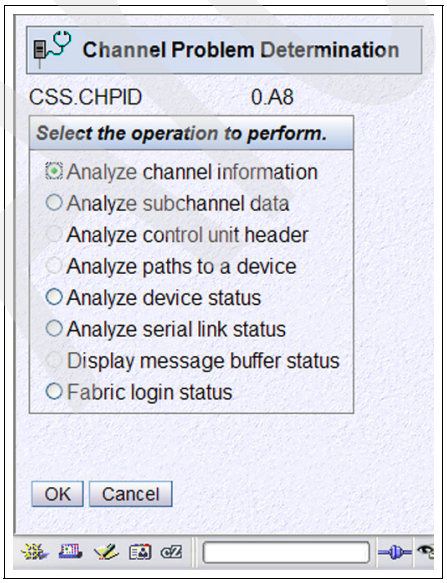

Figure 16 Channel Problem Determination menu

3. Select the Analyze channel information option and click **OK** to open the menu shown in Figure 17.

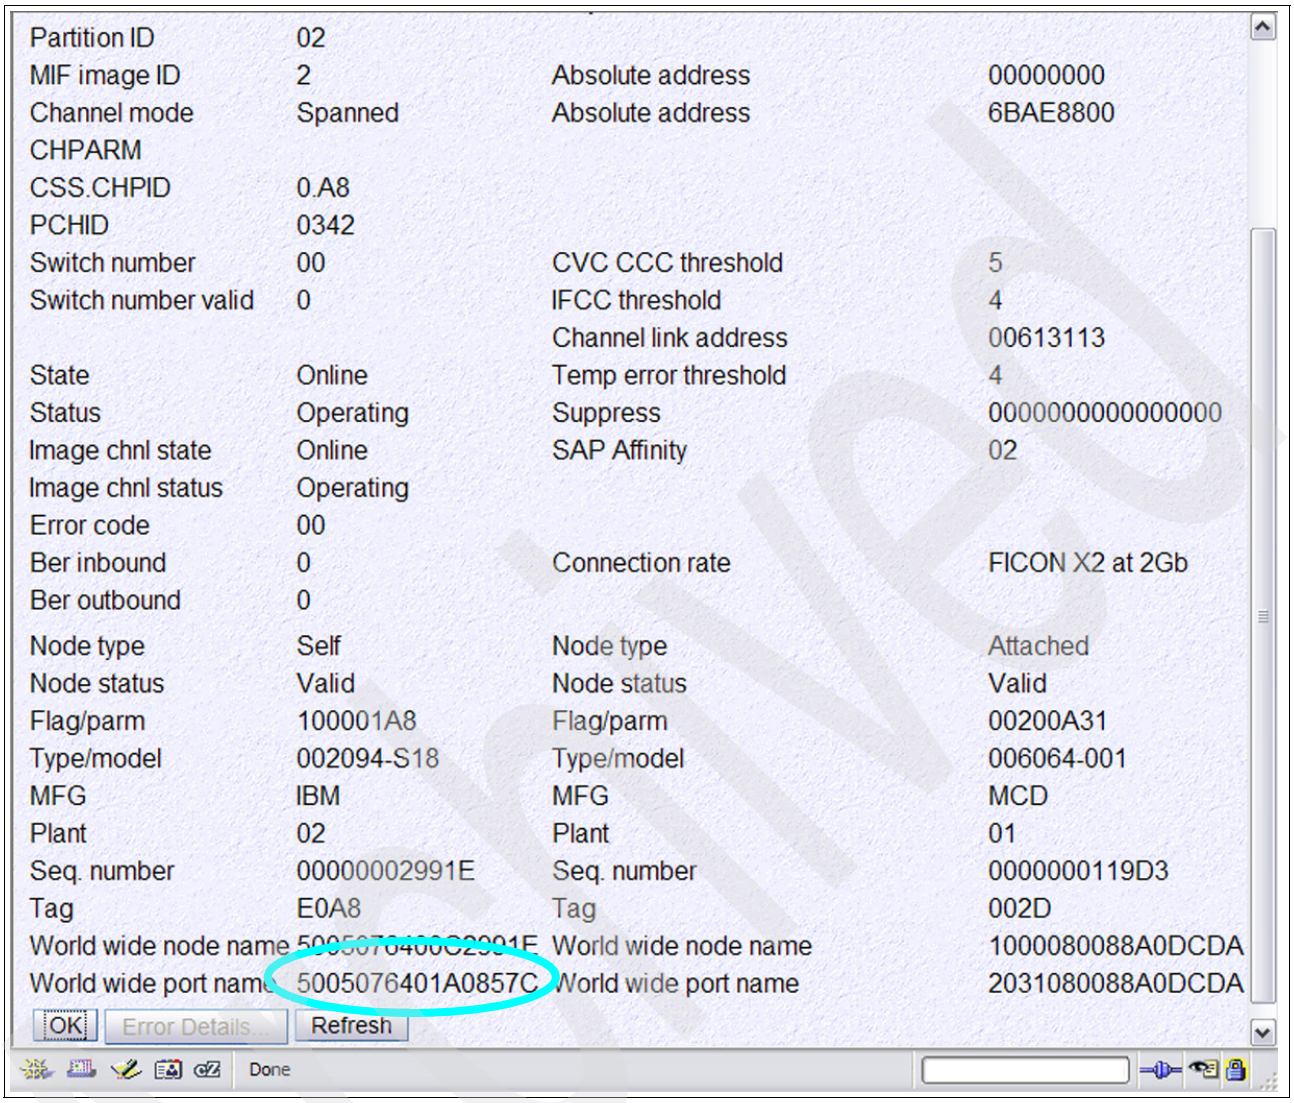

Figure 17 The Analyze Channel Information menu

4. Record the default WWPN for the CHPID (highlighted in the figure).

Configuring fabric security

The NPIV and permanent WWPNs must be defined in the fabric zoning and to the LUN masking on the storage server.

Note: Zoning and LUN masking configuration are specific to the switch and storage server that are used in the fabric. Consult your vendor documentation for details. For a general discussion about zoning and LUN masking, see *Linux on zSeries: Fibre Channel Protocol Implementation Guide*, SG24-6344.

In the examples that follow, the authors used a McData SAN switch and an ESS 800.

Configuring fabric zoning on the switch

To configure fabric zoning on the switch:

1. From the EFCM switch management application, select **Configure** → **Switch Binding** →**Edit Membership List** to open the Switch Binding - Membership List menu in Figure 18.

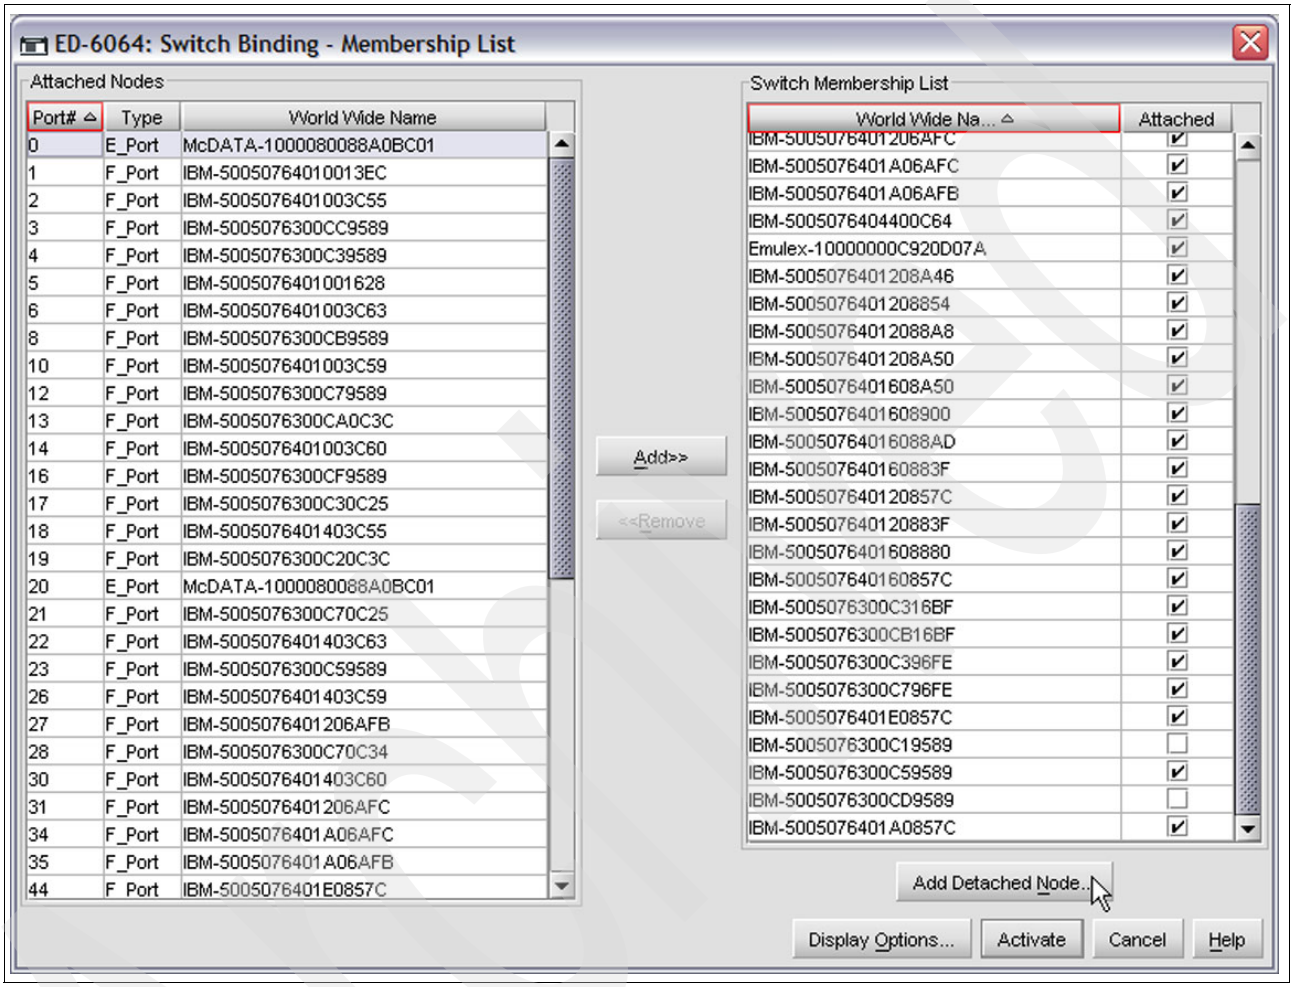

Figure 18 Switch Binding - Membership List menu

2. Click **Add Detached Node** to open the menu shown in Figure 19.

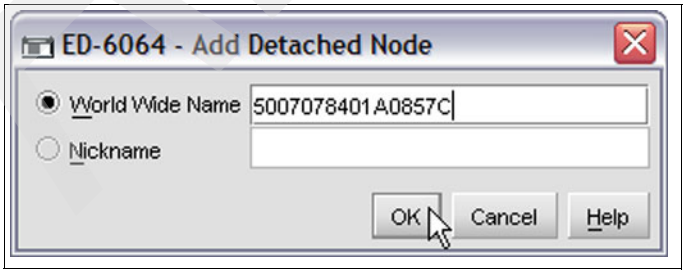

Figure 19 Add Detached Node menu

3. Add the WWPN to the named zone and click **OK**. In this example, the authors added the permanent WWPN to the default zone for the switch.

4. Click **Activate** on the Switch Binding - Membership List menu commit the changes.

Configuring LUN masking on the storage server

LUN masking is configured on the storage server. Access to a specific LUN is granted to a WWPN or group of WWPNs with LUN masking. To configure LUN masking on the ESS, log on to the ESS Specialist application, and start by defining the Linux server to the storage server as follows:

1. Click **Storage Allocation** →**Open System Storage™** →**Modify Host Systems** to open the Modify Host System menu shown in Figure 20.

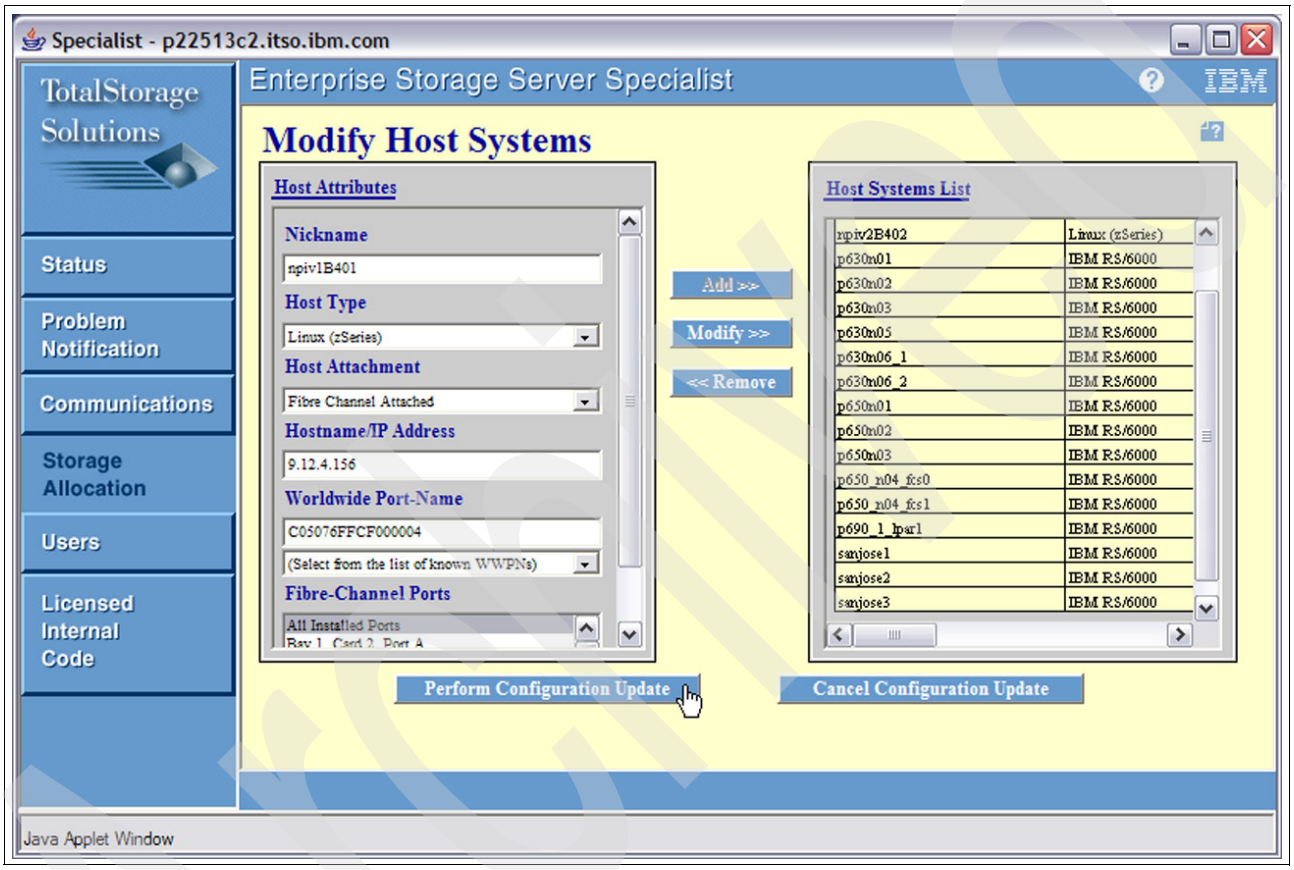

Figure 20 The Modify Host System menu

2. Complete the form in the menu and click **Add** to define the Linux server.
3. Click **Perform Configuration Update** to commit the change.

> **Note:** The authors defined the Linux server using the nickname “npiv1B401”. For the WWPN, we supplied the NPIV WWPN for the B401 device (the device that is used by the NPIV1 Linux guest).

Next, you assign some existing storage volumes to the Linux server. From the Modify Volume Assignments menu for Linux server npiv2B402 (shown in Figure 21 on page 17):

1. Select three LUNs (5100-5102) from the Volume Assignments list.
2. Click **Assign selected volumes to target hosts.**
3. Select **npiv1B401** from the Target Hosts list.

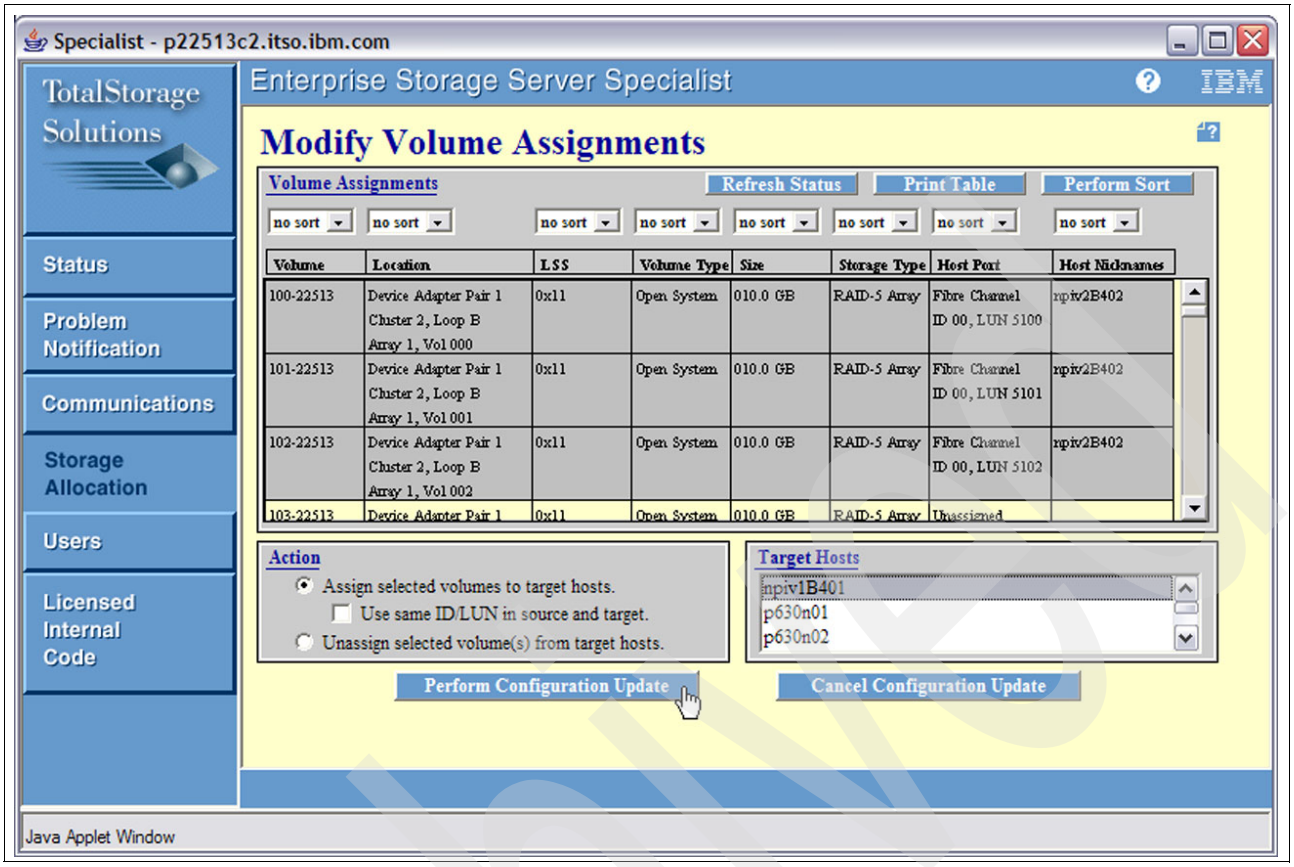

Figure 21 The Modify Volume Assignments menu

4. Click **Perform Configuration Update** to commit the changes.

Finally, you find the WWPN of the HHBA on the ESS that is used to access the LUNs in the Storage Allocation - Graphic View menu shown in Figure 22 on page 18.

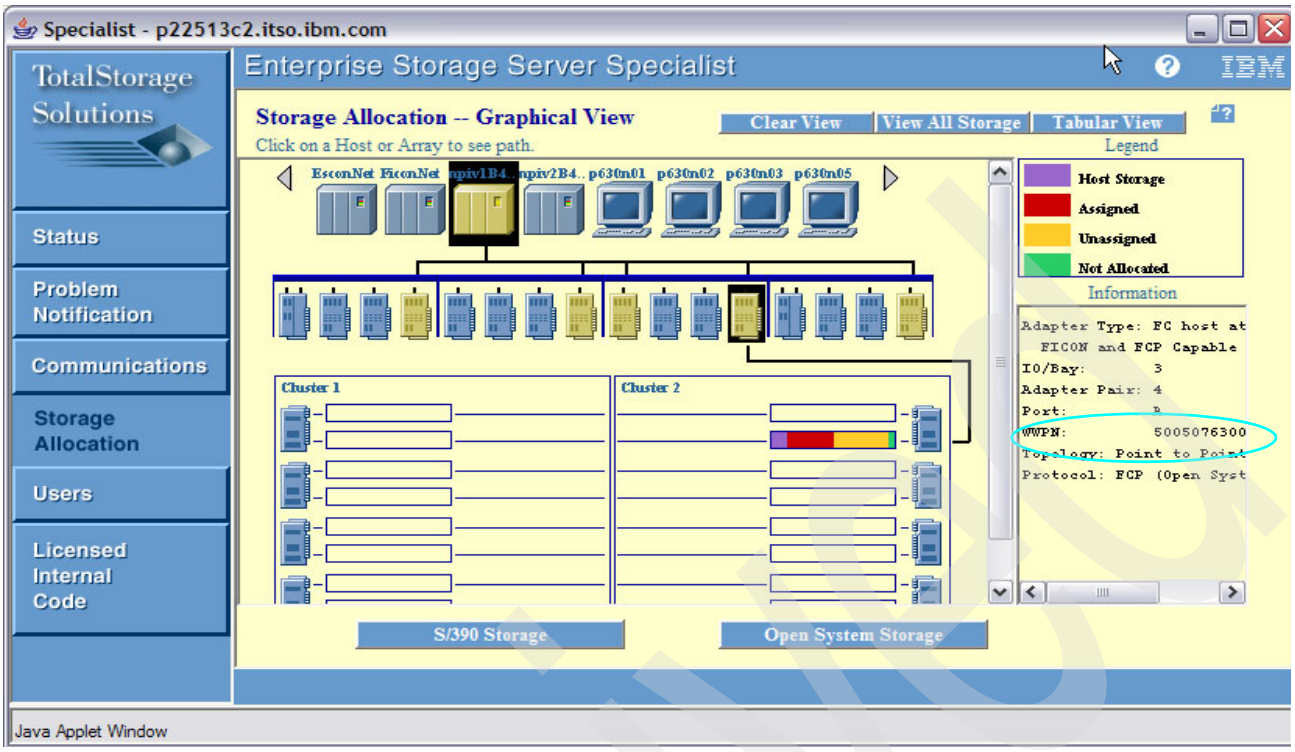

Figure 22 The Storage Allocation - Graphic View menu

Of the 16 ports in the storage server, five are configured for FCP. Any one of these five can be used to access the LUNs. The HBA in bay 3, port 4 connects to the switch. Click this HBA to display the **Information** panel on the right (where the WWPN can be found).

Configuring the Linux server

At this point, the Linux server can be configured to use the NPIV WWPN. With NPIV, each device on the FCP CHPID is assigned a unique WWPN. The WWPN used by a Linux server is determined by the FCP device of the server.

Important: To avoid data corruption, never mount a disk partition in read/write mode to more than one server without proper software support.

You can use the following steps to bring up the disk volume to your Linux images:

1. Attach FCP devices to Linux images from z/VM.
2. Use YaST to setup the Disk Volumes.
3. Verify the NPIV WWPN login.

Attaching the FCP devices to the Linux guest

When running under z/VM, the FCP device must be attached to the virtual machine of the Linux guest. To attach the device dynamically, use the CP ATTACH command. From the MAINT user, the authors attached two FCP devices to two Linux guests (B401 to Linux guest NPIV1 and B402 to Linux guest NPIV2). The commands are illustrated in Figure 23 on page 19.

```
Q CHPID A8
Path A8 online to devices B400 B401 B402 B403 B404 B405 B406 B407
Path A8 online to devices B408 B409 B40A B40B B40C B40D B40E B40F
Path A8 online to devices B410 B411 B412 B413 B414 B415 B416 B417
Path A8 online to devices B418 B419 B41A B41B B41C B41D B41E B41F
Path A8 online to devices B4FC B4FD
Ready; T=0.01/0.01 14:16:27
ATTACH B401 TO NPIV1
FCP B401 ATTACHED TO NPIV1 B401
Ready; T=0.01/0.06 14:16:46
ATTACH B402 TO NPIV2
FCP B402 ATTACHED TO NPIV2 B402
Ready; T=0.01/0.06 14:16:53
Q FCP
FCP  B401 ATTACHED TO NPIV1    B401 CHPID A8
FCP  B402 ATTACHED TO NPIV2    B402 CHPID A8
Ready; T=0.01/0.01 14:16:57
```

Figure 23 Attaching FCP devices to Linux guests

Both Linux guests are running when the FCP devices are attached. Figure 24 shows the messages that appear in the console for guest NPIV1.

```
FCP B401 ATTACHED TO NPIV1 B401
zfcp: adapter 0.0.b401: operational again
Feb 22 14:22:51 npiv1 kernel: crw_info : CRW reports slct=0, oflw=0, chn=0, rsc=
3, anc=1, erc=4, rsid=12
Feb 22 14:22:51 npiv1 kernel: zfcp: adapter 0.0.b401: operational again
zfcp: The adapter 0.0.b401 reported the following characteristics:
WWNN 0x5005076400c2991e, WWPN 0xc05076ffcf000004, S_ID 0x00613128,
adapter version 0x3, LIC version 0x600, FC link speed 2 Gb/s
zfcp: Switched fabric fibrechannel network detected at adapter 0.0.b401.
Feb 22 14:22:52 npiv1 kernel: zfcp: The adapter 0.0.b401 reported the following
characteristics:
Feb 22 14:22:52 npiv1 kernel: WWNN 0x5005076400c2991e, WWPN 0xc05076ffcf000004,
S_ID 0x00613128,
Feb 22 14:22:52 npiv1 kernel: adapter version 0x3, LIC version 0x600, FC link sp
eed 2 Gb/s
Feb 22 14:22:52 npiv1 kernel: zfcp: Switched fabric fibrechannel network detecte
d at adapter 0.0.b401.
```

Figure 24 Console messages from a Linux guest

As expected, the WWPN for the B401 device uses the NPIV WWPN that is assigned by the SE (shown in Figure 14 on page 12). The messages also indicate that the FICON Express 2 adapter is version 0x3 (the minimum required for NPIV). The LIC version (0x600) is also the minimum FCP microcode level required for NPIV.

Note: If the NPIV WWPN is unable to successfully login to the fabric, its WWPN is assigned the value 0x0000000000000000. In this case, check zoning at the switch (see “Configuring fabric zoning on the switch” on page 15). If an “out of resource in fabric” message is shown, the reason might be that the NPIV login limit was set too small on the switch port (see “Configuring NPIV on an individual switch port” on page 7).

Login status can be checked from the Channel Problem Determination menu in the SE (see Figure 16 on page 13). Select **Fabric Login Status** to display the menu shown in Figure 25.

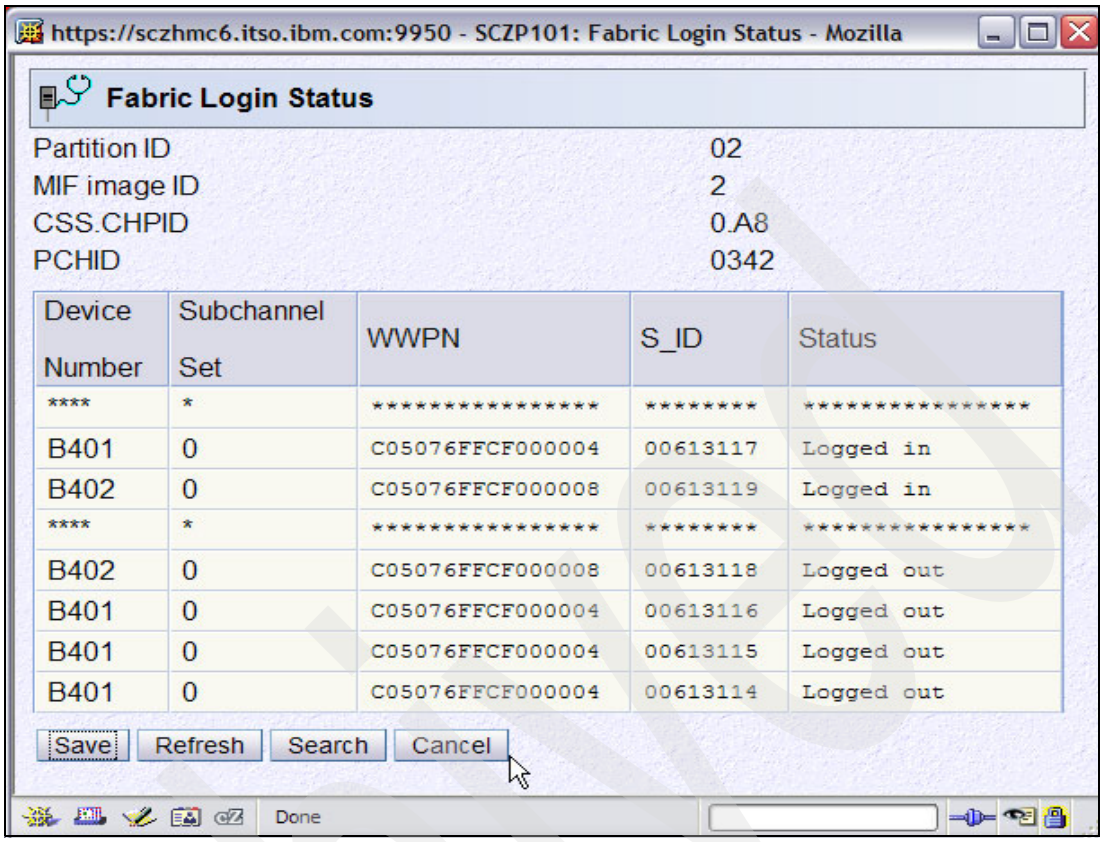

Figure 25 Fabric Login Status menu

If the WWPN is in **Logged in** status but the LUN cannot be accessed, be sure the NPIV WWPN is authorized to access the LUN at the storage server (see “Configuring LUN masking on the storage server” on page 16).

Using YaST to set up the Disk Volumes

Once the device is attached, the FCP disk can be added to the Linux server as follows:

1. From YaST, select **Hardware** → **zFCP** and click **Add** to open the menu shown in Figure 26 on page 21.

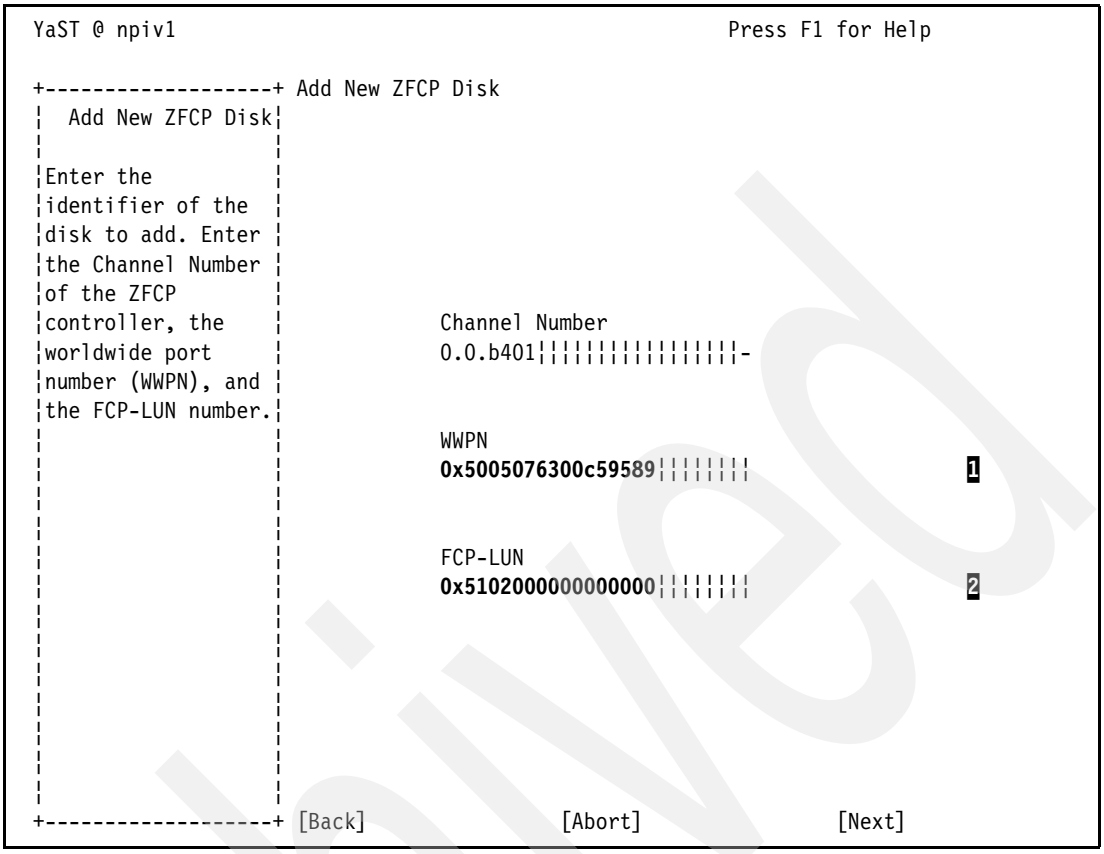

Figure 26 YaST menu to configure FCP disk

2. In this menu, the FCP device (0.0.b401) is automatically detected by the Linux kernel. We can replace it if it is not our intended device. We provided:
   a. The WWPN of the storage (0x5005076300c59589)
   b. The LUN of the SCSI disk (0x5102000000000000)
3. Click **Next** when done to continue.

In Figure 27 on page 22, YaST displays two LUNs that have been added to the Linux guest (0x5100000000000000 and 0x5102000000000000).

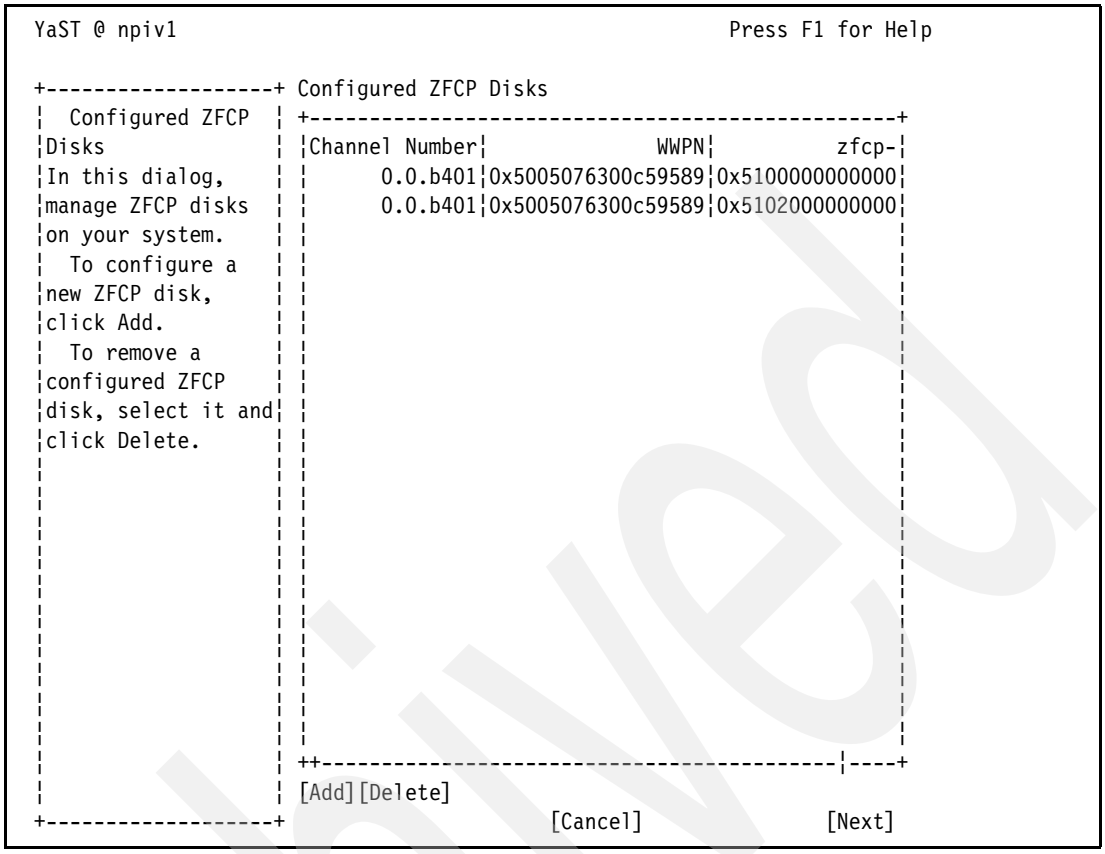

Figure 27 YaST menu for configured FCP devices

We used the **dmesg** command to see that the two LUNs were added as the /dev/sda and /dev/sdb partitions as shown in Figure 28.

```
zfcp: The adapter 0.0.b401 reported the following characteristics:
WWNN 0x5005076400c2991e, WWPN 0xc05076ffcf000004, S_ID 0x00613127,
adapter version 0x3, LIC version 0x600, FC link speed 2 Gb/s
zfcp: Switched fabric fibrechannel network detected at adapter 0.0.b401.
  Vendor: IBM       Model: 2105800           Rev: 3.56
  Type:   Direct-Access                      ANSI SCSI revision: 03
SCSI device sda: 19531264 512-byte hdwr sectors (10000 MB)
SCSI device sda: drive cache: write back
 sda: sda1 sda2
Attached scsi disk sda at scsi0, channel 0, id 1, lun 0
  Vendor: IBM       Model: 2105800           Rev: 3.56
  Type:   Direct-Access                      ANSI SCSI revision: 03
SCSI device sdb: 19531264 512-byte hdwr sectors (10000 MB)
SCSI device sdb: drive cache: write back
 sdb: sdb1
Attached scsi disk sdb at scsi0, channel 0, id 1, lun 1
Attached scsi generic sg0 at scsi0, channel 0, id 1, lun 0,  type 0
Attached scsi generic sg1 at scsi0, channel 0, id 1, lun 1,  type 0
```

Figure 28 Linux kernel messages when adding SCSI devices

Each has a capacity of 10000MB. The kernel recognizes two existing partitions (/dev/sda1 and /dev/sda2) on the /dev/sda device. These are formatted as ext3 partitions.

On the NPIV1 Linux guest, we mounted the /dev/sda1 partition in read/write mode and the /dev/sda2 partition in read-only mode. Next, we created a test text file on the /dev/sda1 partition in Figure 29.

```
# mkdir npiv1a1 npiv1a2
# mount /dev/sda1 npiv1a1
# mount -o ro /dev/sda2 npiv1a2
# echo "test1" > npiv1a1/testa1.txt
# cat npiv1a1/testa1.txt
test1
```

Figure 29 Mounting the SCSI partitions on the NPIV1 Linux guest

On the NPIV2 Linux guest, we mounted the /dev/sda1 partition in read-only mode, and the /dev/sda2 partition in read/write mode. This allows only one host to have write access to a shared partition. In Figure 30, we read the test file that was created in the NPIV1 guest and created a test text file in the /dev/sda2 partition.

```
# mkdir npiv2a1 npiv2a2
# mount -o ro /dev/sda1 npiv2a1
# mount /dev/sda2 npiv2a2
# cat npiv2a1/testa1.txt
test1
# echo "test2" > npiv2a2/testa2.txt
# cat npiv2a2/testa2.txt
test2
```

Figure 30 Mounting the SCSI partitions on the NPIV2 Linux guest

Note: If a shared disk partition is changed by a Linux image with read/write access, the other Linux images that have read-only access to that partition will not be updated until the disk partition is remounted.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Frederick Wong is a Product Engineer at IBM Poughkeepsie. He has nine years of experience with Channel and Connectivity in zSeries. His area of expertise includes 2074, P390, R390, EMIO, FCP and OSA. Before joining IBM, he worked for 10 years at the University of Buffalo, NYSIS as a Senior Application Engineer. He holds MS/BS degree in Electrical Engineering from University of Buffalo. He has written academic papers about superconductivity. He is a member of IEEE.

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

Gerhard Banzhaf, Stefan Mueller
IBM Boeblingen

George P. Kuch, Glenn D. Brooks, Daniel F. Clarke, Leonard J. Dibella
IBM Poughkeepsie

Leslie Geer III
IBM Endicott

Gregory Geisehart, Robert Haimowitz
International Technical Support Ogranization, Poughkeepsie Center

Notices

This information was developed for products and services offered in the U.S.A.

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 give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

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.

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 the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates 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. 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 IBM's application programming interfaces.

This document created or updated on April 6, 2006.

Send us your comments in one of the following ways:

- Use the online **Contact us** review redbook form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYJ Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®	Enterprise Storage Server®	System p™
@server®	FICON®	System z™
Redbooks (logo) ™	IBM®	System z9™
z/VM®	Redbooks™	System Storage™
z9™	System i™	TotalStorage®

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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