Getting started with IBM i™ on an IBM Flex System™ compute node.

Mike Schambureck
(schambur@us.ibm.com)

IBM i Lab Services,Rochester, MN
October 2014

Table of Contents

Abstract

This “Getting started” document provides detailed instructions on using IBM i on a POWER node. It covers prerequisites, supported configurations, preparation for install, hardware and software install, firmware updates and post-install tasks such as backups. The document also contains links to many additional information sources.

1 Overview and Concepts

1.1 IBM PureFlex System

When ordering the IBM PureFlex System with a POWER node as the primary operating system, you will need to choose which version you want: Express, Standard or Enterprise. Each comes with predefined IBM Storwise V7000 or the Flex System V7000 Storage Node® LUNs; zoned Fibre Channel switch or switches; configured Ethernet switch or switches and a preloaded Flex Systems Manager. For more detail on the content and configurations used in these offerings, please see the infoCenter documentation here:
http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp

1.2 Chassis Overview

A chassis and its associated components are like a small data center in a box. There are Ethernet switches, fibre channel switches, a storage node for disks, the compute nodes (both POWER7 and x86) for running VIOS, IBM i and the FSM respectfully. All of these are contained within a 10U space in a rack.

The compute nodes contain sets of adapters, generically called mezzanine adapters, for communicating to the switches (See figure 1 below). The placement of the adapters on the nodes determines which switches that their ports communicate through. The first adapter location on each node is typically used for an Ethernet mezzanine adapter. The first port on the Ethernet adapter is hard wired within the chassis to an internal port on the Ethernet switch in an I/O Bay #1 in the back of the chassis. The second port on the Ethernet adapter is hard wired within the chassis to an internal port on the Ethernet switch in an I/O Bay #2 in the back of the chassis. This second switch gives the node redundant paths in case of a switch failure.

The second adapter location on each node is typically for a Fibre Channel (FC) mezzanine adapter. The first port on the FC adapter is hard wired within the chassis to an internal port on the FC switch in an I/O Bay #3 in the back of the chassis. Similarly, the second port on the FC adapter is hard wired within the chassis to an internal port on the FC switch in an I/O Bay #4 in the back of the chassis. Again, you gain redundancy for the switches.
Figure 1 is showing this relationship between the adapter ports and the switches. ScSE stands for Scalable Switch Elements.

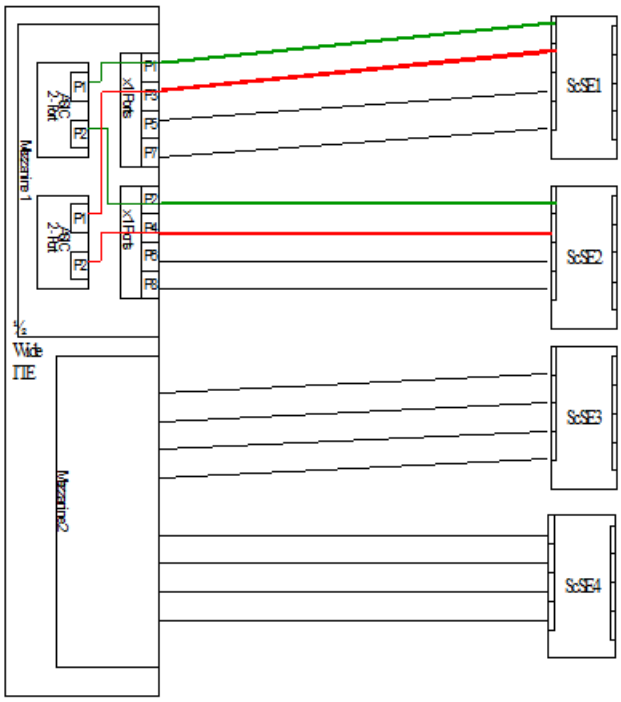

Figure 1: Mezzanine adapter ports and their associated I/O switches.

If a 2 bay node is used, such as the 7895-42X, there are 4 mezzanine adapter locations: locations1 and 3 are used for Ethernet adapters, 2 and 4 are used for FC adapters. Figure 2 shows an example of this.

Figure 2: A view of a 2 bay Power node and its assoicated mezzanine locations.

The chassis Ethernet switches also have external ports that are used to connect to a client's external Ethernet switch(es). The FC switches also have external ports, but for the iEdition do not need to be connected to the client's SAN fabric, although they could be for further storage needs. The Flex System V7000 storage node can have additional expansions cabled from it, but these use Serial-attached SCSI (SAS) cables off of the front of the storage node and do not use the FC switch's external ports.

Finally, there are a pair of redundant chassis management modules (CMMs) that are used as an interface to the other chassis components. The CMM can talk to the service processor on each of the nodes (the Flexible Service Processor (FSP) for the POWER nodes and the Imbedded Management Module (IMM) on the x86 nodes). The FSP is the same as an FSP on a Power server. The CMM can talk to the switch modules too. If you know about BladeCenter chassis's, the CMM is very similar to an Advanced Management Module (AMM) in those chassis's. The Flex Systems Manager (FSM) talks to the CMM to mange the hardware portion of the chassis environment.

1.3 Management interfaces

You can use either the Flex Systems Manager (FSM) or the Hardware Management Console (HMC) for the management of your POWER node. FSM is a hardware appliance: it has its own x86 based node. As of August 2014, the IBM Flex System now supports the IBM Power® Hardware Management Console (HMC). The IBM Power Rack-mounted HMC brings the same, full function PowerVM® management available on POWER7® and POWER7+™ rack servers to the POWER7 and POWER7+ processor-based Flex System compute nodes. The HMC function is available with HMC Version 7 Release 7.7.0 Service Pack 2. The HMC solution requires Power Compute node firmware Version 7.7.3. This guide was written for the FSM interface. Please see chapter 7 of the IBM Flex System p270 Compute Node Planning and Implementation Guide http://www.redbooks.ibm.com/abstracts/sg248166.html?Open for details on the setup and configuration of the HMC to manage your Power nodes.

1.4 Logical Partitioning (LPAR) or Virtual Servers

Similar to other Power® systems, POWER nodes can be partitioned into separate environments called either logical partitions (LPARs) or virtual servers depending upon the management interface being used. An LPAR on FSM is referred to as a **virtual server.** The term LPAR will be used to refer to either term throughout this document.

POWER nodes support the virtual I/O server (VIOS) as a hosting partition and client partitions running IBM i, AIX or Linux. AIX and Linux can also run standalone. This document will only be addressing VIOS hosting IBM i. Any physical adapters installed on the node are owned by VIOS. VIOS then virtualizes disk and optical storage and network resources to the client partitions. An IBM i LPAR on the node does not have direct access to any physical hardware on the node or outside the BladeCenter chassis. Physical adapters cannot be assigned directly to the IBM i partition.

One key function that FSM/HMC brings to POWER nodes is the ability to manage redundant VIO servers. Having redundant VIOS servers allows for the failure or shutdown of one VIOS for maintenance while the redundant VIOS continues to host the IBM i client.

1.5 Operating systems

On top of the virtual server definitions are the operating systems. The currently supported operating systems are the virtual I/O server (VIOS), IBM i, AIX, Power Linux on the Power nodes and Windows® servers; VMWare ® ESX servers; KVM® servers or Linux® servers on the x86 nodes. AIX and Linux can also run standalone on a Power node without the use of a VIOS hosting virtual server, but are then restricted to a single virtual server that uses all of the hardware resources on the node. This document will only be addressing VIOS hosting IBM i.

1.6 Overview of I/O concepts for IBM i on node

The Power node includes sets of adapters for I/O called mezzanine adapters. There are Ethernet and Fibre Channel (FC) mezzanine adapters available. There are no SAS adapters available or planned. A Fibre Channel over Ethernet (FCoE) adapter is also avaialble. Typically for IBM i, the VIOS hosting partition owns these adapters and virtualizes their use to the IBM i client partition. IBM i partitions on a Power node can only use external Fibre Channel storage, there are no SAS adapters available.

The storage is physically connected to VIOS using a supported FC or FCoE mezzanine adapter for the node.
VIOS supports two types of storage virtualization: virtual SCSI (vSCSI) and virtual fibre channel (vFC). The vFC is how VIOS implements N-Port ID Virtualization (NPIV). Each of these will be discussed in more detail in the "Storage management concepts" section of this document. Figure 3 shows this concept.

Figure 3 shows an example Power node environment with two virtual servers as clients of VIOS.

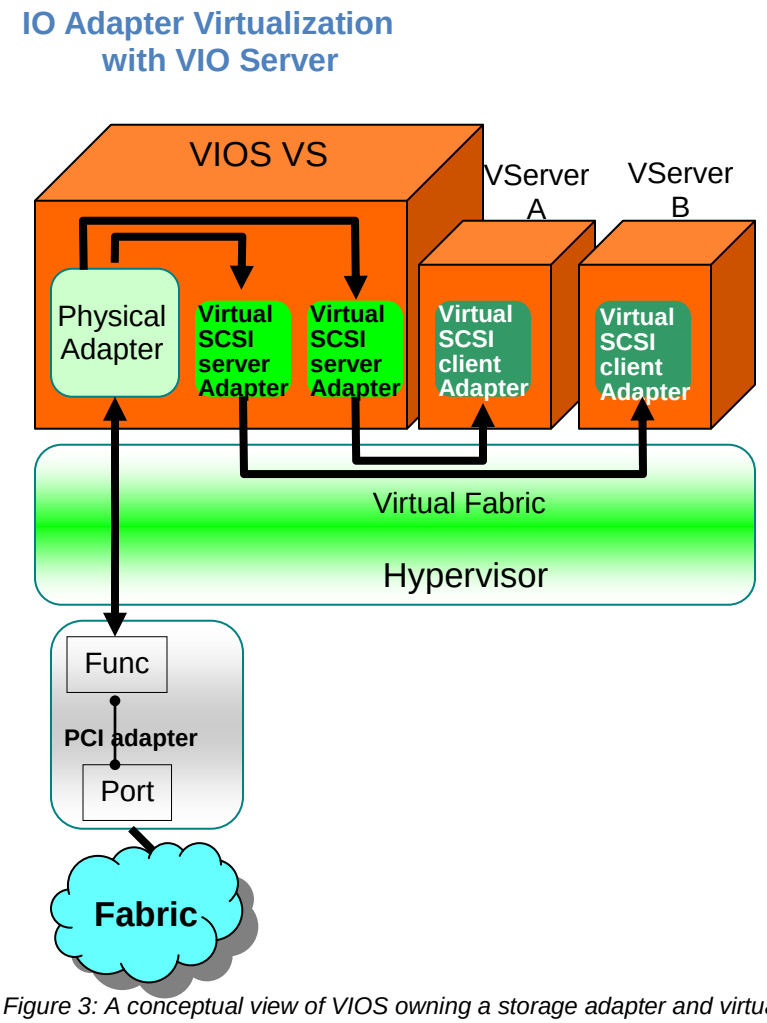

Figure 3: A conceptual view of VIOS owning a storage adapter and virtualizing it to two client virtual servers.

For vSCSI, the FC LUNs are mapped to VIOS and from there they are virtualized to IBM i. On IBM i each LUN appears as a disk drive. For virtual fibre channel adapters, the LUNs are mapped directly to the IBM i through each of the VIOS server, but the LUNS are not seen by VIOS first. IBM i is installed using from a virtual media library managed by VIOS. This repository contains the IBM i DVD iso files for the base operating system and licensed program products. More information on this will come later.

IBM i partitions have Ethernet connectivity also using VIOS virtualization. The ports on the Ethernet mezzanine adapter are owned by the VIOS partition which bridges them to a virtual Ethernet adapter that is associated with the IBM i client partition.

The FSM/HMC supplied console is the only console option. Operations Console (LAN console), thin clients or twinaxial connections are not supported. IBM i partitions use the network virtualization framework explained above for their production TCP/IP traffic. If more than one connection is required, it is recommended that two separate Virtual Ethernet adapters be created for the IBM i partition. The two adapters can then reside on the same Virtual LAN (VLAN) or on separate VLANs. They can connect to the outside network using the same Ethernet adapter port or different ports.

1.7 Review terminology

Terminology	**Description**
Chassis	The enclosure containing the node servers, I/O modules, CMM, power and fan modules.
Chassis management module (CMM)	A control module residing in a special I/O bay in the chassis. The CMM provides browser and command-line interfaces (CLI) into the chassis and interfaces to the other chassis components.
I/O bay	A slot for an I/O module (switch) inside the chassis. A chassis has two sets of high-speed switch bays..
Scalable switch element (ScSE)	A switch residing in the chassis that provides connectivity between the node servers and external I/O. It can be upgraded (scaled) with an activation key to allow more internal and external ports.
Compute Node	A standalone server residing in a node slot in the chassis. The can be POWER or x86 based processors.
Flexible Service Processor (FSP)	The FSP on the POWER node is similar to the FSP on other POWER7 processor-based systems. It contains firmware to manage the hardware on the node; the Power Hypervisor; and Partition Firmware (PFW).
Imbedded Management Module (IMM)	The service processor on an x86 node. It is used to manage the x86 node server for remote access and for firmware updates. It communicates to the CMM as part of chassis management.
System firmware	As with other POWER7 processor-based systems, this is the firmware on the FSP.
I/O expansion card Or mezzanine adapter	Sometimes called a daughter card, this card is an I/O adapter that fits into a PCI Express (PCIe) connector on the node and allows connectivity to external I/O devices through the chassis midplane and I/O modules.
Adapter firmware	The firmware on the I/O expansion cards on the node.
Virtual I/O Server (VIOS)	A software appliance that is located in a logical partition. This software facilitates the sharing of physical I/O resources between client logical partitions

	within the system.
Virtual Ethernet adapter:	A virtual network adapter created in the POWER Hypervisor that is part of an LPAR's hardware resources. On POWER node, IBM i cannot be assigned physical network adapters.
Virtual SCSI adapter:	A virtual storage adapter created in the POWER Hypervisor that is part of an LPAR's hardware resources. On a POWER node, a VSCSI client adapter is created in IBM i and a VSCSI server adapter is created in VIOS for storage virtualization.
Virtual Ethernet bridge:	A VIOS function that allows Layer-2 bridging of a VLAN to an outside physical LAN. It is called a Shared Ethernet Adapter (SEA). It is required on a POWER node to provide networking to IBM i.
Logical Unit Number (LUN):	A logical volume created on a storage area network (SAN) system that appears as a single disk device to a server.
Disk drive module (DDM):	A physical disk unit in a SAN system.
Subsystem Device Driver – Path Control Module (SDDPCM)	A multipath I/O (MPIO) driver installed on top of VIOS for certain storage subsystems – in particular, V7000 and SAN Volume Controller (SVC).
Redundant Disk Array Controller (RDAC)	A MPIO driver for IBM System Storage DS4000® or DS5000, which is included with VIOS.
Serial-attached SCSI (SAS)	A storage access protocol, which is the next generation of the parallel SCSI protocol.
Virtual Fibre Channel adapter	A virtual storage adapter created in VIOS that is part of an LPAR's hardware resources. On a POWER node, a VFC client adapter is created in IBM i and a VFC server adapter is created in VIOS for storage virtualization. This creates virtual World Wide Port Names (WWPNs) for use in the client partition.
Flex Systems Manager (FSM)	x86 node based hardware appliance for managing the components in an PureFlex System chassis.
Host:	Another term for server or physical server, used by FSM, FSM and Systems Director interfaces,
Virtual server:	Another term for a logical partition used by the FSM.

1.8 Plan for necessary IP addresses

You should plan to assign IP addresses to the following components for a minimum IBM i on node configuration. All the addresses below are typically configured on the same subnet.

- **CMM** (this IP address is defaulted to 192.168.70.100)

    - The CMM IP address is a physical LAN IP address. It is used to manage the chassis and node servers
- **IMM** (this IP address is defaulted to 192.168.70.1xx, where xx is the node slot number)
    - This IP address is used to manage the x86 node server.
- **FSP**
    - An IP address is needed on the external LAN to connect to the service processor of the Power node. This is used by FSM to manage the node.
- **I/O module (switch,** these IP addresses are defaulted to 192.168.70.120-123**)**
    - This IP address is used for management of the switch. There can be from 1 to 4 of these modules in a chassis.
- **VIOS**
    - An IP address on the external LAN that is used to connect to the VIOS command line interface.
- **FSM**
    - An IP address on the external LAN that is used to connect to the FSM to manage the chassis and its components
    -
- **IBM i production interface**
    - An IP address on the external LAN that is used for IBM i production network traffic. This address is configured after IBM i is installed.

1.9 Virtual Ethernet concepts

From a network perspective, one significant difference between the POWER node and a standalone server is that the PowerVM hypervisor is virtualizing the Ethernet I/O for the virtual servers with the use of a VLAN-aware virtual Ethernet switch. The client virtual servers communicate to the VIOS virtual servers across similarly configured virtual Ethernet adapters. VIOS then bridges the traffic to a physical port through an object called a shared Ethernet adapter (SEA). Figure 4 shows the Ethernet virtualization concept. The switch shown would be in I/O Bay 1.

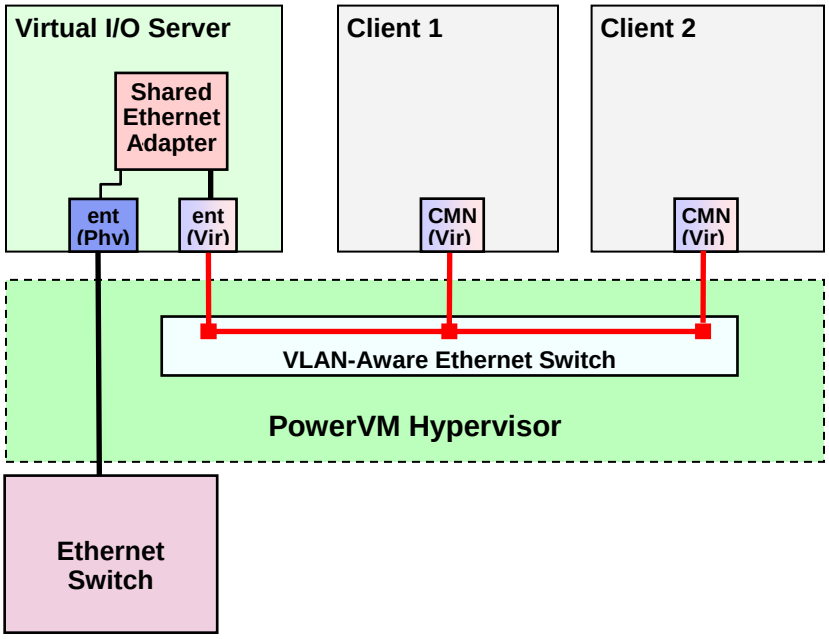

Figure 4: Sample virtual Ethernet network configuration for a basic IBM i installation.

1.9.1 VLAN tags

Ethernet packets can contain tags that allow for isolating the packets from one another. Switches can be configured to continue this traffic isolation as the packets are passed through them. IBM i does not support VLAN tags (at the time of this writing), but VIOS does. If IBM i needs to send tagged packets to match the networking conventions being used there is a way to do this. A virtual Ethernet LAN can be configured for the IBM i and the VIOS that will trigger VIOS to add VLAN tags for the traffic on that VLAN. The VE adapter allows the configuration of up to 20 VLAN tags. The corresponding IBM i VE adapter just needs to specify the port Vlan ID (PVID) value of the VLAN needed. See the "Creating a VIOS virtual server using the FSM" section and the "Create the IBM i Virtual Server using FSM"section for detailed steps.

2 Supported Environments

2.1 Software details

The minimum level of the Virtual IO Server (VIOS) that is supported is 2.2.1.4 which was made available in May 2012.
IBM i release:

- i6.1 is supported, but cannot be ordered preinstalled from IBM manufacturing. You must manually install it yourself.
- i7.1 is supported and can be ordered preinstalled from IBM manufacturing. When ordered in this way, the VIOS supplied media repository will contain the i7.1 microcode, operating system and licensed program products. It will not contain PTFs (fixes), but these can be loaded manually).

2.2 Hardware details

For a list of supported hardware, firmware and software for the IBM i on POWER node environment, see the Flex Interoperability Guide: http://www.redbooks.ibm.com/redpapers/pdfs/redpfsig.pdf

2.2.1 POWER nodes

Use the links below to see current information on the Power nodes.

- IBM PureFlex System Products & Technology, http://www.redbooks.ibm.com/abstracts/sg247984.html?Open
- IBM Flex System p260 and p460 Planning and Implementation Guide, http://www.redbooks.ibm.com/abstracts/sg247989.html?Open
- IBM Flex System p270 Compute Node Planning and Implementation Guide http://www.redbooks.ibm.com/abstracts/sg248166.html?Open
- IBM Flex System p24L, p260 and p460 Compute Nodes http://www.redbooks.ibm.com/abstracts/tips0880.html?Open

3 Chassis installation, configuration and firmware updates

The PureFlex foundation offerings come pre-installed, pre-cabled and pre-configured in a rack, so the following section may not apply. It is included for add-on hardware. The Flex offering is not pre-configured at all, so the follow sections will apply.

3.1 Install the PureFlex System chassis and node hardware

The first step is to install the PureFlex System chassis and node hardware. This might include installing any management modules, power modules, and I/O modules in the chassis. The chassis might have these components already installed if an additional node server is being added to an already functioning chassis. Before installing the nodes in the chassis, any node options must be installed. This might include additional processors, additional memory and I/O mezzanine cards. For POWER node servers, any required mezzanine adapters are installed at this time, depending on the chassis and storage used. Refer to the node and expansion card documentation that came with the option for details on installing each one. After installing the node options, you can install the node in the chassis. Refer to chassis and node documentation for details on how to install the components. After installing all the node options, installing the node into the chassis, and installing the chassis I/O modules, the chassis can be connected to the power outlets.

Note: If you need to add a new node to an existing chassis, consider updating the chassis component's firmware at this time.

3.2 Configure the Chassis Management Module (CMM)

3.2.1 Initial CMM configuration (New chassis install)

On early Express and Standard Foundation single chassis configurations, the CMMs have Ethernet cables plugged into the 7th and 8th Ethernet ports of the EN4093 switch in IO bay 1. For more current configurations the CMMs are not cabled at all. The intention is to plug them into the upstream Ethernet switch(es). For multi-chassis and Enterprise Foundation configurations, the CMMs have Ethernet cables plugged into ports on one of the 8052 Ethernet top of rack (TOR) switches.
Configure the Ethernet adapter on you PC to the 192.168.93.xxx subnet and set the subnet mask to 255.255.255.0. Then cable you laptop to any open port on the switch where the CMM1 is cabled to. These ports will need to be on the management subnet configured on the switch.
You need to perform the following steps on a browser on the computer that is connected to the CMM:

- You can ping the 192.168.93.100 address to confirm communications to the CMM.
- Open a web browser on the computer connected to the CMM. In the address or URL field, type the following: **https://192.168.93.100**. If you cannot communicate to the CMM and you have tried another cable, and rebooted the PC, there is a pin hole reset option at the top of the CMM. Take a paper clip and insert it into the pin hole, push it in and hold for 10 seconds to reset the CMM to its factory defaults. The lights on the CMM should flicker and the fans in the chassis will reset after a few seconds. Resetting the CMM this way returns all configuration values to the factory defaults, including the enabled external ports of the IO modules. This implies that the reset

address will now be https://192.168.70.100, so you will need to put your PC on that address subnet.

- In the Enter Password window that is displayed, type the user name and password. The management module has a default user name of USERID and password of PASSW0RD (where 0 is a zero, not the letter o). Both of the values are case sensitive. If someone else has done the green light check on the chassis, they will have set the new password, else you will be prompted to change the password as this time. Write down the password that you enter in a secure place.
- Select a timeout value on the bottom of the page and click continue.
- A startup wizard runs that can be used for the following steps.

3.2.1.1 New and predefined CMM user IDs

It is recommended that you create a backup user account. To do this:

- From the **Mgmt Module Management** pulldown, select **User Accounts**
- Click on **Create User…**and fill in the ID and password information. Also choose a **Permission Group** for the new user and click **Finish.**

3.2.1.2 IPV6 addressing

On a PureFlex chassis, there is an internal network configured through the CMM. It handles both IPV6 and IPV4 addresses. The Flex System Manager (FSM) can also talk on the IPV6 network in order to manage the chassis. An IBM specific IPV6 prefix (fd8c:215d:178e:c0de:) is added to the link local IPV6 addresses to allow all of the chassis components to communicate. You have to manually add this prefix on each chassis component (see below for IPV4) and finish the IPV6 address with the last 4 octets of the IPV6 link local address shown on each component. The FSM needs to use the same prefix on its IPV6 address assigned to it's eth0 port. So as you go through the IPV4 steps below, also do the IPV6 addressing.

3.2.1.3 Assigning an IP address to I/O modules

With CMM configured the next step is to configure IP addressing for the other components within the chassis. You will need to manually configure the IP addresses for the chassis I/O modules. Instructions in the README files assume that the I/O module you are updating is accessible through the CMM browser UI. To assign an IP address to an I/O module:

- Log into the CMM browser UI with an administrator ID
- From the **Chassis management** pulldown, select **Component IP configuration**
- Select an I/O module by clicking on its **Description.**
- Select the **IPV4** tab.
- Change the **Configuration Method** to **Use Static IP Address** then enter the static address, subnet mask and default gateway IP addresses and click **Apply**.
- Change the IPV6 address using its tab and selecting **IPV6 Static Addressing**. Follow the description given in the "IPV6 addressing" section.
- Then click the OK button to exit.

If you intend to access the I/O modules through the CMM interface, ensure that they are on the

3.2.1.4 Assigning an IP address to the compute node's service processor

The next step is to configure IP addressing for the Power node(s) within the chassis. The Power node has the Flexible Service Processor (FSP) and the xNode has the Imbedded Management Module (IMM). A default value will have been set at the factory, but the client will likely want a different IP assigned. To assign an IP address to the Power node(s):

- Log into the CMM browser UI with an administrator ID
- From the **Chassis management** pulldown, select **Component IP configuration**
- Select a Power node by clicking on its **Description.**

- Select the **IPV4** tab.
- Change the **Configuration Method** to **Use Static IP Address** then enter the static address, subnet mask and default gateway IP addresses and click **Apply**.
- Change the IPV6 address using its tab and selecting **IPV6 Static Addressing**. Follow the description given in the "IPV6 addressing" section.
- Then click the OK button to exit.

3.2.1.5 Assigning an IP address to the storage node's service processor

Your chassis may contain a Flex System V7000 storage node. This is a V7000 that installs in multiple chassis bays. The storage node contains 2 Imbedded Management Modules (IMMs) and 2 controllers/canisters that each requires an IPV4 IP address. When you select this node from the CMM's Component IP configuration screen, you have an extra pulldown to choose which of these addresses you are going to configure. Configure each canister as follows:

- Set the **eth0** address for the IMM address.
- Set the **Service** address for the canister. This can be used to configure the node later.
- You won't be able to configure the Management address from the CMM. It can be set later after running the initial wizard on the V7000. The management IP is for managing both of the canisters. You get to the initial wizard by opening a browser to the URL: https://<service IP address here> You may have to try both canister service IPs to get to the wizard.
- On the V7000 management interface GUI, select the wrench icon and then Network to set the Management IP address. Also set the IPV6 address with the IBM prefix mentioned in the "IPV6 addressing" section.

3.2.1.6 CMM IP configuration

To configure the CMM so that it is accessible on the local network:

- Select the **Mgmt Module Management** pulldown, then click on **Network**
- Enter a host name for the CMM. If the CMM is going to use **DHCP**, choose the option **Try IP DHCP Server. If it fails, use static config** in the **IP address assignment method** drop-down menu
- If the CMM is going to use a static IP address, choose the option **Disable – Use static IP address**, then enter the static address, subnet mask and default gateway IP addresses
- Scroll to the top of the page and click **Apply**.
- Change the IPV6 address using its tab and selecting **IPV6 Static Addressing**. Follow the description given in the "IPV6 addressing" section.
- Then click the OK button to exit.
- A reset of the CMM is required to make the changes take effect.
- On the **Mgmt Module Management** pulldown, select **Restart**. Do a normal restart of the CMM. This takes a few minutes, you can start a continuous **ping <IP address> -t** from a command prompt to see when the CMM is back on line.

3.3 FSM Discovery and access of Power nodes

The FSM is based on an xNode, so as such; it contains an Imbedded Management Module (IMM). The IMM has an IP address, see the "Plan for necessary IP addresses" section. Use a browser and enter in the URL field: https://<IMM address here>. You should see a login prompt. The IMM uses the same UID/Password as the CMM.
Once you are logged in, select **Remote control** and answer the Java prompts until a console window is shown. There is a toolbar across the top of the console window. Look for the power on icon and power on the FSM node. Once booted, you will be presented with the welcome wizard. Your Lab Services or certified business partner can walk you through the wizard.

FSM supports the discovery of POWER nodes, but first you need to manage the chassis that the nodes are in. This is accomplished by the FSM communicating with the CMM. When the FSM starts it automatically discovers the chassis that it is installed in. Next you need to manage that chassis.

From the FSM **Home page**, click the **Manage Chassis…** link. A partially defined chassis should be shown. Select that entry and click **Manage**. This can take a while, as the components of the chassis are discovered and identified. Please be patient.

Discovery is the process which is used to establish connections with resources that can managed. A Power node is one type of resource. Discovery of nodes and virtual servers happens as part of managing the chassis. Next comes a two-step process in the following order: requesting access and running inventory. Each option is available by right clicking on the object and selecting either **Security->Request Access** or **Inventory->Collect Inventory.** Requesting access is an authentication step to ensure you have correct authorization to access that host (a user ID and valid password). Collecting inventory gathers vital product data about the Power node: ie firmware levels, virtual servers configured…

To work with the discovered Power node: from the **Home** tab of the FSM, click on the **Plug-ins** sub tab in the middle of the page. Next search down the list and locate the **Manage Power Systems Resources** link. Click on it. When the Manage Power Systems Resources page comes up, go to the upper right hand pull down and select **Add to my startup page**. The next time you log into the FSM, this tab will automatically be shown.

This page is a lot like the hardware management console (HMC) interface but the **hosts** link is the same as the Managed Server link on the HMC. Left clicking on a blue link takes you one layer lower (from hosts to the virtual servers on that host). The bottom layer is the properties for the object that you clicked on.

3.4 FSM chassis component firmware update process

The intention with PureFlex System is that the Flex System Manager (FSM) will supply a single interface for determining, downloading and installing all chassis component firmware. Assuming that the FSM has connections to the internet, this is done from the home page. The **Plug-ins** tab shows an **Update Manager** that can be used to access the chassis components and IBM fix central to determine which fixes would be applied. If internet access is not available, you can manually download updates from IBM Fix Central and import them into the FSM for applying them at a later time.

4 Storage management concepts

If you are need a better understanding of storage area networking concepts, refer to the **Storage area networks** 101 section of the IBM i Virtualization and Open Storage Read-me First at:
http://www.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

IBM i on all POWER nodes supports Fibre Channel (FC) only. Network Attached Storage (NAS) through iSCSI to VIOS is not supported for IBM i as direct installation disks. NAS storage, and any other storage that is supported by the IBM Storwize V7000 is supported for IBM i through VIOS.
Refer to the System Storage Interoperation Center (SSIC) for the supported combinations of adapters, switches and storage area networks (SANs):
http://www.ibm.com/systems/support/storage/ssic/interoperability.wss

4.1 Storage concepts for IBM i on a POWER node

In the POWER compute node environment, IBM i partitions do not have direct access to any physical I/O hardware on the node, in the chassis or outside the chassis. In PureFlex or Flex System, disk storage is provided by attaching LUNs on a Fibre Channel storage area network (SAN) to VIOS, then directly virtualizing them to IBM i using the FSM interface. DVD access for IBM i installation is provided by either using an external USB DVD or through the VIOS supplied virtual media library. N-port ID Virtualization (NPIV) attached storage, including tape media libraries, are discussed later in the "N-Port ID Virtualization (NPIV)" and "Save and restore with a Fibre Channel-attached tape library" sections.

A *path* to a LUN is a logical connection from the host HBA to the external storage box. The path may include virtual or physical HBAs and may have no switches or multiple switches. IBM i enforces up to 8 active *paths* to the *same* LUN. This is due to the overhead of allocating configuration resources for each path to the IBM i. Here are some ways that paths are created:

1. A direct cable from the host HBA to the external storage HBA will be 1 path.
2. Cabling a host FC adapter ports to a FC switch and zoning any of those adapter ports to a SAN host port will generate a path per world wide node name (WWNN). The IBM i host port will only log into one target port per WWNN even if there are multiple target ports cabled and zoned. The DS8000 is one WWNN where as the SVC and Storwize have 2 WWNNs per controller enclosure. All paths to the DS8000 are considered active where as the SVC or Storwize has one preferred controller that is considered the active path and the other controller will be the passive path. This includes native attached ports or VIOS (NPIV) attached ports.
3. Cabling multiple host ports to a switch and zoning those host ports to the IBM i will generate paths for every host port zoned to the partition. This includes native attached ports or VIOS (NPIV) attached ports.

Remember that the physical path must still be sized to provide adequate throughput. Since you have an 8 or 16 Gb physical port for the NPIV adapter you will have performance problems if too many clients attempt to use the NPIV adapter at the same time.

Note: Each adapter can support up to 128 total paths, but only 64 active paths. This limits the number of devices to 64. The only case where you should have more than 64 paths is when you use SVC or Storwize where you can also have the inactive path configured on the same host port as the active path. The 64 device/active path limit also applies to tape where the control path and tape device each take up a path.

Creating too many paths to your LUNs can potentially inhibit the total number of LUNs that can be allocated to an IBM i partition because there is also a limit of 64 LUNs per FC adapter *port* (direct or native connection) or virtual FC adapter (NPIV) port. NPIV can support up to 64 active and 64 passive LUNs per virtual path.

LUNs are only presented to the host from the owning IO group port.

You can map up to 64 virtual FC (NPIV) adapters from separate IBM i LPARs to the same FC port in VIOS.

This LUN limit only applies to IBM i clients as the limitation is enforced by the IBM i licensed internal code. The limitation of 64 partitions sharing a single FC port is enforced by the FSM/HMC and VIOS, and so that applies to any type of client partition.

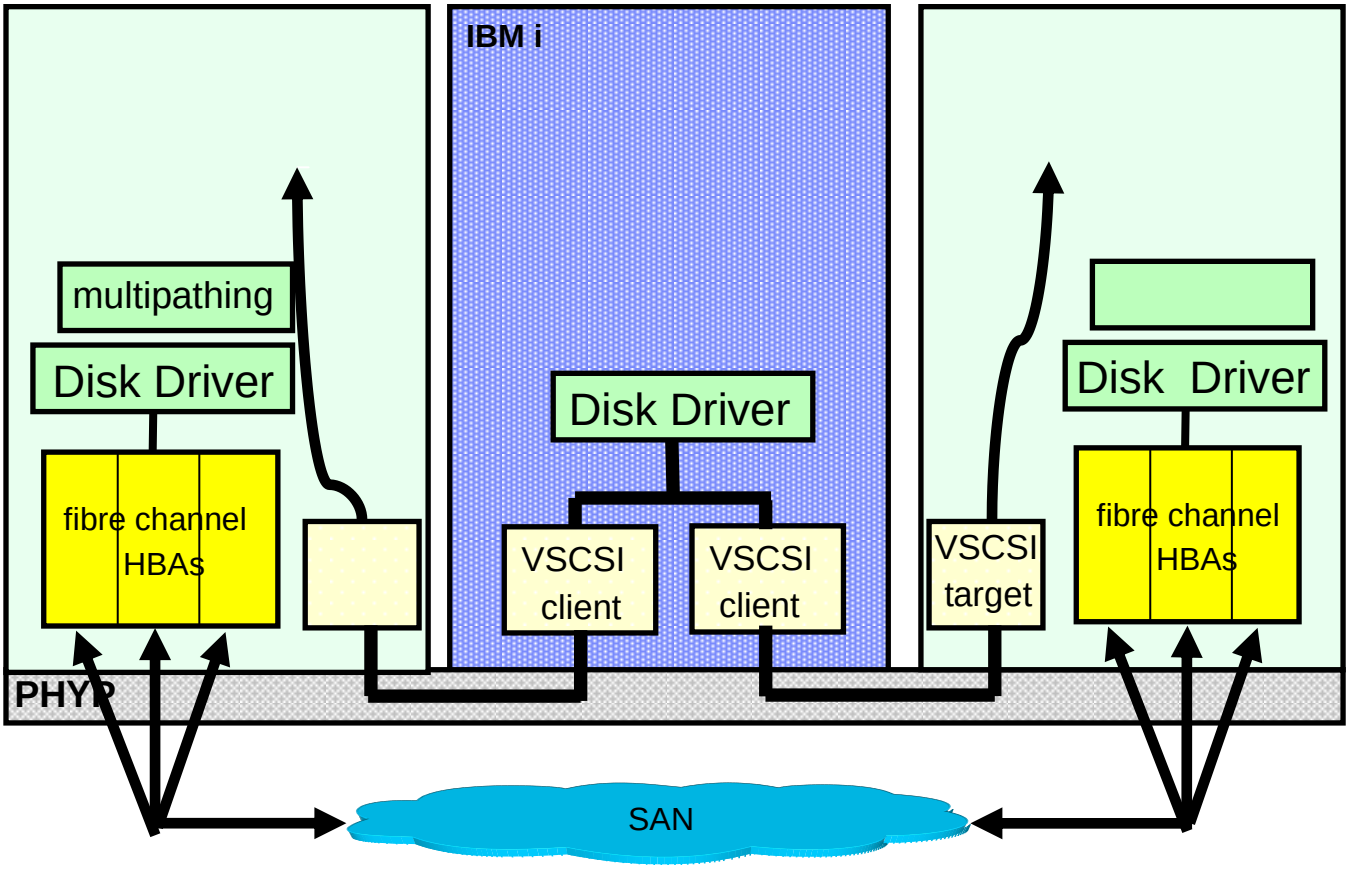

Figure 3 Overview of storage, optical and tape virtualization for IBM i on Power nodes.

An 8 Gb fibre channel (FC) mezzanine adapter is used to connect to FC storage. On a full-wide Power node, two FC adapters can be used to provide separate adapters for VIOS redundancy. The ports on the adapter are hardwired to switches in IO Bays 3 & 4.
Note: Ensure that you look at the combinations of adapters and switch modules for all of the nodes in your chassis. This is especially important when adding a node to an existing chassis.

4.2 PureFlex/Flex Configurations for Storage

When ordering the IBM **PureFlex** configuration with a POWER node as the primary node, you will need to choose which version you want: Express, Standard or Enterprise. Each comes with a predefined IBM Storwize V7000 LUNs and a zoned Fibre Channel switch or switches. For more detail on the content and configurations used in these offerings, please see the infoCenter documentation here: http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.pureflex.doc%2Fp7een_template_landing.html

The **Flex** configuration allows the use of any VIOS supported SANs to host IBM i. It's referred to as a build-to-order configuration. Nothing is preconfigured from IBM manufacturing, including cabling.

When choosing a configuration that uses fibre channel access to the SAN/storage you can choose between Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) technology. More details are contained in the next sections.

4.3 Fibre Channel Concepts

Fibre Channel technology involves compute node adapters and chassis switches of 8 or 16 Gb speeds, transceivers that install in the external switch ports of 8 or 16 Gb speeds and fibre cables for the data transmission. The chassis switch connections may go directly to the SAN or tape media library or they may be cabled to a core SAN switch or switches.

4.4 Fibre Channel over Ethernet Concepts

Fibre Channel over Ethernet technology involves compute node converged adapters and chassis switches of 10 Gb speeds, transceivers that install in the external switch ports of 8, 10 or 16 Gb speeds and typically fibre cables for the data transmission. The FC ports of the chassis switch connections may go directly to the SAN or tape media library or they may be cabled to a core SAN switch(es).

FCoE is also referred to as Fibre Channel over Converged Enhanced Ethernet (FCoCEE).
FCoCEE is an industry standard that allows a single converged adapter in a server to send and receive both Fibre Channel (FC) and Ethernet traffic over a 10-Gb Ethernet (10GbE) fabric. The benefits of this technology are:

- Fewer adapters per server
- Fewer switches and therefore lower energy consumption in the datacenter
- A single converged fabric for both FC and Ethernet traffic, enabling lower infrastructure costs and simplified management

FCoCEE works by encapsulating a FC frame in a new type of Ethernet frame. For a complete general overview of FCoCEE, consult the Fibre Channel Industry Association's Website at http://www.fibrechannel.org/.

Power compute nodes support FCoCEE through a Converged Network Adapter (CNA). Refer to the Flex Interoperability Guide (FIG) for the supported adapter and switch combinations at: http://www.redbooks.ibm.com/redpapers/pdfs/redpfsig.pdf For IBM i the CNA is owned by VIOS, as with other compute node adapters. VIOS then provides disk and network virtualization to any IBM i client partition through the CNA. Note that the converged adapter does not have both FC and Ethernet ports. Instead, the adapter sends and receives both 8Gb FC and 10Gb Ethernet traffic through its two 10GbE ports using FCoCEE frames.
Traffic from the CNA is routed to either:

- An EN4093 10Gb Network switch module in I/O bays 1 and/or 2, and from there to a separate FCoCEE-capable top-of-the-rack (TOR) switch outside the chassis. The TOR switch is capable of routing both types of traffic – 10GbE and 8Gb FC – to separate networks connected to it.
- **Or** to a CN4093 10Gb Converged Network switch module in I/O bays 1 and/or 2 that has both FC and 10 Gb ports to send the FC data to the SAN fabric and to send 10 Gb Ethernet packets out to other external switch(es).

VIOS sees both the Ethernet and FC ports from the CNA as if they are separate adapters. Virtual SCSI (vSCSI) is supported for disk storage. At the time of this writing NPIV is not supported for tape media

libraries or directly to disk storage. NPIV for storage is currently accomplished by putting the CN4093 into NPV mode and using a TOR switch.

The Storage and Network Convergence Using FCoE and iSCSI redbook located here: http://www.redbooks.ibm.com/redbooks/pdfs/sg247986.pdf should be used for education and configuration instructions in the PureFlex environment.

4.5 N-Port ID Virtualization (NPIV)

VIOS uses N_Port ID Virtualization (NPIV) to allow IBM i direct access to a SAN through an NPIV-capable adapter owned by VIOS and an NPIV capable FC switch(es). NPIV is a FC technology that enables a single port on a FC adapter to be presented to the SAN as an N-number of independent ports with different WWPNs. NPIV-capable adapters on IBM Power® nodes allow up to 256 virtual FC ports to be assigned to a single physical FC port. On Power nodes, VIOS always owns and manages the FC adapter.

To leverage NPIV, an IBM i LPAR must have a virtual FC client adapter created, which connects to a virtual FC server adapter in VIOS. However, the virtual FC client adapter does not allow IBM i to access LUNs already assigned to VIOS in the SAN. Instead, the virtual FC server adapter in VIOS is mapped to a FC port on the physical NPIV-capable adapter. This allows the client virtual FC adapter in IBM i direct access to the physical port on the FC adapter, with VIOS having a "pass-through" role, unlike with VSCSI. The logical drives or LUNs in the SAN for IBM i use are not mapped to the WWPNs of the physical ports on the NPIV adapter and it does not become available in VIOS first. Instead, when the virtual FC client adapter in the IBM i LPAR is created, two virtual WWPNs are generated by the PowerVM Hypervisor. The SAN LUNs are zoned directly to the first of the two WWPNs on the virtual FC client adapter in IBM i. The second WWPN is used to facilitate Live Partition Mobility. The PowerVM Hypervisor on a Power node has the default capability to create 32,000 virtual WWPNs. When virtual FC client adapters are deleted, WWPNs are not reused. If all of the default 32,000 WWPNs are used, the client must obtain an enablement code from IBM, which allows the creation of a new set of 32,000 WWPNs.

4.5.1 NPIV disk requirements

- You cannot manage the storage devices by using the SMI-S provider through IBM Flex System Manager.
- NPIV supports multi-disk virtual server disk attachment; if a virtual server has multiple disks, you can attach some of the disks with NPIV and others with vSCSI. The virtual servers can have the following types of paths to the storage on the SAN:
    - A combination of virtual SCSI and NPIV paths
    - Only NPIV paths
    - Each disk that is allocated to the operating system on the virtual server must access its storage through a VIOS virtual SCSI path or a VIOS NPIV path. A disk cannot access its storage through one VIOS virtual SCSI path and another VIOS NPIV path.

Note: NPIV for storage to the Storewize V7000 is supported with either type of chassis FC switch, but on a PureFlex configuration, NPIV is not configured by default. To use VMControl on a Power compute node a Brocade switch is required, either in the chassis (Feature code: FC5022) or as a TOR; the FSM Fixpack 3.1 updates must be installed for all of the components involved; and IBM i 7.1 TR6 PTFs must be applied. i 6.1.1 is not supported with this configuration.

Please refer to this information APAR for NPIV and IBM i 6.1.1 requirements:
http://www.ibm.com/support/docview.wss?uid=nas13b3ed3c69d4b7f25862576b700710198

4.5.2 Support statements and requirements for FC tape libraries

Please refer to this web page for NPIV tape requirements:
https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/IBM%20Remova

ble%20Media%20on%20IBM%20i/page/Fibre%20Channel%20Tape%20library%20devices%20attached%20using%20N%20Port%20ID%20Virtulization%20%28NPIV%29%20with%20VIOS

Note that while the supported FC adapter and switch modules operate at 8Gb/s, the rest of the SAN does not have to operate at that speed. Additional FC switches outside of the chassis can operate at 4Gb/s or 2Gb/s. While this difference in throughput will have an effect on performance, it will not prevent NPIV from working.

4.6 Best practices for PureFlex System and Fibre Channel storage

When configuring LUNs for IBM i (virtualized by VIOS), follow the best practices outlined in chapter 18 of the latest *Performance Capabilities Reference* manual, available here: http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html. Note that some of its recommendations apply only to IBM i using virtual storage outside of the chassis environment. Although this section was written for a POWER node environment, it still pertains to the POWER node environment.

In addition to the guidelines in the *Performance Capabilities Reference* manual, follow these additional recommendations:

- Use Fibre Channel disk drives (and not SATA or FATA) to create the RAID ranks/arrays for production IBM i workloads
- Use 15K RPM drives for medium and heavy I/O IBM i workloads, and 10K RPM drives for low I/O workloads
- When creating a host connection to the WWPN of the Fibre Channel card on the node, specify at most two specific host ports. Do not create the connection so that the Fibre Channel adapter can connect to all host ports on the storage subsystem, which is the default for some subsystems.
- Properly zone the switches between the adapter ports and the SAN host ports.
- When configuring vSCSI LUNs for IBM i within VIOS: as soon as the LUNs that will be virtualized to IBM i have reported in VIOS, change their queue depth to improve performance. Start a Telnet session to VIOS and login as **padmin**. Use the following command for each LUN = hdisk:
    - **lsdev - attr -dev hdisk#,** where # is the hdisk number you want to display.
    - **chdev -perm -attr queue_depth=8 -dev hdisk1#**
        - **Note:** The order of the parameters can be changed, as shown, to facilitate repeating the command and only having to alter the hdisk number.
        - **Note:** Some of the low end SANs might not handle the larger number of concurrent commands as well, which can adversely affect performance.
        - For redundant VIOS servers, each server needs to be able to access the same hdisks, so another attribute needs to be set for this: **reserve_policy=no_reserve** on each hdisk. Add a space between the attribute lists on the command.
    - **lsdev -attr -dev hdisk#**, to validate the change.

- Another parameter that can improve performance is the number of I/O commands to send to the Fibre Channel adapter. The recommended value is 512. Be sure to change the value for all ports of all of the FC adapters:
    - **lsdev -attr -dev fcs#,** where # is the FC adapter number you want to display.
    - **chdev -attr num_cmd_elems=512 –perm –dev fcs#**
    - **lsdev -attr -dev fcs#,** where # is the FC adapter number you want to display to validate the change.

- To improve reliability, enable dynamic tracking and fast fail over for the LUNs virtualized to IBM i. Fast_fail should be done only when multiple paths to the disks exist. Do so for all ports of the all of the adapters:
    - **lsdev -attr -dev fscsi#,** where # is the FC adapter number you want to display.
    - **chdev -attr dyntrk=yes fc_err_recov=fast_fail –perm –dev fscsi#**

- **lsdev -attr -dev fscsi#,** where # is the FC adapter number you want to validate.

Using the **–perm** option in these commands means that the value will be updated only in the VIOS device configuration database (ODM). To make the change effective, **reboot VIOS** when there is downtime available for all client partitions.

5 VIOS and IBM i installation and configuration

On PureFlex configurations, VIOS is always installed on all Power nodes in the order. IBM i can also be preinstalled when ordered as the primary operating system on the primary Power node. This section details the case for a secondary Power node or a Build to Order (BTO) configuration.

5.1 Obtain the VIOS and IBM i installation media and fixes

VIOS is part of IBM PowerVM Editions (formerly known as Advanced Power Virtualization) and is required in the IBM i on POWER node environment. Work with your local sales channel to ensure that PowerVM (Standard or Enterprise Edition) and the latest fix pack are part of the POWER node order. Refer to the Flex System Interoperability Guide (FIG), for current configuration options at:
http://www.redbooks.ibm.com/redpapers/pdfs/redpfsig.pdf

Work with your local sales channel to obtain the IBM i install media and the latest PTF package. The minimum supported release of IBM i is i6.1.1. i6.1 can not be ordered preinstalled from IBM manufacturing, but it is supported with a manual install. PTFs can also be obtained from IBM Fix Central: http://www-912.ibm.com/eserver/support/fixes/

5.2 Install PuTTY (optional)

As mentioned above, FSM is used for both LPAR and virtual resource management in this environment. To install and configure VIOS, and later to save or restore it, a Telnet session to the CMM or to VIOS is required. Use the PuTTY application any time a Telnet session is mentioned in this paper. PuTTY provides better terminal functionality than the Telnet client included with Windows and can be downloaded at no cost the following URL:
http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html

5.3 Installing and configuring a VIOS servers (if needed)

5.3.1 Using an external DVD drive with VIOS and IBM i

VIOS can virtualize an external USB DVD drive to IBM i. The DVD drive becomes physically available only to VIOS. It is then assigned to an IBM partition using the FSM as described in the Moving the DVD drive to another virtual server using FSM section. Note that cd0 is the name of the physical DVD drive as assigned by VIOS.

Note: The USB port on a Power node was only designed for a memory key, so its power capabilities are minimal. The only supported DVD drive must be installed in an IBM 7316 media enclosure as a USB device. Other devices may work, especially if they use an external power source, but they are not officially supported. Better alternatives may be sending VIOS iso files to the FSM for a VIOS installation using the installios command: see the "Installing VIOS from the FSM" section; or sending IBM i iso files to VIOS to add to a media repository: see the "Install IBM i using the VIOS virtual media library (recommended)" section for details.

When the DVD drive must be reassigned to another IBM i LPAR on the same node, the OPTxx device in IBM must first be varied off. Then the cd0 device can be assigned to the second LPAR using the same method described in the Moving the DVD drive to another virtual server using FSM section.

When the DVD drive is assigned to VIOS, you need to perform the following steps:

- SSH to VIOS on the correct node and sign in with **padmin** (as the user ID)

- Type **cfgdev** and press Enter
- Assign the DVD drive to the correct IBM i partition as described in the Moving the DVD drive to another virtual server using FSM section.
- Vary on the OPTxx device in IBM i

VIOS uses a USB 2.0 driver for physical access to the DVD drive. As a result, all client LPARs using the DVD drive as a virtual optical device from VIOS can perform faster reads and writes operations. One benefit of this enhancement is shorter install duration for the IBM i Licensed Internal Code (LIC) and operating environment, as well as Program Temporary Fixes (PTFs).

5.3.2 Writing to DVD-RAM media

VIOS has the capability to write to DVD-RAM media. For IBM i as client of VIOS all DVD-RAM operations that are supported by IBM i are available when using this drive. This allows IBM i to perform small saves and restores (up to 4.7 GB) using this DVD drive.

5.3.3 Powering on the node from FSM

- On the **Home** tab of the FSM, select the **Plug-ins** sub tab.
- Then select the **Manage Power Systems Resources** link a little lower on the page.
- On the **Manage Power Systems Resources** tab on the FSM, select the node from the list of **hosts**
- Right click on the node and select **Operations -> Power On.** Select the **Standby mode** option so that no virtual servers are automatically started on power on.
- View the node's system reference codes from:
    - Right click on the Power node and select Service and Support-> Reference Code history
    - The node should come up to Standby mode, ready to start virtual servers.

5.3.4 Memory recommendations for VIOS and the Power Hypervisor

A minimum of 4 GB of memory is recommended for VIOS with an IBM i client.

You should also plan on leaving at least 512 MB of memory unassigned for use by the system firmware (the POWER Hypervisor). There is no way to view the amount of memory used by the POWER Hypervisor through the FSM.

5.3.5 Processor recommendations for VIOS

For most IBM i on node configurations, it is recommended to configure VIOS with 0.5 processor units and one assigned (or desired) virtual processor. If VIOS is hosting three or more IBM i or other partitions, consider changing its processor allocation to one dedicated processor. Note that you must reboot VIOS to change from shared to dedicated processors.

Note: The Workload Estimator (WLE) can be used to estimate the VIOS that is hosting multiple workloads by using the following link: http://www-304.ibm.com/systems/support/tools/estimator/index.html

5.3.6 Creating a VIOS virtual server using the FSM

Perform the following steps to create a VIOS virtual server profile.

- On the Manage Power Systems Resources page, click the **Hosts** link.
- Right-click on the host on which you want to create the VIOS virtual server and click **System Configuration -> Create Virtual Server**.
- The virtual server ID will default to the next available virtual server number.
- Enter a name for the virtual server.
- Select **Environment: VIOS.** If you plan to use live partition mobility through this VIO server, check the **mover partition** box and then click **Next**.
- Set the desired memory value in GB and click **Next**. 4 GB is the minimum amount recommended.

- Set the desired processor configuration and click **Next**. Half a core is the minimum, but 1 core is recommended.
- Click the check box for the first virtual Ethernet adapter and click the **Edit** button. Set the port VLAN ID (PVID) to 4091 if you are matching the IBM preconfigured PVID. Check the boxes for **required for activation** and **Allow bridging**. The latter allows this adapter to be a bridge to a physical Ethernet port, called a shared Ethernet adapter (SEA) as a later step. If you plan for the client virtual server to have more than one LAN connection, add another virtual Ethernet adapter using a **PVID** of **1** and check the boxes for **required for activation, IEEE802.1Q** (if VLAN tagging is required and also add the vlan tag number that is needed) and **Allow bridging**. Click **OK**.
    - In the case of redundant VIO servers also set the priority field based on the following information: one VIOS has it's priority set to 1, which makes it the **primary** VIOS. The **backup** VIOS **has** to have it's priority set to 2. If you do not do this, you will cause a network storm in your network when the bridging is configured.
- Click **Next**
- The next section is for virtual storage adapters. These are best created once the click virtual server has been defined. Click **Next**.
- On the **Physical IO** page, if you plan to use an external DVD drive choose the **PCI-to-PCI Bridge** adapter as **desired** (to allow sharing). Choose an Ethernet adapter and a storage adapter for the VIOS to own as **required** and click **Next.**
- Review the summary and click **Finish**.

Follow the steps in the "Opening a VIOS console using " section followed by either the "If installing VIOS from DVD" section, "Installing VIOS from the FSM" section or the "If installing VIOS from the NIM" section to proceed with the VIOS install.

5.3.7 Installing a new VIOS

The instructions in this section explain how to install VIOS on the POWER node. VIOS may be pre-installed on the internal SAS drive of your node or on a SAN LUN, but you may have the need to install another VIOS server for redundancy.

The POWER node provides the option for two integrated SAS drives on the node. It is recommended to order both integrated drives on the node, install VIOS on the first SAS drive and mirror the entire install to the second drive. But this is not a requirement. A SAN LUN can also be used for VIOS. In the case of a full-wide node, only one VIOS server can use the SAS drives since there is a only a single controller for them.

5.3.7.1 Installing VIOS from the FSM

The FSM command line interface supports the installios command for installing VIOS from an external DVD drive or from an uploaded iso file (recommended). Use a secure transfer protocol list SCP to send the VIOS DVD iso files to the FSM. The /home/USERID is the default directory where these iso files will be stored. Use the FSM GUI to configure a virtual server for VIOS as noted above.

- Connect a putty session using the SSH option to the IP address for the FSM.
- Signon with the USERID and password are currently configured.
- Enter **installios** and hit enter to be prompted through the command.

You will be shown a list of Power compute nodes like this:
The following objects of type "managed system" were found. Please select one:

- 1. Server-7895-42X-SNxxxxxxx

- Enter the number for the node you want to install VIOS on.

You will be shown a list of VIOS virtual servers that are defined on that node:
The following objects of type "virtual I/O server partition" were found. Please select one:
1. 7895_SNxxxxxxx_VIOS1
2. 7895_SNxxxxxxx_VIOS2

- Enter a number (1-2): 1

You will be shown the VIOS profile, like this:
The following objects of type "profile" were found. Please select one:
1. 7895_SNxxxxxxx_VIOS1
- Enter a number: 1

Now it will ask for the path to the VIOS image to install. This can be the physical DVD (/dev/cdrom), which doesn't require the DVD volume name, or the SCP'd iso files (/home/USERID/<iso file name>:
- Enter the source of the installation images [/dev/cdrom]: /home/USERID/vios2222dvd1.iso

Next enter the IP addressing information that you want to assign to the VIOS server:
- Enter the client's intended IP address: 10.91.4.2
- Enter the client's intended subnet mask: 255.255.255.0
- Enter the client's gateway: 10.91.0.1

You can adjust the Ethernet adapter's speed settings (if you want, these were all defaults):
(Note: To use the adapter's default setting, enter 'default' for speed)
- Enter the client's speed [100]: auto
- Enter the client's duplex [full]: auto
- Enter the numeric VLAN tag priority for the client (0 to 7), 0=none [0]:
- Enter the numeric VLAN tag identifier for the client (0 to 4094), 0=none [0]:

You are prompted to do further network configuration (this example did not do that):
- Would you like to configure the client's network after the installation [yes]/no? no

You are prompted for the FSM's Ethernet adapter to use to connect to the VIOS's Ethernet adapter. This should be the eth1 connected through the chassis Ethernet switch 1:
Please select an adapter you would like to use for this installation.
(WARNING: The client IP address must be reachable through this adapter!
1. eth0
2. eth1 10.91.0.2
3. mgmt0 10.3.0.2
- Enter a number (1-3): 2

Next, the FSM will show a list of Ethernet adapters from the VIOS virtual server:
Retrieving information for available network adapters
This will take several minutes...

The following objects of type "ethernet adapters" were found. Please select one:
1. ent U7895.42X.105D07B-V1-C2-T1 ae4192a05902 /vdevice/l-lan@30000002 n/a virtual
2. ent U7895.42X.105D07B-V1-C3-T1 ae4192a05903 /vdevice/l-lan@30000003 n/a virtual
3. ent U7895.42X.105D07B-V1-C4-T1 ae4192a05904 /vdevice/l-lan@30000004 n/a virtual
4. ent U78AF.001.WZS01WE-P1-C34-L1-T1 0000c9f82c7c /pci@800000020000201/ethernet@0 n/a physical
5. ent U78AF.001.WZS01WE-P1-C34-L1-T2 0000c9f82c7e /pci@800000020000201/ethernet@0,1 n/a physical
6. ent U78AF.001.WZS01WE-P1-C34-L2-T1 0000c9f82c80 /pci@800000020000210/ethernet@0 n/a physical
7. ent U78AF.001.WZS01WE-P1-C34-L2-T2 0000c9f82c82 /pci@800000020000210/ethernet@0,1 n/a physical

Note the terms virtual versus physical, choose the first physical adapter:
- Enter a number (1-7): 4

FSM then recaps the selections you have made:
Here are the values you entered:

managed system = Server-7895-42X-SNxxxxxxx
virtual I/O server partition = 7895_SNxxxxxxx_VIOS1

profile = 7895_SNxxxxxxx_VIOS1
source = /home/USERID/vios2222dvd1.iso
IP address = 10.91.4.2
subnet mask = 255.255.0.0
gateway = 10.91.0.1
speed = auto
duplex = auto
configure network = no
install interface = eth1
ethernet adapters = 00:00:c9:f8:2c:7c

Press enter to proceed or type Ctrl-C to cancel...

- Press enter

The installation should run with a lot of textual output. This is basically a network install manager (NIM) install that is used a lot in the AIX/VIOS world. Keep an eye out for failure messages, especially when the FSM attempts to access the VIOS Ethernet adapter (a single ping command is attempted) . Failures at that point tend to be related to the chassis Ethernet switch configuration.
Scroll through the messages and look for "ioserver nimol: ,info=BOS install 100% complete".
This will be followed by a lot of clean up steps and finally "Successfully unconfigured NIMOL.". You will then be back at the command prompt.
Follow the steps in section "Opening a VIOS console using FSM" to see a login prompt for the VIOS.
Next, follow the steps in section "Completing the VIOS install".

5.3.7.2 Installing VIOS from a local USB DVD device

See the "Using an external DVD drive with VIOS and IBM i" section for supported DVD drives. The USB port on the front of a Power node has an associated PCI-to-PCI Bridge hardware component that needs to be assigned to the new VIOS virtual server in order for the DVD device to seen for the install. Change the VIOS virtual server profile to include this hardware.
Follow the steps in the "Opening a VIOS console using FSM" section and the sections that follow it for opening a console and working with the Systems Management Services (SMS) interface.

5.3.7.3 Installing VIOS from NIM

- First set up the NIM server for providing the VIOS image (not covered in this document)

Follow the steps in the "Opening a VIOS console using FSM" section and the sections that follow it for opening a console and working with the Systems Management Services (SMS) interface.

5.3.8 Starting the VIOS virtual server

To start the new VIOS virtual server in FSM:

- Select the virtual server from the Manage Power Systems Resources virtual server list
- Right click on the virtual server and click **Operations -> Activate -> Profile**

You are prompted with an **Advanced** button, where you can set the boot mode to stop at **SMS** just for this boot. Do that, clicking OK to start the activation and then follow the steps below.

5.3.9 Opening a VIOS console using FSM

To start a console in FSM:

- Select the virtual server from the Manage Power Systems Resources virtual server list
- Right click on the virtual server and click **Operations -> Console Window -> Open Terminal Console**
- Log into the console using the **FSM** user ID and password.
- Log into VIOS using the **padmin** user ID and password.

NOTE: If you have troubles launching this console, check on the **CMM -> Chassis Management -> Compute Nodes:** Select your Power node by clicking its selection box and click the **Settings** button. Next click the **Serial over LAN** tab. Make sure that the **Enable Serial over LAN** check box is unchecked.

5.3.10 Accessing the System Management Services (SMS) menu

- On the console for VIOS, as soon as the PFW screen starts to load, press **1** quickly to enter the System Management Services (SMS) menu
- If you miss the opportunity to press 1, PFW will attempt to boot the partition from the default boot device. If VIOS is not yet installed, the boot will fail and the SMS menu will eventually be displayed.
- On the main SMS menu, click **5.Select Boot Options**
- Click **1.Select Install/Boot Device**

5.3.10.1 If installing VIOS from DVD

- Click **3.CD/DVD**
- Click **6.USB**
- Select Option 1 to choose the **USB CD-ROM**
- Click **2.Normal mode**
- Click **1.Yes** to exit the SMS menus and start the VIOS install. Allow approximately 5 minutes for the install program to load from the media.

5.3.10.2 If installing VIOS from the FSM

- This install does not need to stop at the SMS menu.

5.3.10.3 If installing VIOS from the NIM

- On the main SMS menu, click the menu option **2.Setup Remote IPL (Initial Program Load)**
- Select the port to use (for most scenarios, port 1)
- Click **IPv4**.
- Select **1.BOOTP** as the Network Service
- Select **1.IP Parameters** to fill in the necessary data. Fill in the following information for the IP Parameters by selecting the item number and typing in the data:
    - **Client IP Address:** The IP address you have chosen for this VIOS partition.
    - **Server IP Address:** The NIM server's IP address.
    - **Gateway IP Address:** The gateway IP for this VIOS partition.
    - **Subnet Mask:** The correct subnet mask for the network segment of which VIOS is going to be a part.
- After entering the necessary information correctly, press the Escape key to go back a page and click **3.Ping Test**
- Click **1.Execute Ping Test** to test the connection. If the ping test is successful, press the **M key** to go back to the SMS main menu and continue the install. (If the ping test is not successful, check the network settings on the BladeCenter)
- On the SMS main menu, click **5.Select Boot Options**
- Click **1.Select Install/Boot Device**
- Click **6.Network** for the Device Type
- Click **1.BOOTP** as the Network Service
- Select the port that you just configured above (for most scenarios, port 1)
- Click **2.Normal Mode Boot**
- Click **1.Yes** to start the install (this will exit the SMS menus and the install image will start downloading from the NIM server)

5.3.11 Completing the VIOS install

This section applies to installing from DVD or NIM:

- As soon as the first "Welcome to the Virtual I/O Server" screen has disappeared, enter **1** on the prompt to confirm this terminal as the console and press Enter (the number will not appear on the screen)

- Enter the correct number option for the language you want to use during install
- On the **Welcome to BOS** menu, click **Change/Show Installation Settings and Install**
- Click **option 2** to verify that the correct disk unit is selected for installation. The SAS drive internal to the node is detected as hdisk0, whereas any LUNs attached to VIOS are detected as hdisk1, hdisk2 and so on.
- Click **option 0** twice to begin installation.
- After the installation has been completed and VIOS has been rebooted, log in with the default administrator user ID, **padmin.**
- You will be prompted to enter and verify a password for the **padmin user ID.**
- Hit enter on the prompt to accept licensing **<a>**.
- Enter **license –accept** to accept the VIOS license agreement.

5.3.12 Mirroring of VIOS

Because VIOS provides storage, optical and networking I/O resources to IBM i, any IBM i LPAR on the node depends on VIOS to be operational. Therefore, it is strongly recommended to employ disk protection for the VIOS installation. On a new POWER node where there is a pair of internal SAS drives on the cover of the node itself, the VIOS will can be installed on one of those drives and then mirrored to the other drive. Alternatively, you can install VIOS on a SAN LUN, leveraging the SAN's RAID disk protection. If VIOS is installed entirely on the SAN, it is strongly recommended to use multipath IO (MPIO) to access the LUNs, as discussed in "**Error! Reference source not found.**" section.

Note that if mirroring is used, the correct method to achieve a redundant VIOS installation is to install VIOS on one drive and then mirror to the second. The **incorrect** method is to install VIOS on both drives at the same time. This latter method would not result in two copies of all VIOS files and two bootable drives. Note also that mirroring VIOS does not protect any IBM i or other client partition data; it protects only the system volume group, or storage pool called ***rootvg*** where VIOS is installed.

After installing VIOS, perform the following steps to configure mirroring:

- On FSM select the VIOS virtual server from the **Manage Power Systems Resources** virtual server list, right click on the virtual server and click **Operations -> Console Window -> Open Terminal Console**
- Log into the console using the **FSM** user ID and password.
- Log in with **padmin** as the user ID**.**
- Identify a second available hdisk for the mirrored pair, such as hdisk1 using the **lsdev | grep hdisk** command.
- Add that hdisk to the rootvg with the command, **chsp –add hdisk1.** The command assumes rootvg if a storage pool is not specified
- Enable mirroring using the command, **mirrorios –f –defer hdisk1**. When possible, reboot VIOS to complete the mirroring. For an immediate reboot, omit the **–defer** option

The **mirrorios** command accomplishes three tasks: it creates a copy of all logical volumes (except for the system dump volume) in rootvg on the second hdisk, makes the second hdisk bootable and changes the VIOS bootlist to include it. After completing mirroring mirroring is complete, you can verify that all logical volumes in rootvg have a copy on both hdisks with:

- **lsvg –lv rootvg**. Check for the number 2 in the **PVs** column

You can verify that the VIOS bootlist now includes the second hdisk with:

- **bootlist –mode normal –ls**. The output should include both the hdisks. If both hdisks are not shown on the bootlist, then do a **bootlist -mode normal hdisk0 hdisk1** or whichever hdisks you want to use.

5.3.13 Configure networking in VIOS (if necessary)

If you installed VIOS from a NIM server, basic networking is already configured.

If you installed VIOS from DVD, perform the following steps to configure basic networking in VIOS:

- On FSM select the VIOS virtual server from the **Manage Power Systems Resources** virtual server list, right click on the virtual server and click **Operations -> Console Window -> Open Terminal Console.**
- Log into the console using the **FSM** user ID and password.
- Log into the VIOS using the **padmin** user profile.
- Use the **lsdev | grep en** command to identify the correct network device. In most cases, the first Ethernet adapter port (ent0) is used.
- The network interface that corresponds to **ent0** is **en0**
- To configure the correct **enX** interface: enter **cfgassist**
- Use the arrow keys to move down to the **VIOS TCP/IP Configuration** and press **Enter**
- Select the **enX** port that you want to configure and press **Enter**
- Fill in the host name, IP address, network mask, gateway, domain name and DNS IP address values (as required) and press **Enter**.
- You should see the word **running** in the upper left hand corner. If successful, this will change to **OK**. If unsuccessful, error messages are shown. Press **F3** or **Esc-3** to return to the previous menu and make corrections, or press F10 or **Esc-0** to exit to the command line.
- Test the configuration by pinging the gateway or another address that should be accessible.
- Write down the name of the physical adapter (such as ent0), as you will need it later to configure the virtual Ethernet bridge (as explained in the "Configure the Virtual Ethernet bridge (using FSM) - Express/Standard configurations" section. If you forgot to write it down, enter **lstcpip –num** to find which interface was configured.

Several VIOS commands are available to check, remove or change the networking configuration. For example, to list the current VIOS network configuration, use the following command:

- **lstcpip –stored**

To check that routing is configured correctly, use the following command:

- **lstcpip –routtable**

To remove the VIOS network configuration and start again (this command should be executed on the SOL console for VIOS that the FSM provides), use the following command:

- **rmtcpip –f –all**

  Ensure that you use the –f which forces the removal or the command may hang the session.

Note: Entering the above command from a Telnet or PuTTY session will disconnect the session as they are using the TCP address as the destination for the session. Use the FSM console option instead.

To learn about all available options for these commands and the rest of the VIOS network configuration commands, refer the **Network commands** section of the VIOS command reference in the IBM Systems Information Center at:
http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1_vios_commandslist.htm

5.3.14 FSM update of the Power node firmware (if necessary)

System firmware on a Power node can be updated through the FSM. To view firmware levels and update firmware, select the node from the list of hosts on the **Manage Power Systems Resources** page. Then click **Release Management -> Power Firmware Management** from the Action button. You may have to use the **Gather Target information** button.

To update the firmware from the FSM **Home** page, click on the **Compute Nodes: check and update firmware** link. Select the node to be updated while working through the wizard steps. This assumes internet access from the FSM. If you do not have internet access from the FSM, you can separately download updates from Fix Central and import them into the FSM for installation using the Update Management interface.

5.3.15 Update VIOS (if necessary)

5.3.15.1 Using the VIOS CLI

Updates to the VIOS software are released as fix packs, service packs or rarely for important individual fixes, interim fixes (ifixes). A fix pack is cumulative and contains all fixes from previous fix packs. Each fix pack also includes information on the level of VIOS to which it will update the system. To find your current VIOS level, use the following command in a VIOS console session:

- **ioslevel**

Then compare your current level with the level documented in the latest fix pack. To find the latest VIOS fix pack, visit http://www-933.ibm.com/support/fixcentral/options.

And use the following options:

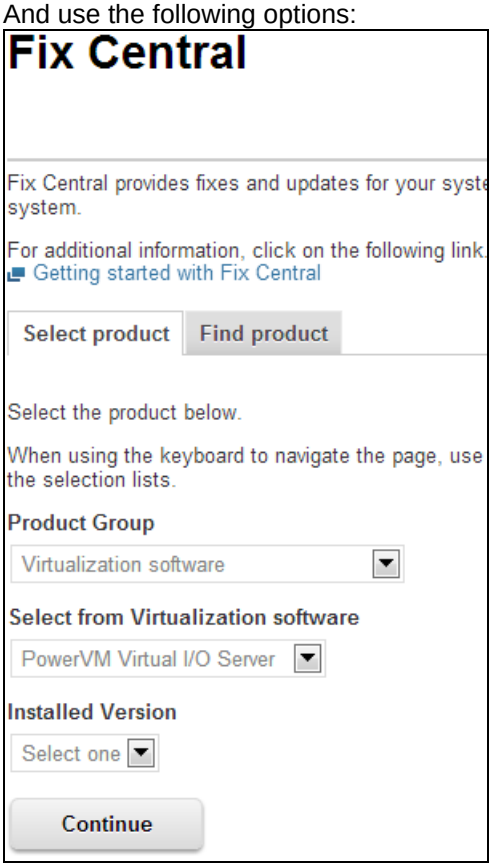

Select the Installed Version that you are running with.
To install a fix pack, service pack or ifix, click on it on the VIOS **Downloads** page above and follow the instructions in the README file for that update. Don't assume the instructions are always the same, they can change between updates.

5.3.15.2 Update VIOS using the FSM

- On the FSM **Manage Power Systems Resources** tab, expand Hosts and right click on the VIOS virtual server.
- Click **Release Management -> Show and Install updates.** This assumes that the FSM has internet access to IBM Fix Central and that VIOS is active.
- If updates exist, you will be shown an option to install them. You need to plan for a VIOS reboot, which implies that the **client virtual servers will need to be shutdown,** unless you are running redundant VIO servers.
- Also look for options to update the firmware on the Mezzanine adapters.

5.3.16 Displaying the current microcode level of the expansion adapters on the node:

Mezzanine firmware levels can be seen using the following command from the VIOS command line:

- **lsfware –all** (list firmware)
- The device names shown can be **fcs0** or similar for FC adapters. Notice a result similar to **fcs0 7710322577107501.03150150505**

5.3.17 Configure the Virtual Ethernet bridge (using FSM) - Express/Standard configurations

As discussed previously, IBM i does not have direct access to any physical network adapters on the POWER node. Instead, it uses a Virtual Ethernet adapter to connect to VIOS and any other virtual servers on the same node over a Virtual LAN (VLAN). VIOS in turn connects to the external LAN by using the Ethernet ports on the expansion card and an Ethernet I/O module in the chassis. The VLAN is bridged over to the external physical LAN using a Virtual Ethernet bridge in VIOS. The Virtual Ethernet bridge associates a VLAN with a physical network port in VIOS. This Layer-2 bridge allows externally routable IP addresses to be assigned to the IBM i partition for regular network communications.

- To configure the Virtual Ethernet bridge:
- From the **Manage Power Systems Resources** page, click on **Hosts** and then **right click** on the Power node you are configuring.
- Click **System Configuration -> Virtual Resources -> Virtual Network Management**.
- On the Virtual Network Management tab, select the VIOS server you want to configure the SEA on and click **Query.**
- On the bottom of the next page there is a **Create SEA** button, click it.
- You are presented with the virtual Ethernet adapter and a physical adapter that would make up the SEA. Click **OK** to create the SEA. **Note**: for redundant VIOS servers, you need an additional virtual Ethernet adapter that is not bridgeable to act as a heartbeat function between the VIO servers. This is called a control channel because it controls when a backup VIO server becomes the primary VIO server.
- By default, only one of a pair of redundant VIOS servers handles ALL of the Ethernet traffic for the client virtual servers. This is configured as *ha_mode=auto* attribute on the shared Ethernet adapter (SEA) command. To better utilize the 10 Gb Ethernet adapters owned by each VIOS server*, ha_mode=sharing* was implemented. The virtual Ethernet traffic handling is negotiated between the VIOS servers as additional LANs are configured, but there is a catch. You have to enable IEEE 802.1Q VLAN tagging on each virtual Ethernet LAN for this to work. A different port VLAN ID (PVID) is not enough. You don't have to specify an explicit VLAN tag, just enable the tagging support on each virtual Ethernet adapter. The FSM interface only supports the ha_mode-auto, you have to use the mkvdev –sea command to use ha_mode=sharing.

5.3.18 Configure the Virtual Ethernet bridge for multiple switches, Enterprise, or multi-chassis configurations

These configurations contain redundant chassis Ethernet switches. There are a few ways to configure these switches and the method you choose affects the way you configure the link aggregation adapter (lnagg) in the VIO server. If your switch and VIOS configurations don't match your communications to, and through, VIOS will not work.

- You can have the switches configured as separate switches for redundancy. The switches are configured for failover: if switch 1 fails, switch 2 takes over. One switch is active, the other is passive until needed. The lnagg mode should be set to standard and the second port, typically ent1 should be specified with the attribute of backup_adapter=ent1. This configuration is referred to as Network Backup Interface (NIB).

- You can configure the switches as one virtual switch. This gives you redundancy and more throughput. This is done with a virtual link aggregation (VLAG) set of commands on the switches. The configurations on both switches have to be nearly identical.
    - The internal switch ports (ie INTA#, where # is the compute node bay number) can be configured as ***statically*** aggregated using ***portchannel*** groups. In a static configuration the VIOS lnagg mode should be ***standard*** and you should not configure a backup adapter.
    - The internal switch ports (ie INTA#, where # is the compute node bay number) can be configured as ***dynamically*** aggregated using link aggregation control protocol (***lacp***) keys. In a dynamic configuration the VIOS lnagg mode should be ***8023ad*** and you should not configure a backup adapter.
    - Please see the EN4093 10Gb Scalable switch Application Guide: http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.networkdevices.doc/00ay513.pdf and the EN4093 command reference guide: http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.networkdevices.doc/88y7943.pdf for more details.

The following is an example of the configuration steps for the VIOS lnagg with **statically aggregated** ports. This assumes that the switch configuration for VLAG and portchannel is already in place on both chassis Ethernet switches.

First, you need to shut off the redundant Internal ports on the second EN4093 Ethernet switch in IO Bay 2. Use the Chassis Management Module's component IP prompt to determine the switch's IP address.

To shut off internal ports on the switch, use an SSH interface to the switch's address and make sure you are in iscli mode. Then run the following:
en
conf t
interface port INTA1-INTB14
shutdown

On VIOS, determine the physical Ethernet ports that are being seen:
lsdev | grep ent This should show you a list of physical and virtual Ethernet ports along with a Shared Ethernet Adapter (SEA). Locate the first 2 physical ports, typically ent0 and ent1.

Now configure a link aggregation between those first two physical ports:
cfgassist -> Devices -> Link Aggregation Adapter -> Add a Link Aggregation Adapter -> then select ent0,ent1 by hitting **F7** on each as the Target ADAPTERS. Add hash_mode=src_dst_port for the ATTRIBUTES field. Then hit enter.
Or use this command:
mkvdev –lnagg ent0,ent1 –attr hash_mode=src_dst_port

This should create another ent# for your link aggregate adapter.

Now make a shared Ethernet adapter (SEA) that bridges the virtual Ethernet adapters (ent4 and ent5) to the link aggregate adapter (ent# created above):
mkvdev –sea ent# -vadapter ent4,ent5 –default ent4 –defaultid 4091

This should create another entY for your SEA

Now assign your IP address to the SEA:
cfgassist -> VIOS TCP configuration -> choose the en# of the link aggregation port -> enter the field parameters for your IP configuration and hit enter.
Or use this command:

mktcpip -hostname <VIOS hostname> -inetaddr <VIOS IP address> -interface **enY** –netmask <subnet mask> –gateway <gateway IP address> –nsrvaddr <DNS server IP address> -nsrvdomain <domain name> –start

Test the IP by pinging another address on the same subnet.

To turn the internal ports back on:
en
conf t
interface port INTA1-INTB14
no shutdown

Note: By default, only one of a pair of redundant VIOS servers handles ALL of the Ethernet traffic for the client virtual servers. This is configured as *ha_mode=auto* attribute on the shared Ethernet adapter (SEA) command. To better utilize the 10 Gb Ethernet adapters owned by each VIOS server*, ha_mode=sharing* was implemented. The virtual Ethernet traffic handling is negotiated between the VIOS servers as additional LANs are configured, but there is a catch. You have to enable IEEE 802.1Q VLAN tagging on each virtual Ethernet LAN for this to work. A different port VLAN ID (PVID) is not enough. You don't have to specify an explicit VLAN tag, just enable the tagging support on each virtual Ethernet adapter. Note: The Edit option will not allow you to set bridging on an active adapter. You must shutdown VIOS or add a new adapter.

5.3.19 MPIO drivers

VIOS includes a basic MPIO driver, which has been the default instead of the RDAC driver since November 2008. The MPIO driver can be used with any storage subsystem that VIOS supports and is included in a default install. In this case, configuration is required only on the storage subsystem in order to connect a single set of LUNs to both of the ports on a FC adapter owned by VIOS.

VIOS also supports the Subsystem Device Driver – Path Control Module (SDDPCM) for certain storage subsystems. The Storwize V7000 San is one of these. To find out whether a particular storage system supports SDD-PCM for VIOS, refer to its interoperability matrix on the SDDPCM Web site: http://www.ibm.com/support/docview.wss?uid=ssg1S4000201 Note that there are separate support statements for AIX and VIOS. If SDDPCM is supported on your storage subsystem for VIOS, download and install the driver following the instructions in the *Multipath Subsystem Device Driver User's Guide* found at that same site.

It is recommended that you install SDDPCM when connecting to the Flex System V7000 storage node. This is not installed by default.

5.3.20 Changing memory configuration for VIOS

A minimum of 4 GB of memory is recommended for VIOS with an IBM i client.
To change the amount of memory that VIOS is using from FSM**:**

- From the **Manage Power Systems Resources** page
- Open the list of virtual servers and right-click the VIOS partition.
- Click **System Configuration -> Manage Virtual Server** and then click the **Memory** tab.
- Change the **assigned memory** and click **Apply.**
- If this is a permanent change, then click the **Tasks** button on the right hand side and select **Save Current Configuration** option. Then select **to overwrite the existing profile** to update the VIOS profile with this same change.
- Close the Manage Virtual server tab.

Note that if the new assigned value is higher than the maximum memory value, a restart of VIOS is required to activate the change. Before restarting VIOS, gracefully shut down any client partitions on the node.

5.1 Redundant VIOS partitions using FSM

FSM adds support for redundant VIOS virtual servers on a POWER node. Each VIOS must have shared access to the same storage LUNs used for the client IBM i virtual server.
Refer to the sections 5.3.2.2 and 4.4.1 in the IBM i Virtualization and Open Storage Read-me First at:
http://www.ibm.com/systems/resources/systems_i_Virtualization_Open_Storage.pdf

The Power 270 compute node supports a Dual VIOS adapter (it's the only one that does) that allows the two local disk drives on the cover of the compute node to be owned separately by two VIO servers. The assumption is that you would later mirror each drive to a LUN also assigned to that VIO server. Ensure that you have the 4 port mezzanine adapters for Ethernet and fibre channel and the associated switch upgrades (as required) to allow each VIO server to meet their IO requirements.

By default, only one of a pair of redundant VIOS servers handles ALL of the Ethernet traffic for the client virtual servers. This is configured as *ha_mode=auto* attribute on the shared Ethernet adapter (SEA) command. To better utilize the 10 Gb Ethernet adapters owned by each VIOS server, *ha_mode=sharing* was implemented. The virtual Ethernet traffic handling is negotiated between the VIOS servers as additional LANs are configured, but there is a catch. You have to enable IEEE 802.1Q VLAN tagging on each virtual Ethernet LAN for this to work. A different port VLAN ID (PVID) is not enough. You don't have to specify an explicit VLAN tag, just enable the tagging support on each virtual Ethernet adapter.

5.2 End to end LUN mapping for VSCSI using FSM

You may need to map the VIOS hdiskx back to their associated LUNs for debug, there is an FSM interface to help do that:

- On the **Manage Power Systems Resources** page, select the Power node
- Right click on it and select to **System Configuration -> Virtual Resources -> Manage Virtual Storage**
- Query the VIOS server you are using.
- Click on the **Physical Storage** tab.
- Slide the columns over to see the full values of the Physical Location Code column. You should see something like this:
- W500507680210976C-**L1**0000 The L1 is the LUN number, or on the Storewize V7000 the SCSI-ID (seen from the logical drive properties on the V7000).
- There is a PTF for IBM i (i7.1 SI45124) that can display the LUN from the IBM i WRKHDWRSC *STG command.

5.3 IBM i install and configuration

FSM defaults the load source and alternate initial program load (IPL) adapters to the VSCSI client adapter in the IBM i partition. That adapter would have the hdisks for IBM i mapped to it and the VIOS media repository (virtual optical drive) mapped to it.

If you choose to use shared processors for IBM i, FSM defaults to assigning 0.1 times the number of processors you select as shared processor units and the whole number of processors you select as virtual processors. For example, if you select four shared processors, FSM initially assigns 0.4 processing units and 4 virtual processors to the partition. Also note, that by default, shared processor partitions are configured as uncapped. When assigning memory during partition creation, you are selecting the assigned or desired amount for the partition. FSM automatically assigns minimum and maximum values. The default processor and memory configuration can be changed by working with the partition properties after creation.

The minimum recommended amount of memory for an IBM i client partition on the POWER node is 1 GB. This is really a minimum value and should be set higher. The actual memory and processor values should be sized individually for each IBM i workload using the Workload Estimator, available at http://www-304.ibm.com/systems/support/tools/estimator/index.html.

5.3.1 Create LUNs for the IBM i partition(s) in PureFlex System

To create additional LUNs for IBM i (virtualized by VIOS) on the Storwize V7000 storage subsystem, follow the instructions in Chapter 5 of the Implementing the IBM Storwize V7000 in IBM Redbooks available at: http://www.redbooks.ibm.com/abstracts/sg247938.html?Open

5.3.2 Create the IBM i Virtual Server using FSM

Perform the following steps to create the IBM i virtual server.

- On the **Manage Power System Resources** page, left click the word **Hosts.**
- Right-click on the Power node on which you want to create the IBM i partition (virtual server) and click **System Configuration -> Create Virtual Server**. The virtual server ID will default to the next available virtual server number.
- Enter a name for the virtual server.
- Select **Environment: IBM i.** If you plan to use partition suspension, check that box and then click **Next**.
- Set the desired memory value in GB and click **Next**.
- Set the desired processor configuration and click **Next**.
- Click the check box for the first virtual Ethernet adapter and click the **Edit** button. Set the port virtual Ethernet (VLAN or PVID) to 1, which should be the default, or to a specific vlan id to match the client's networking conventions, and click **OK**.
- If you need other network connections for this partition add another VE adapter and specify a different PVID value. But ensure that the VIOS VE configuration contains that ID as a tagged vlan.
- Click **Next**
- If you are using virtual SCSI adapters for the LUNs, select **Yes, Automatically manage the virtual storage adapters for this Virtual Server** and then select **Physical Volumes.** VIOS will be queried to see which hdisks are available. The list will be presented to you. Select the hdisks that you want the IBM i partition to use. Click **Next**.
- If you are using virtual Fibre Channel (vFC=NPIV) attached LUNS, select **Yes, Automatically manage the virtual storage adapters for this Virtual Server** and then select **fibre channel**. This will create the vFC adapter pair between the IBM i and the VIOS server.
- If you plan to use an external DVD drive or the VIOS supplied virtual media library, select those devices on the next screen. Multiple iso files can be selected for sequential access. If you are not seeing either of these devices, follow the steps in sections "Using an external DVD drive with VIOS and IBM i" and "Moving the DVD drive to another virtual server using FSM". Once those steps are completed, restart the create virtual server wizard.
- Choose the adapters you want to use for physical I/O (typically none) and click **Next**
- Select the load source adapter (its typically the virtual SCSI adapter), alternate restart device (if any, its typically the virtual SCSI adapter) and the FSM will supply the console and click **Next.**
- Review the summary and click **Finish**.

5.3.3 Increasing the number of virtual adapters in the IBM i partition

The default number of virtual adapters for a client partition is 10. VIOS uses the slots 0 to 10, so these are reserved. The overhead of increasing this value is approximately 1KB of memory per adapter increase used by the hypervisor who manages these. As a rule of thumb, estimate 4-6 virtual adapters per client partition being hosted. Total those up and add a few more, since the value can't be changed while VIOS is active.

5.3.3.1 Using the VIOS CLI

If there is a need for more than 10,slots you can use the **chsyscfg** command to increase this value.

- **chsyscfg -r prof -i "lpar_name=<partition3>,max_virtual_slots=<300>"**

The <values> are variables

5.3.3.2 Using FSM

If there is a need for more than 10 slots, you can use the **FSM** GUI to increase this value by performing the following steps.

- On the **Manage Power Systems Resources** page, click on **Virtual Servers**.
- Select the virtual server for VIOS, right-click on the name and click **System Configuration -> Manage Profile**.
- Select the profile name associated with VIOS by clicking on the name.
- Click the Virtual Adapters tab.
- Change the value in the **Maximum Virtual Adapters** field and click **OK** to accept the change.
- This change will not take affect until the profile is used to activate the VIOS, which implies that the VIOS has to be brought down without an automatic restart.

5.3.4 Creating multiple Virtual SCSI adapters per IBM i partition

IBM i has a restriction of 16 disks (LUNs) per VSCSI adapter. It is possible to create multiple Virtual SCSI client adapters per IBM i partition on a node to allow:

- More than 16 disk and 16 optical devices can be virtualized by VIOS per IBM i partition
- Disk and optical devices can be configured on separate Virtual SCSI adapters

With FSM, you can manually create more adapter pairs. When you have done this, you can use the FSM to assign drives to any additional adapters. Follow the steps in the "Mapping storage to new Virtual SCSI adapters using the FSM" section.

5.3.4.1 Creating Virtual SCSI adapters using FSM

When using an FSM to manage Power nodes, you can create virtual SCSI adapters to use between VIOS and the IBM i client partition. In this section you will see how to create server adapter and client adapter pair.

Create a Client SCSI Adapter

Perform the following steps to create a client SCSI adapter:

- On the Power Systems Resources tab of the FSM, click the Virtual Servers link, then select the checkbox next to the IBM i client you are configuring
- Right click and select **System Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit…** to edit the profile
- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> SCSI Adapter…**
- Note the Adapter ID assigned and **click** the checkbox for "This adapter is required for virtual server activation"
- Select the VIOS server from the Server partition list.
- Click the **System VIOS Information** button to see the currently defined adapters on the server. You need to choose an available adapter number. Then close that window.
- Specify a Server Adapter ID (remember it for the next step) and click **OK** to create the client Virtual SCSI Adapter

Create a Server SCSI Adapter

Perform the following steps to create a server SCSI adapter.

- Select the checkbox next to the VIOS server. Click **Actions-> System Configuration -> Manage Virtual Server**
- Click the **Storage Adapters** tab and then click the **Add** button.
- Specify the Adapter ID you used above and an adapter type of **SCSI**

- Use the list to select the IBM i client partition you want to use
- Specify the adapter ID you created above as the connecting adapter ID and click **OK** to create the server Virtual SCSI Adapter
- These changes are dynamically applied to the VIOS server, so you need to also update the server's profile. On the right hand side of the window, click the **Tasks** button and **Save Current Configuration** option. Then select **to overwrite the existing profile** to update the VIOS profile with these same changes.
- Close the Manage Virtual server tab.

5.3.4.2 Creating Virtual SCSI adapters using the VIOS CLI

If you choose to create additional Virtual SCSI client adapters using the VIOS command line do the following commands:

- Log into VIOS with **padmin** or another administrator user ID.

You need to determine the partition ID VIOS is using. To display the current names and IDs of all existing partitions, use the following command:

- **lssyscfg -r lpar -F "name,lpar_id"**
- If the IBM i partition is not activated, refer to the following example that adds a new Virtual SCSI client adapter in **slot 5** of IBM i partition **"test,"** connecting to a server adapter in the **slot 17** in the partition named **"VIOS**":
    - **chsyscfg -r prof -i "name=test,virtual_scsi_adapters=5/client/1/VIOS/17/1"**
- If the IBM i partition is running on a Power compute node named **P270 Bay2**, refer to the following example that creates a new Virtual SCSI client adapter in **slot 5** of IBM i partition **"test",** connecting to a server adapter in **slot 17** in the partition named **"VIOS**":
    - **chhwres -r virtualio --rsubtype scsi -p test -o a -s 5 -a "adapter_type=client,remote_lpar_name=VIOS,remote_slot_num=17" –m "P270 Bay2"**
    - **Note:** The Power service processor (FSP) will flag an exception: B7005191: A virtual adapter configuration error occurred. between the dynamical add of the client adapter prior to the server adapter being created. There's nothing to be done about it.
- The corresponding server adapter in VIOS is created with this command:
    - **chhwres -r virtualio --rsubtype scsi -p VIOS -o a -s 17 -a "adapter_type=server,remote_lpar_name=test,remote_slot_num=5" –m "P270 Bay2"**
- Notice that there are three variables in the previous commands: the name of the IBM i partition, the new slot for the Virtual SCSI client adapter, and the name of the VIOS partition.

To display which virtual slots have been already been used, by partition and adapter, use:

- **lshwres -r virtualio --rsubtype slot --level slot** (notice the double dashes)
- The type of adapter is shown on the middle of each line. The slot numbers are shown at the beginning of each line.

5.3.4.3 Mapping storage to new Virtual SCSI adapters using the VIOS CLI

After creating a new Virtual SCSI client adapter for IBM i and the associated server adapter for VIOS, you can assign additional LUNs to IBM i by mapping them to the new server adapter in VIOS.

To display LUNs (or other physical volumes) that are available to be assigned to IBM i on the node, use the following list physical volume command:

- **lspv -avail**

To display all existing virtual resource mappings by Virtual SCSI server adapter in VIOS (vhostX) and client partition, as well as any newly created Virtual SCSI server adapters, use the following list mappings command:

- **lsmap -all | more**
- Press the **Space bar** to move forward one screen of data at a time and the **Down Arrow key** for one line at a time. Enter **q** to quit.

Any new Virtual SCSI server adapters will have no resources mapped to them. Assuming that **"hdisk7"** is an available LUN and **"vhost1"** is a newly created Virtual SCSI server adapter, use the following command to make hdisk7 available to IBM i:

- **mkvdev -vadapter vhost1 -vdev hdisk7**

The **lsmap** command above will also show whether a physical DVD drive in the node (typically **cd0**) is already assigned to a client other partition. If so, a **vtoptX** device will exist under a Virtual SCSI server adapter (**vhostX**). Skip this step if the DVD drive is not already assigned to a client partition. To map the DVD drive to a different Virtual SCSI adapter, first delete the correct existing **vtoptX** device (such as **"vtopt0"**) using the following command:

- **rmdev -dev vtopt0**

Next, assign the physical optical device (such as **cd0**) to the IBM i partition using the correct separate Virtual SCSI adapter (such as **"vhost 1"**) using the following command:

- **mkvdev -vdev cd0 -vadapter vhost1**

Note: Only map up to 16 hdisks to a client partition for IBM i. More LUNs are allowed to be mapped to AIX virtual servers, up to 256, but IBM i only supports up to 16. The hdisk mapping may succeed, but the IBM i will only show the first 16 LUNs.

To map a file-backed optical device to a new Virtual SCSI adapter (such as **"vhost 1"**), using the following command:

- **mkvdev -fbo -vadapter vhost1**

5.3.4.4 Mapping storage to new Virtual SCSI adapters using the FSM

- On the **Manage Power Systems Resources** page, select the Power node
- Right click on it and select to **System Configuration -> Virtual Resources -> Manage Virtual Storage**
- Query the VIOS server you are using.
- Click on the **Physical Volumes** tab.
- Select an **hdisk** from the list that is not currently assigned to a virtual server and click on **Modify assignment** button.
- Look near the top of the screen for the **New virtual server assignment pulldown.** Select the virtual server and the correct associated virtual SCSI adapter number (if present) to map the storage to. If only 1 vSCSI adapter is defined in a virtual server, no vSCSI adapter number will be shown.
- Next click **OK.**
- Repeat the process for each hdisk that you want to map.
- When you exit the virtual storage pane, the hdisks will be mapped to the associated adapter.

5.3.4.5 Removing Virtual SCSI adapters using the VIOS CLI

The VIOS command line is also used to remove Virtual SCSI client adapters from an IBM i partition. Note that removing a Virtual SCSI client adapter from IBM i will make any devices it provides access to unavailable. As mentioned above, to check which devices in VIOS are mapped to which Virtual SCSI server adapter, and therefore which client partition, use the following command on the VIOS command line:

- **lsmap –all | more**
- Press the **Spacebar** key to move forward one screen of data at a time and the **Down Arrow** key for one line at a time, and enter **q** to quit **.**

To remove a Virtual SCSI client adapter when the IBM i partition is not activated, refer to the following example which removes the client adapter in **slot 5** of IBM i partition **"test":**

- **chsyscfg -r prof -i "name=test,virtual_scsi_adapters-=5/////"** (note the minus sign before the equal sign)

To remove a Virtual SCSI client adapter when the IBM i partition is running, refer to the following example which removes the client adapter in **slot 5** of IBM i partition **"test":**

- **chhwres -r virtualio --rsubtype scsi -p test -o r -s 5**

5.3.4.6 Removing virtual SCSI adapters using FSM

Note that removing a Virtual SCSI client adapter from IBM i will make any devices it provides access to unavailable.

- Log into FSM with **USERID** as the user ID or another user ID
- On the **Manage Power Systems Resources** tab of the FSM, select the node from the list of hosts and select the checkbox next to the IBM i virtual server you are configuring
- Click **System Configuration -> Manage Profiles**
- Select a profile and then click **Actions->Edit…** to edit the profile

Click the **Virtual Adapters** tab, select the adapter and then click **Actions -> Delete**

5.3.5 Configure an IBM i console connection using PCOMM

IBM i requires a Personal Communications 5250 emulator client application to be used as the console for the operating system. If you happen to have the System i Access application installed on a PC, there is an emulator option in that product that can be used. If you don't have that application on your PC, you can download a version from this website: http://www.ibm.com/software/network/pcomm/ Once this application is installed, start it up to configure the console using the following steps depending which version you are using:

System i Access:

- Follow the instructions starting at ***step 2*** that are listed on this web page: http://www.ibm.com/support/docview.wss?uid=nas137396cfd67d5ef5886256f01000bda50
    - Use the FSM IP address in place of the HMC IP address if you are using an FSM but still use the userid of Q#HMC as shown.

Personal Communications:

- **Start Programs->IBM Personal Communications->Start or Configure sessions.**
- Click the **New Session** button.
- Type of host is **iSeries,** then click the **Link Parameters** button.
- For the **workstation ID** use the value: **Q#HMC.**
- On the **Primary** host name or IP address field, enter the **IP address of the FSM**. Change the **Port** field on the right hand side to **2300.** Then click **OK.**
- Then click **OK** again.
- A 5250 emulation window will come up and wait for connection from the IBM i.

5.3.6 Install IBM i using the a Network install

Starting in IBM® i 6.1, you have the ability to perform complete installations and upgrades of IBM® i partitions, in addition to loading PTFs or restoring data via a virtual optical image on an image server, all without having to manually FTP the data across the network to the individual partitions. This change allows for the install of Licensed Internal Code, operating system, licensed programs and PTFs across IBM® i partitions from a single image source. This capability is available through PTFs beginning in i 6.1. The assumption here is that you have successfully saved a valid IBM i image to a file server that can be accessed from your Power node.

The steps to configure and install across a network connection are described in this document: ftp://public.dhe.ibm.com/systems/support/power/i/nfs_optical_upgrade.pdf

5.3.7 Install IBM i using the VIOS virtual media library (recommended)

VIOS comes with a virtual media library that is preloaded with iso files for the preferred operating system that was ordered. This library can be shared across multiple partitions doing installs or PTF loads concurrently. Multiple iso files can be selected for sequential access.
From FSM

- On the **Manage Power Systems Resources** page, select the Power node

- Right click on it and select to **System Configuration -> Virtual Resources -> Manage Virtual Storage**
- Query the VIOS server you are using.
- Click on the **Optical media** tab.
- Select the iso image or images that you want to use and then click **Modify Partition assignment** to assign them to your partition. If you choose more than one image to be mounted, you will need to choose the specific image to use during the install process.

5.3.7.1 Adding IBM i iso files to the media repository

There are a few ways to add iso files to the media repository:

- FTP the iso files to the VIO server.
  - These will end up in the /home/padmin directory. Just don't do too many or you will run out of space in the VIOS file system and cause command failures.
  - From the FSM:
    - On the **Manage Power Systems Resources** page, select the Power node
    - Right click on it and select to **System Configuration -> Virtual Resources -> Manage Virtual Storage**
    - Query the VIOS server you are using.
    - Click on the **Virtual Optical media** tab. Add media from file system. Use short names without the .iso value and specify the path as /home/padmin/< iso file name here>.
  - Remove the files from /home/padmin using rm /home/padmin/<iso file name here>
- From physical DVD's and a supported DVD drive connected to the Power node.
  - From the FSM:
    - On the **Manage Power Systems Resources** page, select the Power node
    - Right click on it and select to **System Configuration -> Virtual Resources -> Manage Virtual Storage**
    - Query the VIOS server that owns the PCI-to-PCI Bridge adapter (this adapter has access to the USB port).
    - Click on the **Virtual Optical media** tab. From the Actions button, select **Add media** from physical **cd0** device. Use short names without the .iso value and specify the path as /dev/cd0/< iso file name here>.

5.3.8 Installing IBM i using an external USB DVD from FSM

- Place the IBM i installation media in the DVD-ROM drive..
- Refer to the "Moving the DVD drive to another virtual server using FSM" section and perform the necessary steps to assign the DVD-ROM to the node for this virtual server.
- In FSM, right-click the IBM i virtual server name and click **Operations -> Activate -> Profile** to activate the virtual server.
- Select a profile you want to use (such as the default profile) and select the **Open a terminal window or console session** check box. Then click **Advanced**.
- On the Advanced tab, select **Manual** from the Keylock position list and **D-mode** from the Boot mode list. Then click **OK**.
- The welcome page will open and show the status of the virtual server including SRC information.

Note that during installation, the load source and other disk units will be initialized before they are included in the system or other ASPs by IBM i. This initialization time will vary depending on the type of storage used.

5.3.9 Installing IBM i using a Fibre Channel tape library

A SAVSYS tape backup can be used for installing a new IBM i virtual server, but there are some PTFs that are needed on the source system before the SAVSYS is done. Use the link in the "Support statements and requirements for FC tape libraries" section and look closely for the correct operating

system release level and then the **alternate IPL PTFs** for that release. Note that those PTFs **must** be **permanently applied** to the source system before the SAVSYS is done. Then look through the steps in the "Performing a D-mode IPL from a FC tape library using FSM" section.

5.3.10 Mirroring in IBM i

With SAN storage, it is not necessary to configure disk protection within the IBM i operating system. Data integrity is provided by the SAN, where each LUN made available to IBM i through VIOS is created from multiple physical drives in a RAID array. However, unlike LUNs physically that are attached to IBM i, LUNs that are virtualized by VIOS will not appear in IBM i as parity-protected. Instead, LUNs from a SAN are recognized in IBM i as unprotected DHxxx devices.

5.3.11 Post-install tasks and considerations

5.3.11.1 Install IBM i PTFs (if necessary)

Refer to the "Fixes concepts and terms" topic in the IBM i Information Center for the steps to install any required PTFs at the following URL:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzam8/rzam8fixconcepts1.htm
Refer to the IBM i *Recommended Fixes* website to get a list of the latest recommended PTFs at:
http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixes.

Starting in IBM® i 6.1, you have the ability to perform complete installations and upgrades of IBM® i partitions, in addition to loading PTFs or restoring data via a virtual optical image on an image server, all without having to manually FTP the data across the network to the individual partitions. This change allows for the install of Licensed Internal Code, operating system, licensed programs and PTFs across IBM® i partitions from a single image source. This capability is available through PTFs beginning in i 6.1. The assumption here is that you have successfully saved a valid IBM i image to a file server that can be accessed from your Power node.

The steps to configure and install across a network connection are described in this document:
ftp://public.dhe.ibm.com/systems/support/power/i/nfs_optical_upgrade.pdf

5.3.11.2 Configure IBM i networking

Refer to the Ethernet topic in the IBM i Information Center to configure IBM i networking for production at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm. Note that while any virtual Ethernet adapter available in IBM i (as a CMNxx device) can be configured, only those on VLANs bridged with a virtual Ethernet bridge by the VIOS server can communicate with the external LAN.

5.3.11.3 Configure Electronic Customer Support (ECS) over LAN

There is no physical modem available to IBM i on POWER node, so ECS over LAN should be configured. Refer to the *Setting up a connection to IBM* topic in the IBM i Information Center at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm

5.3.11.4 How to display the IBM i partition System Reference Code (SRC) history from FSM

- View the node's system reference codes from the **Manage Power Resources.**
- Right click on the Power node and select **Service and Support-> Reference Code history.**

5.3.12 Moving the DVD drive to another virtual server using FSM

Use the following steps:

- From the client IBM i partition, vary off the OPTxx device.

- On the FSM Manger Power Resources page, expand **Hosts**. Select the node and then select the IBM i virtual server that is using the DVD.
- Right click on the name and click **System Configuration -> Manage Virtual Server.**
- Click the **Media Devices** link on the left hand side of the pane.
- The DVD drive is named **cdX.** Select it, then click the **Remove** to remove it from the virtual server.
- Repeat the steps for the target virtual server using the **Add** option, then vary on the OPTxx device there.

5.4 Creating the virtual media library on a Power node

It is not recommended to move the virtual media library from one node to another because this leaves the source node's VIOS's internal database in a bad state.

On the Storwize V7000:

- Open a browser and enter the IP address of the V7000 cluster management interface: use an https for a secure session.
- Sign on to the V7000 as superuser
- Navigate to the Volumes pane select the **New Volume** button.
- Create either a Generic (allocated to full capacity) or a Thin-Provisioned volume. Select the **mdiskgrp#** to get the space. Supply a volume name and the size you feel you will need (the default media repository is 400 GB, but not all of it is allocated) and then click **Create and Map to Host**.
- The volume will be created and you will be prompted for the host to map it to.
- Select the correct host from the drop down menu. Verify the selections and then click **Apply**.

On the new VIOS that now owns the new LUN:

- Open a PuTTy session to the VIOS address using the SSH option.
- Sign on to the command line interface as padmin.
- Issue the 'lsdev | grep hdisk' command to see which hdisks are currently known to this VIOS. Note the hdisk# numbers. Be aware of gaps in the numbers as one of these may be used.
  - VIOS will not automatically see the new LUN without a reboot or without running the cfgdev command.
- Issue the 'cfgdev' command to discover the newly mapped LUN.
- Issue the 'lsdev | grep hdisk' command to locate new drive. Note the drive number (hdisk#).

On the FSM Manage Power System Resources tab:

- Select the Power node where the VIOS media repository is to be located:
- Right click on the Power node -> System Configuration -> Virtual Configuration -> Virtual Storage Management.
- When the window comes up, select the VIOS and click query.
- You need to create a storage pool on the hdisk that represents the new volume. When the window comes up, select the **Storage Pools** tab. Then click the **Create Storage pool** button.
- Give the storage pool a name, ie mediaRep and a Storage pool type of **Volume Group-Based**. Select the **hdisk#** from the list and click **OK**.
- Next click the virtual Optical Media tab, then click the **Create** library button. Select the **storage pool** that you just created and specify a **size** just short of the amount free as shown in next to the storage pool name. Click **OK**.
- When the media library is created, the **Actions** pull down will have an option to **Add media**.
  - You can ftp in **binary mode** .iso files to the VIOS /home/padmin directory.
  - If you have a DVD drive attached and seen by VIOS, you can add media through that.
  - Or you can create blank media to be used as a target to hold data.
  - Its recommended to use shorter names.
- Once you have the media loaded. You can select one of the iso files and from the **Actions** box select **Modify Virtual Server assignment** and choose the client virtual server to assign it to.
- When the window comes up, check the box next to the media, and click OK.
- This will add the virtual media assignment to the IBM i virtual server.

On the IBM i client virtual server that will be using the media repository:

- Vary on the virtual optical device using option 1 from the wrkcfgsts *dev opt* display.

6 Backup and restore

6.1 Overview of backup and restore for IBM i on a POWER node

The save and restore process for IBM i on a POWER node is not similar to that on other Power servers. Clients can only use a Fibre Channel-attached tape library with LTO drives. The tape drive from the Fibre Channel tape library is recognized in the IBM i LPAR and follows the same naming convention as in any other IBM i environment. This allows an administrator to use native IBM i save and restore commands or Backup, Recovery and Media Services (BRMS) as part of the same backup and recovery strategy employed on other Power servers.

The physical adapters connecting to the tape devices are owned by VIOS. The Fibre Channel-attached tape library solution uses NPIV, so that a tape drive from the library is directly mapped to IBM i in the SAN, with VIOS managing the physical FC adapter in a passthrough method.

The following sections describe the technical details, support statements and implementation steps for backup and restore.

6.2 Save and restore with a Fibre Channel-attached tape library - Technical overview

NPIV attached Fibre Channel (FC) tape libraries are supported for save and restore option for IBM i on a POWER node. This capability allows an administrator to perform backup and recovery operations using the same IBM i native or BRMS commands as on any other IBM i LPAR attached to a tape library. Advanced capabilities such as automatic media change and hardware encryption of tape media are supported. Encrypted tape volumes can then be read and used for restore on a non-node Power server connected to a tape library.

VIOS uses N_Port ID Virtualization (NPIV) to allow IBM i direct access to a tape library in the SAN through an NPIV-capable adapter owned by VIOS. NPIV is a Fibre Channel technology that enables a single port on a FC adapter to be presented to the SAN as an N-number of independent ports with different World-wide Port Names (WWPNs). NPIV-capable adapters on Power servers and nodes allow up to 256 virtual FC ports to be assigned to a single physical FC port. On Power servers and nodes, VIOS always owns and manages the FC adapter.

To leverage NPIV, an IBM i LPAR must have a Virtual Fibre Channel (virtual FC) client adapter created, which connects to a virtual FC server adapter in VIOS, similar to the VSCSI model. However, the virtual FC client adapter does not allow IBM i to access LUNs or tape drives already assigned to VIOS in the SAN. Instead, the virtual FC server adapter in VIOS is mapped to a FC port on the physical NPIV-capable adapter. This allows the client virtual FC adapter in IBM i direct access to the physical port on the FC adapter, with VIOS having a “pass-through” role, unlike with VSCSI. The tape drive in the FC tape library for IBM i use is not mapped to the WWPNs of the physical ports on the NPIV adapter and it does not become available in VIOS first. Instead, when the virtual FC client adapter in the IBM i LPAR is created, two virtual WWPNs are generated by the PowerVM Hypervisor. The tape drive in the library is zoned directly to the first of the two WWPNs on the virtual FC client adapter in IBM i. The second WWPN is used to facilitate Live Partition Mobility, although only AIX and Linux LPARs can leverage it at the moment. The PowerVM Hypervisor on a Power server or node has the default capability to create 32,000 virtual WWPNs. When virtual FC client adapters are deleted, WWPNs are not reused. If all of the 32,000 WWPNs are used, the client must obtain an enablement code from IBM, which allows the creation of a new set of 32,000 WWPNs.

Zoning a tape drive in a FC tape library directly to a WWPN on the virtual FC client adapter in IBM i has two important advantages:

- It allows simpler SAN zoning and I/O virtualization configuration on the POWER node. The tape drive in the library does not have to be made available to VIOS first and then assigned to IBM i; instead, it is mapped to IBM i in the SAN
- From the perspective of both the SAN device (the tape drive in the library) and IBM i, an NPIV connection is identical to a direct FC connection, which is otherwise not possible on POWER nodes. NPIV enables IBM i to recognize and use all characteristics of the FC device, instead of using it as a generic tape drive virtualized by VIOS

Figure 5 presents an overview of accessing a FC tape library through NPIV for IBM i on a POWER node:

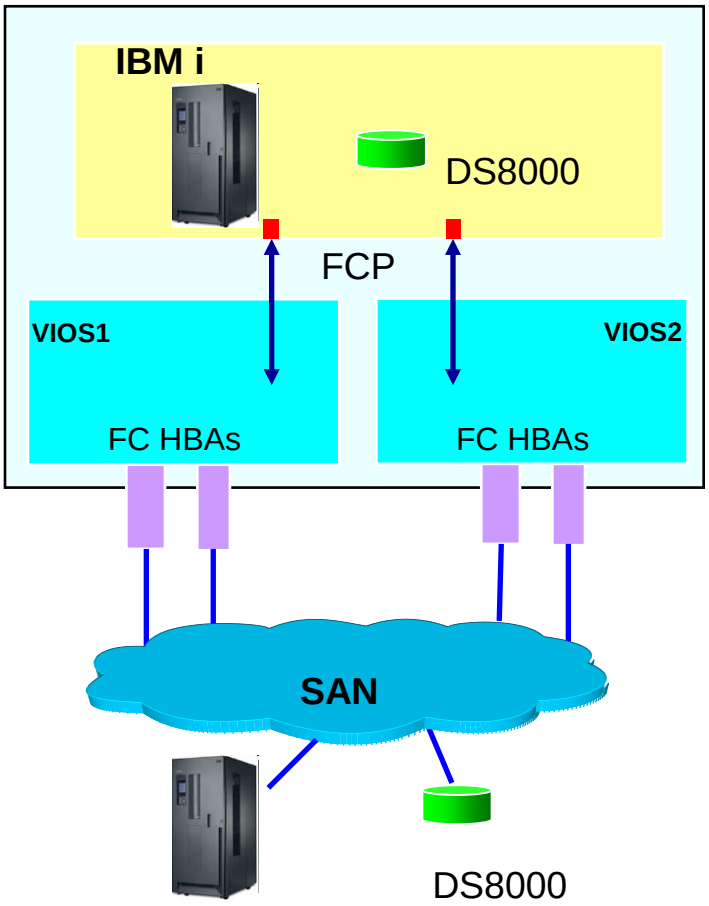

Figure 5: An overview of accessing a FC tape library through NPIV for IBM i on a POWER processor-based node.

6.2.1 Support statements and requirements for FC tape libraries

Refer to the following APAR for the IBM i PTFs required and for D-Mode limitations for NPIV virtualization of Fibre Channel libraries and a current list of supported devices at: http://www-912.ibm.com/n_dir/nas4apar.nsf/c79815e083182fec862564c00079d117/02327e067599839e86257659003c6d66?OpenDocument&Highlight=2,ii14526.
Note: that the NPIV PTFs must be **permanently applied** on the partition **before** a SAVSYS is done in order to allow D Mode IPL from that tape.

Note that while all supported FC adapters and most supported FC switch modules operate at 8Gbps, the rest of the SAN does not have to operate at that speed. Additional FC switches outside of the chassis

can operate at 4Gbps or 2Gbps. While this difference in throughput will have an effect on performance, it will not prevent NPIV from working. Similarly, only the first FC switch connected to the NPIV adapter must be NPIV-enabled. For this solution, that first FC switch is one of the supported switch modules (explained in this section).

IBM i hosted by VIOS supports NPIV attachment of DS8000/DS5100/53000 logical drives. These logical drives must be formatted as IBM i 520 byte sector drives. As with other support statements, this paper reflects the current support and its online master version is updated as IBM's support position changes.

6.2.2 Creating the virtual FC configuration using FSM

Before the tape library can be mapped to IBM i in the SAN, the virtual FC server and client adapters must be created in VIOS and IBM i and the IBM i LPAR must be associated with the physical FC adapter. Perform the following steps define the virtual FC adapter:

- On the **Manage Power Systems Resources** page:
- Expand **Virtual Servers** and right-click the IBM i client virtual server you are configuring.
- Click **System Configuration -> Manage Profile. Note:** Do **not** use the Manage Virtual Server option as this creates a dynamic set of WWPNs. Later when you update the profile you will generate yet another set of WWPNs and they will not match.
- Select the profile to update and select **Actions -> Edit.**
- Click the **Virtual Adapters** tab.
- Click the **Actions -> Create Virtual Adapter -> Fibre Channel Adapter.**
- Select the box for making this adapter a **required adapter.**
- Select the correct **VIOS server partition** that is hosting this client partition.
- Click **OK**.
- This creates the virtual FC adapter and the virtual WWPNs that are now assigned to that client virtual server. If you go back into the IBM i partition, you should see the virtual WWPNs that were generated. You can use these for zoning the tape to the IBM i client partition.
- Perform the steps in the "Making the tape library available to IBM i" section.

6.2.3 Making the tape library available to IBM i

You need to associate the virtual WWPNs generated from the virtual Fibre Channel adapter configuration (from the prior section) with a fibre channel adapter port owned by the VIO server. This allows the VIO server to pass the storage traffic through its physical connection.
From FSM

- On the **Manage Power Systems Resources** page, select the Power node
- Right click on it and select to **System Configuration -> Virtual Resources -> Manage Virtual Storage**
- Query the VIOS server you are using.
- Click on the **Virtual Fibre Channel** tab. A list of fibre channel ports will be shown as VIOS sees them (ie fcs# are the resource names that VIOS assigns to the ports). The port you choose is tied to specific FC switches in the back of the chassis. For instance, fcs0 goes to switch bay 3 and fcs1 goes to switch bay 4. So choose a port to the switch that you have run a fibre cable from your tape device to its switch port(s). If multiple tape devices are to be used, consider using different switches, if available, for the different tape device cables to balance the IO traffic.
- Select the **fcs#** port associated with the correct switch (discussed above) and click **Actions->Modify virtual server connections…**
- Select the checkbox next to the partition that you want to associate the tape device and click OK.

Once this assignment is made, you can use either the IBM i client profile or virtual server interface to log into the virtual WWPN to facilitate the zoning on the FC switch.

- On the **Manage Power Systems Resources** page, select the **virtual servers** link
- Right click on the IBM i virtual server and select to **System Configuration ->Virtual server or Manage Profile.**

- Click the **Virtual Adapters** tab and then click on the **virtual Fibre Channel** adapter to see its properties.
- Select the **Fetch** button (if present, based on the interface you are using) and then the **Login** button,.
- This will bring the FC port on line and it should be seen on the FC switch for zoning as described below.

Note: You can also use the change port login (**chnportlogin**) command from the FSM CLI to bring the virtual FC ports online for the SAN fabric to detect. This allows the SAN zoning to be done easier. Refer to the HMC commands at:
http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha1/usingthehmcremotecommandline.htm

After configuring the virutal FC configurations is complete, making a tape drive within the FC tape library available to IBM i requires the correct zoning in the FC switch module in the chassis and any other switches to which the tape library might be connected. This process is identical to connecting a FC tape library directly over the SAN fabric to IBM i on a Power server without NPIV.

Use both of the virtual WWPNs from the IBM i client virtual server profile and refer to the Redbooks *Implementing Tape in i5/OS* at: http://www.redbooks.ibm.com/abstracts/sg247440.html?Open to perform the correct zoning.

There is a web interface on the tape media library where you need to enable control paths from each device that you want IBM i to be able to work with. Selecting **Enable** creates the control paths.

IBM i cannot dynamically detect these control paths. To detect the control paths, you need to re-IPL (reboot) the virtual I/O adapter (IOA).

First determine which virtual IOA has been created for the virtual fibre channel adapters. To do this, enter a WRKHDWRSC *STG command and check for a 6B25 (virtual FC) adapter. Note the IOP/IOA name.

Next, use STRSST and start a service function (1) -> hardware service manager (7) -> Logical Hardware Resources (2) – System bus resources (1). Enter a 255 in the System bus(es) to work with field and hit enter. This is the virtual adapter bus. Locate the virtual IOA from above and enter a 6 for I/O debug, then option 4 to IPL the IOA. Use F3 to exit SST.

Return to the WRKHDWRSC *STG and use an option 9 to see the tape devices under the virtual FC IOA. With auto configuration turned on, a new tape device(s) should show up under WRKCFGSTS *DEV TAP*.

6.2.4 Performing an IBM i save

In IBM i, the FC tape library is recognized with the correct type and model, such as 3584-032 for TS3500, and is assigned a standard TAPMLBxx device name. For instance, each LTO4 tape drive within the library mapped to the IBM i LPAR is recognized as a 3580-004 device with a TAPxx name.
The tape library connected through NPIV supports all standard and advanced library functions, as well as both native save restore and BRMS commands. Use the standard backup procedure or refer to the IBM i Information Center's **Backup and recovery** topic at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm
IBM i can also use a DVD writer, if it available. See the "Writing to DVD-RAM media" section.

6.2.5 Performing an IBM i restore

IBM i can read from any supported tape media depending upon the model tape device used in the FC tape library and can perform both full-system and partial restores. You can use the standard restore procedure or refer to the IBM i Information Center's **Backup and recovery** topic at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm.

6.2.6 Performing a D-mode IPL from a FC tape library using FSM

It is possible to perform a full-system save in an IBM i partition on a different system and then use that save for a D-mode IPL and full-system restore on a node with access to a FC tape library. To perform a D-mode IPL from a tape library, use the following steps:
Note: Make sure that you have the requisite PTFs **permanently** applied on the system **prior** to the save system.

- Verify that the tape drive is powered on and the tape is loaded
- Verify that the correct tape drive in the library is zoned to one of the WWPNs on the virtual FC client adapter in the IBM i LPAR
- Verify that the correct virtual FC client adapter is selected as an alternate IPL resource for the IBM i partition:
- On the **Manage Power Systems Resources** page:
- Expand **Virtual Servers** and right-click the IBM i client virtual server you are configuring.
- Click **System Configuration -> Manage Profile.**
- On the **General partition properties** tab, verify that the virtual FC client adapter with the correct WWPNs is selected from the **Alternate restart adapter** list.
- On the same tab, verify that the IBM i partition is configured for a D-mode manual IPL.
- Close the partition properties window.
- Select the IBM i partition and click **Operations -> Activate -> Profile.**

6.3 Backup and restore of VIOS

For a "bare metal" type of restore of the entire POWER node, backups of VIOS and the LPAR configuration are also required. Saving VIOS, the LPAR configuration and IBM i on the node can be performed in any order. During a restore, VIOS and the LPAR configuration are restored first, followed by IBM i as described earlier in this section. VIOS can be backed up to DVDs or to a Network Install Manager (NIM server). For more information on NIM servers and backing up to them, see the Redbook: **NIM from A to Z in AIX 5L** at http://www.redbooks.ibm.com/abstracts/sg247296.html

6.3.1 Backing up VIOS to DVDs

The VIOS can be backed up to DVDs using the following steps.
The client (hosted) partitions are do not have to be inactive.
Ensure a blank tape is loaded in the tape device you is using for the backup.

- Telnet or PuTTY into VIOS. Log in using the **padmin** user ID.
- Backup the partition profile data by typing the following command:
  - **viosbr –backup <–file <directory and file path here>**
    If a location is not specified with the -file option, the file is placed in the default location /home/padmin/cfgbackups.

Alternatively, the partition profile can be backed up through FSM using the following steps:

- Log into the FSM.
- On the **Manage Power Systems Resources** page, click on POWER node you want to work with.
- Right click and select **System Configuration -> Manage Virtual Server Data -> Backup**.
- Provide a backup file name and click **OK**.
- As soon as the operation is complete, a confirmation window is shone.
- Find the DVD device name by typing the following command from the VIOS CLI: **lsdev –type optical**
  If the device has a status of *Defined*, type the following command, with ***name*** as the DVD device name: **cfgdev –dev *name***
- Check the current physical block size for the DVD device using the following command, with ***name*** as the tape device name: **lsdev –dev *name* –attr**
- Perform the VIOS backup by using the following command:
  - **backupios -cd /dev/cd0 -cdformat**

The resulting DVDs are a bootable VIOS backup. If you need to create multi-volume discs because the image does not fit on one disc, the backupios command gives instructions for disk replacement and removal until all the volumes have been created.

Alternatively, the partition profile can be backed up through the FSM interface using the following steps:

- On the FSM welcome page, expand **Hosts**. Select the node server and then click **System Configuration -> Manage Virtual Server Data -> Backup.**
- Enter a file name for the backup when prompted.
- Click **Ok.**

6.3.2 Restoring VIOS from a DVD backup

The steps to restore the VIOS from a DVD backup are described in this section. Ensure that the backup DVD media is loaded in the DVD device before proceeding.

Perform the following steps to restore VIOS.

- Log into the FSM.
- On the **Manage Power Systems Resources** page, click on POWER node you want to work with.
- Open a console for VIOS on the FSM using the steps in section “Opening a VIOS console using FSM”.
- Power on or restart the POWER node: Right click on the node and select **Operations -> Power On.**
- On the console, when the Partition Firmware displays its initial screen (a series of IBM’s scroll across the console), press **1** to enter the SMS menu. If you miss the prompt to press 1, the partition will attempt to boot from the default device, which might boot back into VIOS, if there is still a bootable image installed on the default device (probably the internal hard drive). Otherwise, the boot will fail and the SMS menu will eventually be displayed.

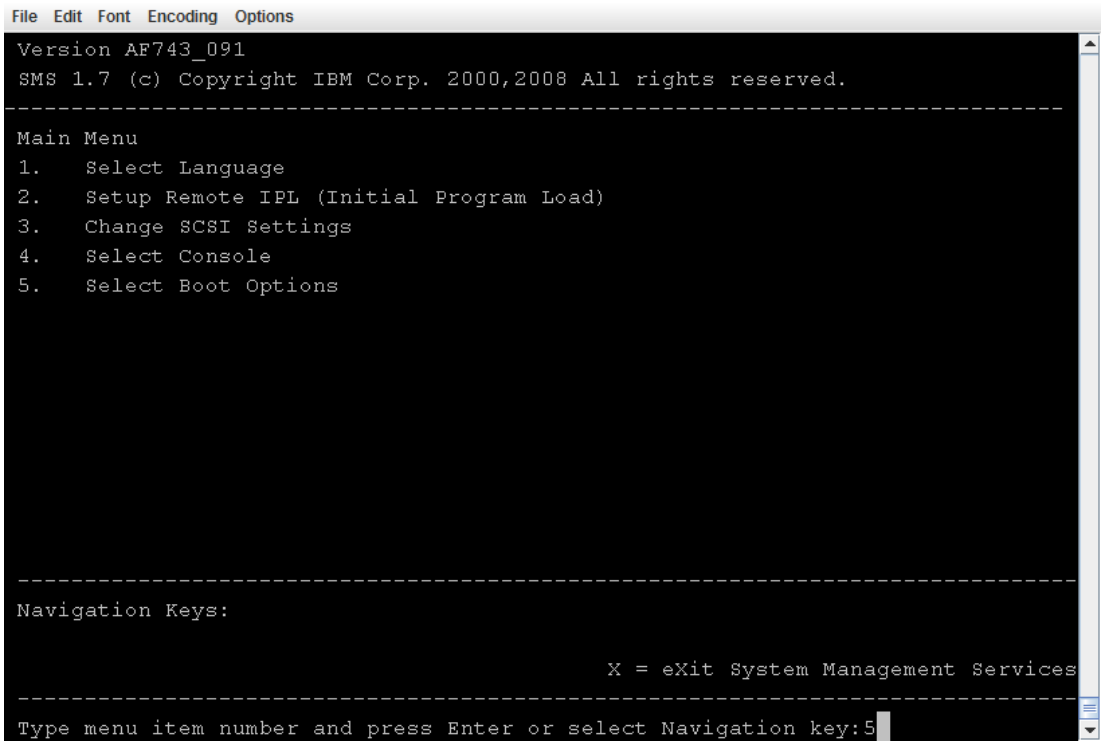

- On the SMS **Main Menu**, type **5** to select the **Select Boot Options** and press **Enter**.

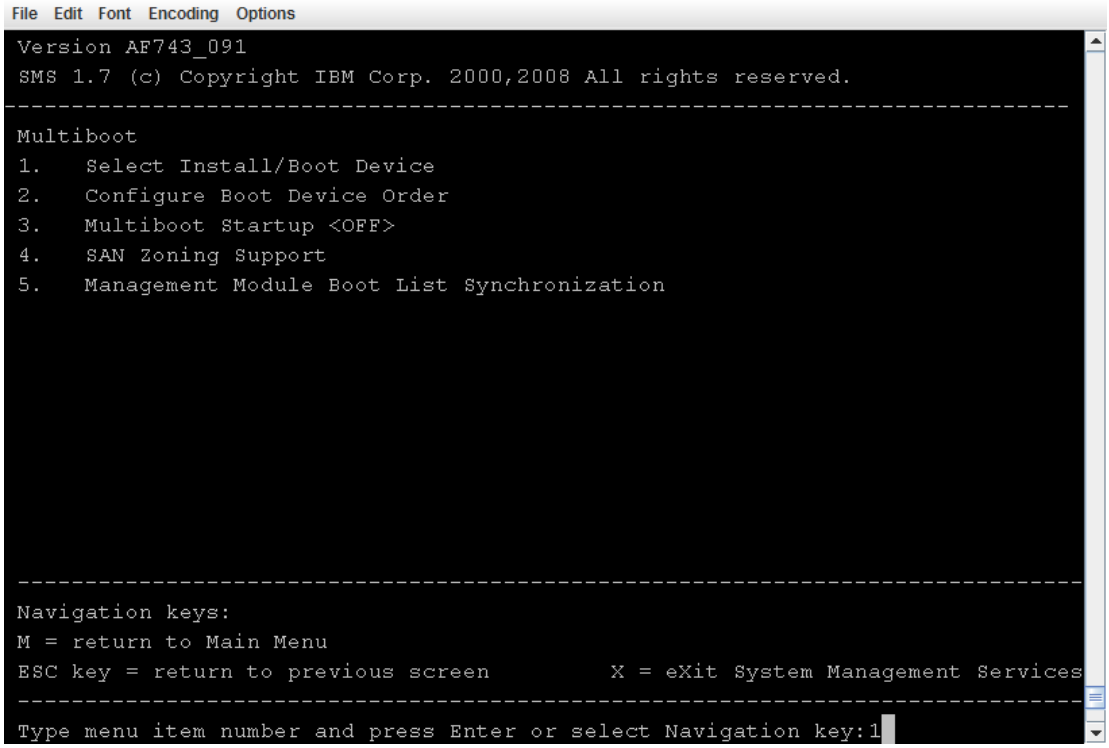

- On the **Multiboot** menu, type **1** to select the **Select Install/Boot Device** and press **Enter**.

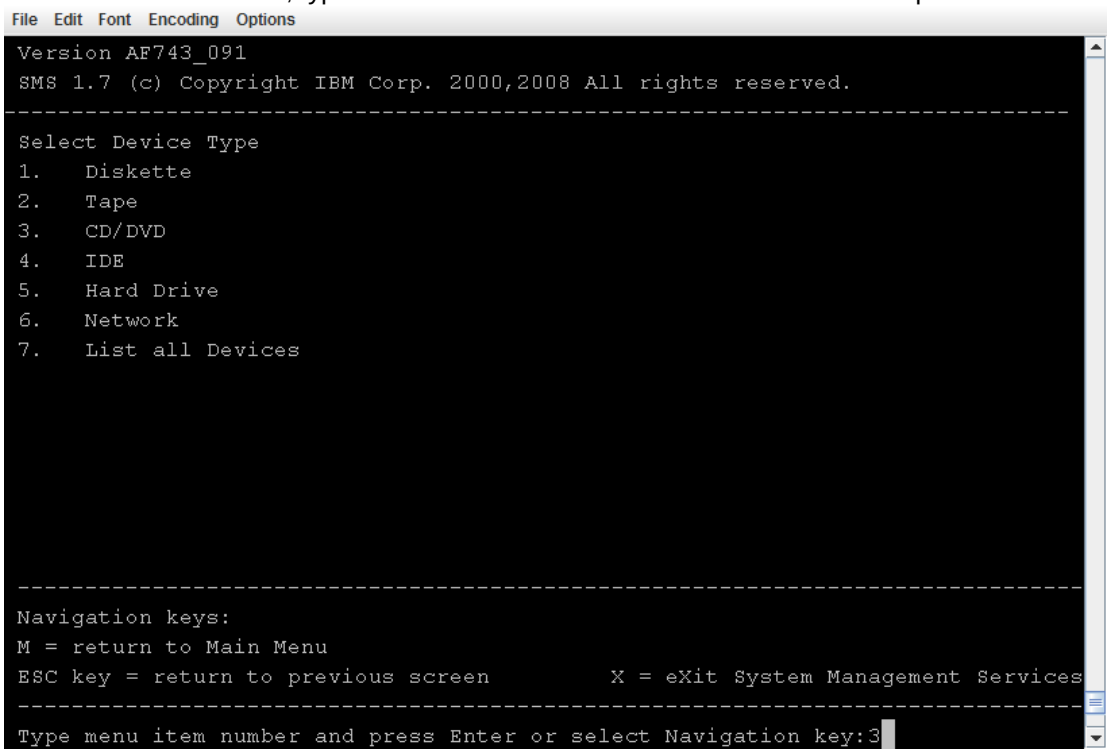

- On the **Select Device Type** menu, type **3** to select **CD/DVD** and press **Enter**.

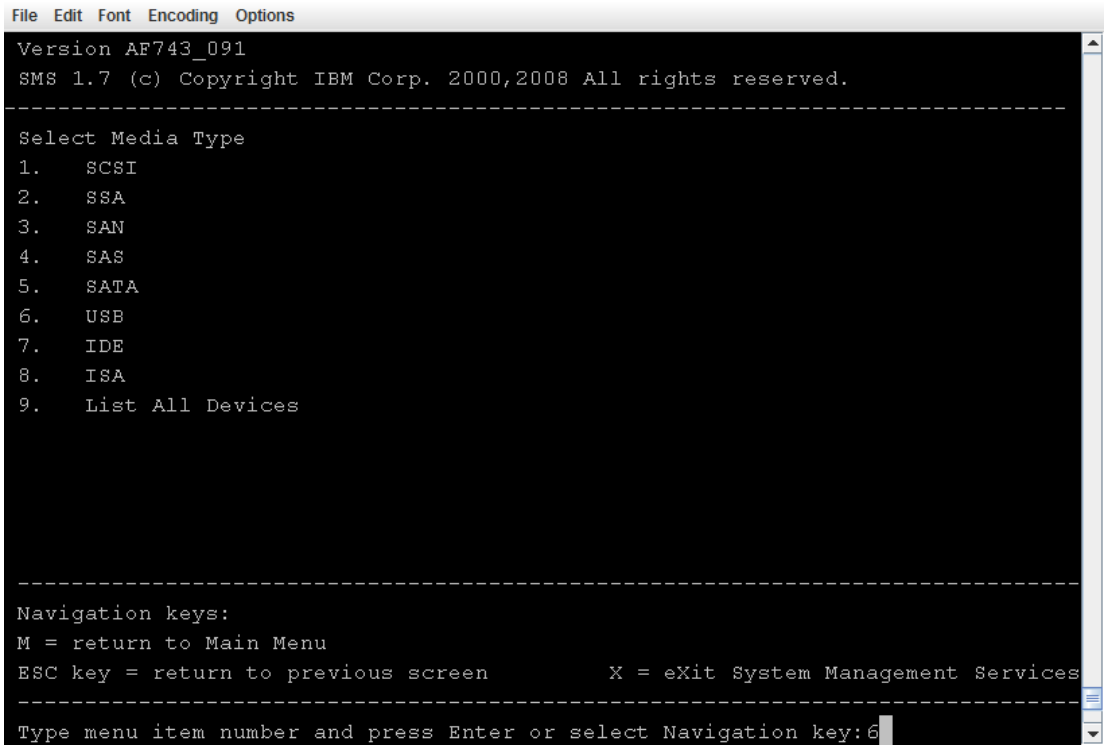

- On the **Select Media Type** menu, type **6** to select **USB** and press **Enter**.

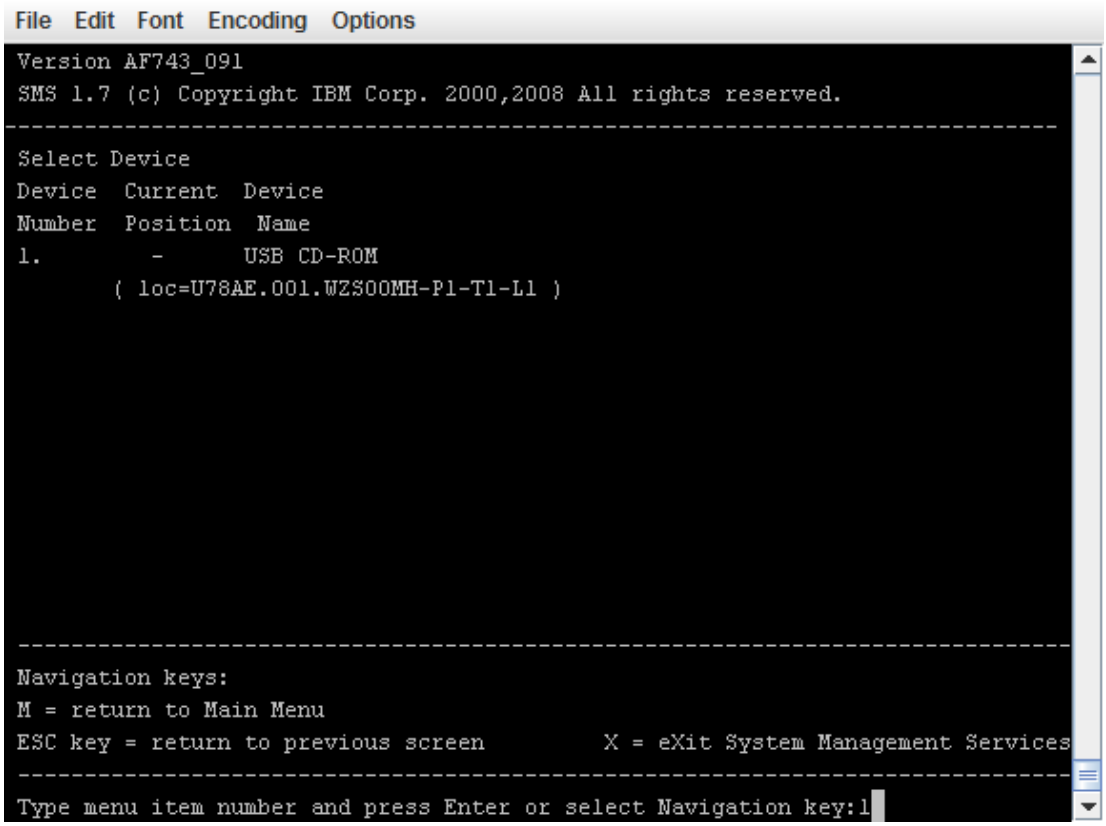

- On the **Select Device** menu, type the device number for the tape device with the backup media loaded and press **Enter**.

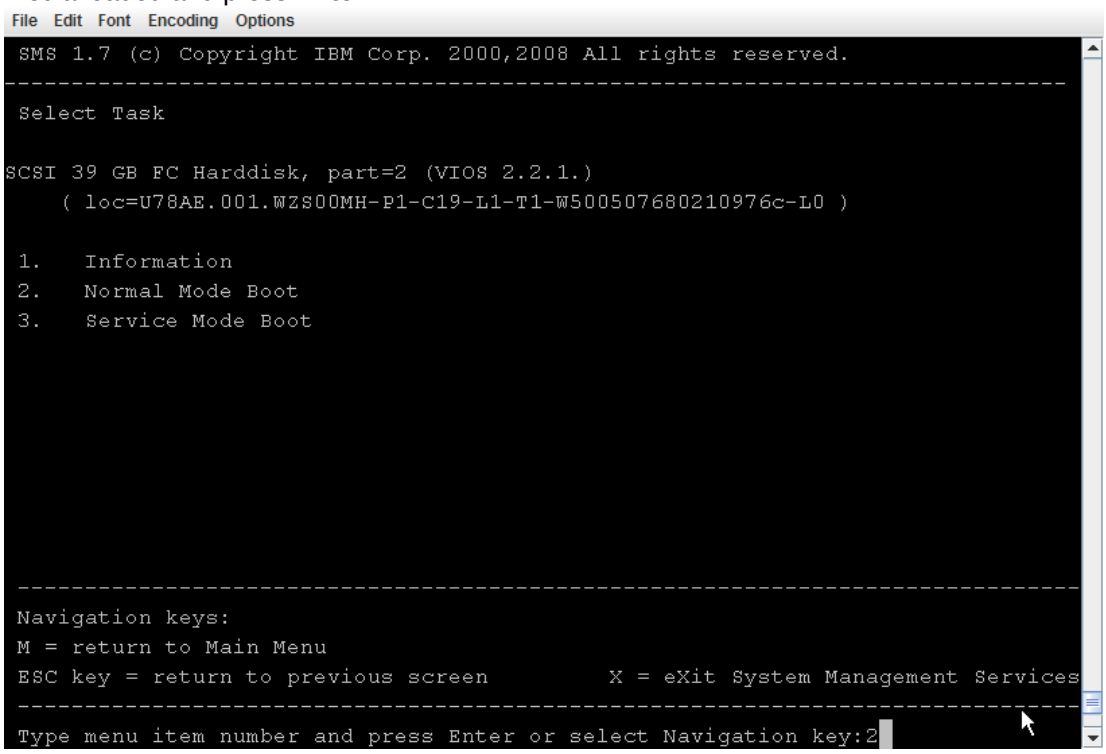

- On the **Select Task** menu, type **2** for **Normal Boot Mode** and press **Enter**.

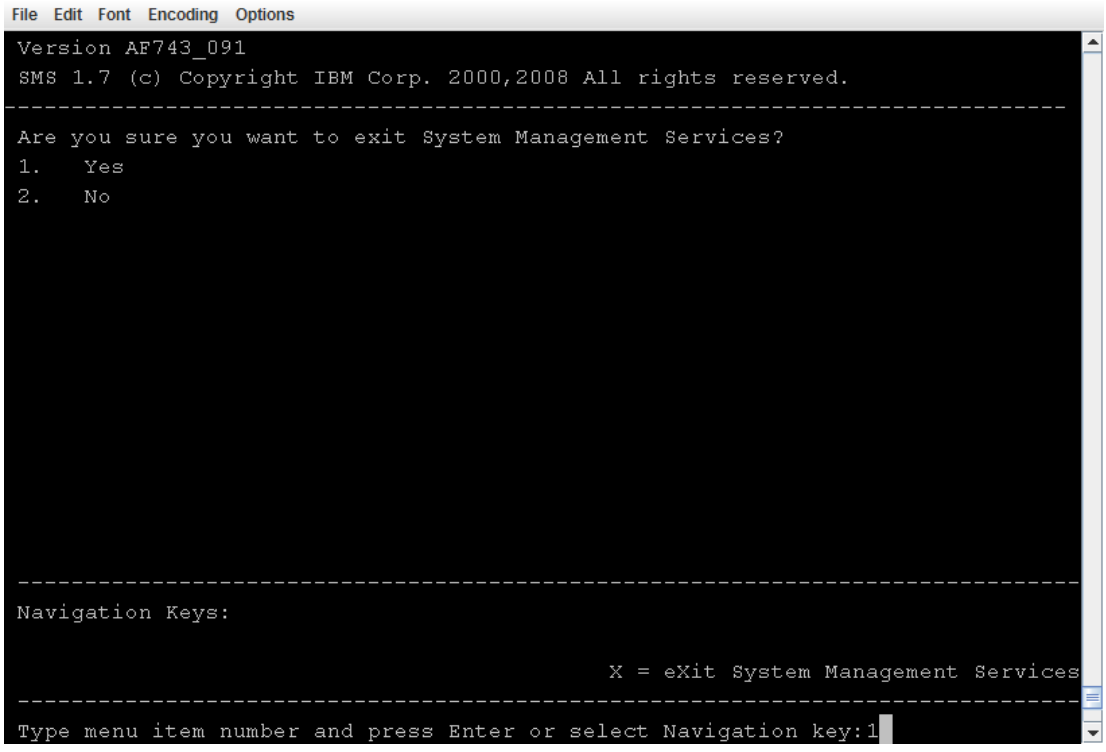

- Type **1** and press **Enter** to confirm that you want to exit the SMS menus. It will take several minutes to boot from the tape.

```
******* Please define the System Console. *******

Type a 2 and press Enter to use this terminal as the
  system console.
Pour definir ce terminal comme console systeme, appuyez
  sur 2 puis sur Entree.
Taste 2 und anschliessend die Eingabetaste druecken, um
  diese Datenstation als Systemkonsole zu verwenden.
Premere il tasto 2 ed Invio per usare questo terminal
  come console.
Escriba 2 y pulse Intro para utilizar esta terminal como
  consola del sistema.
Escriviu 1 2 i premeu Intro per utilitzar aquest
  terminal com a consola del sistema.
Digite um 2 e pressione Enter para utilizar este terminal
  como console do sistema.
```

- As soon as you are past the **Welcome to the Virtual I/O Server** screen, a screen with the **Please define the System Console** message is displayed . Type **2** and press **Enter** as the screen directs. Note that these keystrokes will not appear on the screen.

```
                              Error Warning

The data file indicated that all logical volumes should
be created exactly as they were before, but the disks
are not the same as they were on the source system.


To reboot the system, press reset.

>>> 1   Continue with Install









    88  Help ?

>>> Choice [1]: 1
```

- You may see a warning screen indicating the disks on the system do not match the disks on the source (of the backup) system. Type **1** and press **Enter** to continue with the installation.

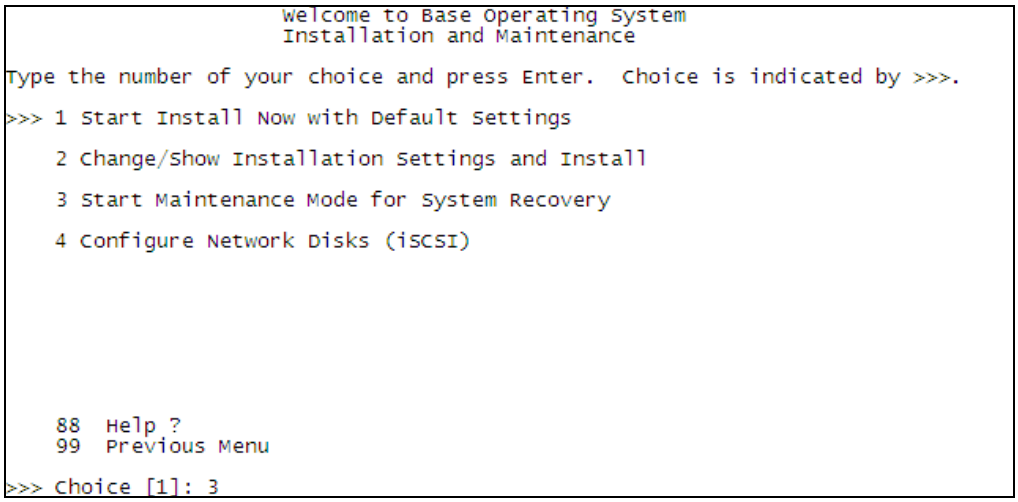

- On the **Welcome to Base Operating System – Installation and Maintenance** menu, type **3** to select the **Start Maintenance Mode for System Recovery** option and press **Enter**.

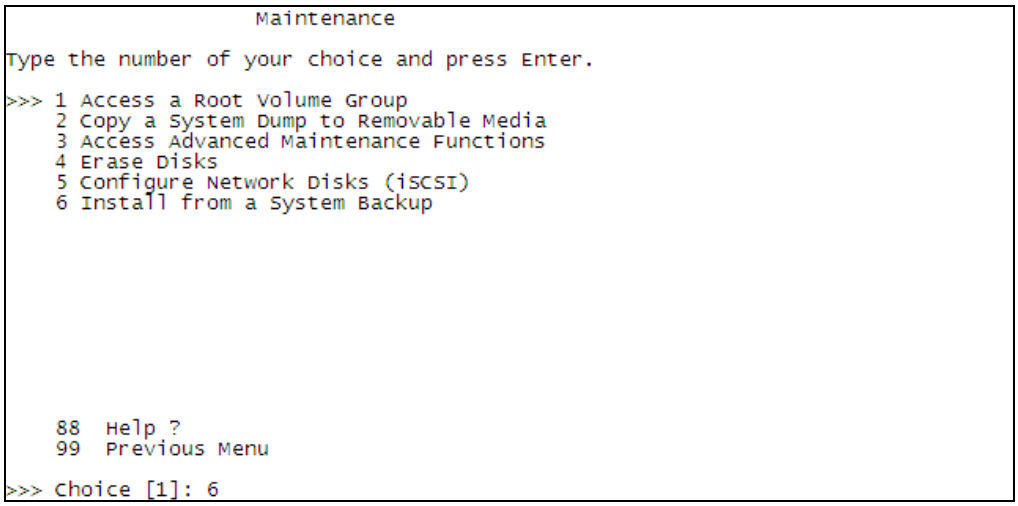

- On the **Maintenance** menu, type **6** to select the **Install from a System Backup** and press **Enter**.

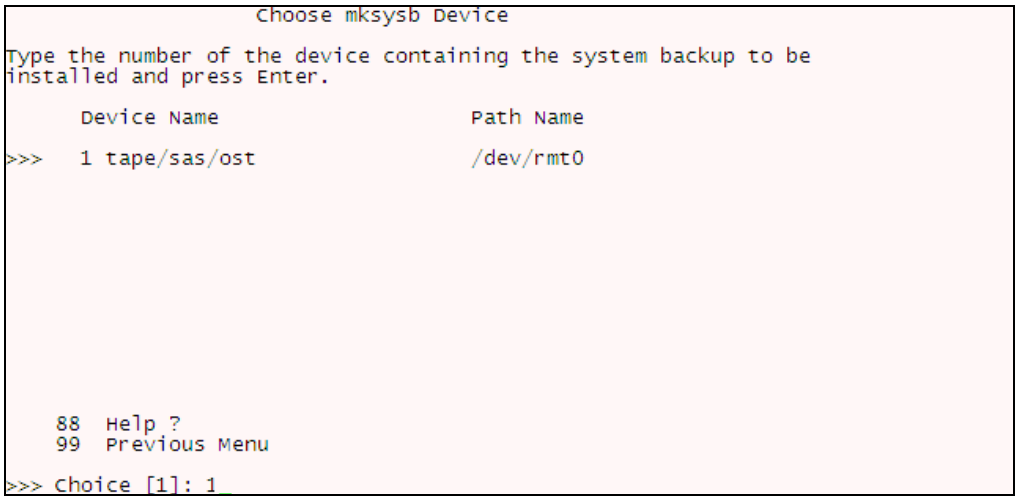

- On the **Choose mksysb Device** menu, type the number of the device with the backup tape mounted and press **Enter**. The restore will now begin. A progress screen is displayed until the restore is complete.

As soon as the restore is complete, the VIOS will reboot from the just completed restore image.

The partition (virtual server) data can be restored through the FSM (if saved through the FSM) using the following steps:

- Log into the FSM.
- On the **Manage Power Systems Resources** page, click on POWER node you want to work with.
- Right click and select **System Configuration -> Manage Virtual Server Data -> Restore**.
- Select the backup file you want to restore.
- Select the backup you want to perform from the following types:
  - **Full Restore** – Restore from the backup file and don't merge with the existing FSM virtual server data.
  - **Backup Priority** – Merge the backup data with the current FSM virtual server data but prioritize the backup file contents if there are conflicts between it and current FSM data.
  - **Host Priority** – Merge the backup data with the current FSM virtual server data but prioritize the FSM contents if there are conflicts between the current FSM data and the backup data.
  - Click **OK.**

7 Additional resources

7.1 PureFlex System documentation

We now have a new portal on our web site for the Product Guides, book, and papers:
http://www.redbooks.ibm.com/portals/puresystems

Redbooks and Redpapers - online in draft format

- IBM PureFlex System Products & Technology, http://www.redbooks.ibm.com/abstracts/sg247984.html?Open
- IBM Flex System p260 and p460 Planning and Implementation Guide, http://www.redbooks.ibm.com/abstracts/sg247989.html?Open
- IBM Flex System Networking in an Enterprise Data Center, http://www.redbooks.ibm.com/abstracts/redp4834.html?Open

xREF:

- xREF: IBM BladeCenter and System x Reference, http://www.redbooks.ibm.com/xref

Product Guides:

All available individually or from our Product Guide landing page:
http://www.redbooks.ibm.com/portals/puresystems?Open&page=pgbycat

Chassis and compute nodes

- IBM Flex System Enterprise Chassis
- IBM Flex System Manager
- IBM Flex System p260 and p460 Compute Node

Product Guides - Switches and pass-thru modules

- IBM Flex System EN2092 1Gb Ethernet Scalable Switch
- IBM Flex System EN4091 10Gb Ethernet Pass-thru Module
- IBM Flex System Fabric EN4093 10Gb Scalable Switch
- IBM Flex System FC3171 8Gb SAN Switch and Passthru
- IBM Flex System FC5022 16Gb SAN Scalable Switch and FC5022 24-port 16Gb ESB SAN Scalable Switch

Product Guides - Adapter cards

- IBM Flex System CN4054 10Gb Virtual Fabric Adapter and EN4054 4-port 10Gb Ethernet Adapter
- IBM Flex System EN2024 4-port 1Gb Ethernet Adapter
- IBM Flex System EN4132 2-port 10Gb Ethernet Adapter
- IBM Flex System FC3052 2-port 8Gb FC Adapter
- IBM Flex System FC3172 2-port 8Gb FC Adapter
- IBM Flex System FC5022 2-port 16Gb FC Adapter
- ServeRAID M5115 SAS/SATA Controller for IBM Flex System

7.2 Storage

- Implementing the IBM Storwize V7000 in IBM Redbooks available at: http://www.redbooks.ibm.com/abstracts/sg247938.html?Open
- *Performance Capabilities Reference* manual (chapter 6 for best practices for IBM i and VIOS or chapter 18 for POWER blades): http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html
- SDDPCM driver: http://www.ibm.com/support/docview.wss?uid=ssg1S4000201

7.3 VIOS

- *PowerVM Editions Operations Guide*: http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/iphdx/abstract_pvme_guide.htm.
- *IBM Power Virtualization: Introduction and Configuration* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247940.html
- VIOS command reference: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1_vios_commandslist.htm
- VIOS Datasheet: http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html

7.4 IBM i

- IBM i installation: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahc/installprocess.htm
- IBM i PTF installation: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahc/installptfsforlicandos.htm
- Latest recommended IBM i PTFs: http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixesIBM Fix Central: http://www-912.ibm.com/eserver/support/fixes/
- IBM Systems Workload Estimator: http://www-304.ibm.com/systems/support/tools/estimator/index.html
- IBM i networking: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahg/rzahgicnet2.htm
- ECS over LAN: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzaji/rzajiecs.htm
- IBM i virtual client partitions topics in the IBM i Information Center: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahc/rzahcblade5osvirtualpartitionsoverview.htm
- Backup and recovery Information Center topic: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm
- *Implementing Tape in i5/OS* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247440.html?Open.
- APAR for PTFs required for NPIV virtualization of Fibre Channel tape libraries: http://www-912.ibm.com/n_dir/nas4apar.nsf/c79815e083182fec862564c00079d117/02327e067599839e86257659003c6d66?OpenDocument&Highlight=2,ii14526.

8 Trademarks and disclaimers

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.
IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.
Other company, product, or service names may be trademarks or service marks of others.
Information is provided "AS IS" without warranty of any kind.
All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.
Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.
All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.
Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.
Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogrCMMing in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.
Photographs shown are of engineering prototypes. Changes may be incorporated in production models.
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.