IBM

IBM PowerVM Getting Started Guide

- Step-by-step virtualization configuration to the first partition
- Single and dual VIOS setups using three common management interfaces
- Advanced configuration of a dual Virtual I/O Server setup

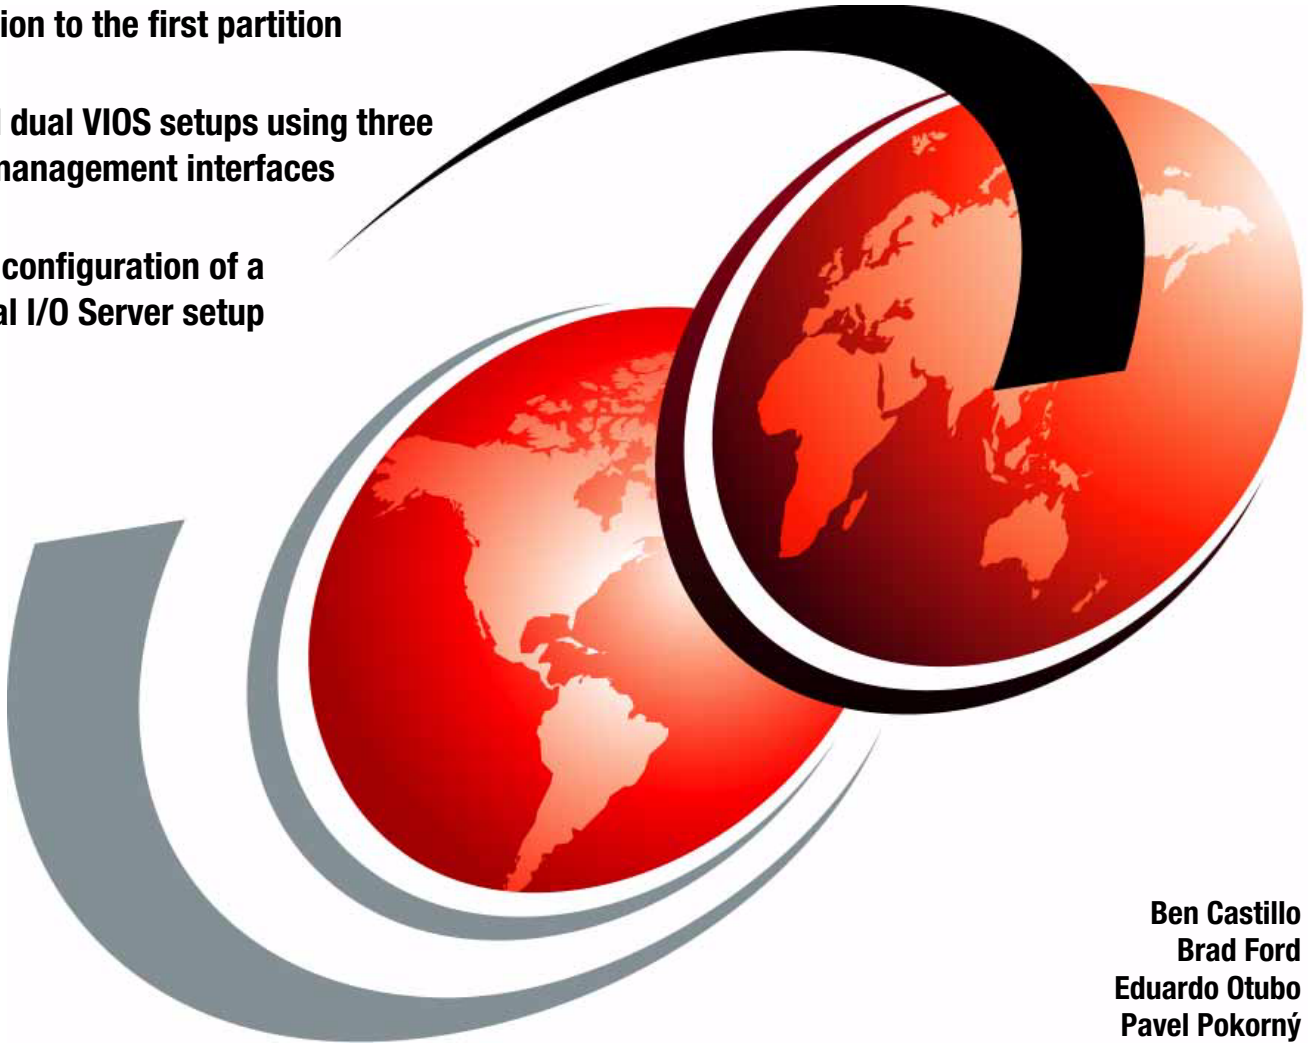

Ben Castillo
Brad Ford
Eduardo Otubo
Pavel Pokorný

Redpaper

ibm.com/redbooks

International Technical Support Organization

IBM PowerVM Getting Started Guide

February 2012

REDP-4815-00

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

First Edition (February 2012)

This edition applies to IBM Virtual I/O Server, versions 2.2.0 and 2.2.1; IBM Systems Director Management Console, version 6.7.4.0; and IBM Hardware Management Console.

Contents

Notices v
Trademarks vi

Preface vii
The team who wrote this paper vii
Now you can become a published author, too! viii
Comments welcome ix
Stay connected to IBM Redbooks ix

Chapter 1. Introduction to PowerVM 1
1.1 Overview 2
1.2 Planning 3
1.3 Terminology differences 6
1.4 Prerequisites 6

Chapter 2. Configuring PowerVM with Integrated Virtualization Manager 7
2.1 Setting up a single VIOS using IVM 8
2.1.1 Installing a VIOS 8
2.1.2 Creating a partition for the client operating system 11
2.1.3 Configuring a VIOS for a client network 12
2.1.4 Configuring a VIOS for client storage 12
2.1.5 Installing a client operating system 14
2.2 Setting up a dual VIOS with IVM 15
2.3 Setting up an N_Port ID Virtualization Fibre Channel with IVM 15

Chapter 3. Configuring PowerVM with the Hardware Management Console 17
3.1 Setting up a single VIOS using an HMC 18
3.1.1 Creating a VIOS partition profile 18
3.1.2 Installing a VIOS 23
3.1.3 Configuring a VIOS partition 26
3.1.4 Creating a logical partition profile for a client operating system 30
3.2 Setting up a dual VIOS using HMC 32
3.2.1 Creating dual VIOS partition profiles 32
3.2.2 Installing a VIOS 34
3.2.3 Configuring VIOS partitions for a dual setup 34
3.2.4 Creating a logical partition profile for the client operating system 39
3.3 Setting up a virtual Fibre Channel using HMC 40
3.4 Adding additional client partitions 44

Chapter 4. Configuring PowerVM with the IBM Systems Director Management Console 45
4.1 Implementing a dual VIOS setup using the SDMC 46
4.1.1 Creating the virtual servers for VIOS1 and VIOS2 46
4.1.2 Installing VIOS1 and VIOS2 49
4.1.3 Configuring the TCP/IP stack in VIOS1 and VIOS2 49
4.1.4 Creating the SEA failover configuration by using the SDMC 50
4.1.5 Configuring storage devices 51
4.1.6 Creating a virtual server for a client operating system 53
4.1.7 Installing a client operating system 54

4.1.8 Configuring virtual Fibre Channel adapters using the SDMC 54
4.2 Setting up a single VIOS using the SDMC 56
4.2.1 Creating a VIOS virtual server 57
4.2.2 Installing a VIOS 59
4.2.3 Configuring a VIOS 60
4.2.4 Creating a virtual server for a client operating system 64
4.2.5 Installing a client operating system 65
4.3 Setting up a dual VIOS using the SDMC 65
4.3.1 Creating a second VIOS virtual server 65
4.3.2 Installing a second VIOS using NIM 67
4.3.3 Configuring a second VIOS 69
4.4 Setting up a virtual Fibre Channel using the SDMC 74
4.4.1 Configuring a client virtual server for NPIV 74
4.4.2 Configuring a VIOS for NPIV 75
4.4.3 Configuring a second VIOS for NPIV 77

Chapter 5. Advanced configuration 79
5.1 Adapter ID numbering scheme 80
5.2 Partition numbering 81
5.3 VIOS partition and system redundancy 81
5.4 Advanced VIOS network setup 82
5.4.1 Using IEEE 802.3ad Link Aggregation 82
5.4.2 Enabling IEEE 802.1Q VLAN tagging 83
5.4.3 Multiple SEA configuration on VIOS 84
5.4.4 General network considerations 84
5.5 Advanced storage connectivity 85
5.6 Shared processor pools 87
5.7 Live Partition Mobility 87
5.8 Active Memory Sharing 87
5.9 Active Memory Deduplication 87
5.10 Shared storage pools 88

Related publications 89
IBM Redbooks 89
Online resources 89
Help from IBM 89

Notices

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US 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 http://www.ibm.com/legal/copytrade.shtml

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

Active Memory™
AIX®
BladeCenter®
GPFS™
IBM®
POWER Hypervisor™
Power Systems™
POWER6®
POWER7®
PowerHA®
PowerVM®
POWER®
Redbooks®
Redpaper™
Redbooks (logo) ®
System i®
System p5®
System Storage®
Tivoli®

The following terms are trademarks of other companies:

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

Preface

IBM® PowerVM® virtualization technology is a combination of hardware and software that supports and manages virtual environments on IBM POWER5, POWER5+, POWER6®, and POWER7® processor-based systems. These systems are available on IBM Power Systems™ and IBM BladeCenter® servers as optional editions, and are supported by the IBM AIX®, IBM i, and Linux operating systems. With this set of comprehensive systems technologies and services, you can aggregate and manage resources with a consolidated, logical view.

By deploying PowerVM virtualization and IBM Power Systems, you can take advantage of the following benefits:

- Lower energy costs through server consolidation
- Reduced cost of your existing infrastructure
- Better management of the growth, complexity, and risk of your infrastructure

This IBM Redpaper™ publication is a quick start guide to help you install and configure a complete PowerVM virtualization solution on IBM Power Systems. It highlights how to use the following management console interfaces to configure PowerVM:

- Integrated Virtualization Manager (IVM)
- Hardware Management Console (HMC)
- Systems Director Management Console (SDMC)

This paper also highlights advanced configuration of a dual Virtual I/O Server setup.

This paper targets new customers who need assistance with quickly and easily installing, configuring, and starting a new PowerVM server in a virtualized environment.

The team who wrote this paper

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

Ben Castillo is a Technical Specialist working for the Emerging Technology Delivery team, part of IBM Integrated Technology Delivery (ITD) in Australia. He has been working for IBM since 2000. Ben's expertise consists of PowerVM, Power Systems, AIX, IBM PowerHA®, and IBM General Parallel File System (IBM GPFS™). He is also working on AIX global server builds for IBM ITD. Ben is an IBM Certified Advanced Technical Expert - IBM System p5®.

Brad Ford is an IBM Certified Technical Specialist and President of i400 Technology Inc., an IBM Advanced Business Partner in Salt Lake City, Utah, specializing in IBM Power Systems and IBM i. He has worked in the computer industry since 1973, specializing in hardware and operating systems. Brad writes the certification exams for IBM i and Power Systems.

Eduardo Otubo is a Software Engineer in Brazil. He has 10 years of experience with Linux and Open Source. He has worked at IBM STG Linux Technology Center in Hortolandia, Brazil, for four years. Eduardo holds a degree in computer science from Universidade Estadual Paulista "Julio de Mesquita Filho." His areas of expertise include Linux development for virtualization on Power.

Pavel Pokorný is an IT Consultant at GC System a.s., an IBM Business Partner in the Czech Republic. He has seven years of experience in the Information Technology field. His areas of expertise include PowerVM, Power Systems, AIX, Data Protection, IBM System Storage®, and Storage Area Network. Pavel holds a master degree in Information Technology from the University of Defence in Brno, Czech Republic.

The project that produced this publication was managed by Scott Vetter, Certified Executive Project Manager, for the ITSO in Austin, Texas.

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

David Bennin
Richard M. Conway
Brian King
Ann Lund
Linda Robinson
Alfred Schwab
Don S. Spangler
IBM US

Nicolas Guerin
IBM France

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Introduction to PowerVM

Businesses are turning to IBM PowerVM virtualization to consolidate multiple workloads onto fewer systems, to increase server utilization, and to reduce cost. PowerVM provides a secure and scalable virtualization environment for AIX, IBM i, and Linux applications that is built on the advanced reliability, availability, and serviceability features and the leading performance of the IBM Power platform.

This chapter provides an overview of the key PowerVM concepts, a planning model, and preferred practices to follow. It includes the following sections:

- Overview
- Planning
- Terminology differences
- Prerequisites

1.1 Overview

With a PowerVM system, you can immediately install and configure virtual machines and have a fully functional logical partition (LPAR). This paper highlights how to use the following management console interfaces to install and configure PowerVM step by step:

- Integrated Virtualization Manager (IVM) in Chapter 2, “Configuring PowerVM with Integrated Virtualization Manager” on page 7)
- Hardware Management Console (HMC) in Chapter 3, “Configuring PowerVM with the Hardware Management Console” on page 17
- IBM Systems Director Management Console (SDMC) in Chapter 4, “Configuring PowerVM with the IBM Systems Director Management Console” on page 45

These chapters use a cookbook style approach that includes similar steps to accomplish the same task.

Logical partition: The term *logical partition* is used as a generic term in this document. Other terms that are used include guest partition, partitions, and virtual servers. All of these terms refer to virtualized guest servers that run their own operating systems (OS).

IBM BladeCenter: This paper does not address IBM BladeCenter.

All three management console interfaces manage virtualization on IBM Power Systems. Table 1-1 shows how these interfaces differ in managing virtualization.

Table 1-1 Feature comparison of management console interfaces to manage virtualization

	IVM	HMC	SDMC
Included in PowerVM	✓		
Manage Power Blades	✓		✓
Manage more than one server		✓	✓
Hardware monitoring	✓	✓	✓
Service agent call home	✓	✓	✓
Graphical interface	✓	✓	✓
Requires a separate server to run on		✓	✓
Run on virtualized environments			✓
Advanced PowerVM features		✓	✓
High-end servers		✓	✓ (hardware appliance only)
Low-end and midrange servers	✓	✓	✓
Server families support	POWER5/POWER5+: ✓ POWER6/POWER6+: ✓ POWER7: ✓	POWER5/POWER5+: ✓ POWER6/POWER6+: ✓ POWER7: ✓	POWER5/POWER5+: POWER6/POWER6+: ✓ POWER7: ✓
Redundant setup		✓	✓

Chapters 2 - 4 include the following basic tasks, but they vary in order and complexity from one managing system to another:

- This paper guides you through all the installation and configuration from scratch. You can factory reset your server if you prefer, no previous configurations needed.

  Important: Before performing a factory reset, back up all of your data.

- Depending on the case, install one or two Virtual I/O Servers (VIOS). A redundant VIOS is only supported by HMC and SDMC.
- Configure your network and storage. This procedure might require information from your network administrator, storage administrator, or both.
- Create a client LPAR.

Each chapter guides you step by step to achieve a fully functional PowerVM solution with one LPAR ready for use.

1.2 Planning

During the development of this paper, a unique server was used for each of the management interfaces. Figure 1-1 shows the model, type, and management interface used.

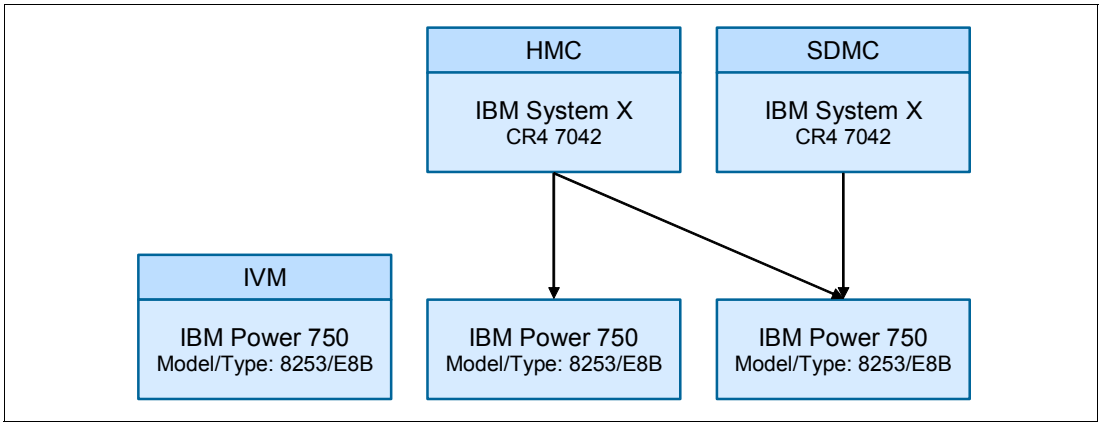

Figure 1-1 Hardware schema used in this paper

Before you start to configure your environment, complete the following planning tasks:

- Check the firmware levels on the server and the HMC or SDMC.
- Decide whether to use Logical Volume Mirroring (LVM) in AIX LPARs or Multipath I/O (MPIO.) The examples in this paper use MPIO.

  MPIO: MPIO is a fault-tolerance and performance enhancement technique where more than one path is used between a system and its storage devices.

- Make sure that your Fibre Channel switches and adapters are N_Port ID Virtualization (NPIV) capable.

  NPIV: *N_Port ID Virtualization* is a subset of Fibre Channel standard that PowerVM uses to virtualize Fibre Channel adapters.

- Make sure that your network is properly configured.
- Check the firewall rules on the HMC or SDMC.
- Plan how much processor and memory to assign to the VIOS for best performance.
- Plan the slot numbering scheme of the VIOS virtual adapter. This paper uses the scheme shown in Figure 1-2. SDMC offers automatic handling of slot allocation.

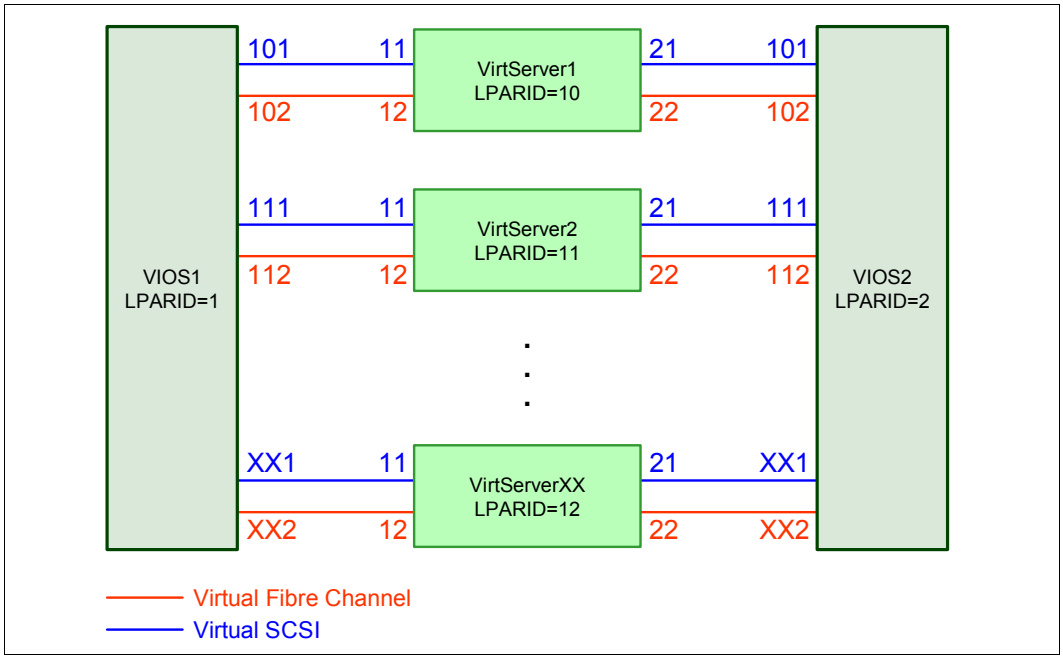

Figure 1-2 Virtual adapter slot numbering used in this paper

- Plan for two VIOS. Use the dual VIOS architecture so that you can have serviceability and scalability.

> **Dual VIOS architecture:** The dual VIOS architecture is available only when using the HMC or SDMC as managers. You cannot use dual VIOS with IVM.

The dual VIOS setup offers serviceability to a PowerVM environment on the managed system. It also provides added redundancy and load balancing of client network and storage.

The mechanisms involved in setting up a dual VIOS configuration use Shared Ethernet Adapter (SEA) failover for network and MPIO by using shared drives on the VIOS partitions for client storage. Other mechanisms can be employed, but SEA failover for networks and MPIO for storage require less configuration on the client partitions.

> **SEA:** A SEA is a VIOS component that bridges a physical Ethernet adapter and one or more virtual Ethernet adapters. For more information, see the IBM Power Systems Hardware Information Center at the following address, and search for *POWER6 Shared Ethernet Adapters*:
>
> `http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha8/hwicwelcome.htm`

SEA failover and MPIO allow for serviceability, redundancy, and load balancing with the VIOS partitions. One VIOS can act as a primary VIOS for networks and can be a standby for

storage. The other VIOS can act as a standby for networks and can be the primary VIOS for storage. The flexibility afforded by using a dual-VIOS setup caters to a range of client requirements.

Table 1-2 shows the adapter allocation for VIOS1 (illustrated in Figure 1-2). This table describes the relationship between the virtual client adapter ID and the client adapter IDs of the virtual servers.

Table 1-2 VIOS1 adapter ID allocation

Virtual adapter	Server adapter ID	VLAN ID	Server adapter slot	Client partition or virtual server	Client adapter ID	Client adapter slot
Virtual Ethernet	2 (used default allocation)	1 (used default allocation)	C2	All virtual servers		
Virtual Ethernet[a]	3 (used default allocation)	99 (default for SDMC only)	C3			
Virtual Ethernet[b]		1		VirtServer1	2	C2
Virtual VSCSI	101		C101	VirtServer1	11	C11
Virtual fiber	102		C102	VirtServer1	12	C12
Virtual Ethernet[b]		1		VirtServer2	2	C2
Virtual VSCSI	111		C111	VirtServer2	11	C11
Virtual fiber	112		C112	VirtServer2	12	C12

a. Use this virtual Ethernet adapter as the control channel adapter (SEA failover adapter).
b. This client virtual Ethernet adapter is not associated with a VIOS. The VLAN ID configured on the adapter is the link to the SEA configuration.

Similarly Table 1-3 describes the adapter ID allocation for VIOS2 and its relationship to the client adapter IDs of the virtual servers.

Table 1-3 VIOS2 adapter ID allocation

Virtual adapter	Server adapter ID	VLAN ID	Server adapter slot	Client partition or virtual server	Client adapter ID	Client adapter slot
Virtual Ethernet	2 (used default allocation)	1 (used default allocation)	C2			
Virtual Ethernet[a]	3 (used default allocation)	99 (default for SDMC only)	C3			
Virtual VSCSI	101		C101	VirtServer1	21	C21
Virtual fiber	102		C102	VirtServer1	22	C22
Virtual VSCSI	111		C111	VirtServer2	21	C21
Virtual fiber	112		C112	VirtServer2	22	C22

a. Use this virtual Ethernet adapter as the control channel adapter (SEA failover adapter).

1.3 Terminology differences

IVM, HMC, and SDMC use IBM Power Systems terminology, which can differ from x86 terminology. Table 1-4 lists the terms used in Power Systems environments, maps them to similar x86 terms, and provides a definition for the terms.

Table 1-4 Power Systems and x86 terms

Power Systems terms	x86 term or concept	Definition
Managed system	Server or system	A physical server, including physical processors, memory, and I/O resources that are often virtualized into virtual servers, which are also known as *client LPARs*.
Management partition	Virtual machine, virtual server, management operating system, VMWare Service Console, or KVM Host partition	The LPAR that controls all of the physical I/O resources on the server and provides the user interface from which to manage all of the client LPARs within the server. The LPAR in which IVM is installed.
Client LPAR, LPAR, or virtual server	Virtual machine or virtual server	The collection of virtual or physical processor, memory, and I/O resources defined to run the client OS and its workload.
IBM POWER® Hypervisor™	x86 hypervisor	The underlying software of VIOS that enables the sharing of physical I/O resources between client LPARs within the server. In IVM environments, the terms *Virtual I/O Server* and *Integrated Virtualization Manager* are sometimes used interchangeably.

1.4 Prerequisites

Verify the following prerequisites to get as close to an ideal scenario as possible:

- Ensure that your HMC or SDMC (the hardware or the virtual appliance) is configured, up, and running.
- Ensure that your HMC or SDMC is connected to the HMC port of the new server. Use either a private network or a direct cable connection.
- Ensure that TCP port 657 is open between the HMC or SDMC and the virtual server to enable dynamic LPAR functions.
- Ensure that you properly assigned IP addresses for the HMC and SDMC.
- Ensure that IBM Power Server is ready to power on.
- Ensure that all your equipment is connected to 802.3ad capable network switches with link aggregation enabled. For more information, see Chapter 5, “Advanced configuration” on page 79.
- Ensure that the Fibre Channel fabrics are redundant. For more information, see Chapter 5, “Advanced configuration” on page 79.
- Ensure that the Ethernet network switches are redundant.
- Ensure that SAN storage for virtual servers (logical partitions) is ready to be provisioned.

2

Configuring PowerVM with Integrated Virtualization Manager

IBM developed the Integrated Virtualization Manager (IVM) as a server management solution. This solution performs a subset of the Hardware Management Console (HMC) and IBM Systems Director Management Console (SDMC) features for a single server, avoiding the need for a dedicated HMC or SDMC server. IVM manages a single stand-alone server. A second server managed by IVM has its own instance of IVM installed. With the subset of HMC and SDMC server functions in IVM, an administrator can quickly set up a server. IVM is integrated within the Virtual I/O Server (VIOS) product, which services I/O, memory, and processor virtualization in IBM Power Systems.

Many environments need small partitioned systems, either for testing reasons or for specific requirements, for which the HMC and SDMC solutions are not ideal. A sample situation is where small partitioned systems cannot share a common HMC or SDMC because they are in multiple locations.

IVM is a simplified hardware management solution that inherits most of the HMC features. It manages a single server, avoiding the need for an independent personal computer. With the IVM solution, an administrator can take advantage of reduced system setup time and easier hardware management, at a lower cost.

This chapter explains how to set up PowerVM using IVM. It includes the following sections:

- Setting up a single VIOS using IVM
- Setting up a dual VIOS with IVM
- Setting up an N_Port ID Virtualization Fibre Channel with IVM

2.1 Setting up a single VIOS using IVM

This section explains how to install a single instance of the VIOS and IVM on the server. As a requirement of this installation, the server must not be attached to an HMC or SDMC. If an HMC or SDMC is attached, the server must be reset to the manufacturing default.

To reset the server to the manufacturing default configuration, see the IBM Power Systems Hardware Information Center at the following address, and search for *manufacturing default configuration for POWER6*:

http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha8/hwicwelcome.htm

See also 2.1, “Reset to Manufacturing Default Configuration” of *Integrated Virtualization Manager on IBM System p5*, REDP-4061.

When not using the HMC or the SDMC, VIOS takes control of all the hardware resources. You do not need to create a specific partition for the VIOS. When VIOS is installed using the default settings, it installs on the first internal disk controller of the server and onto the first disk on that controller. IVM is part of VIOS and is activated when VIOS is installed without an HMC or SDMC.

2.1.1 Installing a VIOS

This section is the only section on using the IVM where you cannot use the graphical browser interface. You must use a text-based interface to install and configure a VIOS. The initial installation and configuration of VIOS requires a serial ASCII console (a physical ASCII terminal or a suitable terminal emulator) and cross-over cable connected to one of the two serial ports on the server. These ports might be a DB9 or an RJ45 connector.

To install a VIOS on the server:

1. Cable between the PC and the serial port (S1 or S2) on the server.
2. Configure the ASCII console to communicate based on the following settings:
   - 19200 baud
   - 8 data bits
   - No parity
   - 1 stop bit
   - Xon/Xoff protocol
3. Ensure that the server is in the normal boot mode as indicated on the operator panel (1 N V=N T).
4. Press the white power button on the front panel to power on the server.
5. On the ASCII console, if presented with options to set this panel as the active console, press the keys indicated on the panel.
6. If prompted about the license agreements or software maintenance terms, accept the agreements or terms.

7. After the first menu selection panel (Figure 2-1) is displayed, as soon as you see the word `keyboard` the bottom of the panel, press the 1 key. If you delay, the system attempts to start any operating system (OS) that might be loaded on the server.

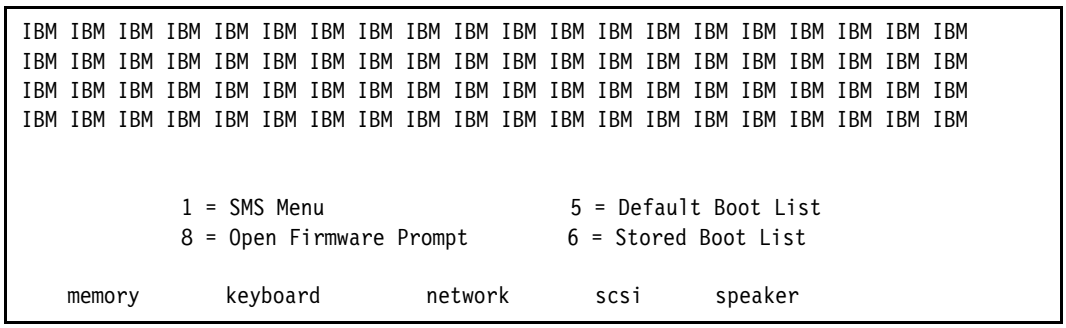

Figure 2-1 SMS startup panel

8. Select the language if necessary.
9. Enter the Service Processor password for the admin user account. The default password is `admin`. If the default password does not work, nor do you have the `admin` password, contact hardware support to walk through the sign-on process with an IBM service representative profile.
10. Insert the VIOS installation media in the CD or DVD drive.
11. To start the media from the CD or DVD drive, complete these steps:
    a. Type `5` for Boot Options.
    b. Type `1` for Install/Boot Device.
    c. Type `3` for CD/DVD.
    d. Type `9` for List All Devices.
    e. Select the correct CD or DVD device from the list (probably the last device at the bottom of the list).
    f. Select the media type from the list.
12. From the SMS menu, select **Normal Mode Boot**, and then select **Exit**.
13. Select the console number, and then press Enter.
14. Select the preferred language.
15. On the Installation and Maintenance menu, select option **1** to start with default settings. For the other panels, select the default options.

    A progress panel shows `Approximate% Complete` and `Elapsed Time`. This installation takes between 15 minutes and an hour to complete.

Accepting the license agreements

When installation is complete, sign on and accept the license agreements:

1. Sign on as `padmin`.
2. When prompted, change the password to a secure password.
3. If prompted to accept the license agreement or to accept the software maintenance agreement, accept these agreements and continue.
4. After receiving the `$` prompt, to accept the license agreement, enter the following command:

```
license -accept
```

Attaching the VIOS to the Ethernet and configuring TCP/IP

To attach VIOS to the external Ethernet and to configure TCP/IP:

1. List all Ethernet adapter ports by entering the following command:

   ```
   lsdev -vpd | grep ent
   ```

   For this installation, we plugged the Ethernet cable into the top port of the 4-port Ethernet card in slot C4 of the central electronics complex (CEC). Example 2-1 shows the `lsdev` listing.

   Example 2-1 Output from lsdev -vpd | grep ent command

   ```
   ent0 U78A0.001.DNWKGPB-P1-C6-T1 4-Port 10/100/1000 Base-TX
   ent1 U78A0.001.DNWKGPB-P1-C6-T2 4-Port 10/100/1000 Base-TX
   ent2 U78A0.001.DNWKGPB-P1-C6-T3 4-Port 10/100/1000 Base-TX
   ent3 U78A0.001.DNWKGPB-P1-C6-T4 4-Port 10/100/1000 Base-TX
   ent4 U78A0.001.DNWKGPB-P1-C4-T1 4-Port 10/100/1000 Base-TX
   ent5 U78A0.001.DNWKGPB-P1-C4-T2 4-Port 10/100/1000 Base-TX
   ent6 U78A0.001.DNWKGPB-P1-C4-T3 4-Port 10/100/1000 Base-TX
   ent7 U78A0.001.DNWKGPB-P1-C4-T4 4-Port 10/100/1000 Base-TX
   ```

   In Example 2-1, the top port (T1) of the Ethernet card in slot 4 (C4) of the CEC drawer (P1, serial number DNWKGPB) is assigned to ent4.

2. Enter the **`cfgassist`** command, and then select **VIOS TCP/IP Configuration**. Select the appropriate en*n* interface related to the adapter port chosen previously. In this case, the interface is en4, which is related to adapter port ent4.

   ent*n*: Each ent*n* has an associated en*n* and et*n* (where *n* is the same number). In the example of ent4, en4 and et4 are all related to the same Ethernet port on the card. Always use the en*n* entry to assign TCP/IP addresses.

3. In the VIOS TCP/IP Configuration panel (Figure 2-2), enter TPC/IP configuration values for the VIOS connectivity.

   ```
                          VIOS TCP/IP Configuration

   Type or select values in entry fields.
   Press Enter AFTER making all desired changes.

   [Entry Fields]
   *Hostname                                          [Vios]
   *Internet ADDRESS (dotted decimal)                 [172.16.22.10]
   Network MASK (dotted decimal)                      [255.255.252.0]
   *Network INTERFACE                                 en4
   Default Gateway (dotted decimal)                   [172.16.20.1]
   NAMESERVER
        Internet ADDRESS (dotted decimal)             [ ]
        DOMAIN Name                                   [ ]
   Cable Type                                         [tp]
   ```

   Figure 2-2 VIOS TCP/IP Configuration panel

   Initializing the Ethernet might take a few minutes to complete.

4. Ping the Internet address set for VIOS from your PC. In this example, the address is 172.16.22.10.

5. Open a browser on your PC, and connect to following address:

   `HTTPS://internet-address`

   In this example, we use the following Internet address:

   `HTTPS://172.16.22.10`

   Browser connection: Before you can proceed, you must get the browser to connect. Not all Windows browsers work with IVM. Use Microsoft Internet Explorer version 8 or earlier or Mozilla Firefox 3.6 or earlier. You must also enable pop-up windows in the browser.

Complete steps 6 and 7 from a browser window.

6. Using a browser, sign on to VIOS using the `padmin` profile and the password set previously.
7. To check for updates to VIOS, from the Service Management section in the left panel, click **Updates**.

VIOS is now installed and ready for client partitions. At this time, VIOS owns all the hardware in the server. VIOS can supply virtual adapters to the various client partitions, or it can give up control of a hardware adapter or port for assignment to the client partition.

2.1.2 Creating a partition for the client operating system

For consistency with the sections on using the HMC and SDMC to create partitions (virtual servers), you create one partition without storage and network bridging. You add these features to the partition in 2.1.3, “Configuring a VIOS for a client network” on page 12, and 2.1.4, “Configuring a VIOS for client storage” on page 12.

To create the partition:

1. In the left pane, click **View/Modify Partitions**. In the right pane, from the Partition Details section, click **Create Partition**.
2. Select the next number for the partition ID, enter a name for the partition, and then select the OS environment for the partition. Click **Next**.
3. Enter an appropriate amount of memory for the partition. Click **Next**.
4. Enter the amount of processing power for the partition, and then select shared or dedicated processor mode. Click **Next**.

   Amount of memory and processing power: The amount of memory and processing power to assign to the client partition depend on the available resources in the server and on the anticipated workload of the client partition.

5. If installing AIX or Linux, clear the Host Ethernet Adapter ports in the upper half of the panel. From the pull-down list, select the first Ethernet adapter for the partition, and then select the virtual Ethernet this adapter will be on. In most cases, the default values suffice. You can ignore the warning message that the Ethernet is not bridged at this time. You bridge these ports later when you assign them to a partition.

   Click **Next**.

6. For the storage assignment, click **None**. You add disk later. Click **Next**.

   Blank storage selection panel: If the storage selection panel is blank, you might be using a browser that is not supported by the IVM. Use another browser or an earlier version of your browser.

7. If installing AIX or Linux, skip the virtual fiber connections. Click **Next**.
8. Confirm that one virtual optical device (CD or DVD) is selected. Click **Next**.
9. In the final summary panel, which shows the settings for this partition, if everything is correct, click **Finish**. VIOS finishes creating the partition environment.

2.1.3 Configuring a VIOS for a client network

Before you can assign virtual Ethernet to the individual partitions, you must enable VIOS to bridge the physical virtual Ethernet ports:

1. From the left pane, click **View/Modify Host Ethernet Adapters**.
2. Select the physical Ethernet port, and then click the **Properties** button (above the listing).
3. Select the **Allow virtual Ethernet bridging** box, and verify the performance characteristics in the lower part of the panel. Change the setting as needed, and click **OK**.

   The VIOS internal Ethernet is now enabled to bridge the physical Ethernet and the virtual Ethernet ports. It is also ready to connect to the partition.
4. From the left pane, click **View/Modify Virtual Ethernet**. You see the four virtual Ethernets and which partitions are connected to which virtual Ethernet. Notice the red asterisks (*) next to the virtual Ethernet on the first partition. They indicate that virtual Ethernets are enabled for bridging.
5. Select the **Virtual Ethernet Bridge** tab at the top of the panel. The Virtual Ethernet Bridge tab shows which physical Ethernet port is connected to which virtual Ethernet ID.
6. From the Virtual Ethernet ID 1 pull-down list, select the appropriate physical interface to bridge. Notice the card slot and port numbers that match the physical card and port on the back of the server unit.
7. Click **Apply**.

Repeat these steps to bridge and configure additional virtual Ethernet IDs if needed.

2.1.4 Configuring a VIOS for client storage

Now decide how to assign storage (virtual or physical) for this partition.

Assigning a virtual disk

You are going to create a storage pool, assign physical disks to the storage pool, create virtual disks from the storage pool, and assign the virtual disks to the individual partition.

Important: For better performance, separate the storage pool for client partitions from the rootvg storage pool.

Consider disk protection (RAID 5 or RAID 6) at the physical disk layer. If you are going to assign internal disk, before creating virtual disks, use the **`diagmenu`** commands from the character-based VIOS console to create the RAID array. For detailed instructions, see 4.2.5, “SCSI RAID adapter use” in *Integrated Virtualization Manager on IBM System p5*, REDP-4061.

After you set up disk protection, continue creating the virtual disks:

1. From the left pane, click **View/Modify Virtual Storage**.
2. At the top of the right pane, click the **Storage Pool** tab.

3. On the **Create Storage Pool** tab, complete the following steps:
   a. Enter a name for the storage pool.
   b. Accept the default setting of **Logical volume based**.
   c. Optional: Leave the **Assign as default storage pool** check box selected to make logical volume creation easier later.
   d. Select as many of the physical volumes from the listing at the bottom of the panel as needed to size the storage pool.
   e. When you are finished, click **OK**.

   Error message: You might receive an error message about the disk being previously assigned to another storage pool. Correct the assignment if you made a mistake, or select the **Force physical volume addition** check box at the bottom of the panel to continue.

4. Back on the **Storage Pool** tab, view the size of the storage pool that was created. This pool is now ready to divide into individual virtual disks.

Dividing the pool into virtual disks

Performance hint: When dividing the storage pool into virtual disks, guest partitions operate differently if they are presented various numbers of disks and sizes. For example, IBM i performs better if presented more than four disk units (virtual or physical).

To divide the pool into individual virtual disks:

1. From the left pane, click **View/Modify Virtual Storage**.
2. On the **Virtual Disks** tab, click **Create Virtual Disk**.
3. Select a naming convention for the virtual disk names.
4. In the Virtual disk name window, complete the following steps:
   a. Enter the first name.
   b. Select the appropriate storage pool to get the storage from, and then select the amount of storage to assign to the virtual disk. Be careful to change the MB to GB as appropriate for your size selection.

      Important: Sizing the disk depends on the OS that you want to install and the anticipated workload. For example, the first virtual disk must meet the minimum disk size for the OS that is being loaded. For size requirements, see the release notes for the operating system you are installing.

      Disk assignment: You can assign this single virtual disk to your partition at this time, or create all the virtual disks and assign them to the partition at one time later.

   c. Click **OK**.

   When you are finished, the first virtual disk is displayed in the Virtual Disk panel.

Repeat steps 1 through 4 to create additional virtual disks for your partition at this time.

Assigning the virtual disk to the LPAR

To assign the virtual disk to the LPAR:

1. From the left pane, click **View/modify Virtual Storage**.
2. On the **Virtual Disk** tab, complete these steps:
   a. Select all the virtual disk units to be assigned.
   b. Click the **Modify partition assignment** button.
   c. Select the partition from the pull-down list, and click **OK**.

You are done assigning virtual storage to the client partition and can now install the client OS as explained in 2.1.5, “Installing a client operating system” on page 14.

Assigning a physical disk

Assigning the physically attached disk is similar to assigning physical disk to a storage pool. To assign the physical disks:

1. From the left pane, click **View/Modify Virtual Storage**.
2. From the right pane, on the **Physical Volumes** tab, complete the following steps:
   a. Select the physical volumes to be assigned.
   b. Click the **Modify partition assignment** button.
   c. From the pull-down list, select the partition.
   d. Click **OK**.

Continue with the installation of the client OS in the next section.

2.1.5 Installing a client operating system

Now you install the client OS on the partition. These steps vary depending on whether you are installing AIX, IBM i, or Linux.

First decide whether you plan to install the client OS from the physical media in the CD or DVD drive of the server or from the image files that are uploaded into the VIOS environment. In either case, the VIOS does not switch to the next CD or DVD image when the client OS requests it. You must manually change the CD or DVD, or change the file assigned to the virtual optical drive.

The following example explains how to install IBM i from physical media. You must have IBM System i® Access for Windows installed on a PC to configure the LAN console to access the partition when installing IBM i.

To begin the installation, assign the physical CD or DVD drive to the partition:

1. In the left pane, click **View/Modify Virtual Storage**.
2. In the right pane, on the **Optical/Tape** tab, complete these steps:
   a. If necessary, expand the **Physical Optical Devices** section.
   b. Select the **cd0** device.
   c. Click the **Modify partition assignment** button.
   d. From the pull-down list, select the partition.
   e. Click **OK**.

   The physical CD or DVD drive in the server now belongs to that partition.

3. Select the IPL type for the IBM i partition, and verify the other partition settings:
   a. In the left pane, click the **View/Modify Partitions**.
   b. In the right pane, check the partition. Complete these steps:
      i. From the **More-Tasks** pull-down list, select **properties**.
      ii. Change the IPL type to **D (IPL from CD/DVD)**, and change the keylock position to **Manual**.
4. Place the I_Base_01 CD in the CD/DVD drive of the server. Click **OK** at the bottom of the panel.
5. Select the partition again, and use the **Activate** button to start the partition IPL.

Progress and reference codes: For IBM i, if the partition reaches the C600-4031 reference code, the partition is operating normally and looks for the LAN console session.

If the IBM i partition reaches reference code A600-5008, the partition was unsuccessful in contacting the console session. Therefore, you must troubleshoot the LAN console connectivity. Make sure that you bridged the correct VLAN ports and that the LAN console PC is on the same subnet as the bridged Ethernet port.

After you reach the language selection panel on the console, the installation of IBM i proceeds the same as installing on a stand-alone server. Continue with the Dedicated Service Tools functions to add the disk to the auxiliary storage pool (ASP) and loading the OS.

You have now installed and configured VIOS and at lease one client partition. The following sections expand on this basic installation with more advanced features.

2.2 Setting up a dual VIOS with IVM

VIOS does not support dual VIO installations on the same physical server without the use of either an HMC or SDMC.

2.3 Setting up an N_Port ID Virtualization Fibre Channel with IVM

Fiber connected storage can be assigned to partitions in the following ways:

- Assigning the physical storage attached to them to a particular partition (see “Assigning a physical disk” on page 14)
- Creating virtual Fibre Channel adapters and assigning the virtual adapter to the partition

Preferred practice: Use the internal disk for the installation of VIOS, mirroring the rootvg volumes. Use external SAN storage for the installation of client OSs. This approach positions the client partitions to use partition mobility later.

To configure N_Port ID Virtualization (NPIV) attached storage, create the virtual fiber adapters to generate the worldwide port name (WWPN) to allow the configuration and assignment of the storage. To configure the virtual fiber adapters:

1. In the left pane, click **View/Modify Partitions**.
2. Select the partition.
3. From the **More Tasks** pull-down list, select **Properties**.
4. On the **Storage** tab, complete these steps:
   a. Expand the **Virtual Fiber Channel** section.
   b. If an interface is not shown, click **Add** to create the first interface. Select the first interface listed (listed as **Automatically Generated**), and select the correct physical port from the pull-down list.
   c. Click **OK** to complete the generation of the WWPNs for this interface.
5. Return to the partition storage properties (steps 1 - 3) to view the WWPNs. Record these numbers to configure the fiber attached storage.

After the OS is installed and the NPIV attached storage is provisioned, directly assign the storage to the OS of the partition. VIOS is unaware of the storage. Use the normal procedures to add newly attached storage to the OS (AIM, IBM i, or Linux.)

After you finish the installation using the IVM, you can increase the reliability, availability, and serviceability (RAS) of the configuration by applying the information in Chapter 5, “Advanced configuration” on page 79.

3

Configuring PowerVM with the Hardware Management Console

The Hardware Management Console (HMC) started managing POWER4 systems at a dedicated partition level where administrators assigned whole processors, memory, and physical adapters to partitions. The HMC has progressed with the Power Systems virtualization from the early days of virtualization to its current form with micropartitioning, virtualization of adapters, processors, and memory.

One of the most notable benefits that HMC has over the Integrated Virtualization Manager (IVM) is that multiple HMCs can manage a single Power System and a single HMC can manage multiple Power Systems. For enterprise Power Systems, use a redundant HMC setup. Table 1-1 on page 2 outlines other differences.

This chapter highlights the use of the HMC to set up a virtualized environment on a Power System. It explains step-by-step how to set up a single Virtual I/O Server (VIOS) partition environment and how to set up a dual VIOS partition and using N_Port ID Virtualization (NPIV).

This chapter includes the following sections:

- Setting up a single VIOS using an HMC
- Setting up a dual VIOS using HMC
- Setting up a virtual Fibre Channel using HMC
- Adding additional client partitions

Before you begin this chapter, you must meet the following prerequisites:

- Be familiar with the information in Chapter 1, “Introduction to PowerVM” on page 1, particularly 1.4, “Prerequisites” on page 6.
- Install and configure HMC with the latest applicable updates.
- Connect your system and ensure that it is visible to the HMC.
- Ensure that your system status is in the *Standby* or Operating state.

3.1 Setting up a single VIOS using an HMC

You can use a single VIOS in your managed server environment for a basic setup of PowerVM with an HMC.

For the exercise in this chapter, the following adapters are cabled on the system for the deployment of a single VIOS setup:

- The first port of `U78A0.001.DNWHZS4-P1-C2 Quad 10/100/1000 Base-TX PCI-Express Adapter`
- The first port of `U5802.001.0087356-P1-C2 Fibre Channel Serial Bus for virtual SCSI`
- The second port of `U5802.001.0087356-P1-C2 Fibre Channel Serial Bus` for virtual fiber (NPIV)

For placement of your cards, see the *Adapter Placement Guide* for the system you are working on.

The basic network setup in a single VIOS environment uses the Shared Ethernet Adapter (SEA), where a single physical Ethernet adapter is bridged to a virtual Ethernet adapter. The virtual Ethernet adapter configured for bridging is set for one VLAN, VLAN ID 1, to be used on the client partitions. The virtual adapter IDs for the virtual Ethernet adapters use the default adapter IDs allocated by the HMC.

The basic storage setup in a single VIOS environment uses Multipath I/O (MPIO) where the disks are configured on the VIOS partition and mapped to a client partition by using virtual SCSI. The virtual adapter IDs for virtual SCSI use the scheme outlined in Table 1-2 on page 5.

This chapter demonstrates how to use user-defined adapter IDs and HMC-derived adapter IDs (both acceptable options). You might have your own adapter ID scheme or use the default adapter ID allocations provided by the HMC.

For more information about adapter ID numbering, see 5.1, “Adapter ID numbering scheme” on page 80.

3.1.1 Creating a VIOS partition profile

To create a VIOS partition profile:

1. Log on to the HMC as `hscroot` or with a user ID that has similar access.
2. Select **Systems Management** → **Servers** to open the Servers panel.

3. In the Servers panel, complete these steps:

   a. Select your managed system (the first check box in the first column on the right sides of the table (circled in Figure 3-1)).

   b. In the Name field for your managed system, click the » button. Then select **Configuration** → **Create Logical Partition** → **VIO Server**.

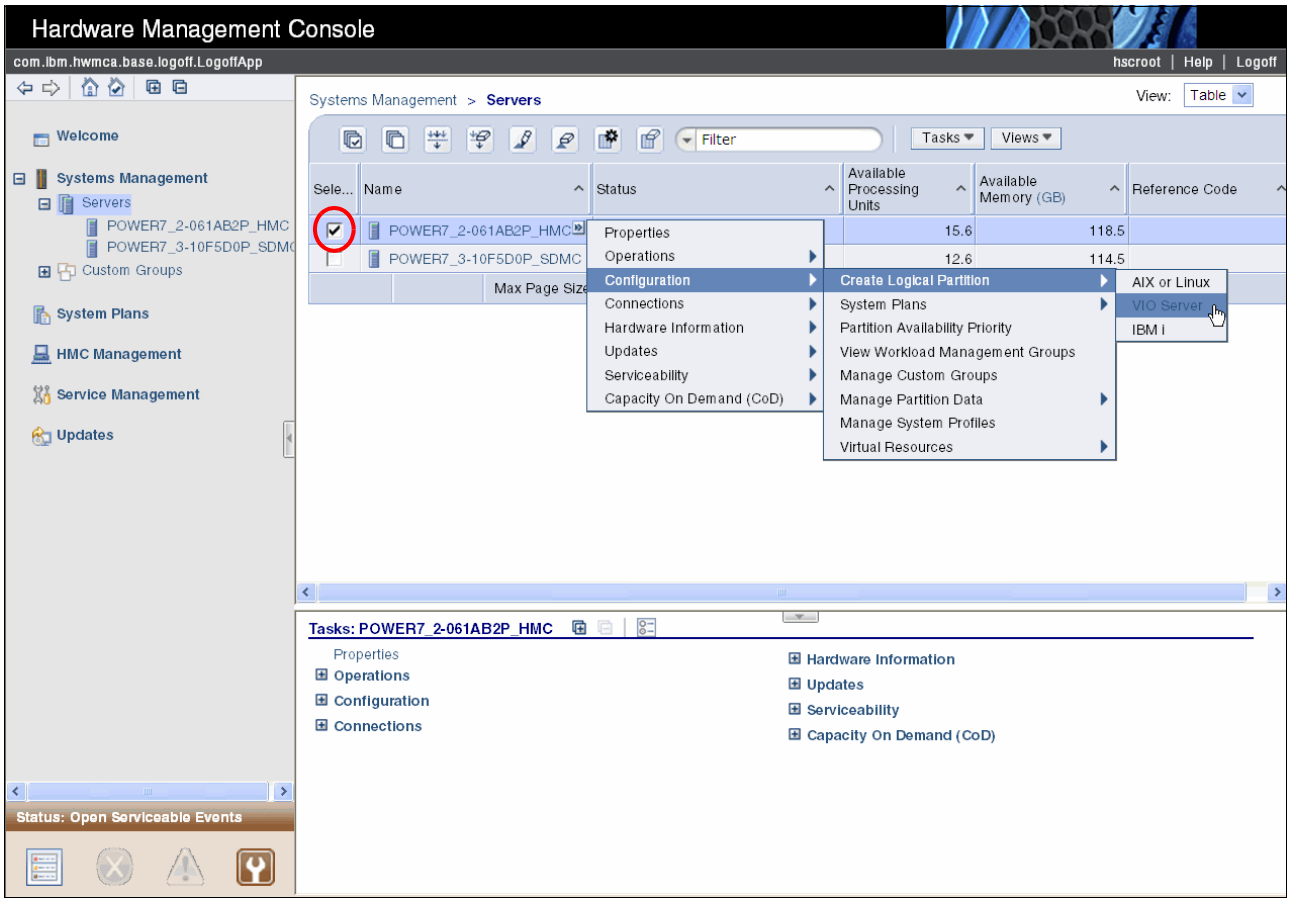

Figure 3-1 HMC window showing the system

4. In the Create Partition panel, specify the name of your partition. In this example, we enter `VIOS1`. Then click **Next**.

5. In the Partition Profile panel, enter your profile name. In this example, we enter `Normal`. Then click **Next**.

6. In the Processors panel, verify that **Shared** is selected. Then click **Next**.

7. In the Processor Settings panel, complete the following settings:

   a. For Desired processing units, enter `0.2`.
   b. For Maximum processing units, enter `10`.
   c. For Desired virtual processors, enter `2`.
   d. For Desired maximum processors, enter `10`.
   e. Select the **Uncapped** check box.
   f. Update Weight setting to `192`.

Processor settings: The processor settings allow for the lowest utilization setting for the VIOS of 0.2 (Desired processing units), but it is scalable up to two processing units (Desired virtual processors) if necessary. The higher weight gives the VIOS priority over the other logical partitions (LPARs). For more information, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

8. In the Memory Settings panel, complete the following fields:
   a. For Minimum Memory, enter `1 GB`.
   b. For Desired Memory, enter `4 GB`.
   c. For Maximum Memory, enter `8 GB`.
9. In the I/O panel, complete the following steps:
   a. Select the following check boxes:
      - **The RAID or SAS controller where the internal disks are attached to (disk controllers for the VIOS internal drives)**
      - **The Ethernet adapter (newer adapters are described as PCI-to-PCI Bridge) where it has been cabled to the network**
      - **The Fibre Channel adapter attached to the SAN fabric**
   b. Click **Add as desired**.
   c. Click **Next**.

   Figure 3-2 shows the selected adapters.

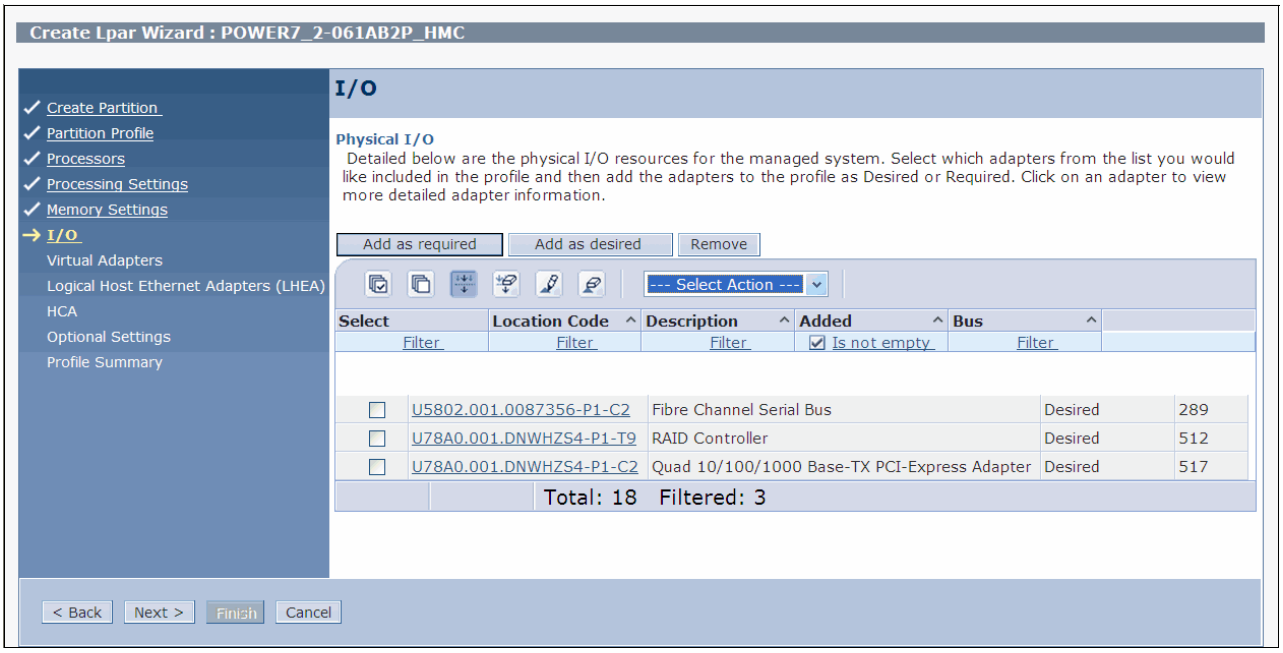

Figure 3-2 I/O adapter selection

10. In the Virtual Adapters panel, update the Maximum virtual adapters setting to `1000`.

Adapter numbering scheme: You can plan your own adapter numbering scheme. You must set the Maximum virtual adapters setting in the Virtual Adapters panel to allow for your numbering scheme. The maximum setting is 65535. The higher the setting is, the more memory the managed system reserves to manage the adapters.

11. Create a virtual Ethernet adapter for Ethernet bridging. In the Virtual Adapters panel (Figure 3-3), complete these steps:

    a. Select **Actions** → **Create Virtual Adapter** → **Ethernet Adapter**.

    b. In the Create Virtual Ethernet Adapter panel, select the **Use this adapter for Ethernet bridging** check box and click **OK**.

    The virtual Ethernet adapter is created and is shown in the Virtual Adapters panel.

> **Default settings:** When creating the virtual Ethernet adapter, we accepted the default settings for Adapter ID, Port VLAN ID, and Ethernet Bridging Priority (Trunk Priority). These settings are customizable for a range of planning designs or standards.

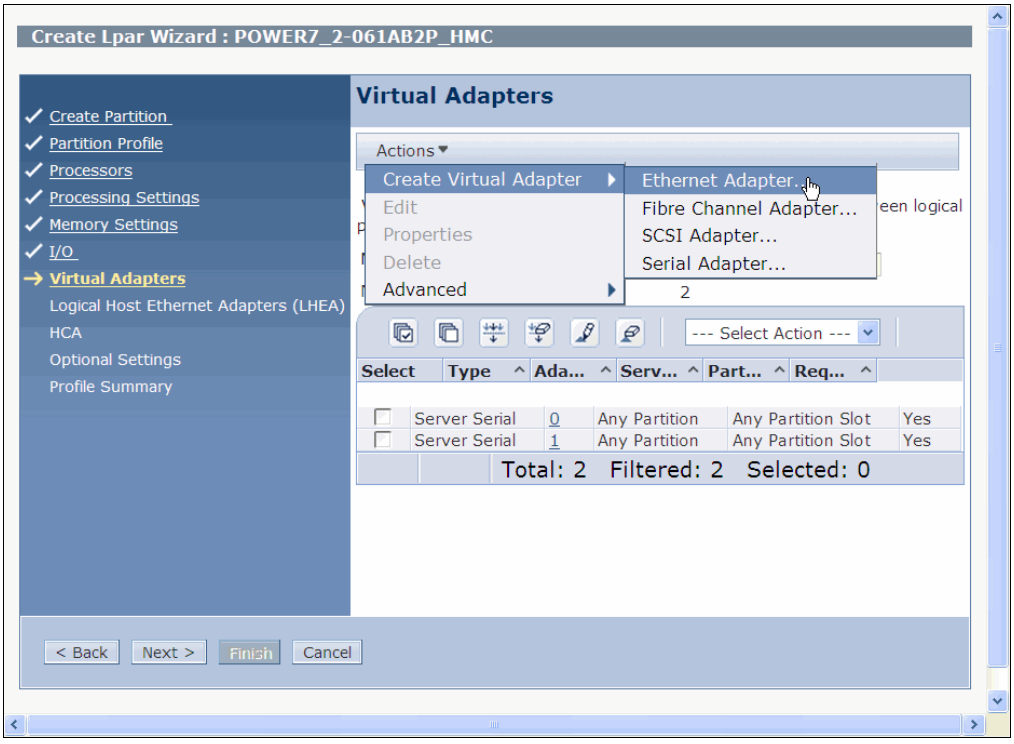

Figure 3-3 Creating the virtual Ethernet adapter

12. Depending on your VIOS partition setup, choose one of the following options:

    – For a single VIOS partition setup, skip to step 13.

    – For dual VIOS partition setup, continue to create a virtual Ethernet adapter for SEA failover, in the Virtual Adapters panel:

        i. Select **Actions** → **Create Virtual Adapter** → **Ethernet Adapter**.

        ii. In the Create Virtual Ethernet Adapter panel, update the Port Virtual Ethernet value to 99.

        iii. Click **OK**. The virtual Ethernet adapter is created and is shown in the Virtual Adapters panel.

13. To create the virtual SCSI adapter, complete these steps:

    a. In the Virtual Adapters panel, select **Actions** → **Create Virtual Adapter** → **SCSI Adapter**.

    b. In the next panel, complete the following steps:

        i. Select the **Only selected client partition can connect** check box.
        ii. For Adapter, enter 101.
        iii. For Client partition, enter 10.
        iv. For Client adapter ID, enter 11.
        v. Click **OK** to accept settings.

> **Information:** For the client partition, we begin at partition ID 10 (reserving partition IDs 2 - 9 for future VIOS or infrastructure servers). For the adapter ID, we chose 101 as a numbering scheme to denote the partition and virtual device 1. For the Client adapter ID, we chose 11 as the first disk adapter for the client partition.

14. In the Virtual Adapter panel (Figure 3-4), which shows the virtual adapters that you created, click **Next**.

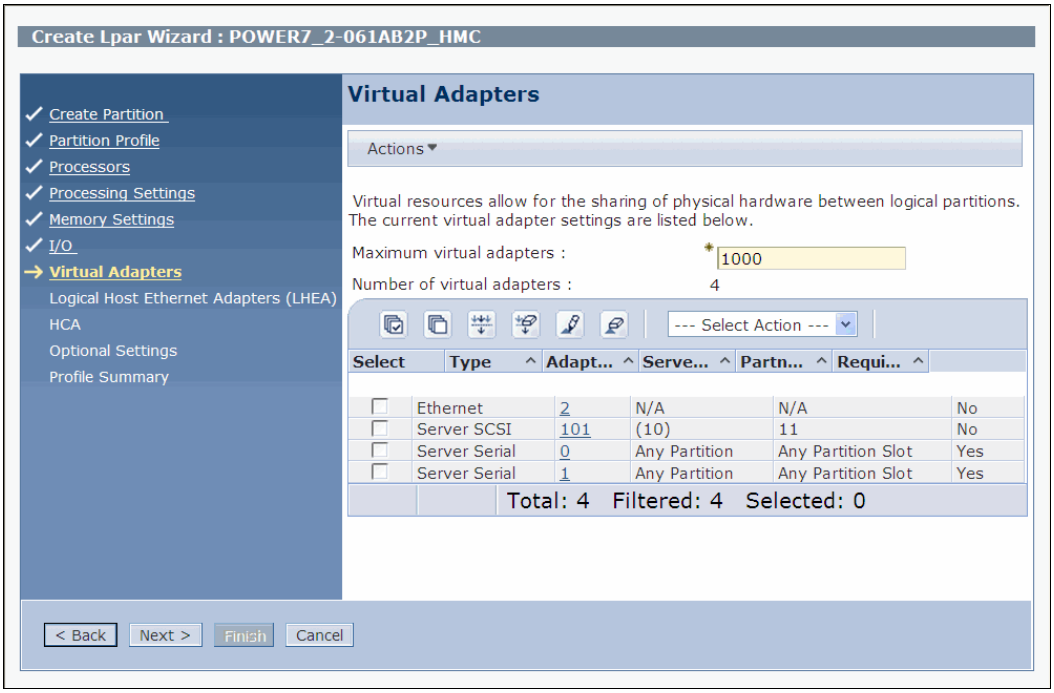

Figure 3-4 Virtual Adapters panel with virtual Ethernet and virtual SCSI adapter defined

15. For the remaining panels, click **Next** until you reach the Profile Summary panel.
16. In the Profile Summary panel, verify your settings, and then click **Finish**.
17. Click your managed system to view the VIOS partition profile you created.

You have now created an LPAR (virtual server) for the VIOS installation.

3.1.2 Installing a VIOS

You can install a VIOS by using the following methods:

- DVD
- The HMC using the **installios** command
- Network Installation Manager (NIM)

Installing a VIOS by using a DVD

To install a VIOS by using a DVD:

1. Select **Systems Management** → **Servers** → ***your system***.
2. In the Servers panel, complete these steps:
   a. Select the VIOS partition check box.
   b. In the Name field for your partition, click the button. Then select **Operations** → **Activate** → **profile**.
3. In the Activate Logical Partition panel, click **Advanced**.
4. In the Advanced options panel, from the Boot mode list, select **SMS**, and then click **OK**.
5. In the Activate Logical Partition panel, click **OK** to activate the VIOS partition.
6. Open a terminal window to the VIOS partition. Observe the VIOS partition being booted into the SMS Main Menu.
7. Continue with steps 10 - 15 on page 9 in 2.1.1, “Installing a VIOS” on page 8. Then complete the steps in “Accepting the license agreements” on page 9.

The VIOS is ready to be configured for client network and storage service.

For client storage, you can extract the worldwide port name (WWPN) from the Fibre Channel adapter interface and give it to the storage area network (SAN) administrator for zoning. You can use the following command to extract the WWPN:

```
lsdev -dev fcs0 -vpd | grep "Network Address"
```

Example 3-1 shows the WWPN for Fibre Channel adapter port `fcs0`. To obtain the WWPN for fcs1, run the command as shown, but replace `fsc0` with `fcs1`.

Example 3-1 WWPN of fcs0 Fibre Channel adapter port

```
$ lsdev -dev fcs0 -vpd | grep "Network Address"
        Network Address.............10000000C99FC3F6
```

Installing a VIOS by using HMC

To use the **installios** command on the HMC, you need the following TCP/IP details for the VIOS partition:

- VIOS TCP/IP address
- The subnet mask of the VIOS network
- The VIOS network gateway TCP/IP address

To install a VIOS by using the **installios** command on the HMC console command-line interface (CLI):

1. Insert the VIOS media into the HMC DVD drive. (If multiple media are available, insert the first DVD.)
2. Log on to the HMC CLI with an ASCII terminal emulator (SSH to the TCP/IP address of the HMC).
3. At the command prompt, enter **installios**.
4. Select your system from the list of systems connected to the HMC.
5. Select the VIOS partition on which you are conducting the installation.
6. Select the VIOS partition profile.
7. Press Enter to accept `/dev/cdrom` as the default source of the media.
8. Enter the VIOS TCP/IP address.
9. Enter the VIOS subnet mask.
10. Enter the VIOS TCP/IP gateway.
11. For the VIOS adapter speed, type `auto`.
12. For the VIOS adapter duplex, type `auto`.
13. To not configure the TCP/IP address on the VIOS after installation, type `no`.
14. Select the open TCP/IP address of the HMC.

    Adapters: At least two adapters are shown with their TCP/IP addresses. One address is for the HMC open network. The other address is the private network to the Flexible Service Processor (FSP) port of your system.

15. After the HMC retrieves the Ethernet adapter details based on the VIOS partition profile configuration, select the Ethernet adapter port that is cabled in 3.1, “Setting up a single VIOS using an HMC” on page 18.
16. Press Enter to accept `en_US` as the language and the locale defaults.

    Alternative: If en_US is not your default language and locale, enter the language and locale that you regularly use.

17. In the window that shows the details that you selected, press Enter.
18. Review the License Agreement details. At the end of the License Agreement window, type `Y` to accept the agreement.
19. If the installation media spans multiple DVDs: When prompted, change DVDs. Then type `c` to continue.

    Using the details that you provided, the HMC uploads the software from the installation media to a local file system within the HMC. NIM on Linux (NIMOL) features on HMC are used to network boot the VIOS partition and network install the VIOS software.
20. Open a terminal window to the VIOS partition.
21. After the VIOS installation is completed and the VIOS partition prompts you to log in, enter the `padmin` user ID.
22. When prompted, change the password to a secure password.
23. To accept the VIOS software maintenance terms and conditions, type `a`.

24. To accept the VIOS license agreement, enter the following command:

```
license -accept
```

25. To list the physical Fibre Channel adapters on the VIOS, enter the **lsnports** command.

    Example 3-2 shows the Fibre Channel adapter ports configured on VIOS1. As explained in 3.1, “Setting up a single VIOS using an HMC” on page 18, the first port (T1) is planned for virtual SCSI. The second port (T2) is planned for virtual Fibre Channel, as explained later in this chapter.

Example 3-2 Fibre Channel adapter port listing on VIOS1

```
$ lsnports
name             physloc                        fabric tports aports swwpns  awwpns
fcs0             U5802.001.0087356-P1-C2-T1          1     64     64   2048    2046
fcs1             U5802.001.0087356-P1-C2-T2          1     64     64   2048    2048
```

26. For client storage, extract the WWPN from the Fibre Channel adapter interface and given to the SAN administrator for zoning by using the following command:

    ```
    lsdev -dev fcsX -vpd | grep "Network Address"
    ```

    Example 3-3 shows the WWPN for Fibre Channel adapter port fcs0. To obtain the WWPN for fcs1, run the command but replace `fsc0` with `fcs1`.

Example 3-3 WWPN for fcs0 Fibre Channel Adapter port

```
$ lsdev -dev fcs0 -vpd | grep "Network Address"
        Network Address.............10000000C99FC3F6
```

You can now configure the VIOS for client network and storage service.

Installing a VIOS by using NIM

For information about NIM and using it to install the VIOS software and other client partition software, see *NIM from A to Z in AIX 5L*, SG24-7296.

A VIOS NIM installation entails the following main steps:

1. Register the VIOS partition as a NIM machine object.
2. Create an installation NIM resource object for the VIOS software.
3. Create a NIM resource object to start using the VIOS software.
4. Allocate the NIM resources you created to your VIOS partition NIM machine object.
5. On the HMC, boot the VIOS partition using the network interface.

If a NIM server is not available and you want to use NIM to build a PowerVM environment on your system, complete these steps:

1. Build the VIOS partition using either DVD or the **installios** command.
2. Build the first client partition as an AIX NIM server.
3. If you plan to build a second VIOS partition, build the second VIOS using NIM.
4. Deploy any Linux or AIX client partitions by using NIM.

3.1.3 Configuring a VIOS partition

You can configure a VIOS partition for a client network or for client storage.

Configuring a VIOS for a client network

To configure a VIOS for a client network:

1. Open a console session and log on to the VIOS terminal window.
2. List the Ethernet devices configured on the VIOS to show the logical name relationship to the physical device details:

   ```
   lsdev -vpd | grep ent
   ```

   Example 3-4 lists the VIOS Ethernet devices.

Example 3-4 Listing of VIOS Ethernet devices

```
$ lsdev -vpd | grep ent
  ent4 U8233.E8B.061AB2P-V1-C2-T1 Virtual I/O Ethernet Adapter (l-lan)
  ent0 U78A0.001.DNWHZS4-P1-C2-T1 4-Port 10/100/1000 Base-TX PCI-Express Adapter
  ent1 U78A0.001.DNWHZS4-P1-C2-T2 4-Port 10/100/1000 Base-TX PCI-Express Adapter
  ent2 U78A0.001.DNWHZS4-P1-C2-T3 4-Port 10/100/1000 Base-TX PCI-Express Adapter
  ent3 U78A0.001.DNWHZS4-P1-C2-T4 4-Port 10/100/1000 Base-TX PCI-Express Adapter
```

In Example 3-4 on page 26, ent0 (`U78A0.001.DNWHZS4-P1-C2-T1`) is the physical Ethernet adapter port that is cabled. The `U78A0.001.DNWHZS4-P1-C2` Ethernet adapter is the adapter selected in Figure 3-2 on page 20. Adapter ent4 (`U8233.E8B.061AB2P-V1-C2-T1`) is the virtual Ethernet adapter shown in Figure 3-4 on page 22.

Virtual Ethernet adapter: For the virtual Ethernet adapter `U8233.E8B.061AB2P-V1-C2-T1`, the V in V1 indicates that it is a virtual adapter, and C2 indicates that it is a slot with adapter ID 2 as shown in step 14 on page 22.

If you plan to use 802.3ad Link Aggregation, your respective adapters must be cabled and the network switch ports must be configured for 802.3ad Link Aggregation. To create the Link Aggregation adapter, enter the following command:

```
mkvdev -lnagg <entX> <entY> -attr mode=8023ad
```

Alternatively, create the adapter by using the **`cfgassist`** command:

a. On a command line, enter **`cfgassist`**.

b. Select **Devices** → **Link Aggregation Adapter** → **Add a Link Aggregation Adapter**.

c. In the Target Adapters field, enter the physical network adapters (spaces between each physical network adapter).

d. In the ATTRIBUTES field, enter `mode=8023ad`.

List all physical Ethernet adapters and EtherChannel adapters available for creating an SEA:

```
ledev -type ent4sea
```

3. Create the SEA, which bridges the physical adapter and the virtual adapter:

   ```
   mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 1
   ```

   Where:

   ent0 The physical adapter in step 2 (use the EtherChannel adapter if one is available for the SEA configuration).

   ent4 The virtual adapter found in step 2.

   1 The port VLAN ID of ent4 where you accepted the default Port VLAN ID allocation.

   Example 3-5 shows creating the SEA virtual network devices, where:

   ent5 An Ethernet network adapter device.
   en5 A standard Ethernet network interface where TCP/IP addresses are assigned.
   et5 An Institute of Electrical and Electronics Engineers (IEEE) 802.3 Ethernet network interface.

   Example 3-5 Creating an SEA interface

   ```
   $ mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 1
   ent5 Available
   en5
   et5
   ```

4. Configure the TCP/IP connection for the VIOS with details provided by the network administrator.

   ```
   mktcpip -hostname vios1 -interface en5 -inetaddr 172.16.22.15 -netmask
   255.255.252.0 -gateway 172.16.20.1
   ```

   Where:

   Network IP address 172.16.22.15
   Network subnet 255.255.252.0
   Network gateway 172.16.20.1

   Alternatively, configure the connection by using the **cfgassist** command:

   a. On a command line, enter **cfgassist**.
   b. Select **VIOS TCP/IP Configuration**.
   c. Select **en5**, which is the SEA interface created in step 3. Then press Enter.
   d. Enter the TCP/IP details in listed previously in step 4 after the **mktcpip** command (see “Where:”).

   Interface and port details: Interface en5 is the SEA created in step 3 on page 27. Alternatively, an additional virtual adapter can be created for the VIOS remote connection, or another physical adapter can be used (must be cabled) for the TCP/IP remote connection.

   TCP and UDP port 657 must be open between the HMC and the VIOS, which is required for dynamic LPAR (DLPAR; using the Resource Monitoring and Control (RMC) protocol).

Configuring a VIOS for a client storage

You might encounter a scenario where the SAN administrator allocated a group of disks but only one disk is planned for the client partition. For this exercise, the SAN administrator allocated the disk with logical unit number (LUN) ID `60:0a:0b:80:00:11:46:32:00:00:92:78:4e:c5:0f:0b`.

To configure and map the disks to the client partition:

1. List any Fibre Channel adapter SCSI protocol devices that are configured on the VIOS:

   ```
   lsdev | grep fscsi
   ```

   In Example 3-6, `fscsi0` and `fscsi1` are the Fibre Channel adapter SCSI protocol devices configured on VIOS1. Their attributes must be updated to allow for dynamic tracking and fast failover (applicable for a multiple Fibre Channel adapter VIOS).

   Example 3-6 List of Fibre Channel adapter SCSI protocol devices on VIOS1

   ```
   $ lsdev | grep fscsi
   fscsi0          Available   FC SCSI I/O Controller Protocol Device
   fscsi1          Available   FC SCSI I/O Controller Protocol Device
   ```

2. Update the device attributes of the Fibre Channel adapter SCSI protocol listed in step 1 to enable dynamic tracking and fast failover:

   ```
   chdev -dev fscsi0 -attr dyntrk=yes fc_err_recov=fast_fail
   chdev -dev fscsi1 -attr dyntrk=yes fc_err_recov=fast_fail
   ```

   Tip for a busy Fibre Channel adapter SCSI protocol: If the Fibre Channel adapter SCSI protocol device is busy, append the **`-perm`** flag to the command to update the VIOS database only as shown in the following example:

   ```
   chdev -dev fscsi0 -attr dyntrk=yes fc_err_recov=fast_fail -perm
   ```

   The attributes are not applied to the device until the VIOS is rebooted.

   Tip for the fast_fail and dyntrk settings: The `fast_fail` and `dyntrk` settings are useful in a setup with multiple Fibre Channel adapters or dual VIOS. With the `fast_fail` setting, the I/O can immediately fail if the adapter detects link events such as a lost link between the Fibre Channel adapter and the SAN switch port. With the `dyntrk` setting, the VIOS tolerates cabling changes in the SAN.

3. To configure the disks on the VIOS, enter **`cfgdev`**.

4. List the disks on the VIOS partition, and see the disk type:

   ```
   lsdev -type disk
   ```

   In Example 3-7, VIOS1 lists two internal SAS disks and six DS4800 disks.

   Example 3-7 List of disks with the disk type on the VIOS1

   ```
   $ lsdev -type disk
   name            status      description
   hdisk0          Available   SAS Disk Drive
   hdisk1          Available   SAS Disk Drive
   hdisk2          Available   MPIO DS4800 Disk
   hdisk3          Available   MPIO DS4800 Disk
   hdisk4          Available   MPIO DS4800 Disk
   hdisk5          Available   MPIO DS4800 Disk
   hdisk6          Available   MPIO DS4800 Disk
   hdisk7          Available   MPIO DS4800 Disk
   ```

5. To confirm the SAN LUN ID on VIOS1, enter the following command for each disk listed in step 4 until the correct disk is found with LUN ID provided by the SAN administrator:

   ```
   lsdev -dev hdiskX -attr | grep -i -E "reserve|unique_id"
   ```

Example 3-8 shows the hdisk, which the SAN administrator assigned. Also, the SCSI reserve policy is set with `single_path`. You must update this setting with no SCSI reserve locks. The LUN ID is embedded in the unique_id string for hdisk6, beginning with the sixth character.

Example 3-8 Disk attributes of hdisk6

```
$ lsdev -dev hdisk6 -attr | grep -E "unique_id|reserve"
reserve_policy  single_path Reserve Policy                  True
unique_id       3E213600A0B8000114632000092784EC50F0B0F1815      FAStT03IBMfcp Unique
device identifier         False
```

Additional disk information: Disks using EMC PowerPath drivers, IBM Subsystem Device Driver Path Control Module (SDDPCM) drivers, and IBM Subsystem Device Drivers (SDDs) also have their LUN IDs embedded in the unique_id string. Use their supplied commands to view the LUN IDs in a more readable format. To obtain the disks complete with LUN IDs, see the driver manual.

EMC disks are displayed with hdiskpowerX notation, and SDD disks are displayed with a vpathX notation. Use their disk notations with the **`lsdev`** command sequence instead of hdisk.

Other disk subsystems can use different fields to set their SCSI reserve locks. Use the **`lsdev`** command sequence without the pipe to grep, as in the following example:

```
lsdev -dev sampledisk -attr
```

6. Deactivate the SCSI reserve lock on the disk, in this case hdisk6:

   ```
   chdev -dev hdisk6 -attr reserve_policy=no_policy
   ```

   Disks using SDDPCM and SDD drivers: Ignore this step if the disks are using SDDPCM and SDD drivers because the SCSI reserve locks are already deactivated. For EMC disks and disks using native MPIO, you must deactivate the SCSI reserve locks.

   The SCSI reserve lock attribute differs among the disk subsystems. The IBM System Storage SCSI reserve lock attribute is `reserve_policy` as shown in Example 3-8. The attribute on EMC disk subsystem is `reserve_lock`.

   If you are unsure of the allowable value to use to deactivate the SCSI reserve lock, use the following command to provide a list of allowable values:

   ```
   lsdev -dev hdisk6 -range reserve_policy
   ```

7. Determine the virtual adapter name of the virtual SCSI adapter created in step 13 on page 22:

   ```
   lsdev -vpd | grep "Virtual SCSI"
   ```

   In Example 3-9, the virtual SCSI adapter with server adapter ID C101 is vhost0 to use in the next step.

   Example 3-9 List of virtual SCSI devices

   ```
   $ lsdev -vpd | grep "Virtual SCSI"
     vhost0           U8233.E8B.061AB2P-V1-C101 Virtual SCSI Server Adapter
   ```

8. Use the MPIO setup to map a whole LUN to client the operating system (OS) partitions. To map hdisk6 to CLIENT1, enter the following command:

   ```
   mkvdev -vdev hdisk6 -vadapter vhost0
   ```

   Where:

   hdisk6 The disk in step 5 on page 28.

   vhost0 The virtual server SCSI adapter with adapter ID 101 created for CLIENT1, in step 6.

   In Example 3-10, the Virtual Target Device (VTD) `vtscsi0` is created.

   Example 3-10 Creating a disk mapping to a client partition

   ```
   $ mkvdev -vdev hdisk6 -vadapter vhost0
   vtscsi0 Available
   ```

9. Check the mapped devices to vhost0:

   ```
   lsmap -vadapter vhost0
   ```

   In Example 3-11, the `vhost0` virtual SCSI adapter shows one disk mapped, where hdisk6 is mapped to the `vtscsi0` device.

   Example 3-11 vhost0 disk mapping

   ```
   $ lsmap -vadapter vhost0
   SVSA            Physloc                                        Client Partition ID
   --------------- ---------------------------------------------- ------------------
   vhost0          U8233.E8B.061AB2P-V1-C101                      0x0000000a

   VTD                   vtscsi0
   Status                Available
   LUN                   0x8100000000000000
   Backing device        hdisk6
   Physloc               U5802.001.0087356-P1-C2-T1-W202200A0B811A662-L5000000000000
   Mirrored              false
   ```

3.1.4 Creating a logical partition profile for a client operating system

In this section, you create a client partition using the following settings:

- Partition ID: `10`
- Partition Name: `CLIENT1`
- Profile Name: `Normal`
- Processing Units: `0.2`
- Virtual Processors: `4`
- Virtual Ethernet adapter using the default adapter ID allocation: `2`
- Virtual SCSI adapter using adapter ID: `11`

The following steps explain how to create an AIX or Linux partition. However, you can also use these steps to create an IBM i partition.

To begin, navigate to the Servers panel and continue with the following steps:

1. Select the check box for your managed system.
2. In the Name field for your managed system, click the button. Then select **Configuration** → **Create Logical Partition** → **AIX or Linux**.
3. In the Create Partition panel, specify the ID and name of your partition. In this example, for Partition ID, we enter `10`. For Partition name, we enter `CLIENT1`. Then click **Next**.

4. In the Partition Profile panel, enter your Profile name. In this example, we enter `Normal`. Then click **Next**.
5. In the Processors panel, ensure that the **Shared** option is selected, and then click **Next**.
6. In the Processor Settings panel, complete the following steps:
   a. For Desired processing units, enter `0.4`.
   b. For Maximum processing units, enter `10`.
   c. For Desired virtual processors, enter `4`.
   d. For Desired maximum processors, enter `10`.
   e. Select the **Uncapped** check box.
7. In the Memory Settings panel, complete the following fields:
   a. For Minimum Memory, enter `1 GB`.
   b. For Desired Memory, enter `16 GB`.
   c. For Maximum Memory, enter `24 GB`.
8. In the I/O panel, click **Next**.
9. In the Virtual Adapters panel, for the Maximum virtual adapters setting, enter `50`.
10. To create virtual Ethernet adapters, in the Virtual Adapters panel:
    a. Select **Actions** → **Create Virtual Adapter** → **Ethernet Adapter**.
    b. In the Create Virtual Ethernet Adapter panel, click **OK**.

    The virtual Ethernet adapter is created and displayed in the Virtual Adapters panel.
11. To create the virtual SCSI adapter, in the Virtual Adapters panel:
    a. Select **Actions** → **Create Virtual Adapter** → **SCSI Adapter**.
    b. In the Create Virtual SCSI Adapter panel, complete the following steps:
       i. Select the **Only selected client partition can connect** check box.
       ii. For Adapter, enter `11`.
       iii. For Server partition, enter `1`.
       iv. For Server adapter ID, enter `101`.
       v. Click **OK** to accept the settings.

    The virtual SCSI adapter is created and is displayed in the Virtual Adapters panel.
12. Depending on your VIOS partition setup, choose one of the following options:
    – For a single VIOS partition setup, skip to step 13.
    – For dual VIOS partition setup, create an additional virtual SCSI adapter to map to VIOS2 virtual server SCSI adapter. To begin, select **Actions** → **Create Virtual Adapter** → **SCSI Adapter**. Then in the next panel, complete the following steps:
       i. Select the **Only selected client partition can connect** check box.
       ii. For Adapter, enter `21`.
       iii. For Server partition, enter `2`.
       iv. For Server adapter ID, enter `101`.
       v. Click **OK** to accept the settings.

       The virtual SCSI adapter is created and is displayed in the Virtual Adapters panel.
13. In the Virtual Adapter panel, which shows the virtual adapters you created, click **Next**.
14. For the remaining panels, click **Next** until you reach the Profile Summary panel.
15. In the Profile Summary panel, click **Finish**.
16. Click your managed system to view the partition profile you created.

3.2 Setting up a dual VIOS using HMC

The benefit of a dual VIOS setup is that it promotes redundancy, accessibility, and serviceability (RAS). It also offers load balancing capabilities for MPIO and for multiple SEA configuration setups.

A single VIOS setup and a dual VIOS setup have the following differences:

- The additional VIOS partition
- The additional virtual Ethernet adapter used as the SEA Control Channel adapter per VIOS
- Setting the trunk priority on the virtual Ethernet adapters used for bridging to physical adapters in an SEA configuration

  The trunk priority determines which virtual adapter is to be the primary adapter in an SEA failover setup.

On the client partition, no additional virtual Ethernet adapters are required. However, an additional client virtual SCSI adapter is required because of the one-to-one mapping to the server virtual SCSI adapter.

Tip: Each VIOS is configured with a server virtual SCSI adapter for the client partition.

If you are planning to update from a single VIOS setup to a dual VIO setup, you can use DLPAR operations to avoid interfering with the operations of a deployed VIOS or client partition.

The dual VIOS setup expands on the single VIOS setup. The following steps to create a dual VIOS setup use the adapter ID allocation in Table 1-2 on page 5 for VIOS1 and in Table 1-3 on page 5 for VIOS2.

If you completed the single VIOS setup and are looking to change to a dual VIOS setup, keep the adapter IDs as consistent as possible.

You might have your own adapter ID scheme that you want to use, or you can use the default adapter ID allocations.

3.2.1 Creating dual VIOS partition profiles

Follow these steps to create partition profiles for VIOS1 and VIOS2:

1. Depending on whether the VIOS1 partition exists, choose one of the following options:
   - If the VIOS1 partition does not exist, skip to step 4 on page 33.
   - If the VIOS1 partition exists, update the VIOS1 partition profile to add the virtual Ethernet adapter for SEA failover:
     i. By using the DLPAR method, select the check box of the client partition.
     ii. In the Name field for your client partition, click the button. Then select **Dynamic Logical Partitioning** → **Virtual Adapters**.
     iii. Select **Actions** → **Create Virtual Adapter** → **Ethernet Adapter**.
     iv. In the Create Virtual Ethernet Adapter panel, change the Port Virtual Ethernet field to 99. Then click **OK**.

The virtual Ethernet adapter is created with Adapter ID 3 and is displayed in the Virtual Adapters panel with virtual Ethernet for bridging with Adapter ID 2 and server virtual SCSI adapter with Adapter ID 101 (Figure 3-5).

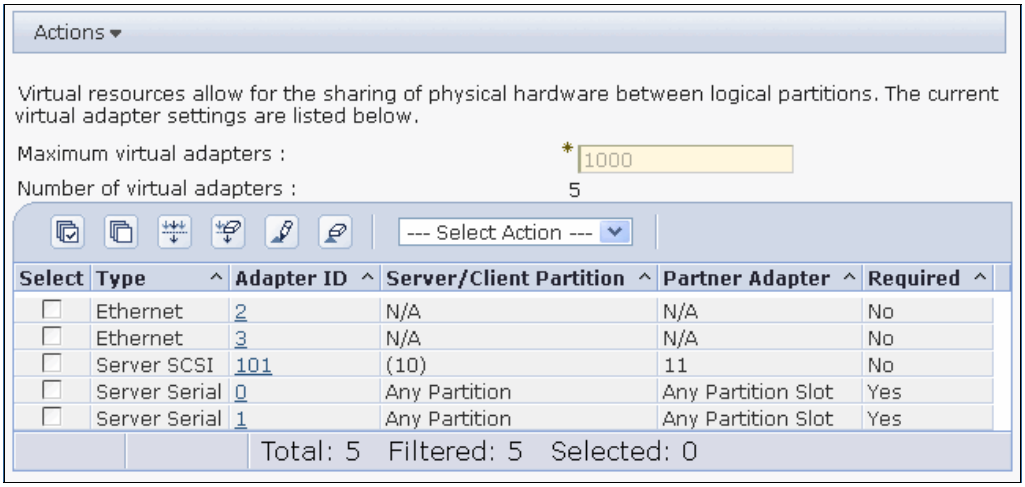

Figure 3-5 Virtual Adapters panel showing virtual adapters created

v. Click **OK** to dynamically add the virtual Ethernet adapter.

2. To save the DLPAR updates to VIOS1 to its profile, click the button at the end of the client partition, and select **Configuration** → **Save Current Configuration**.

DLPAR failure: DLPAR relies on RMC connectivity between the HMC and VIOS1. If DLPAR fails, use steps a - h on page 40 as a reference to create the virtual Ethernet adapter for the SEA failover.

3. In the Save Partition Configuration panel, click **OK**. Go to step 5 to create the VIOS2 partition profile.
4. To create the VIOS1partition profile, follow the steps in 3.1.1, “Creating a VIOS partition profile” on page 18.
5. To create the VIOS2 partition profile, follow the steps in 3.1.1, “Creating a VIOS partition profile” on page 18.

Important changes to VIOS2: VIOS2 requires the following changes:

- Ensure that the priority of the virtual Ethernet adapters used for bridging is different. By default, VIOS1 is created with a priority of 1. For VIOS2, when the virtual Ethernet adapter used for bridging is created in step 11 on page 21, the priority is set to 2.
- For VIOS1 and VIOS2, ensure that the virtual Ethernet adapters used for SEA failover in step 12 on page 21 are created with the same Port VLAN ID. This step is essential for inter-VIOS communication.
- For VIOS2, ensure the virtual SCSI adapter in step 13 on page 22 is created with a different client adapter ID than VIOS1 by using the following settings:
  - For Adapter, enter 101.
  - For Client partition, enter 10.
  - For Client adapter ID, enter 22.

Requirements for using virtual Ethernet adapters as SEA failover adapters: For virtual Ethernet adapters to be used as SEA failover adapters, keep in mind the following requirements:

- Ensure that the Port Virtual Ethernet ID (also known as VLAN ID) is consistent for VIOS1 and VIOS2.
- Ensure that the Port Virtual Ethernet ID is not a known VLAN ID in the network.
- Ensure that the virtual Ethernet adapter is not configured for IEEE 802.1Q.
- Ensure that the virtual Ethernet adapter is not bridged to a physical adapter.

VIOS1 and VIOS2 are now displayed in your system server listing. Their partition profiles are ready to use for the installation process.

3.2.2 Installing a VIOS

The two VIOS partitions can be installed by using the options in 3.1.2, “Installing a VIOS” on page 23.

3.2.3 Configuring VIOS partitions for a dual setup

In this section, you configure the client network and client storage partitions for a dual VIOS configuration.

Tip: If you installed VIOS partition that is not yet configured, skip to “Configuring the VIOS for a client network” on page 35.

Updating the VIOS for a dual configuration

If you configured VIOS1 as a single VIOS setup and you need to update it to serve in a dual VIOS setup, complete these steps:

1. Log on to the VIOS1.
2. To configure the virtual Ethernet adapter created in step – on page 32 on VIOS1, enter **`cfgdev`**.
3. List the virtual Ethernet devices configured on VIOS1:

   ```
   lsdev -vpd | grep Virtual | grep ent
   ```

   Example 3-12 shows the output of the command on VIOS1. Device ent6 is the SEA failover virtual Ethernet adapter created in step – on page 32.

   Example 3-12 VIOS1 virtual Ethernet adapters

   ```
   $ lsdev -vpd | grep Virtual | grep ent
     ent6             U8233.E8B.061AB2P-V1-C3-T1 Virtual I/O Ethernet Adapter (l-lan)
     ent4             U8233.E8B.061AB2P-V1-C2-T1 Virtual I/O Ethernet Adapter (l-lan)
   ```

4. Query the ent device that is the SEA:

   ```
   lsdev -type sea
   ```

Example 3-13 shows the output for VIOS1.

Example 3-13 List of the SEAs configured on VIOS1

```
$ lsdev -type sea
name             status      description
ent5             Available   Shared Ethernet Adapter
```

5. Update the SEA to add the SEA failover functions:

   ```
   chdev -dev ent5 -attr ctl_chan=ent6 ha_mode=auto
   ```

   Where:

   ent5 The SEA in step 4.
   ent6 The SEA failover virtual Ethernet Adapter in step Example 3-12.

Configuring the VIOS for a client network

To configure the VIOS2 partition:

1. Log on to the VIOS terminal window.
2. List the Ethernet devices configured on the VIOS showing the logical name relationship to the physical device details:

   ```
   lsdev -vpd | grep ent
   ```

   Example 3-14 shows the listing of Ethernet devices that you see after running the **lsdev** command.

   Example 3-14 Listing of Ethernet devices on the VIOS

   ```
   $ lsdev -vpd | grep ent
     ent5 U8233.E8B.061AB2P-V2-C3-T1 Virtual I/O Ethernet Adapter (l-lan)
     ent4 U8233.E8B.061AB2P-V2-C2-T1 Virtual I/O Ethernet Adapter (l-lan)
     ent0 U78A0.001.DNWHZS4-P1-C3-T1 4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent1 U78A0.001.DNWHZS4-P1-C3-T2 4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent2 U78A0.001.DNWHZS4-P1-C3-T3 4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent3 U78A0.001.DNWHZS4-P1-C3-T4 4-Port 10/100/1000 Base-TX PCI-Express Adapter
   ```

   In Example 3-14, note the following explanation:

   ent0 (U78A0.001.DNWHZS4-P1-C2-T1)
   The physical Ethernet adapter port cabled.

   U78A0.001.DNWHZS4-P1-C2
   The Ethernet adapter selected in Figure 3-2 on page 20.

   ent4 (U8233.E8B.061AB2P-V1-C2-T1)
   The virtual Ethernet adapter shown in Figure 3-5 on page 33.

   ent5 (U8233.E8B.061AB2P-V2-C3-T1)
   The virtual Ethernet adapter also shown in Figure 3-5 on page 33.

   If you plan to use 802.3ad Link Aggregation, the adapters must be cabled and the network switch ports must be configured for 802.3ad Link Aggregation. To create the Link Aggregation adapter, enter the following command:

   ```
   mkvdev -lnagg <entX> <entY> -attr mode=8023ad
   ```

   Alternatively, use the **cfgassist** command as explained in the following steps:

   a. On the command line, enter **cfgassist**.

   b. Select **Devices** → **Link Aggregation Adapter** → **Add a Link Aggregation Adapter**.

c. In the Target Adapters field, enter the physical network adapters (with spaces between each physical network adapter).

d. In the ATTRIBUTES field, enter `mode=8023ad`.

List all physical Ethernet adapters and EtherChannel adapters that are available for creating an SEA:

```
lsdev -type ent4sea
```

3. Create the SEA, which bridges the physical adapter and the virtual adapter:

```
mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 1 -attr ctl_chan=ent5
ha_mode=auto
```

Where:

ent0 The physical adapter in step 2 on page 35. (Use an EtherChannel adapter if one was created for the SEA configuration.)

ent4 The bridging virtual adapter in step 2 on page 35.

1 The port VLAN ID of ent4.

ent5 The SEA failover virtual adapter in step 2 on page 35.

Example 3-15 shows the SEA virtual network devices that are created, where:

ent6 An Ethernet network adapter device.
en6 A standard Ethernet network interface where TCP/IP addresses are assigned.
et6 An IEEE 802.3 Ethernet network interface.

Example 3-15 Creating an SEA interface

```
$ mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 1 -attr ctl_chan=ent5
ha_mode=auto
ent6 Available
en6
et6
```

4. Configure the TCP/IP connection for the VIOS with details provided by the network administrator:

```
mktcpip -hostname vios1 -interface en6 -inetaddr 172.16.22.15 -netmask
255.255.252.0 -gateway 172.16.20.1
```

Alternatively, you can use the **`cfgassist`** command:

a. Enter **`cfgassist`**.

b. Select **VIOS TCP/IP Configuration**.

c. Select **en5**, which is the SEA interface created in step 3 on page 27, and then click **Enter**.

d. Enter the TCP/IP details shown in the following list.

Regardless of the method you choose, in this example, we use the following details:

– Network IP address: 172.16.22.15
– Network subnet: 255.255.252.0
– Network gateway: 172.16.20.1

Configuring a VIOS for client storage

You might encounter a scenario where the SAN administrator allocated a group of disks, but only one of the disks is planned for the client partition. For this exercise, the SAN administrator allocated the disks with the LUN ID `60:0a:0b:80:00:11:46:32:00:00:92:78:4e:c5:0f:0b`.

To configure and map the disks to the target client partition:

1. List any Fibre Channel adapter SCSI protocol devices that are configured on the VIOS:

   ```
   lsdev | grep fscsi
   ```

   In Example 3-16, `fscsi0` and `fscsi1` are the Fibre Channel adapter SCSI protocol devices configured on VIOS1. Their attributes are updated to allow for dynamic tracking and fast failover (applicable for a multiple Fibre Channel adapter VIOS).

   Example 3-16 List of Fibre Channel adapter SCSI protocol devices on VIOS1

   ```
   $ lsdev | grep fscsi
   fscsi0          Available   FC SCSI I/O Controller Protocol Device
   fscsi1          Available   FC SCSI I/O Controller Protocol Device
   ```

2. Update the Fibre Channel adapter SCSI protocol device attributes listed in step 5 on page 35 to enable dynamic tracking and fast failover:

   ```
   chdev -dev fscsi0 -attr dyntrk=yes fc_err_recov=fast_fail
   chdev -dev fscsi1 -attr dyntrk=yes fc_err_recov=fast_fail
   ```

   Busy SCSI protocol device: If the SCSI protocol device of the Fibre Channel adapter is busy, append the **`-perm`** flag to the command, as shown in the following example, to update the VIOS database only. The attributes are not applied to the device until the VIOS is rebooted.

   ```
   chdev -dev fscsi0 -attr dyntrk=yes fc_err_recov=fast_fail -perm
   ```

3. To configure the disks on the VIOS, enter **`cfgdev`**.
4. List the disks on the VIOS partition and show the disk type:

   ```
   lsdev -type disk
   ```

   In Example 3-17, VIOS1 lists two internal SAS disks and six DS4800 disks.

   Example 3-17 List disks with their type on VIOS1

   ```
   $ lsdev -type disk
   name            status      description
   hdisk0          Available   SAS Disk Drive
   hdisk1          Available   SAS Disk Drive
   hdisk2          Available   MPIO DS4800 Disk
   hdisk3          Available   MPIO DS4800 Disk
   hdisk4          Available   MPIO DS4800 Disk
   hdisk5          Available   MPIO DS4800 Disk
   hdisk6          Available   MPIO DS4800 Disk
   hdisk7          Available   MPIO DS4800 Disk
   ```

5. To confirm the SAN LUN ID on VIOS1, enter the following command for each disk listed in step 2 on page 35 until the correct disk is found with LUN ID provided by the SAN administrator:

   ```
   lsdev -dev hdiskX -attr | grep -i -E "reserve|unique_id"
   ```

Example 3-18 shows the hdisk that the SAN administrator assigned. Also, the SCSI reserve policy was set with `single_path`, which must be updated with no SCSI reserve locks. The LUN ID is embedded in the `unique_id` string for hdisk6.

Example 3-18 Disk attributes of hdisk6

```
$ lsdev -dev hdisk6 -attr | grep -E "unique_id|reserve"
reserve_policy  single_path Reserve Policy                   True
unique_id       3E213600A0B8000114632000092784EC50F0B0F1815      FAStT03IBMfcp Unique
device identifier         False
```

Additional information: Disks using EMC PowerPath, IBM SDDPCM, and IBM SDD drivers also have their LUN IDS embedded in the unique_id string. Use their supplied commands to show the LUN IDS in a more readable format. To obtain the disks complete with LUN IDs, see their manuals.

EMC disks are displayed with the hdiskpowerX notation, and SDD disks are displayed with vpathX notation. Use these disk notations with the **`lsdev`** command sequence instead of hdisk.

Other disks subsystems might use different fields to set their SCSI reserve locks. Use the **`lsdev`** command sequence without the pipe to grep as in the following example:

```
lsdev -dev sampledisk -attr
```

6. For users with EMC disks and disks using native MPIO: Deactivate the SCSI reserve lock on the disk, which is hdisk6 in this example:

   ```
   chdev -dev hdisk6 -attr reserve_policy=no_policy
   ```

   Disks using SDDPCM and SDD drivers: Ignore the following step if the disks are using SDDPCM and SDD drivers because the SCSI reserve locks are already deactivated.

   SCSI reserve lock: The SCSI reserve lock attribute differs among disk subsystems. The IBM System Storage SCSI reserve lock attribute is `reserve_policy` as displayed in Example 3-8. The attribute on EMC disk subsystem is `reserve_lock`.

   If you are unsure of the allowable value to use to deactivate the SCSI reserve lock, the following command provides a list of allowable values:

   ```
   lsdev -dev hdisk6 -range reserve_policy
   ```

7. Determine the virtual adapter name of the virtual SCSI adapter in step 13 on page 22:

   ```
   lsdev -vpd | grep "Virtual SCSI"
   ```

   In Example 3-19, the virtual SCSI adapter with server adapter ID C101 is vhost0 to use in the next step.

   Example 3-19 List of virtual SCSI devices

   ```
   $ lsdev -vpd | grep "Virtual SCSI"
     vhost0           U8233.E8B.061AB2P-V1-C101 Virtual SCSI Server Adapter
   ```

   The MPIO setup is used to map whole LUNS to client OS partitions.

8. To map hdisk hdisk6 to CLIENT1, enter:

   ```
   mkvdev -vdev hdisk6 -vadapter vhost0
   ```

Where:

hdisk6 The disk found in step 5 on page 28.
vhost0 The virtual server SCSI adapter found in step 6 on page 29.

In Example 3-20, the virtual target device (VTD) `vtscsi0` is created.

Example 3-20 Creating a disk mapping to a client partition

```
$ mkvdev -vdev hdisk6 -vadapter vhost0
vtscsi0 Available
```

9. To check mapped devices to vhost0, enter:

   `lsmap -vadapter vhost0`

   In Example 3-21, the `vhost0` virtual SCSI adapter shows one disk mapped, where hdisk6 is mapped to the `vtscsi0` virtual target device.

Example 3-21 Disk mapping for vhost0

```
$ lsmap -vadapter vhost0
SVSA            Physloc                                     Client Partition ID
--------------- ------------------------------------------- ------------------
vhost0          U8233.E8B.061AB2P-V1-C101                   0x0000000a

VTD                   vtscsi0
Status                Available
LUN                   0x8100000000000000
Backing device        hdisk6
Physloc               U5802.001.0087356-P1-C2-T1-W202200A0B811A662-L5000000000000
Mirrored              false
```

10. Repeat step 1 on page 35 through step 9, to configure the VIOS2 partition. For step 1 on page 35, ensure that you log on to the VIOS2 terminal window.

3.2.4 Creating a logical partition profile for the client operating system

Create the client partition profile as explained in 3.1.4, “Creating a logical partition profile for a client operating system” on page 30.

Alternatively, if the client partition profile exists and you want to configure an additional virtual SCSI adapter, you can choose one of the following methods:

- Add the virtual SCSI adapter by using DLPAR. Then save the current configuration, overwriting the current profile. (The client partition must be running and have RMC connectivity to the HMC.)
  a. Select your client partition.
  b. Click the button at the end of your client partition, and then select **Dynamic Logical Partitioning** → **Virtual Adapters**.
  c. Select **Actions** → **Create Virtual Adapter** → **SCSI Adapter**.
  d. In the Create Virtual SCSI Adapter panel, complete the following steps:
     i. Select the **Only selected client partition can connect** check box.
     ii. For Adapter, type 22.
     iii. For Server partition, type 2.
     iv. For Server adapter ID, type 101.
     v. Click **OK** to accept the settings.

e. Click **OK** to dynamically add the virtual SCSI adapter.

f. Click the button at the end of your client partition, and then select **Configuration** → **Save Current Configuration**.

g. In the Save Partition Configuration panel, click **OK**.

h. Click **Yes** to confirm the save.

- Update the client partition profile to add the additional virtual SCSI adapter. Then shut down the client partition (if it is running), and activate the client partition.

Important: Shutting down the client partition and then activating it causes the client partition to re-read its profile. A partition restart does not re-read the partition profile.

a. Click the button at the end of your client partition, and then select **Configuration** → **Manage Profiles**.

b. Click the profile to update.

c. Click the **Virtual Adapters** tab

d. Select **Actions** → **Create Virtual Adapter** → **SCSI Adapter**.

e. In the Create Virtual SCSI Adapter panel, complete the following steps:

i. Select the **Only selected client partition can connect** check box.
ii. In the Adapter field, type `22`.
iii. In the Server partition field, type `2`.
iv. In the Server adapter ID field, type `101`.
v. Click **OK** to accept settings.

f. Click **OK** to save the profile.

g. Run the shutdown command on the client partition.

h. After the client partition is displayed with the `Not Activated` state, activate the client partition.

3.3 Setting up a virtual Fibre Channel using HMC

With Virtual Fibre Channel, disks can be assigned directly to the client partitions from the SAN storage system. With virtual SCSI, the disks are assigned to the VIOS partition before they are mapped to a virtual SCSI adapter.

For a client partition operating system disk, use virtual SCSI, and for the data, use virtual Fibre Channel. You want to use virtual SCSI for client partition OS disks for the following reasons:

- When using Fibre Channel LUNs for client storage, an advantage is ensuring that they are visible from the VIOS before assigning them to a client. After the client partition is loaded, the mapped LUNs can be checked from the client. By using virtual Fibre Channel, the LUNs cannot be determined until the client partition is loaded from an installation source.
- Operating systems, such as AIX and Linux, have their kernels running in memory. If serious SAN issues are experienced, the VIOS first detects the problem and severs the link to the client partition. The client partition halts abruptly, reducing any risk to data corruption. With OSs using virtual Fibre Channel or physical Fibre Channel, the partition remains running for some time. During that time, the client partition is susceptible to data corruption.
- Operating system disks using virtual SCSI do not rely on external device drivers, but OS disks using virtual Fibre Channel do rely on them. When upgrading the external device drivers, the client partitions must follow special procedures.

A server virtual Fibre Channel adapter has a one-to-one relationship with a client virtual Fibre Channel adapter.

For a single VIOS setup, the VIOS is configured with one virtual Fibre Channel adapter, and the client partition is configured with one virtual Fibre Channel adapter mapped to each other.

For a dual VIOS setup, each VIOS is configured with one virtual Fibre Channel adapter, and the client partition is configured with two virtual Fibre Channel adapters mapped to each VIOS virtual Fibre Channel adapter.

You can add the virtual Fibre Channel adapters to existing partition profiles similarly to the virtual SCSI adapters as explained in 3.2.4, “Creating a logical partition profile for the client operating system” on page 39.

Use the DLPAR method as explained previously to add the server virtual Fibre Channel adapter to the VIOS and to add the corresponding client virtual Fibre Channel adapter to the client partition.

Updating a VIOS partition profile

To update a VIOS partition profile:

1. On the HMC, select the VIOS1 partition.
2. Click the button at the end of your client partition, and then select **Dynamic Logical Partitioning** → **Virtual Adapters** (Figure 3-6).

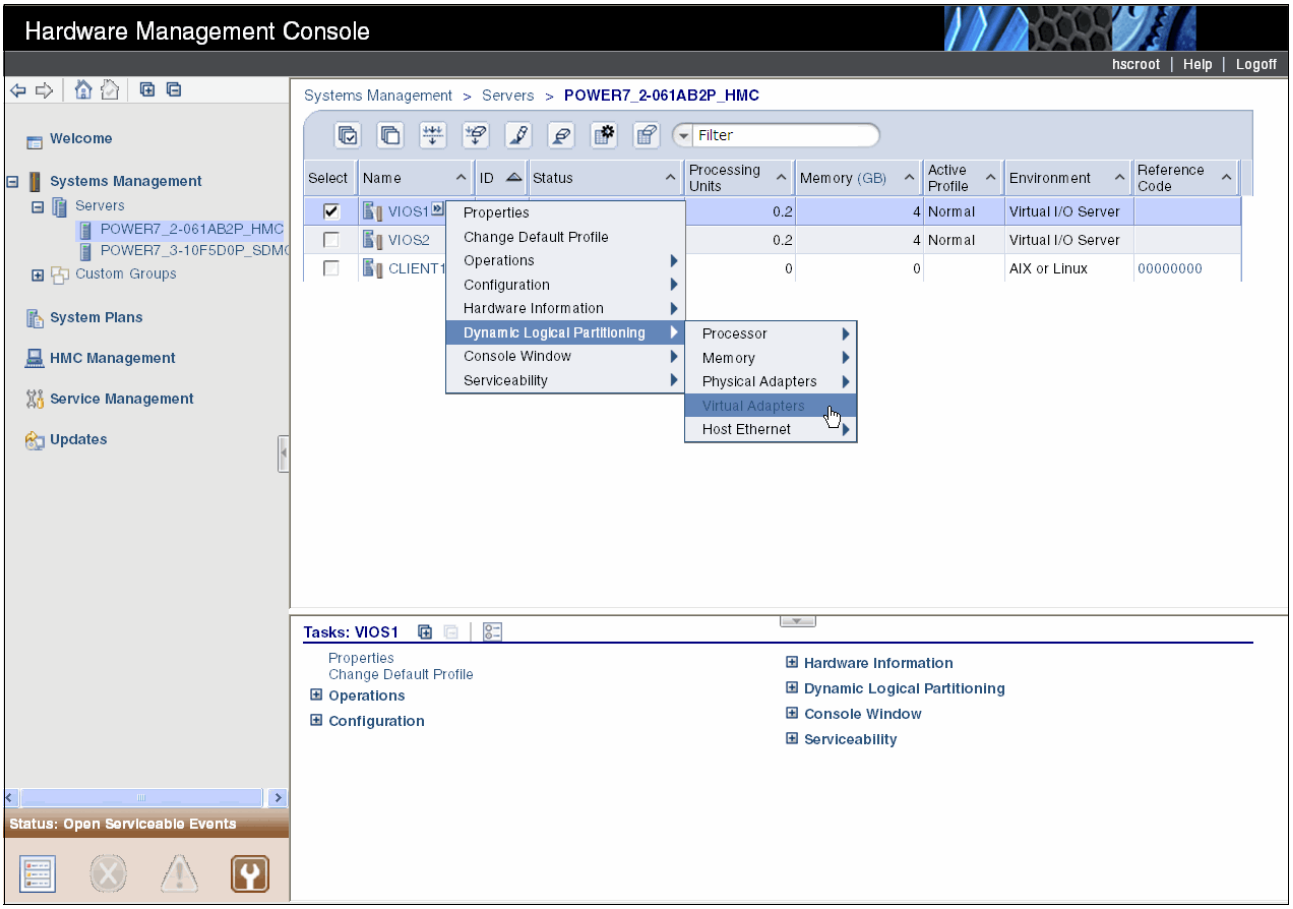

Figure 3-6 Creating a virtual adapter using DLPAR

3. Select **Actions** → **Create Virtual Adapter** → **Fibre Channel Adapter**.
4. In the Create Virtual Fibre Channel Adapter panel, complete the following steps:
   a. In the Adapter field, type 102.
   b. In the Client partition field, type 10.
   c. In the Client adapter ID field, type 12.
   d. Click **OK** to accept the settings.
5. Click **OK** to dynamically add the virtual Fibre Channel adapter.
6. Click the » button at the end of your client partition, and then select **Configuration** → **Save Current Configuration** (Figure 3-7).

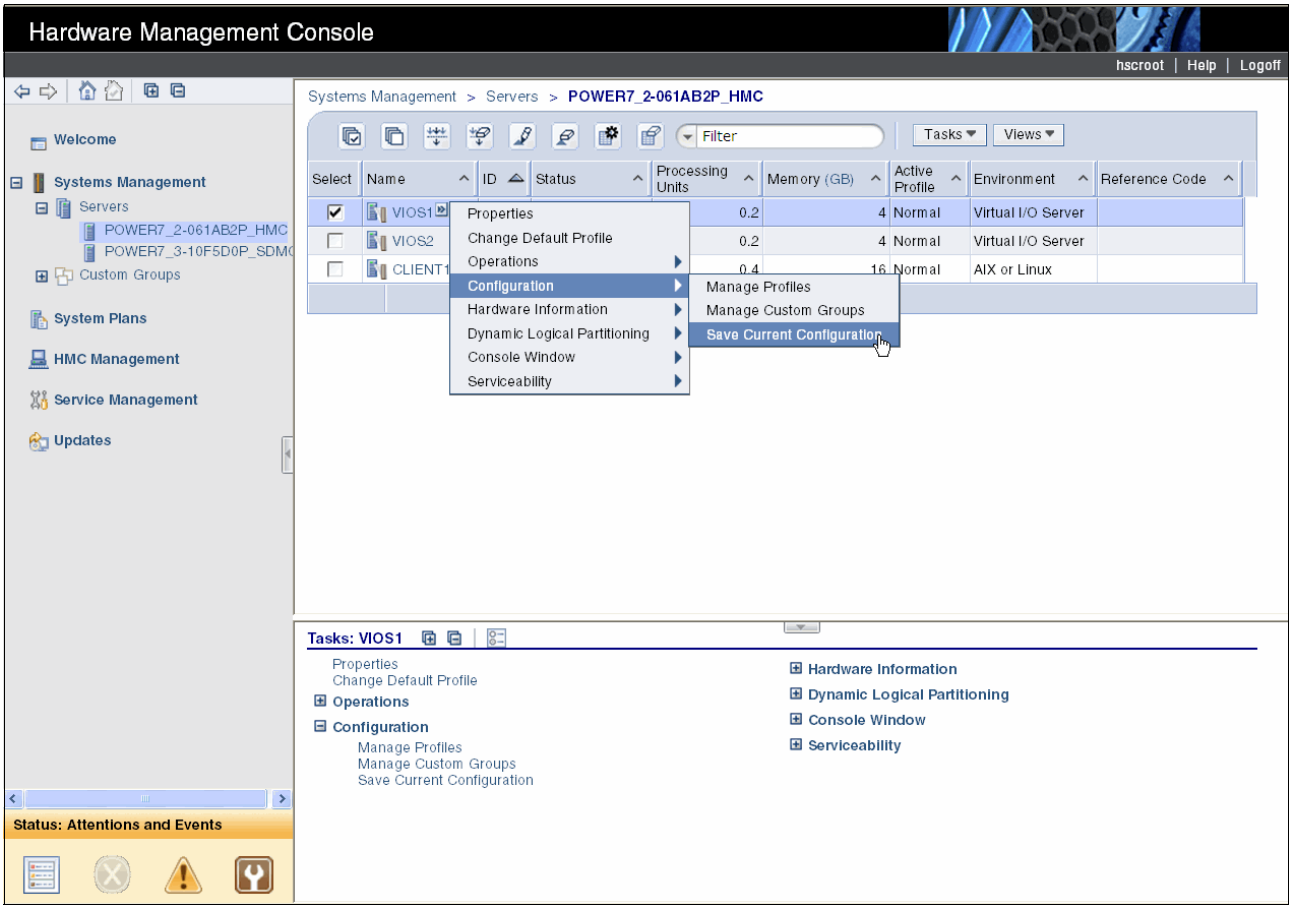

Figure 3-7 Saving the current configuration to a profile

7. In the Save Partition Configuration panel, click **OK**.
8. Click **Yes** to confirm the save, overwriting the existing profile.

Updating a client partition profile

To update a client partition profile:

1. On the HMC, select the client partition.
2. Click the » button at the end of your client partition, and then select **Dynamic Logical Partitioning** → **Virtual Adapters**.
3. Select **Actions** → **Create Virtual Adapter** → **Fibre Channel Adapter**.

4. In the Create Virtual Fibre Channel Adapter panel, complete the following steps:
   a. In the Adapter field, type 12.
   b. In the Server partition field, type 1.
   c. In the Server adapter ID field, type 102.
   d. Click **OK** to accept settings
5. In the Virtual Adapters panel, which shows the virtual Fibre Channel adapters that were created (Figure 3-8), click **OK** to dynamically add the virtual Fibre Channel adapter.

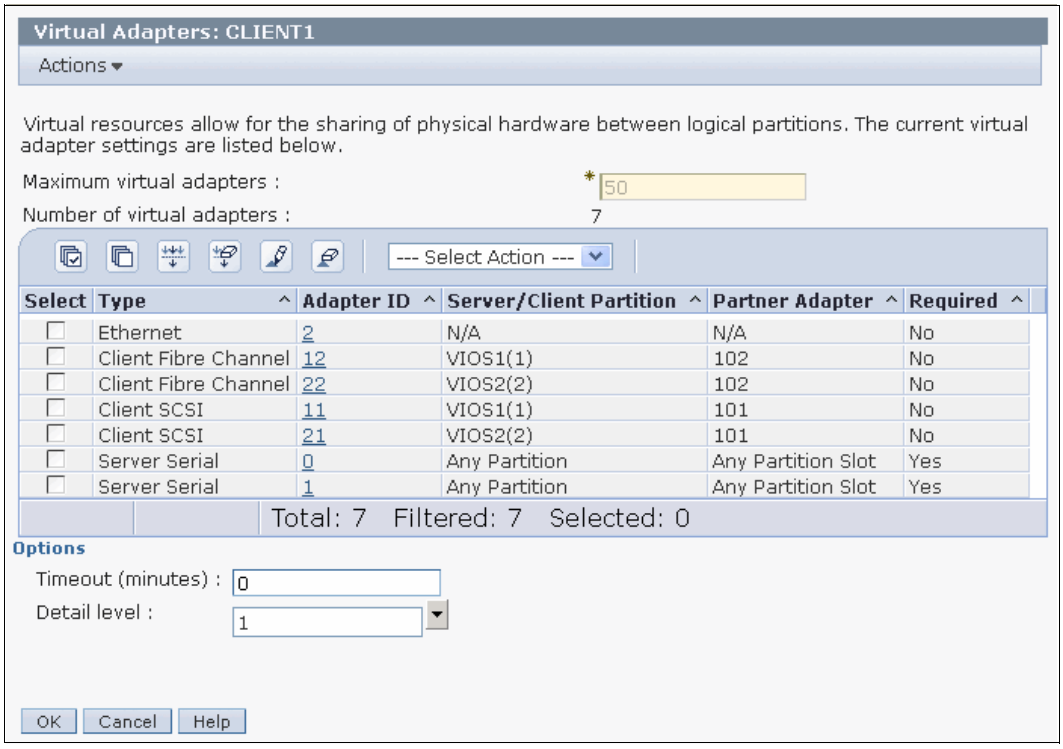

Figure 3-8 Client Virtual Adapters panel

6. Click the » button at the end of your client partition, and then select **Configuration** → **Save Current Configuration**.
7. In the Save Partition Configuration panel, click **OK**.
8. Click **Yes** to confirm the save.

Mapping virtual Fibre Channel adapters on the VIOS

To map virtual Fibre Channel adapters on the VIOS:

1. Log on to the VIOS partition.
2. To list the physical Fibre Channel adapters on the VIOS, enter the **`lsnports`** command.

   In Example 3-22, `fcs0` is the first adapter port zoned for allocation of disks for virtual SCSI. Also `fcs1` is the second adapter port zoned for NPIV, which is planned in 3.1, “Setting up a single VIOS using an HMC” on page 18.

Example 3-22 List of physical Fibre Channel adapter ports on the VIOS

```
$ lsnports
name             physloc                        fabric tports aports swwpns  awwpns
fcs0             U5802.001.0087356-P1-C2-T1          1     64     64   2048    2046
fcs1             U5802.001.0087356-P1-C2-T2          1     64     64   2048    2048
```

3. To configure the virtual Fibre Channel adapter that was added by using DLPAR in step 4 on page 42, enter the **cfgdev** command.
4. List the virtual Fibre Channel adapters:

   ```
   lsdev -vpd | grep vfchost
   ```

   In Example 3-23, one virtual Fibre Channel adapter is listed with an adapter slot ID of 102 (C102) created in step 4 on page 42, called *vfchost0*.

   Example 3-23 List of virtual Fibre Channel adapters on the VIOS

   ```
   $ lsdev -vpd | grep vfchost
     vfchost0         U8233.E8B.061AB2P-V1-C102 Virtual FC Server Adapter
   ```

5. Map the client virtual Fibre Channel adapter to the physical Fibre Channel adapter that is zoned for NPIV:

   ```
   vfcmap -vadapter vfchost0 -fcp fcs1
   ```

   Tip: You can map multiple client virtual Fibre Channel adapters to a physical Fibre Channel adapter port. Up to 64 client virtual Fibre Channel adapters can be active at one time per physical Fibre Channel adapter.

6. Verify the virtual Fibre Channel mapping for vfchost0:

   ```
   lsmap -vadapter vfchost0 -npiv
   ```

   Alternatively, list all virtual Fibre Channel mapping on the VIOS:

   ```
   lsmap -all -npiv
   ```

7. If you have a dual VIOS setup, repeat step 1 on page 43 through step 5 for VIOS2. Ensure that the client partition adapter IDs are unique.

3.4 Adding additional client partitions

If you have additional client partitions to add, complete the following steps:

1. Create the virtual VSCI and virtual Fibre Channel adapters on the VIOS, as explained in step 1 on page 41 through step 8 on page 42. Repeat this procedure for the other VIOS if you have a dual VIOS setup.
2. Create the client partition profile, as explained in 3.1.4, “Creating a logical partition profile for a client operating system” on page 30.

   Reminder: Add virtual Fibre Channel adapters if they are required.

3. Configure the VIOS virtual SCSI adapter mapping to the client virtual SCSI adapter, as explained in step 3 on page 37 through step 9 on page 39. Repeat this procedure for the other VIOS if you have a dual VIOS setup.
4. Configure the VIOS virtual Fibre Channel adapter mapping to the client virtual Fibre Channel adapter, as explained in step 1 on page 43 through step 7. Repeat this procedure for the other VIOS if you have a dual VIOS setup.

4

Configuring PowerVM with the IBM Systems Director Management Console

With the IBM Systems Director Management Console (SDMC), system administrators can manage IBM Power System servers and BladeCenter blades. The SDMC organizes tasks in a single panel that simplifies views of systems and day-to-day tasks. The SDMC can also be integrated into the administrative framework of IBM Systems Director.

This chapter includes the following sections:

- Implementing a dual VIOS setup using the SDMC

  This section explains how to create a dual VIOS configuration for PowerVM and how to have the SDMC manage slot allocations for all the virtual adapters.

- Setting up a single VIOS using the SDMC

  This section explains how to set up PowerVM and a virtual server that are prepared for installing the client operating system. It follows the scheme for virtual adapter ID numbers as explained in 4.1, “Implementing a dual VIOS setup using the SDMC” on page 46.

  Virtual server: Virtual server is the SDMC term for a virtualized guest server running its own operating system.

- Setting up a dual VIOS using the SDMC

  This section explains how to set up PowerVM with an additional VIOS for added redundancy, accessibility, and serviceability (RAS).

- Setting up a virtual Fibre Channel using the SDMC

  This section explains how to set up PowerVM with virtual Fibre Channel functionality, by configuring virtual Fibre Channel adapters on the virtual server.

4.1 Implementing a dual VIOS setup using the SDMC

You can create a PowerVM dual Virtual I/O Server (VIOS) configuration by using the SDMC to automatically handle the allocation of slot numbers for virtual adapters. The final configuration uses the following functions of PowerVM:

- Virtual SCSI adapters to provide the disk for the virtual server OS
- Virtual Ethernet adapters to provide LAN connectivity to a client OS partition
- A highly available virtual Ethernet connection using Shared Ethernet Adapter (SEA) failover functions
- Virtual Fibre Channel adapters to provide disk for a virtual server (for data storage)

All the virtual SCSI adapters and virtual Fibre Channel adapters are dual to both VIOSs, enabling the use of multiple paths to the virtualized disks in the client OS.

Implementing a guided dual VIOS setup using the SDMC entails the following tasks:

1. Creating the virtual servers for VIOS1 and VIOS2
2. Installing VIOS1 and VIOS2
3. Configuring the TCP/IP stack in VIOS1 and VIOS2
4. Creating the SEA failover configuration by using the SDMC
5. Configuring storage devices
6. Creating a virtual server for a client operating system.
7. Installing a client operating system.
8. Configuring virtual Fibre Channel adapters using the SDMC.

Before you begin the tasks in this chapter, you must complete the following tasks:

1. Verify the prerequisites in Chapter 1, “Introduction to PowerVM” on page 1.
2. Verify that the SDMC is installed and configured.
3. Verify that the host is already discovered and visible to the SDMC.
4. Verify that your host is in a *Started* state.

4.1.1 Creating the virtual servers for VIOS1 and VIOS2

In this section, you create two virtual servers for the dual VIOS configuration, one virtual server for VIOS1 and the other virtual server for VIOS2.

Before you create the virtual servers, complete these steps:

1. Log in to the SDMC environment.
2. If the Home panel is not displayed automatically, in the left pane, click **Home**.
3. From the Home panel, locate the host on which to create the virtual server. Select the check box of the host. Then select **Actions** → **System Configuration** → **Create Virtual Server** (Figure 4-1 on page 47).

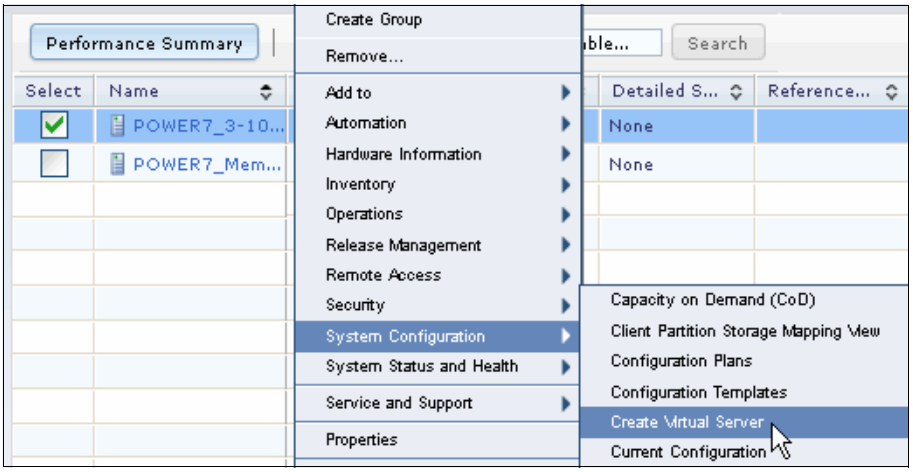

Figure 4-1 Starting the Create Virtual Server task

To create the virtual server for VIOS1:

1. In the Name panel, complete the following steps:
   a. In the Virtual Server name field, enter a name. We enter `VIOS1` in this example.
   b. For Environment, select VIOS.
   c. For the other fields, accept the default values.
   d. Click **Next**.
2. In the Memory panel, select the **Dedicated for Memory Mode** check box (if shown), and enter an appropriate amount of memory in the Assigned memory field. Use 4 GB of memory.

   Amount of memory: The amount of memory your VIOS needs depends on the functions of VIOS that you will use. Start with 4 GB of memory, and then periodically monitor the memory usage on VIOS.

   Click **Next**.
3. In the Processor panel, complete these steps:
   a. Select **Shared for Processing Mode**.
   b. In the Assigned Processors field, type `1` for a single shared processor.

      Assigned Processors field: Start with one shared processor, and then periodically monitor the processor on VIOS.

   c. Click **Next**.

   Virtual Ethernet adapters: By default, the wizard creates two virtual Ethernet adapters. The first virtual Ethernet adapter uses adapter ID 2 and VLAN ID 1. The second virtual Ethernet adapter uses adapter ID 3 and VLAN ID 99. The second virtual Ethernet adapter is used for a control channel between two VIOSs in dual VIOS configurations and the Shared Ethernet Adapter failover configuration. For more information about control channel and dual VIOS configuration for virtual Ethernet, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

4. In the Ethernet panel, complete the following steps:
   a. Expand **Virtual Ethernet**.
   b. Select the check box to the left of the first adapter (ID 2), and then click **Edit**.
   c. In the Virtual Ethernet - Modify Adapter panel, complete these steps:
      i. Select the **Use this adapter for Ethernet bridging** check box.
      ii. In the Priority field, enter `1`.

      For more information about the priorities, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.
      iii. Click **OK** to confirm changes to the first Ethernet adapter.
   d. Back in the Ethernet panel, click **Next**.
5. In the Virtual Storage Adapter panel, click **Next**. Because the client virtual servers are added and assigned storage, the console automatically creates the virtual SCSI or virtual Fibre Channel server adapters.
6. In the Physical I/O Adapters panel, select the check boxes to the left of the Location Code and Description of the adapter that you need. These physical adapters and controllers are used later to virtualize devices to the virtual server for the client OS. To use all functions as explained in this paper, you must make the following selections:
   – One SAS or SCSI disk controller (controller for internal disk drives)
   – One Ethernet adapter (Ethernet adapter for connection to LAN)
   – One Fibre Channel adapter (Fibre Channel adapter for connection to SAN and a virtual Fibre Channel configuration)

   In our case, we selected the following physical adapters:

   ```
   U78A0.001.DNWKF81-P1-T9 RAID Controller
   U5802.001.RCH8497-P1-C7 Quad 10/100/1000 Base-TX PCI-Express Adapter
   U5802.001.RCH8497-P1-C3 Fibre Channel Serial Bus
   ```

   The RAID Controller that we selected also has the physical CD or DVD drive connected.

   Verifying the adapters: Check the physical location codes of your adapters and verify that you are using the correct adapters for your virtual server. Sometimes the description can be misleading. For example, the PCI-to-PCI bridge can be the Ethernet adapter.

   Click **Next**.
7. In the Summary panel, verify the information, and then click **Finish** to confirm the creation of the virtual server.

To create the virtual server for VIOS2, follow step 1 on page 47 - step 7, but change the following values:

- In step 1, change the name for the virtual server to `VIOS2`. For Environment, make sure that you choose **VIOS**.
- In step 4, change the priority of the virtual Ethernet adapter to `2`. Also, select the **Use this adapter for Ethernet bridging** check box.
- In step 6, select the different adapters and controller. We selected the following adapters:

  ```
  U5802.001.RCH8497-P1-C2 PCI-E SAS Controller
  U5802.001.RCH8497-P1-C6 Quad 10/100/1000 Base-TX PCI-Express Adapter
  U5802.001.RCH8497-P1-C5 Fibre Channel Serial Bus
  ```

For a more detailed description of the options when creating a virtual server, see *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860.

4.1.2 Installing VIOS1 and VIOS2

You can choose from three ways to install VIOS1 and VIOS2 into the previously created virtual servers:

- Install VIOS from VIOS DVD media by using the allocated physical DVD drive.
- Install VIOS from the SDMC command line by using the **`installios`** command. The **`installios`** command has a progressive wizard that guides you through the installation process.
- Install VIOS by using the IBM Network Installation Manager (NIM).

To install VIOS1, follow the steps in 4.2.2, “Installing a VIOS” on page 59. To install VIOS2, follow the steps in 4.3.2, “Installing a second VIOS using NIM” on page 67. You can also use the **`installios`** command from the SDMC command line to install both VIOS1 and VIOS2.

4.1.3 Configuring the TCP/IP stack in VIOS1 and VIOS2

Now you assign management IP addresses to VIOS1 and VIOS2. You do this configuration from the terminal console of the VIOS virtual servers.

Before you configure TCP/IP, follow these steps:

1. Log on to the SDMC environment to open the console to the VIOS1 virtual server.
2. From the home page, locate the host that contains the VIOS1 virtual server, and select the name of the host.
3. In the Resource Explorer window that opens, select the check box to the left of the virtual server VIOS1, and then select **Actions** → **Operations** → **Console Window** → **Open Terminal Console**.
4. Enter the password for the SDMC account.

To configure TCP/IP on VIOS1, complete the following steps on the VIOS1 console:

1. Find the device name of the physical Ethernet adapter ports:

   ```
   lsdev -vpd | grep ent | grep -v Virtual
   ```

   Example 4-1 shows the listing.

   Example 4-1 Listing of physical Ethernet adapter ports on VIOS

   ```
   $ lsdev -vpd | grep ent | grep -vi Virtual
     Model Implementation: Multiple Processor, PCI bus
     ent0  U5802.001.RCH8497-P1-C7-T1  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent1  U5802.001.RCH8497-P1-C7-T2  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent2  U5802.001.RCH8497-P1-C7-T3  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent3  U5802.001.RCH8497-P1-C7-T4  4-Port 10/100/1000 Base-TX PCI-Express Adapter
   ```

   Select the correct Ethernet port from the listed ports that is to be used for the LAN connection and has the Ethernet cable plugged in. The interface device name of this physical adapter port is used in the next step. In this example, it is `en0`.

2. Enter the **`cfgassist`** command, and select **VIOS TCP/IP Configuration**. Then select the appropriate interface device name from the previous step.

3. In the VIOS TCP/IP Configuration panel, enter TPC/IP configuration values for VIOS connectivity. For these values, consult your network administrator. Figure 2-2 on page 10 shows the TPC/IP configuration values for this example. After entering the needed values for TCP/IP configuration press Enter.
4. When you see the output `Command: OK`, press F10. Alternatively, press Esc+0.

To configure TCP/IP on VIOS2, follow the same steps on VIOS2 console, but in step 3, change the IP configuration.

SSH: From this point onward, you can use Secure Shell (SSH) to connect to VIOS1 and VIOS2.

4.1.4 Creating the SEA failover configuration by using the SDMC

Next, you set up the SEA failover configuration. In VIOS terminology, the *SEA* is a bridge that connects the physical Ethernet adapter to the virtual Ethernet adapter. In normal operations, the virtual Ethernet connection for the client virtual server is provided by using VIOS1. The failover capability enables the automatic failover to VIOS2 if you need to restart VIOS1, such as after a VIOS update.

Before you create the SEA failover configuration, complete these steps:

1. Log in to the SDMC environment.
2. If the home page is not automatically displayed, click **Home** in the left pane.
3. From the home page, locate the host on which to configure the SEA failover configuration. (This host is the same one that was used in the previous chapters.)
4. Select the check box to the left of the host, and then select **Actions** → **System Configuration** → **Virtual Resources** → **Virtual Network Management** to open the Virtual Network Management panel.

You create the SEA configuration on both VIOS1 and VIOS2. To create an SEA failover configuration for your host:

1. To start the configuration, click **Create SEA**.
2. As shown in Figure 4-2, for VIOS, select **VIOS1(1)**. Select the **Configure network failover** check box. Enter the physical adapter. Then click **OK**.

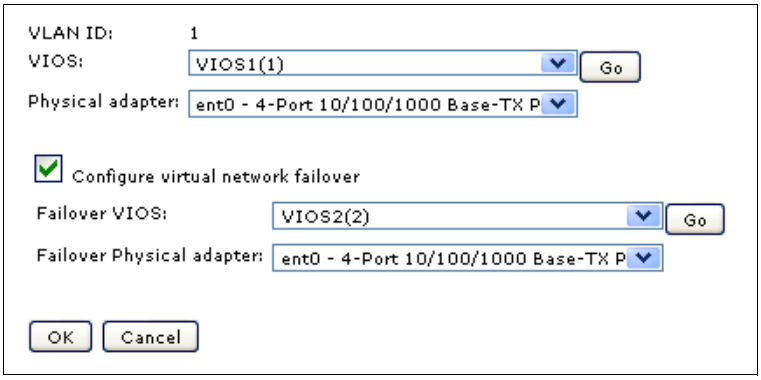

Figure 4-2 The SEA failover configuration using the SDMC

The SDMC automatically creates the SEAs on both VIOS1 and VIOS2. The SDMC also configures the control channel as a part of this step. The virtual Ethernet adapter with the highest VLAN ID is used for the SEA control channel.

3. In the Virtual Network Management panel, confirm the created SEAs. Two SEAs are created, each with a different priority as shown in Figure 4-3.

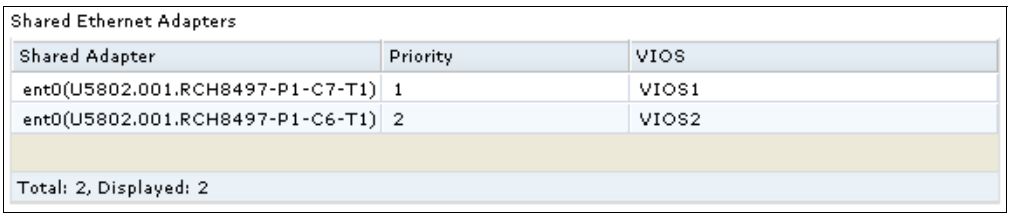
Shared Ethernet Adapters

Shared Adapter	Priority	VIOS
ent0(U5802.001.RCH8497-P1-C7-T1)	1	VIOS1
ent0(U5802.001.RCH8497-P1-C6-T1)	2	VIOS2

Total: 2, Displayed: 2

Figure 4-3 View of the SEA from the SDMC

4.1.5 Configuring storage devices

Now you configure the storage devices on VIOS1 and VIOS2. SAN logical unit number (LUN) is used for installation of the client OS. First you must connect the VIOS virtual server to the SAN and then provide the necessary information to the storage administrator. After the storage administrator provisions the needed SAN LUN, map this SAN LUN over the virtual SCSI adapter to virtual server for the client OS using the SDMC.

To attach VIOS1 to a SAN and configure storage, complete the following steps on the VIOS console:

1. Find the Fibre Channel adapters owned by VIOS1:

   ```
   lsdev -vpd | grep fcs
   ```

   The number of Fibre Channel adapters can vary. You receive a list similar to the following example:

   ```
   fcs0  U5802.001.RCH8497-P1-C3-T1  8Gb PCI Express Dual Port FC Adapter
   fcs1  U5802.001.RCH8497-P1-C3-T2  8Gb PCI Express Dual Port FC Adapter
   ```

   In our case, we used the Fibre Channel port `fcs0` for LUN masking the SAN LUNs for the installation device of the client OS.

2. Find the worldwide port name (WWPN) address for the fcs0 device:

   ```
   lsdev -dev fcs0 -vpd | grep Address
   ```

   Your output is similar to the following example:

   ```
   Network Address.............10000000C9E3AB56
   ```

3. Repeat the step 1 and 2 on VIOS2.

4. Provide the location codes and the WWPN addresses from the previous steps for both VIOS1 and VIOS2 to your storage administrator.

   At this time, your storage administrator provisions your SAN LUN. The storage administrator makes the LUN masking so that VIOS1 and VIOS2 both see the same SAN LUN. The storage administrator also gives you the SAN LUN ID of the disk for client OS installation. For this exercise, the SAN administrator allocated disk with LUN ID 60:0a:0b:80:00:11:46:32:00:00:92:75:4e:c5:0e:78 and size 25 GB.

5. After the storage administrator provisions the storage, run the **`cfgdev`** command on the VIOS1 and VIOS2 command line to discover any new devices.

Before SAN LUN can be virtualized and provisioned to the virtual server for the client OS, you must change the behavior of locking the SCSI reservation of the physical disk (here SAN LUN). You do not want VIOS to lock the SCSI reservation (to be prepared for dual VIOS configuration). To change the behavior of locking SCSI reservations, you must complete the following steps on both VIOS1 and VIOS2:

1. Log on to the VIOS console and list the physical disks attached to your VIOS:

   ```
   lsdev -type disk
   ```

   Example 4-2 shows the listing. In the output, you can see four internal disks (SAS hdisk0 to hdisk3) and six external disks from the IBM DS4800 (MPIO hdisk4 to hdisk9).

   Example 4-2 Listing of physical disk devices on VIOS

   ```
   $ lsdev -type disk
   hdisk0           Available   SAS Disk Drive
   hdisk1           Available   SAS Disk Drive
   hdisk2           Available   SAS Disk Drive
   hdisk3           Available   SAS Disk Drive
   hdisk4           Available   IBM MPIO DS4800 Array Disk
   hdisk5           Available   IBM MPIO DS4800 Array Disk
   hdisk6           Available   IBM MPIO DS4800 Array Disk
   hdisk7           Available   IBM MPIO DS4800 Array Disk
   hdisk8           Available   IBM MPIO DS4800 Array Disk
   hdisk9           Available   IBM MPIO DS4800 Array Disk
   ```

2. Confirm the SAN LUN ID on VIOS:

   ```
   lsdev -dev hdisk4 -attr | grep unique_id
   ```

   Example 4-3 shows the output with the LUN ID highlighted.

   Example 4-3 Listing disk LUN ID on VIOS

   ```
   $ lsdev -dev hdisk4 -attr | grep unique_id
   unique_id  3E213600A0B8000114632000092754EC50E780F1815  FAStT03IBMfcp PCM
   ```

   Runt the same **lsdev** command again to find the physical disk with the correct LUN ID you received from the storage administrator. The device name for the provisioned external disk used in the following steps is `hdisk4`.

3. Change the behavior of locking SCSI reservations:

   ```
   chdev -dev hdisk4 -attr reserve_policy=no_reserve
   ```

Complete steps 1 - 3 on both VIOS1 and VIOS2. Make a note of the correct device names of the SAN LUN on both VIOS1 and VIOS2. Use these device names in Chapter 5, “Advanced configuration” on page 79, to virtualize them to the virtual server for the client OS.

Device names: You might notice that the names of the devices on VIOS1 are not the same as the device names on VIOS2. The reason for this difference might be that VIOS2 has a different number of internal disks.

The reserve_policy attribute name: If you use a storage subsystem from a different vendor, the `reserve_policy` attribute can have a different name. For example if you use EMC PowerPath drivers to connect LUNs from the EMC storage subsystem, you must use the `reserve_lock` attribute and the value `no`.

Same disk configuration: Make the disk configuration of both VIOS1 and VOS2 the same. This approach makes management of dual VIOS configurations easier and less prone to administrator mistakes.

4.1.6 Creating a virtual server for a client operating system

Next, you create the virtual server for the client OS. The created virtual server for the client OS will be fully virtualized and have the virtual SCSI adapters (most likely for OS disk) and a virtual Ethernet adapter to connect to the LAN. The Fibre Channel adapters (for data disks) will be added in following sections.

Before you create the virtual server for the client OS, complete the following steps:

1. Log on to the SDMC environment.
2. From the home page, locate the host on which the virtual server for the client OS will be created. This host is the same host used in 4.1.1, “Creating the virtual servers for VIOS1 and VIOS2” on page 46.
3. Select the check box to the left of the host, and then select **Actions** → **System Configuration** → **Create Virtual Server**.

To create the virtual server for the client OS:

1. In the Name panel, complete the following steps:
   a. For Virtual Server name, type `VirtServer1`.
   b. For Environment, select **AIX/Linux**.
   c. Click **Next**.
2. In the Memory panel, complete the following steps:
   a. Select the **Dedicated for Memory Mode** check box (if present).
   b. In the Assigned memory field, enter an appropriate amount of memory for this virtual server.
   c. Click **Next**.
3. In the Processor panel, complete the following steps:
   a. Select **Shared for Processing Mode**.
   b. In the Assigned Processors field, enter `1`. You can change this value so that it reflects your needs.
   c. Click **Next**.

   By default, the Ethernet wizard creates two virtual Ethernet adapters. Only the first virtual Ethernet adapter (with VLAN ID 1) is used for a network connectivity.
4. Select the check box to the left of the second adapter (ID 3), and then click **Delete**.

5. If the Storage selection panel opens, select **Yes, Automatically manage the virtual storage adapters for this Virtual Server**. You can provision the virtual disks, physical volumes, or virtualize the Fibre Channel adapters here. Select the check box to the left of the physical volumes. Click **Next**.

   Virtual Fibre Channel adapter configuration: The virtual Fibre Channel adapters will be configured in Chapter 5, “Advanced configuration” on page 79.

6. In the Physical Volumes part of the panel, select the physical disk to virtualize to the virtual server for the client OS. These disks are the same disks that you changed the SCSI reservation policy on in 4.1.5, “Configuring storage devices” on page 51. You can also check the Physical Location Code column to find the correct physical disk.

   Important: Make sure that you select the appropriate physical disk on both VIOS1 and VIOS2.

7. In the Optical devices panel, on the **Physical Optical Devices** tab, select the check box to the left of cd0 to virtualize the physical DVD drive to the virtual server for the client OS. Click **Next**.
8. In the Physical I/O Adapters panel, do not select any physical I/O adapters. The client OS is installed on the disk that is connected by using the virtual SCSI adapter, and all other devices are virtualized.
9. In the Summary panel, verify the information, and then click **Finish** to confirm creation of the virtual server.

4.1.7 Installing a client operating system

The virtual server for the client OS is now fully virtualized and ready for installing a client OS. The virtual server for the client OS now owns the following components:

- The virtual disk for OS installation (the provisioned SAN LUN)
- The virtual CD/DVD drive that can be used to install the client OS
- The virtual Ethernet adapter for connection to your LAN

An installation of a client OS is not within the scope of this document. To install a client OS, follow installation instructions provided by the manufacturer.

4.1.8 Configuring virtual Fibre Channel adapters using the SDMC

Next, you configure the virtual Fibre Channel for the virtual server for the client OS. This function uses a portion of the Fibre Channel standard called *N_Port ID Virtualization* (NPIV). The virtual Fibre Channel adapters have the WWPN generated automatically by the SDMC.

From this point, we assume that the client OS is already installed and TCP/IP connectivity exists between the SDMC and the client OS. With this assumption, you can add the second virtual SCSI adapter dynamically. This function of PowerVM is called *dynamic logical partitioning* (DLPAR).

Before you create virtual Fibre Channel adapters for the virtual server for the client OS, complete these steps:

1. Log in to the SDMC environment.
2. From the home page, locate the host that contains the VirtServer1 virtual server. Click the host name.
3. In the Resource Explorer window that opens, select the check box to the left of the virtual server (VirtServer1), and then select **Actions** → **System Configuration** → **Manage Virtual Server**.

To create virtual Fibre Channel adapters for the virtual server for the client OS:

1. From the left pane, click **Storage Devices**.
2. Under Fibre Channel, click **Add**.
3. In the Add Fibre Channel panel, which shows the physical Fibre Channel adapters that support NPIV, select the physical Fibre Channel adapter that you want to virtualize to the virtual server for the client OS. In our case, we selected the physical Fibre Channel adapter with device name `fcs1` for both VIOS1 and VIOS2. Click **OK**.

   Physical Fibre Channel adapter fcs0: The physical Fibre Channel adapter with device name `fcs0` was already used in 4.1.5, “Configuring storage devices” on page 51 to provision the SAN LUN.

4. Click **Apply**.

Now update the configuration profiles of VirtServer1, VIOS1, and VIOS2. To update the profile on VirtServer1 Virtual Server, log on to the SDMC environment. Then complete the following steps:

1. From the home page, locate the host that contains the VirtServer1 virtual server, and then click the name of the host.
2. In the Resource Explorer window that opens, select the check box to the left of the virtual server VirtServer1, and then select **Actions** → **System Configuration** → **Save Current Configuration**.
3. Select the **Overwrite existing profile** check box, and then select the **OriginalProfile** profile.
4. Click **OK**.
5. In the Save Profile window, click **Yes**.

Repeat steps 1- 5 for VIOS1 and VIOS2 to update their configuration profiles.

You now have a running virtual server with the following virtualized configurations:

- One virtual processor from the shared processor pool (can be adjusted dynamically to meet your needs)
- 4 GB of memory (can be adjusted dynamically to meet your needs)
- One virtual Ethernet adapter with high-available failover mode
- Two virtual SCSI adapters for the OS disk

  This disk uses two paths: one path to VIOS1 and a second path to VIOS2.
- Two virtual Fibre Channel adapters, likely for connecting the SAN LUNs for data

  Each virtual Fibre Channel adapter is provided by a separate VIOS.

Example 4-4 shows devices from the virtual server running the AIX OS.

Example 4-4 List of the virtual devices from AIX

```
# lsdev -Cc adapter
ent0   Available Virtual I/O Ethernet Adapter (l-lan)
fcs0   Available C5-T1 Virtual Fibre Channel Client Adapter
fcs1   Available C6-T1 Virtual Fibre Channel Client Adapter
vsa0   Available LPAR Virtual Serial Adapter
vscsi0 Available Virtual SCSI Client Adapter
vscsi1 Available Virtual SCSI Client Adapter
# lsdev -Cc disk
hdisk0 Available Virtual SCSI Disk Drive
# lspath
Enabled hdisk0 vscsi0
Enabled hdisk0 vscsi1
```

4.2 Setting up a single VIOS using the SDMC

Now you create a PowerVM VIOS configuration using the SDMC. In this section, you configure one VIOS and one virtual server for the client OS. You add the second VIOS in 4.3, “Setting up a dual VIOS using the SDMC” on page 65. This section configures the virtual server for the client OS with the following components:

- Virtual processor from the shared processor pool
- Appropriate amount of the memory
- Virtual disk for OS installation
- Virtual CD or DVD drive that can be used to install the client OS
- Virtual Ethernet adapter for connection to your LAN

Before you begin, you must complete the following tasks:

1. Check the prerequisites in Chapter 1, “Introduction to PowerVM” on page 1.
2. Verify that the SDMC is installed and configured.
3. Verify that the host is already discovered and visible to the SDMC.
4. Verify that your host is in a *Started* state.

In this section, you use commands that show the locations of the hardware cards and adapters. These location codes look similar to this example: `DNWKGPB-P1-C4-T1`. In the location code, the number after the `C` character indicates the card slot, the number after the `T` character is the number of the port on the card. The portion preceding the `P` character is the serial number of the drawer. In this example, `DNWKGPB-P1-C4-T1` refers to the first port on a card in slot C4 of the drawer with serial number DNWKGPB. This location code provides information that you need to find the correct card and to plug cables into the correct ports.

Drawer serial number: You can find the serial number of the drawer in the front of the drawer under the plastic cover.

4.2.1 Creating a VIOS virtual server

Before you create the virtual server for running VIOS using the SDMC environment, you must complete the following steps:

1. Log in to the SDMC environment.
2. If the home page is not automatically displayed, click **Home** from the left pane.
3. On the home page, locate the host on which to create the virtual server.
4. Select the check box to the left of the host, and then select **Actions** → **System Configuration** → **Create Virtual Server** (Figure 4-4).

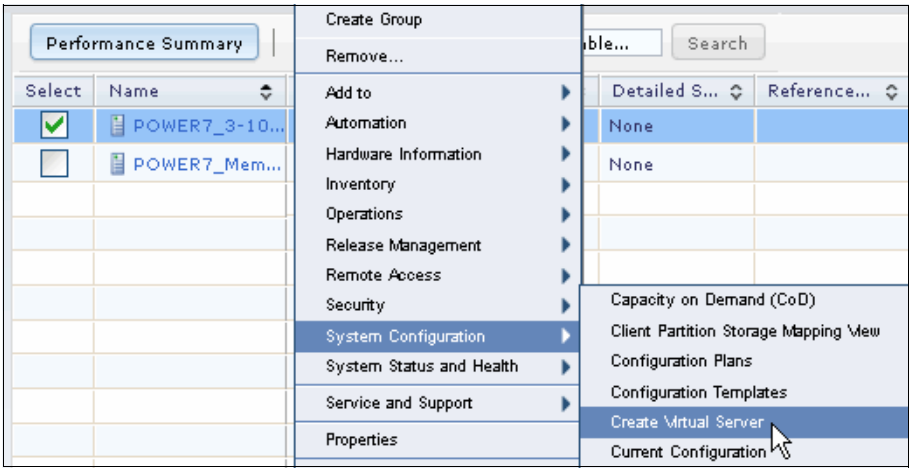

Figure 4-4 Accessing the option to create the virtual server in the SDMC

To create the virtual server for VIOS1:

1. In the Name panel, complete the following steps:
   a. In the Virtual Server name field, enter a name. In this example, we enter `VIOS1`.
   b. For Virtual server ID, enter an ID. We follow the naming convention shown in 1.2, “Planning” on page 3. The default value is the next partition number that is available.
   c. For Environment, select **VIOS**.
   d. Click **Next**.
2. In the Memory panel, complete these steps:
   a. Select the **Dedicated for Memory Mode** check box (if present).
   b. In the Assigned memory field, enter 4 GB of memory or an appropriate amount of memory.

   Amount of memory: The amount of memory your VIOS needs depends on the functions of VIOS that you will use. Start with 4 GB of memory, and then periodically monitor the memory usage on VIOS.

   c. Click **Next**.

3. In the Processor panel, complete the following steps;
   a. Select **Shared for Processing Mode**.
   b. In the Assigned Processors field, type 1 for a single shared processor (from the Shared Processor Pool **DefaultPool(0)**).
   c. Click **Next**.

   Assigned Uncapped Processing Units: In background, the value of Assigned Uncapped Processing Units is 0.1 by default. Start with one shared processor, and then periodically monitor the processor on the VIOS.

4. In the Ethernet panel, complete the following steps:
   a. Expand **Virtual Ethernet**.

      By default, the wizard creates two virtual Ethernet adapters. The first virtual Ethernet adapter uses adapter ID 2 and VLAN ID 1. The second virtual Ethernet adapter uses adapter ID 3 and VLAN ID 99. The second virtual Ethernet adapter is used as a control channel between two VIOSs in a dual VIOS configuration and is not used in a single VIOS configuration. For more information about a control channel and dual VIOS configuration for virtual Ethernet, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.
   b. Select the check box to the left of the first adapter (ID 2), and then click **Edit**.
   c. In the Virtual Ethernet - Modify Adapter panel, complete these steps:
      i. Select the **Use this adapter for Ethernet bridging** check box.
      ii. In the Priority field, type 1.

         For information about the priorities, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.
      iii. Click **OK** to confirm changes to first Ethernet adapter.
   d. Back in the Ethernet panel, click **Next**.
5. In the Virtual Storage Adapter panel, complete the following steps:
   a. In the Maximum number of virtual adapters field, type 200.
   b. Click **Create Adapter**.
   c. In the Create Virtual Adapter panel, complete the following steps:
      i. Adapter ID:101
      ii. Adapter type:SCSI
      iii. Connecting Virtual Server ID:10
      iv. Connecting adapter ID:11

         For the numbers for virtual adapters and the virtual server ID, see the naming convention in 1.2, “Planning” on page 3.
      v. Click **OK**.
   d. Back in the Virtual Storage Adapter panel, click **Next**.
6. In the Physical I/O Adapters panel, select the check boxes to the left of the Location Code and Description of the adapter that you need. To use all the functions as explained in this paper, you must select the following components:
   – One SAS or SCSI disk controller (controller for internal disk drives)
   – One Ethernet adapter (Ethernet adapter for connection to a LAN)

– One Fibre Channel adapter (Fibre Channel adapter for connection to a SAN and a virtual Fibre Channel configuration)

Tip: If the busy icon called “Working” seems to hang, click another tab, and then come back to the previous window.

In our case, we selected the following physical adapters:

```
U78A0.001.DNWKF81-P1-T9 RAID Controller
U5802.001.RCH8497-P1-C7 Quad 10/100/1000 Base-TX PCI-Express Adapter
U5802.001.RCH8497-P1-C3 Fibre Channel Serial Bus
```

The RAID controller that is selected also has the physical CD or DVD drive connected.

Physical location codes: Check the physical location codes of your adapters and be sure to use the correct adapters for your virtual server. Sometimes the description can be misleading. For example, the PCI-to-PCI bridge can be the Ethernet adapter.

In the Physical I/O Adapters panel, click **Next**.

7. In the Summary panel, verify the information, and then click **Finish** to confirm the creation of the virtual server.

For a more detailed description of the options available from the virtual server creation wizard, see *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860.

4.2.2 Installing a VIOS

You can install a VIOS into the previously created virtual server by using one of the following three methods:

- Install VIOS from VIOS DVD media by using the allocated physical DVD drive.
- Install VIOS from the SDMC command line by using the **`installios`** command. The **`installios`** command has a progressive wizard that guides you through the installation process.
- Install VIOS using the IBM Network Installation Manager.

In this section, you use the VIOS DVD media to install VIOS. Before you install VIOS, insert the VIOS installation media into the CD or DVD drive of the system. Then complete these steps:

1. Log on to the SDMC.
2. From the home page, locate the host on which the virtual server for VIOS was created and click its name.
3. Select the check box to the left of the virtual server name that you created (VIOS1), and then select **Actions** → **Operations** → **Activate** → **Profile**.
4. In the Activate Virtual Server: VIOS1 panel, click **Advanced**.
5. Change Boot mode to **SMS**, and then click **OK**.
6. Select the **Open a terminal window or console session** check box, and then click **OK**.
7. When the terminal console for VIOS1 opens, enter your SDMC user ID and password to open the console.

In the terminal console window, install VIOS:

1. If you are prompted with options to set this console as the active console, press the appropriate key indicated in the panel.
2. Type `5` for Select Boot Options.
3. Type `1` for Select Install/Boot Device.
4. Type `7` for List All Devices.
5. Find the CD-ROM device in the list. You might need to type `N` to scroll down. Record the number of the device, and press Enter.
6. Type `2` for Normal Bode Boot, and then type 1 for Yes to exit the SMS menu.
7. Select the console number, and then press Enter.
8. Select the preferred language. To select English, press Enter.
9. When prompted with the Installation and Maintenance menu, type `2` for Change/Show Installation Settings and Install to open installation settings panel.
10. Type `1` for Disk(s) where you want to install option to select the target installation device (the target installation device is marked with three closing angle brackets (>>>)). Usually this device is the first physical disk device. Therefore, you can leave the default value.
11. Type `99` for Previous Menu.
12. Type `5` for Select Edition to choose the correct PowerVM edition.
13. Type `0` for Install with the settings listed above to start the installation. A progress panel shows the `Approximate % Complete` and `Elapsed Time`.
14. Insert volume 2 when prompted by the installation process, and press Enter. This installation takes between 15 minutes and an hour to complete.
15. When VIOS1 first opens, log in with the `padmin` user name.
16. When prompted by the VIOS, change the password and accept the software maintenance terms. After you change the password and agree to the license, enter the following command:

    ```
    license -accept
    ```

4.2.3 Configuring a VIOS

Now you configure a virtual Ethernet and virtual storage (virtual SCSI) for the virtual server for the client OS.

Configuring a VIOS for a client network

In this section, you connect the virtual server for a client OS to a LAN network. You also assign the management IP address to VIOS.

Complete the following steps on the VIOS1 console:

1. To open the console to VIOS1 virtual server, log on to the SDMC environment.
2. From the home page, locate the host that contains the VIOS1 virtual server, and then click the name of the host. This host is the same host used in 4.2.1, “Creating a VIOS virtual server” on page 57.
3. In the Resource Explorer window that opens, select the check box to the left of the virtual server, VIOS1, and then select **Actions** → **Operations** → **Console Window** → **Open Terminal Console**.

4. Find the device name for the physical Ethernet adapter ports:

   ```
   lsdev -type ent4sea
   ```

   To also find the physical locations, enter the following command (Example 4-5):

   ```
   lsdev -vpd | grep ent | grep -v Virtual
   ```

   Example 4-5 Listing of physical Ethernet adapter ports on VIOS

   ```
   $ lsdev -vpd | grep ent | grep -vi Virtual
     Model Implementation: Multiple Processor, PCI bus
     ent0  U5802.001.RCH8497-P1-C7-T1  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent1  U5802.001.RCH8497-P1-C7-T2  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent2  U5802.001.RCH8497-P1-C7-T3  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent3  U5802.001.RCH8497-P1-C7-T4  4-Port 10/100/1000 Base-TX PCI-Express Adapter
   ```

   From the listed ports, select the correct Ethernet port that is to be used for the LAN connection and that has the Ethernet cable plugged in. In this example, we select device `ent0`. This physical adapter port device name is used in the following steps as a value for the **`-sea`** attribute.

5. Find the device name for the virtual Ethernet adapter port with adapter ID 2:

   ```
   lsdev -vpd | grep ent | grep Virtual | grep C2
   ```

   The following example shows the output:

   ```
   ent4  U8233.E8B.10F5D0P-V1-C2-T1  Virtual I/O Ethernet Adapter
   ```

   Command value C2: The value `C2` used in the command in this step is related to adapter ID 2 of the virtual Ethernet adapter created in 4.2.1, “Creating a VIOS virtual server” on page 57. You can also find this ID and the slot number in Table 1-2 on page 5.

   The device name from the output is the virtual Ethernet adapter port. In this example, the name is device `ent4`. This virtual adapter port device name is used in the following step as a value for the **`-vadapter`** and **`-default`** attributes.

   The virtual Ethernet adapter in this step must use VLAN ID 1. Confirm the VLAN ID by using the following command:

   ```
   entstat -all ent4 | grep "Port VLAN ID"
   ```

   VLAN ID 1 is confirmed by the following output:

   ```
   Port VLAN ID: 1
   ```

6. Create a virtual bridge or an SEA in VIOS terminology (Example 4-6):

   ```
   mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 199
   ```

   Example 4-6 Creating an SEA on VIOS

   ```
   $ mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 199
              main:  86 Recived SEA events bytes 164
   ent6 Available
   en6
   et6
   ```

   In this example, the SEA device name is `ent6`. Make a note of the name of the created device. This SEA device name is required in the first part of “Configuring a second VIOS for a client network” on page 69 for changing the attributes of the SEA on VIOS1.

   Tip: The SEA is bridging a virtual and physical network using VIOS.

7. Run the **cfgassist** command, and then select the **VIOS TCP/IP Configuration** option. Select the appropriate interface device name from the previous step, which is **en6** in this example.
8. In the VIOS TCP/IP Configuration panel, enter the TPC/IP configuration values for VIOS connectivity. For these values, consult your network administrator. Figure 2-2 on page 10 shows TPC/IP configuration values used in this example. After you enter the required values for TCP/IP configuration, press Enter.
9. When you see the output `Command: OK`, press F10, or press Esc+0.

SSH: From this point onward, you can use SSH to connect to VIOS1.

Configuring a VIOS for client storage

In this section, the disk is virtualized and provisioned to the virtual server for the client OS. You use SAN LUN to install the client OS. You connect the VIOS virtual server to the SAN, and then you provide the necessary information to the storage administrator.

To attach VIOS to a SAN and configure storage follow these steps on VIOS1 console:

1. To find the Fibre Channel adapters owned by VIOS enter the **lsdev -vpd | grep fcs** command.

   The number of Fibre Channel adapters can vary. You see a list similar to the following example:

   ```
   fcs0  U5802.001.RCH8497-P1-C3-T1  8Gb PCI Express Dual Port FC Adapter
   fcs1  U5802.001.RCH8497-P1-C3-T2  8Gb PCI Express Dual Port FC Adapter
   ```

2. Find the WWPN address for the fcs0 device:

   ```
   lsdev -dev fcs0 -vpd | grep Address
   ```

   This command has the following output:

   ```
   Network Address.............10000000C9E3AB56
   ```

3. Provide the location code and the WWPN address from the previous steps to your storage administrator.

   At this time, your storage administrator provisions the necessary SAN LUN. The storage administrator also gives you the SAN LUN ID of the disk for the client OS installation. For this exercise, the SAN administrator allocated disk with LUN ID 60:0a:0b:80:00:11:46:32:00:00:92:75:4e:c5:0e:78 and a size of 25 GB.

4. After the storage administrator provisions the storage, find any new devices by running the **cfgdev** command.
5. List the physical disks attached to your VIOS:

   ```
   lsdev -type disk
   ```

   In the output shown in Example 4-7, you can see four internal disks (SAS hdisk0 to hdisk3) and six external disks from the IBM DS4800 (MPIO hdisk4 to hdisk9).

Example 4-7 Listing of physical disk devices on VIOS

```
$ lsdev -type disk
hdisk0          Available   SAS Disk Drive
hdisk1          Available   SAS Disk Drive
hdisk2          Available   SAS Disk Drive
hdisk3          Available   SAS Disk Drive
hdisk4          Available   IBM MPIO DS4800 Array Disk
hdisk5          Available   IBM MPIO DS4800 Array Disk
```

```
hdisk6           Available   IBM MPIO DS4800 Array Disk
hdisk7           Available   IBM MPIO DS4800 Array Disk
hdisk8           Available   IBM MPIO DS4800 Array Disk
hdisk9           Available   IBM MPIO DS4800 Array Disk
```

6. Confirm the SAN LUN ID on VIOS:

   `lsdev -dev hdisk4 -attr | grep unique_id`

   Example 4-8 shows the output with the LUN ID highlighted in bold.

   Example 4-8 Listing disk LUN ID on VIOS

```
$ lsdev -dev hdisk4 -attr | grep unique_id
unique_id  3E213600A0B8000114632000092754EC50E780F1815  FAStT03IBMfcp PCM
```

   If you have more disk devices, repeat this **lsdev** command to find the physical disk with the correct LUN ID you received from the storage administrator. The device name for the external disk used in the next steps is `hdisk4`.

7. If do not plan to use a dual VIOS configuration skip this step. In this step, you change the behavior of locking the SCSI reservation on the physical disk device. You do not want VIOS to lock the SCSI reservation (to be prepared for dual VIOS configuration).

   To change the behavior of locking SCSI reservations, enter the following command:

   `chdev -dev hdisk4 -attr reserve_policy=no_reserve`

   The reserve_policy attribute name: If you use a storage subsystem from a different vendor, the `reserve_policy` attribute can have a different name. For example, if you use EMC PowerPath drivers to connect LUNs from the EMC storage subsystem, you must use the `reserve_lock` attribute and the value `no`.

Mapping the SAN LUN

After the storage administrator provisions the necessary SAN LUN, map this SAN LUN over the virtual SCSI adapter to the virtual server for the client OS. You can do this mapping by using the SDMC interface or by using the VIOS command line. Use the SDMC to create the mapping of the physical disk to the virtual SCSI adapter. To create this mapping using the VIOS command, see “Configuring a second VIOS for client storage” on page 72.

To create the mapping:

1. Log on to the SDMC.
2. From the home page, locate the host that contains the VIOS1 virtual server. Select the check box to the left of the host that contains the VIOS1, and then select **Actions** → **System Configuration** → **Virtual Resources** → **Virtual Storage Management**.
3. In VIOS/SSP, select **VIOS1**, and then click **Query**.
4. Click **Physical Volumes**.
5. Select the check box to the left of the hdisk4 physical disk. This device name was in the previous steps.
6. Click **Modify assignment**.
7. In the New virtual server assignment panel, select the **VirtServer1(10)**, and then click **OK**.
8. Click **Physical Volumes**.
9. Select the check box to the left of the cd0 physical DVD drive.
10. Click **Modify assignment**.

11. In the New virtual server assignment panel, select **VirtServer1(10)**, and then click **OK**.
12. Click **Close**.

4.2.4 Creating a virtual server for a client operating system

Now you create the virtual server for the client OS. The created virtual server for the client OS is fully virtualized and has the virtual SCSI adapter and the virtual Ethernet adapters. To add the virtual Fibre Channel adapters, see 4.4, “Setting up a virtual Fibre Channel using the SDMC” on page 74.

Before you begin, complete the following steps:

1. Log on to the SDMC environment.
2. From the home page, locate the host on which the virtual server for the client OS will be created. This host is the same one used in 4.2.1, “Creating a VIOS virtual server” on page 57.
3. Select the check box to the left of the host, and then select **Actions** → **System Configuration** → **Create Virtual Server**.

To create the virtual server for the client OS.

1. In the Name panel, complete the following steps:
   a. In the Virtual Server name field, type `VirtServer1`.
   b. In the Virtual server ID field, type `10`. The default value is the next available ID.
   c. In the Environment field, select **AIX/Linux**.
   d. Click **Next**.
2. In the Memory panel, complete these steps:
   a. Select the **Dedicated for Memory Mode** check box (if present).
   b. In the Assigned memory field, enter an appropriate amount of memory for this virtual server.
   c. Click **Next**.
3. In the Processor panel, complete these steps:
   a. Select **Shared for Processing Mode**.
   b. In the Assigned Processors field, type `1` or a value that reflects your needs.
   c. Click **Next**.

   By default, the Ethernet wizard creates two virtual Ethernet adapters. Only the first virtual Ethernet adapter (with VLAN ID 1) is used for a network connectivity.
4. Select the check box to the left of the second adapter (ID 3), and then click **Delete**.
5. If the Storage selection panel opens, select the **No, I want to manage the virtual storage adapters for this Virtual Server** check box.
6. In the Virtual Storage Adapter panel, complete the following steps:
   a. In Maximum number of virtual adapters field, type `30`.
   b. Click **Create Adapter**, and then complete the following steps:
      i. In the Adapter ID field, type `11`.
      ii. In the Adapter type field, type `SCSI`.
      iii. In the Connecting Virtual Server ID field, type `VIOS (1)`.
      iv. In the Connecting adapter ID field, type `101`.
      v. Click **OK**.
   c. Click **Next**.

7. In the Physical I/O Adapters panel, do not select any physical I/O adapters.

   The client OS is installed on the disk that is connected by using the virtual SCSI adapter. The virtual Fibre Channel adapters are added in 4.4, “Setting up a virtual Fibre Channel using the SDMC” on page 74.

   Click **Next**.

8. In the Summary panel, verify the information, and then click **Finish** to confirm creation of the virtual server.

4.2.5 Installing a client operating system

The virtual server for the client OS is now fully virtualized and ready for installing a client OS. The virtual server for the client OS now has the following components:

- The virtual disk for OS installation
- The virtual CD or DVD drive that can be used to install the client OS
- The virtual Ethernet adapter for connection to your LAN

An installation of a client OS is not within the scope of this publication. To install a client OS, follow instructions provided by the installation guide for your OS.

4.3 Setting up a dual VIOS using the SDMC

In this section, you add a second VIOS to the configuration explained in 4.2, “Setting up a single VIOS using the SDMC” on page 56. By adding a second VIOS, you can take advantage of better redundancy, accessibility, and serviceability (RAS). This configuration also increases the I/O bandwidth to your virtual servers for client OS.

4.3.1 Creating a second VIOS virtual server

To start the dual VIOS configuration, you must create a virtual server for the second VIOS (VIOS2). Before you create a virtual server for VIOS2, complete these steps:

1. Log on to the SDMC environment.
2. From the home page, locate the host on which the virtual server for VIOS2 will be created. This host is the same host used in 4.2.1, “Creating a VIOS virtual server” on page 57.
3. Select the check box to the left of the host, and then select **Actions** → **System Configuration** → **Create Virtual Server**.

Now create a virtual server for the second VIOS (VIOS2):

1. In the Name panel, complete the following steps:
   a. In the Virtual Server name field, type `VIOS2`.
   b. In the Virtual server ID field, type `2`.
   c. In the Environment field, type `VIOS`.
   d. Click **Next**.

2. In the Memory panel, select the **Dedicated for Memory Mode** check box (if shown). In the Assigned memory field, enter 4 GB of memory or an appropriate amount of memory.

   Amount of memory: The amount of memory your VIOS needs depends on the functions of VIOS that you will use. Start with 4 GB of memory, and then periodically monitor the memory usage on VIOS.

   Click **Next**.

3. In the Processor panel, complete the following steps;
   a. Select **Shared for Processing Mode**.
   b. In the Assigned Processors field, type `1` for a single shared processor.
   c. Click **Next**.

   Assigned Uncapped Processing Units: In background, the value of Assigned Uncapped Processing Units is 0.1 by default. Start with one shared processor, and then periodically monitor the processor on the VIOS.

4. In the Ethernet panel, complete these steps:
   a. Expand **Virtual Ethernet**.

      By default, the wizard creates two virtual Ethernet adapters. The first virtual Ethernet adapter uses adapter ID 2 and VLAN ID 1. The second virtual Ethernet adapter uses adapter ID 3 and VLAN ID 99. The second virtual Ethernet adapter is used for control channel between two VIOSs in the dual VIOS configuration. For more information about the control channel and the dual VIOS configuration for virtual Ethernet, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.
   b. Select the check box to the left of the first adapter (ID 2), and then click **Edit**.
   c. In the Virtual Ethernet - Modify Adapter panel, complete the following steps:
      i. Select the **Use this adapter for Ethernet bridging** check box.
      ii. In the Priority field, type `2`.

         For an explanation of the priorities, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.
      iii. Click **OK** to confirm changes to the first Ethernet adapter.
   d. In the Ethernet panel, click **Next**.

5. In the Virtual Storage Adapter panel, complete the following steps:
   a. In the Maximum number of virtual adapters field, type `200`.
   b. Click **Create Adapter**.
   c. In the Create Virtual Adapter panel, complete the following steps:
      i. In the Adapter ID field, type 101.
      ii. In the Adapter type field, type `SCSI`.
      iii. In the Connecting Virtual Server ID field, type `10`, or select **VirtServer1 (10)**.
      iv. In the Connecting adapter ID field, type `21`.
      v. Click **OK**.
   d. Back in the Virtual Storage Adapter panel, click **Next**.

6. In the Physical I/O Adapters panel, select the check boxes to the left of the Location Code and Description of the required adapters. To use all the functions described in this paper, you must make the following selections:
   - One SAS or SCSI disk controller (controller for internal disk drives)
   - One Ethernet adapter (Ethernet adapter for connection to LAN)
   - One Fibre Channel adapter (Fibre Channel adapter for connection to SAN)

   In this example, we selected the following physical adapters:

   ```
   U5802.001.RCH8497-P1-C2  PCI-E SAS Controller
   U5802.001.RCH8497-P1-C6  Quad 10/100/1000 Base-TX PCI-Express Adapter
   U5802.001.RCH8497-P1-C5  Fibre Channel Serial Bus
   ```

   In the Physical I/O Adapters panel, click **Next**.
7. In the Summary panel, verify the information, and then click **Finish** to confirm the creation of the virtual server.

For a more detailed description of the options that are available, see *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860.

4.3.2 Installing a second VIOS using NIM

In this section, you use NIM to install the VIOS. You can still use the **`installios`** command as explained in 4.2.2, “Installing a VIOS” on page 59.

Before you install VIOS2, prepare your NIM environment to enable a boot over Ethernet and install the VIOS `mksysb` image. Because the installation of the second VIOS is over Ethernet, as a prerequisite, you must have functional network connectivity between the NIM server and the virtual server for the second VIOS. A detailed description of NIM preparation is outside the scope of this publication. This section provides only the basic steps to be done on the NIM server side.

You must prepare the following resources to install a VIOS using NIM:

- The NIM `mksysb` resource of the VIOS installation media, which is the `mksysb` image of VIOS

  This resource can be generated from the `mksysb` image copied from the VIOS installation media.

  The mksysb file for VIOS 1.5.1 or later: If using VIOS 1.5.1 or later, the `mksysb` file can be split into two parts. In this case, combine the two `mksysb` images to create one file.

- The NIM Shared Product Object Tree (SPOT) resource

  This resource is created from the NIM VIOS `mksysb` resource.

- The NIM bosinst_data resource

  The bosinst_data resource is defined by using the `bosinst.data` file that is copied from the VIOS installation media.

You must add the second VIOS (VIOS2) to the NIM environment as a stand-alone machine. Then initialize a VIOS2 installation from NIM by using the prepared NIM stand-alone machine and NIM resources.

Important: In the Install the Base Operating System on the Standalone Clients panel, ensure that the Remain NIM client after install attribute field is set to `NO`. With this setting, NIM does not set up the TCP/IP configuration on a newly installed VIOS so that you can create an SEA in this VIOS.

Now all the necessary resources are prepared for your NIM environment, and an installation of the second VIOS is initialized from NIM.

For a detailed explanation of how to prepare NIM to install VIOS, see the NIM installation and backup of the VIOS Technote at:

`https://www.ibm.com/support/docview.wss?uid=isg3T1011386#4`

Before you install VIOS2 into the virtual server created in 4.3.1, “Creating a second VIOS virtual server” on page 65, complete these steps:

1. Log on to the SDMC.
2. From the SDMC home page, locate the host on which the virtual server for VIOS2 was created, and click its name.
3. Select the check box to the left of the virtual server name VIOS2, and then select **Actions** → **Operations** → **Activate** → **Profile**.
4. In the Activate Virtual server: VIOS2 panel, click **Advanced**.
   a. Change Boot mode to **SMS**, and click **OK**.
   b. Select the **Open a terminal window or console session** check box, and click **OK**.
5. In the terminal console for VIOS2 that opens, enter your SDMC user ID and password to open the terminal console

To install a second VIOS:

1. If prompted with options to set this console as the active console, press the key indicated in the panel.
2. Type `2` for Setup Remote IPL (Initial Program Load).
3. Select the number of the port that is connected to the Ethernet switch and the subnet used during installation. In this example, we type 3 for Port 1.
4. Type `1` for IPv4 - Address Format 123.231.111.222.
5. Type `1` for BOOTP.
6. Type `1` for IP Parameters.
7. Enter the TCP/IP configuration parameters. We used these parameters:

```
1.   Client IP Address                  [172.16.22.13]
2.   Server IP Address                  [172.16.20.40]
3.   Gateway IP Address                 [172.16.20.1]
4.   Subnet Mask                        [255.255.252.0]
```

   The Server TCP/IP Address address is the TCP/IP address of your NIM server.
8. Press the ESC key.
9. Type `3` for Ping Test.
10. Type `1` for Execute Ping Test. You see the following message:

```
.-----------------.
|  Ping  Success. |
`-----------------'
```

11. Press any key to continue, and then press the ESC key five times to go to the Main Menu.
12. From the Main Menu, type `5` for Select Boot Options.
13. Type `1` for Select Install/Boot Device.
14. Type `6` for Network.
15. Type `1` for BOOTP.
16. In the Select Device panel, select the number of the port that is connected to the switch and subnet that are used during the installation. In this example, we typed `3` for Port 1:
17. Type `1` for Normal Mode Boot, and then type `1` for Yes to leave the SMS menu and start the installation.
18. When prompted to define the system console, type `1`. The number that you type might be different for your installation. Press Enter.
19. To confirm English as the language to use during the installation, press Enter.
20. From the Installation and Maintenance menu, type `2` for Change/Show Installation Settings and Install to open the installation settings panel.
21. Type `1` for Disk(s) where you want to install to select the target installation device. Usually this device is the first physical disk device. Therefore, you can accept the default. After you select the target installation device, type `0` for Continue with choices indicated above to return to the main menu.
22. Type `5` for Select Edition to choose the PowerVM edition.
23. Start the installation by typing `0` for Install with the settings listed above. A progress panel shows `Approximate % Complete` and `Elapsed Time`. This installation takes between 15 minutes and one hour to complete.
24. When VIOS2 first starts, log in with the `padmin` user name.
25. When prompted by VIOS, change the password and accept software maintenance terms. After you change the password and agree to the license, enter the following command:

```
license -accept
```

4.3.3 Configuring a second VIOS

Now you configure the second VIOS to enable the Ethernet adapter failover in the virtual server for the client OS. During normal operations, the client virtual server has the Ethernet adapter provisioned by VIOS1. The Ethernet adapter in the client virtual server automatically switches to VIOS2 if VIOS1 is not operating, such as during a restart or normal maintenance of VIOS1.

Configuring a second VIOS for a client network

Before you configure the second VIOS for a client network, you must change the attributes of the SEA on the first VIOS (VIOS1) to support the SEA failover.

Important: Make sure that you are logged on the first VIOS, which in this example is VIOS1.

To change the appropriate attributes, log on to VIOS1, and complete the following steps:

1. Open the console for VIOS1 using the SDMC. For details about this process, open the terminal console by using the SDMC is described in “Configuring a VIOS for a client network” on page 60.
2. Find the device name for the virtual port that will function as a control channel

   ```
   lsdev -vpd | grep ent | grep C3
   ```

   The command produces the following output:

   ```
   ent5  U8233.E8B.10F5D0P-V1-C3-T1  Virtual I/O Ethernet Adapter
   ```

   This adapter is the second virtual Ethernet adapter with adapter ID 3 that was created by default in 4.2.1, “Creating a VIOS virtual server” on page 57. This device name is used in the next step in the ctl_chan attribute. In this example, the device name is `ent5`.

   The virtual Ethernet adapter used in this step must use VLAN ID 99. Confirm the VLAN ID:

   ```
   entstat -all ent5 | grep "Port VLAN ID"
   ```

   VLAN ID 99 is confirmed by the following output:

   ```
   Port VLAN ID:    99
   ```

3. Change the attributes of the SEA on VIOS1:

   ```
   chdev -dev ent6 -attr ha_mode=auto ctl_chan=ent5
   ```

   In this command, the **`-dev`** attribute contains the SEA device name from “Configuring a VIOS for a client network” on page 60. To confirm the attributes of the SEA on VIOS1, enter the following command:

   ```
   lsdev -dev ent6 -attr
   ```

Now configure the virtual Ethernet bridge (known as the SEA) on the second VIOS (VIOS2) and also configure the management TCP/IP address for the second VIOS. Follow these steps in VIOS2 console:

Important: Make sure that you are logged on the second VIOS, which in this example is VIOS2.

1. Find the device names for the physical Ethernet adapter ports:

   ```
   lsdev -vpd | grep ent | grep -v Virtual
   ```

   Select one of the listed ports (Example 4-9) that is used for a LAN connection and has an Ethernet cable plugged in. In this case, the device is `ent0`. This physical adapter port device name is used in the next steps as the value for the **`-sea`** attribute.

   Example 4-9 Listing of physical Ethernet adapter ports on VIOS

   ```
   $ lsdev -vpd | grep ent | grep -v Virtual
     Model Implementation: Multiple Processor, PCI bus
     ent0  U5802.001.RCH8497-P1-C6-T1  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent1  U5802.001.RCH8497-P1-C6-T2  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent2  U5802.001.RCH8497-P1-C6-T3  4-Port 10/100/1000 Base-TX PCI-Express Adapter
     ent3  U5802.001.RCH8497-P1-C6-T4  4-Port 10/100/1000 Base-TX PCI-Express Adapter
   ```

2. Find the device name for the virtual port:

   ```
   lsdev -vpd | grep ent | grep C2
   ```

   This command has the following output:

   ```
   ent4  U8233.E8B.10F5D0P-V2-C2-T1  Virtual I/O Ethernet Adapter (l-lan)
   ```

Command value C2: C2 used in the previous command is related to the number of the adapter created in 4.3.1, “Creating a second VIOS virtual server” on page 65.

The device name from the output is the virtual Ethernet adapter port. In our case, the device is `ent4`. This virtual adapter port device name is used in the next step as a value for the **`-vadapter`** and **`-default`** attributes.

The virtual port device name found in this step uses VLAN ID 1. Confirm the VLAN ID by using the following command:

```
entstat -all ent4 | grep "Port VLAN ID"
```

VLAN ID 1 is confirmed by the output:

```
Port VLAN ID: 1
```

3. Find the device name for the virtual port that functions as a control channel in the output of the following command:

   ```
   lsdev -vpd | grep ent | grep C3
   ```

   This command has the following output:

   ```
   ent5  U8233.E8B.10F5D0P-V2-C3-T1  Virtual I/O Ethernet Adapter (l-lan)
   ```

   This adapter is the second virtual Ethernet adapter with adapter ID 3 that was created by default in 4.3.1, “Creating a second VIOS virtual server” on page 65. This device name is used in the next step in the `ctl_chan` attribute. In this example, the device name is `ent5`.

   The virtual Ethernet adapter in this step must use VLAN ID 99. Confirm the VLAN ID by using the following command:

   ```
   entstat -all ent5 | grep "Port VLAN ID"
   ```

   VLAN ID 99 is confirmed by the following output:

   ```
   Port VLAN ID:    99
   ```

4. Create a virtual bridge or an SEA in VIOS terminology (Example 4-10):

   ```
   mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 199 -attr ha_mode=auto
   ctl_chan=ent5
   ```

   Example 4-10 Creating an SEA on a second VIOS

   ```
   $ mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 199 -attr
   ha_mode=auto ctl_chan=ent5
              main:  86 Recived SEA events bytes 164
   ent6 Available
   en6
   et6
   ```

   Note the name of the created SEA and interface. In this example, the device name of the interface is `en6`.

Important: Mismatching SEA and the SEA failover can cause broadcast storms to occur in the network and effect the network stability. For more information, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

5. Run the **`cfgassist`** command, and then select **VIOS TCP/IP Configuration**. Select the appropriate interface device name from the previous step. In this example, we selected **en6**.

6. In the VIOS TCP/IP Configuration panel, enter TPC/IP configuration values for VIOS2 connectivity. For these values, consult your network administrator. See Figure 2-2 on page 10 for the TPC/IP configuration values in this example. After entering the necessary values for TCP/IP configuration, press Enter.
7. When you see the output `Command: OK`, press F10 or press ESC+0.

SSH: From this point onward, you can use SSH to connect to VIOS2.

From the Ethernet point of view, the virtual server for the client OS is already prepared for the dual VIOS configuration. You do not need to make changes to the virtual server for the client OS.

Configuring a second VIOS for client storage

Configuring the second VIOS to provision the same disk to the client virtual server provides the client OS access to the same disk over two paths. You must add another virtual SCSI adapter to your virtual server for the client OS (VirtServer1) to use multipathing for the disk provided by both VIOSs.

From this point, we assume that the client OS is already installed and TCP/IP connectivity exists between the SDMC and the client OS. With this assumption, you can add the second virtual SCSI adapter dynamically. This function of PowerVM is called *dynamic logical partitioning*.

Important: If the OS in the client virtual server is not already installed and no TCP/IP connectivity is established, you must change the configuration of virtual server profile and restart the virtual server. Rebooting just the client OS in the virtual server is not enough. For more information about this configuration, see *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860.

To add the second virtual SCSI adapter to your virtual server VirtServer1:

1. Change the virtual server for the client OS dynamically:
   a. Log in to the SDMC environment.
   b. From the home page, locate the host that contains VirtServer1 virtual server. This host is the same host used in 4.2.4, “Creating a virtual server for a client operating system” on page 64. Click the host name.
   c. In the Resource Explorer window that opens, select the check box to the left of the Virtual Server (VirtServer1), and then click **Actions** → **System Configuration** → **Manage Virtual Server**.
   d. Click **Storage Adapters** from the left menu.
   e. Click **Add**.
   f. In the Create Virtual Storage Adapter window, complete these steps:
      i. In the Adapter Id field, type `21`.
      ii. In the Adapter type field, select **SCSI**.
      iii. In the Connecting virtual server field, type `VIOS2(2)`.
      iv. In the Connecting adapter ID field, type `101`.
      v. Click **OK**.
   g. Click **Apply** to dynamically add the virtual SCSI adapter to the virtual server.

2. Update the configuration profile of VirtServer1:
   a. Log on to the SDMC environment.
   b. From the home page, locate the host that contains the VirtServer1 virtual server, and select the name of the host.
   c. In the Resource Explorer window that opens, select the check box to the left of virtual server VirtServer1, and then select **Actions** → **System Configuration** → **Save Current Configuration**.
   d. Select the **Overwrite existing profile** check box, and then select the **OriginalProfile** profile.
   e. Click **OK**.
   f. In the Save Profile window, click **Yes**.

Now configure the second VIOS (VIOS2) to provision the disk to the virtual server for the client OS. To attach the second VIOS to a SAN and configure the storage, complete the following steps in VIOS2 console:

1. Provide Fibre Channel card location codes and their WWPN addresses to your storage administrator. For the steps to find the location codes and WWPN addresses, see “Configuring a VIOS for client storage” on page 62.

   At this time, your storage administrator provides the same SAN LUN (and its LUN ID) that was provisioned and used in “Configuring a VIOS for client storage” on page 62.
2. After the storage administrator completes the provisioning, run the **`cfgdev`** command to find the new devices.
3. List the physical disks attached to your VIOS:

   ```
   lsdev -type disk
   ```

   Example 4-11 shows out system output from the **`lsdev`** command. In the output, you can see six internal disks and six external disks from the IBM DS4800 storage subsystem. Make sure that you find the correct physical disk device names as explained in “Configuring a VIOS for client storage” on page 62. In this example, the physical disk with LUN ID 60:0a:0b:80:00:11:46:32:00:00:92:75:4e:c5:0e:78 has the device name `hdisk6`. This device name is used in the following steps.

Example 4-11 Listing physical disks on the VIOS

```
$ lsdev -type disk
hdisk0           Available   SAS RAID 0 SSD Array
hdisk1           Available   SAS RAID 0 SSD Array
hdisk2           Available   SAS RAID 0 SSD Array
hdisk3           Available   SAS RAID 0 SSD Array
hdisk4           Available   SAS RAID 0 SSD Array
hdisk5           Available   SAS RAID 0 SSD Array
hdisk6           Available   IBM MPIO DS4800 Array Disk
hdisk7           Available   IBM MPIO DS4800 Array Disk
hdisk8           Available   IBM MPIO DS4800 Array Disk
hdisk9           Available   IBM MPIO DS4800 Array Disk
hdisk10          Available   IBM MPIO DS4800 Array Disk
hdisk11          Available   IBM MPIO DS4800 Array Disk
```

Device names: You might notice that the names of devices on VIOS1 are not the same as the device names on VIOS2. The reason for this difference is that VIOS2 has more internal disks, and the external disk has a higher disk number.

Same disk configuration: Make the disk configuration of both VIOS1 and VOS2 the same. This approach makes management of a dual VIOS configuration easier and less prone to administrator mistakes.

4. Change the behavior of locking SCSI reservations:

   ```
   chdev -dev hdisk6 -attr reserve_policy=no_reserve
   ```

5. Find the device name for the virtual adapter connected to virtual server for the client OS:

   ```
   lsdev -vpd | grep vhost | grep C101
   ```

   `C101` is the slot number from the Table 1-3 on page 5. In this example, this command produces the following output:

   ```
   vhost0  U8233.E8B.10F5D0P-V1-C101  Virtual SCSI Server Adapter
   ```

   The device name for the virtual adapter is used in the next step. In this example, the device name is `vhost0`.

6. Map the external disk to the virtual server for the client OS:

   ```
   mkvdev -vdev hdisk6 -vadapter vhost0
   ```

4.4 Setting up a virtual Fibre Channel using the SDMC

This section explains how to virtualize Fibre Channel adapters on both VIOSs and how to provide virtual Fibre Channel adapters to the client virtual server. This function uses a method of the NPIV Fibre Channel standard.

4.4.1 Configuring a client virtual server for NPIV

You add virtual Fibre Channel adapters to the VirtServer1 virtual server for the client OS. From this point, we assume that the client OS is already installed and TCP/IP connectivity exists between the SDMC and the client OS. With this assumption, you can add the virtual Fibre Channel adapters dynamically. This function of PowerVM is called *dynamic logical partitioning*.

Before you begin, complete these steps:

1. Log in to the SDMC environment.
2. From the home page, locate the host that contains the VirtServer1 virtual server, which is the same host used in 4.2.4, “Creating a virtual server for a client operating system” on page 64. Click the name of the host.
3. In the Resource Explorer window that opens, select the check box to the left of the VirtServer1 virtual server name, and then select **Actions** → **System Configuration** → **Manage Virtual Server**.

Now add the virtual Fibre Channel adapters to the VirtServer1 virtual server:

1. Click **Storage Adapters** from the left pane.
2. Click **Add**.

3. In the Create Virtual Storage Adapter window, complete the following steps:
   a. In the Adapter ID field, type `12`.
   b. In Adapter type field, select **Fibre Channel**.
   c. In the Connecting virtual server field, type `VIOS1(1)`.
   d. In the Connecting adapter ID field, type `102`.
   e. Click **Add**.
4. Click **Add**.
5. In the Create Virtual Storage Adapter window, complete the following steps:
   a. In the Adapter Id field, type `22`.
   b. In the Adapter type field, select **Fibre Channel**.
   c. In the Connecting virtual server, type `VIOS2(2)`.
   d. In the Connecting adapter ID field, type `102`.
   e. Click **Add**.
6. Click **Apply** to dynamically add virtual Fibre Channel adapters to VirtServer1.

Now update the configuration profiles of the VirtServer1 virtual server. To update the profile of VirtServer1 virtual server, log in to the SDMC environment:

1. From the home page, locate the host that contains the VirtServer1 virtual server, and click the name of the host.
2. In the Resource Explorer window that opens, select the check box to the left of the virtual server VirtServer1, and then select **Actions** → **System Configuration** → **Save Current Configuration**.
3. Select the **Overwrite existing profile** check box, and then select the **OriginalProfile** profile.
4. Click **OK**.
5. In the Save Profile window, click **Yes**.

4.4.2 Configuring a VIOS for NPIV

Next you add virtual Fibre Channel adapters to the first VIOS (VIOS1). Before you begin, complete the following steps:

1. Log in to the SDMC environment.
2. From the home page, locate the host that contains the VIOS1 virtual server, which is the same host used in 4.2.1, “Creating a VIOS virtual server” on page 57. Click the host name.
3. In the Resource Explorer window that opens, select the check box to the left of the VIOS1 virtual server name, and then select **Actions** → **System Configuration** → **Manage Virtual Server**.

Add the virtual Fibre Channel adapters to VIOS1:

1. Click **Storage Adapters** in the left panel.
2. Click **Add**.
3. In the Create Virtual Storage Adapter window, complete the following steps:
   a. In the Adapter Id, type `102`.
   b. In the Adapter type field, select **Fibre Channel**.
   c. In the Connecting virtual server field, type `VirtServer1(10)`.
   d. In the Connecting adapter ID field, type `12`.
   e. Click **Add**.

4. Click **Apply** to dynamically add the virtual Fibre Channel adapter to VIOS1.
5. On the VIOS1 command line, run the **cfgdev** command to check for newly added devices.
6. On the VIOS1 command line, list virtual Fibre Channel adapters as shown in Example 4-12:

   ```
   lsdev -type adapter | grep "Virtual FC"
   ```

   Example 4-12 Listing virtual Fibre Channel adapters on VIOS1

   ```
   $ lsdev -type adapter | grep "Virtual FC"
   vfchost0         Available   Virtual FC Server Adapter
   ```

   The device name vfchost0 is used in the following steps as the **-vadapter** attribute.
7. List the physical Fibre Channel ports and NPIV attributes by using the **lsnports** command (Example 4-13).

   Example 4-13 Listing the NPIV capable Fibre Channel ports on VIOS1

   ```
   $ lsnports
   name  physloc                        fabric tports aports swwpns  awwpns
   fcs0  U5802.001.RCH8497-P1-C3-T1          1     64     64   2048    2046
   fcs1  U5802.001.RCH8497-P1-C3-T2          1     64     64   2048    2048
   ```

   NPIV capable ports have a value of 1 in the fabric column. For Fibre Channel virtualization, select the physical port with the device name fcs1. This device name is used in the following steps to create the mapping. The physical port, fcs0, was used for a SAN LUN masking in “Configuring a VIOS for client storage” on page 62.
8. Create the virtual Fibre Channel adapter to physical Fibre Channel adapter mapping. You can perform this mapping by using the SDMC interface or the VIOS command line. Use the SDMC interface to create the mapping between the virtual Fibre Channel adapter and the physical Fibre Channel adapter. Creation of this mapping using the VIOS command is explained in 4.4.3, “Configuring a second VIOS for NPIV” on page 77.

   To create the mapping by using the SDMC:

   a. Log on to the SDMC:
   b. From the home page, locate the host that contains the VIOS1 virtual server.
   c. Select the check box to the left of the host that contains the VIOS1, and then select **Actions** → **System Configuration** → **Virtual Resources** → **Virtual Storage Management**.
   d. In VIOS/SSP section, select **VIOS1**, and click **Query**.
   e. Click **Virtual Fibre Channel**.
   f. Select the check box to the left of the **fcs1** physical Fibre Channel port. This device name was in previous steps.
   g. Click **Modify virtual server connections**.
   h. Select the check box to the left of the VirtServer1 virtual server name.
   i. Click **OK**.

Now update configuration profiles of the VIOS1 virtual server:

1. Log on to the SDMC environment.
2. From the home page, locate the host that contains the VIOS1 virtual server.
3. Click the name of the host.

4. In the Resource Explorer window that opens, select the check box to the left of the virtual server VIOS1, and then select **Actions** → **System Configuration** → **Save Current Configuration**.
5. Select the **Overwrite existing profile** check box, and then select the **OriginalProfile** profile.
6. In the Confirm window, click **OK**.
7. In the Save Profile window, click **Yes**.

4.4.3 Configuring a second VIOS for NPIV

Next you add the virtual Fibre Channel adapter to second VIOS (VIOS2). Before you begin this task, complete the following steps:

1. Log in to the SDMC environment.
2. From the home page, locate the host that contains the VIOS2 virtual server, which is the same host used in 4.3.1, “Creating a second VIOS virtual server” on page 65. Click the host name.
3. In the Resource Explorer window that opens, select the check box to the left of the VIOS2 virtual server name, and then select **Actions** → **System Configuration** → **Manage Virtual Server**.

Now add the virtual Fibre Channel adapters to VIOS2:

1. Click **Storage Adapters** from the left pane.
2. Click **Add**.
3. In the Create Virtual Storage Adapter window, complete the following steps:
   a. In the Adapter Id field, type `102`.
   b. In the Adapter type field, select **Fibre Channel**.
   c. In the Connecting virtual server field, type `VirtServer1(10)`.
   d. In the Connecting adapter ID field, type `22`.
   e. Click **Add**.
4. Click **Apply** to dynamically add the virtual Fibre Channel adapter to VIOS2.
5. On the VIOS2 command line, run the **`cfgdev`** command to check for newly added devices.
6. On the VIOS2 command line, list the virtual Fibre Channel adapters as shown in Example 4-14:

   ```
   lsdev -type adapter | grep "Virtual FC"
   ```

   Device name `vfchost0` is used in the following steps as a **`-vadapter`** attribute.

Example 4-14 Listing virtual Fibre Channel adapters on VIOS2

```
$ lsdev -type adapter | grep "Virtual FC"
vfchost0         Available   Virtual FC Server Adapter
```

7. List the physical Fibre Channel ports and NPIV attributes by using the **lsnports** command as shown in Example 4-15.

Example 4-15 Listing NPIV capable Fibre Channel ports on VIOS2

```
$ lsnports
name  physloc                         fabric tports aports swwpns  awwpns
fcs0  U5802.001.RCH8497-P1-C5-T1          1     64     64   2048    2046
fcs1  U5802.001.RCH8497-P1-C5-T2          1     64     64   2048    2048
```

NPIV-capable ports have a value of 1 in the fabric column. For Fibre Channel virtualization, select the physical port with the device name fcs1. The physical port fcs0 was used for the SAN LUN masking in “Configuring a second VIOS for client storage” on page 72.

8. Create the Fibre Channel virtualization:

```
vfcmap -vadapter vfchost0 -fcp fcs1
```

9. Verify the virtual Fibre Channel mapping:

```
lsmap -all -npiv
```

Example 4-16 shows the complete listing. The status of the virtual Fibre Channel adapter must be LOGGED_IN.

Checking for new devices: Make sure that the client OS in virtual server VirtServer1 checks for new devices after adding devices in 4.4.1, “Configuring a client virtual server for NPIV” on page 74. In AIX, use the **cfgmgr** command to check for newly added devices.

Example 4-16 Listing virtual Fibre Channel mapping on VIOS

```
$ lsmap -all -npiv
Name          Physloc                            ClntID ClntName       ClntOS
------------- ---------------------------------- ------ -------------- -------
vfchost0      U8233.E8B.10F5D0P-V2-C102              10 VirtServer1    AIX

Status:LOGGED_IN
FC name:fcs1                    FC loc code:U5802.001.RCH8497-P1-C5-T2
Ports logged in:1
Flags:a<LOGGED_IN,STRIP_MERGE>
VFC client name:fcs1            VFC client DRC:U8233.E8B.10F5D0P-V10-C22
```

Now update configuration profiles of VIOS2 Virtual Server. To update the profile of the VIOS2 Virtual Server log on to the SDMC environment and:

1. From the home page, locate the host that contains the VIOS2 virtual server, and click the name of the host.
2. In the Resource Explorer window that opens, select the check box to the left of the virtual server VIOS2, and then select **Actions** → **System Configuration** → **Save Current Configuration**.
3. Select the **Overwrite existing profile** check box, and then select the **OriginalProfile** profile.
4. In the Confirm window, click **OK**.
5. In the Save Profile window, click **Yes**.

5

Advanced configuration

This chapter describes additional configurations to a dual Virtual I/O Server (VIOS) setup and highlights other advanced configuration practices. The advanced setup addresses performance concerns over the single and dual VIOS setup.

This chapter includes the following sections:

- Adapter ID numbering scheme
- Partition numbering
- VIOS partition and system redundancy
- Advanced VIOS network setup
- Advanced storage connectivity
- Shared processor pools
- Live Partition Mobility
- Active Memory Sharing
- Active Memory Deduplication
- Shared storage pools

5.1 Adapter ID numbering scheme

Chapter 3, “Configuring PowerVM with the Hardware Management Console” on page 17 and Chapter 4, “Configuring PowerVM with the IBM Systems Director Management Console” on page 45 showed a combination of default adapter IDs and specific adapter IDs. Because this paper is a quick start guide, default adapter IDs were chosen.

If you are dealing with a large virtual server (client partition) environment, devise an adapter ID numbering scheme. It makes managing your environment easier, particularly when you have multiple Shared Ethernet Adapters (SEAs), a host of virtual SCSI adapters, and virtual Fibre Channel adapters to list. Then configure them back to a network, virtual server, or VIOS.

Table 5-1 shows an example (for one VIOS) of how adapter numbering schemes are used for complex environments. This table shows a two-SEA environment, where one SEA is hosting a VLAN ID of 20, and the other SEA is hosting a group of VLANs using 802.1Q. The table shows three virtual servers, where VirtServer1 has two network interfaces and the C11 adapter hosting VLAN 845. The VirtServer2 C11 adapter is hosting VLAN 865, and the VirtServer3 C11 adapter is hosting VLAN 861.

Each virtual server is configured with one virtual SCSI adapter and two virtual Fibre Channel adapters from VIOS1. A dual VIOS setup shows additional virtual SCSI and virtual Fibre Channel Adapters referencing unique adapter ID slots.

Table 5-1 VIOS1 adapter ID allocation schema example

Virtual adapter	Server adapter ID	VLAN ID	Additional VLANs	Server adapter slot	Client partition or virtual server	Client adapter ID	Client adapter slot
Virtual Ethernet	20	20	No additional	C20			
Virtual Ethernet	21	1	845 865 861	C21			
Virtual Ethernet	22	98	No additional	C22			
Virtual Ethernet	23	99	No additional	C23			
Virtual Ethernet		20		C10	VirtServer1	10	C10
Virtual Ethernet		845		C11	VirtServer1	11	C11
Virtual SCSI	101			C101	VirtServer1	21	C21
Virtual Fibre	103			C103	VirtServer1	23	C23
Virtual Fibre	105			C105	VirtServer1	25	C25
Virtual Ethernet		20		C10	VirtServer1	10	C10
Virtual Ethernet		865		C11	VirtServer1	11	C11

Virtual adapter	Server adapter ID	VLAN ID	Additional VLANs	Server adapter slot	Client partition or virtual server	Client adapter ID	Client adapter slot
Virtual SCSI	111			C111	VirtServer2	21	C21
Virtual Fibre	113			C113	VirtServer2	23	C23
Virtual Fibre	115			C115	VirtServer2	25	C25
Virtual Ethernet		20		C10	VirtServer1	10	C10
Virtual Ethernet		861		C11	VirtServer1	11	C11
Virtual SCSI	121			C121	VirtServer3	21	C21
Virtual Fibre	123			C122	VirtServer3	23	C23
Virtual Fibre	125			C123	VirtServer3	25	C25

5.2 Partition numbering

The VIOS partitions are planned as partition ID 1 and 2, and the virtual servers start at partition ID 10, as explained in the following chapters:

- Chapter 3, “Configuring PowerVM with the Hardware Management Console” on page 17
- Chapter 4, “Configuring PowerVM with the IBM Systems Director Management Console” on page 45

If you have additional infrastructure partitions, number them 3 - 9 so that they are grouped and viewed with the VIOS partitions. Such additional infrastructure partitions might include an IBM Network Installation Manager (NIM) virtual server, IBM Tivoli® Storage Manager virtual server, or additional VIOS partitions.

When maintenance is done on a system, the VIOS partitions are the first partitions to start and the last partitions to shut down.

5.3 VIOS partition and system redundancy

With a redundant VIOS partition setup, you can perform system maintenance on one VIOS partition without affecting the virtual servers. System maintenance might entail a reboot, update, upgrade, or reinstallation of a VIOS partition.

On a VIOS partition level, redundancy can be added by introducing additional physical adapters. For networks, combining additional Ethernet adapters with 802.3ad Link Aggregation provides redundancy and throughput. For storage, using multiple Fibre Channel adapters with Multipath I/O (MPIO) or supported MPIO device drivers provides the necessary redundancy, throughout, and load balancing. The MPIO device drivers include PowerPath for EMC storage and Subsystem Device Driver Path Control Module (SDDPCM) for enterprise IBM System Storage.

Redundancy can be applied to a system which spans multiple I/O drawers and multiple central electronics complex (CEC) servers. Separate CEC loops can be created as explained in the product manual of the system.

If you have a system that spans multiple I/O drawers and CECs, allocate the adapters from different I/O drawers or CECs to a VIOS partition for a highly redundant setup.

5.4 Advanced VIOS network setup

In a VIOS partition, you can apply redundancy to physical network connections by combining multiple physical connections to form a logical Ethernet device. To achieve this redundancy, you use the network communication protocol of Institute of Electrical and Electronics Engineers (IEEE) 802.3ad Link Aggregation (known to network administrators as *Link Aggregation Control Protocol (LACP)*.

Another network protocol that you can apply is IEEE 802.1Q VLAN tagging (also known as *VLAN trunking*). Traditionally, one network port hosted one VLAN. With 802.1Q, now multiple networks can be hosted on a single network port.

For more information about the two network protocols, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

5.4.1 Using IEEE 802.3ad Link Aggregation

Figure 5-1 shows the configuration of an SEA failover setup using dual VIOS with additional Ethernet adapters.

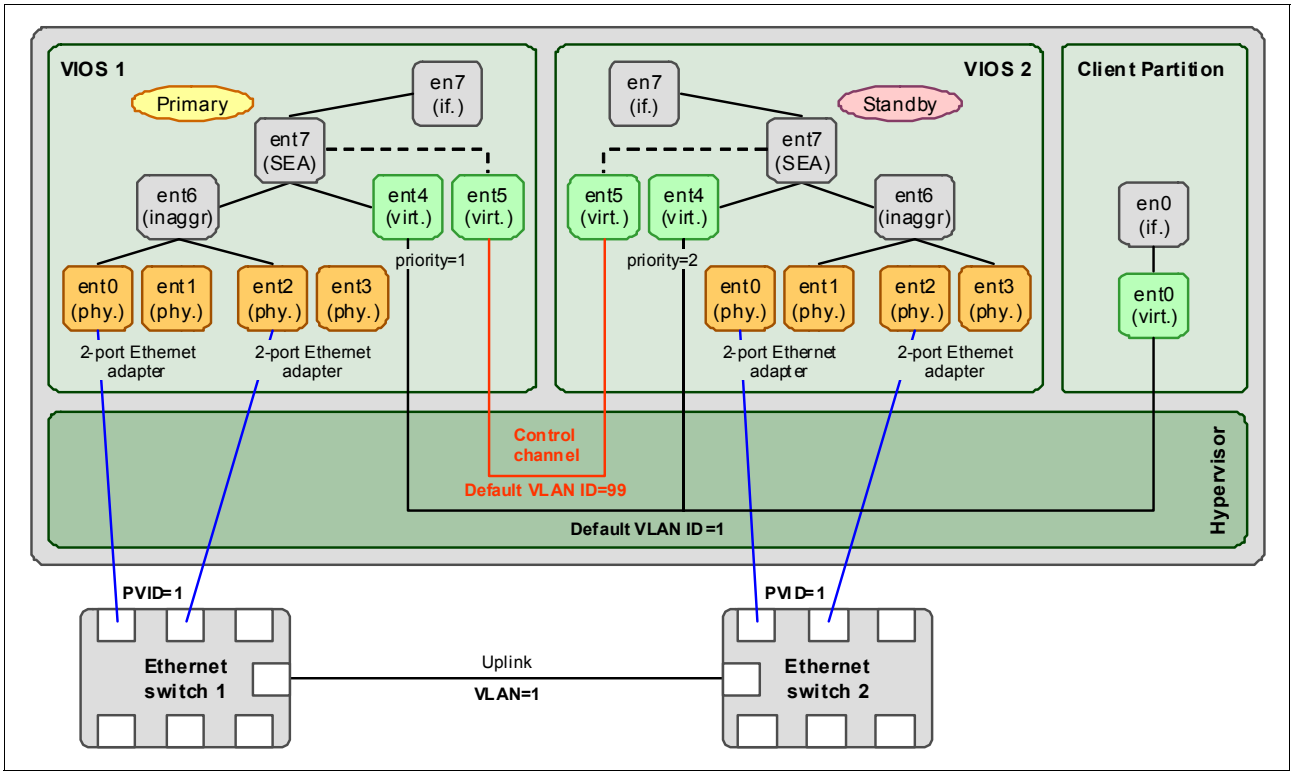

Figure 5-1 Redundant SEA setup

Each VIOS partition is configured with two physical Ethernet ports on different Ethernet adapters. If an Ethernet adapter fails, the external link is still active with the port of the other Ethernet adapter. The two Ethernet adapter ports are linked together to form a logical EtherChannel adapter.

EtherChannel: In this section, EtherChannel is used to describe a link aggregated adapter on VIOS or AIX. In this regard, EtherChannel is not the same as CISCO EtherChannel.

Figure 5-1 on page 82 also shows each VIOS partition connected to a network switch. Linked adapters in the 802.3ad Link Aggregation configuration do not support spanning across multiple network switches if the network switches are in a nonvirtualized state.

To configure a logical EtherChannel adapter, configure 802.3ad Link Aggregation settings on the VIOS partition and on the network switches. Activate Portfast on the network switches to allow faster failover time.

To create a Link Aggregation adapter between physical Ethernet adapters entX and entY, use the following command syntax:

```
mkvdev -lnagg <entX> <entY> -attr mode=8023ad
```

802.3ad setting: You must set the 802.3ad setting on the VIOS partition side and network switch end. The Link Aggregation adapter is nonresponsive if the setting is set on either side only.

5.4.2 Enabling IEEE 802.1Q VLAN tagging

Enable 802.1Q VLAN tagging on the virtual Ethernet adapters that are used for bridging to the physical Ethernet adapters:

1. Open the Virtual Adapters window from the manage profiles window of the VIOS partition, the dynamic logical partition (DLPAR) option, or during VIOS partition profile creation.
2. Depending on whether bridging is required, choose one of the following options:
   - Skip to step 3 if virtual Ethernet adapter for bridging is not present in the Virtual Adapters window.
   - If bridging is required, select the virtual Ethernet adapter, and then select **Actions** → **Edit**. Go to step 4.
3. Select **Actions** → **Create Virtual Adapter** → **Ethernet Adapter**.
4. In the Create Virtual Ethernet Adapter window, select the **IEEE 802.1q compatible adapter** option. Then enter the VLAN ID one at a time, clicking the **Add** button after each VLAN ID.
5. In the Virtual Adapters window, click **OK** to update the virtual Ethernet adapter details.
6. In the Virtual Adapters window, click **OK** to save the information.

5.4.3 Multiple SEA configuration on VIOS

The VIOS partition is not restricted to one SEA. It can host multiple SEAs in these situations:

- A company security policy might advise a separation of VLANs so that one SEA hosts secure networks and another SEA hosts unsecure networks.
- A company might advise a separation of production, testing, and development networks connecting to specific SEA configurations.

5.4.4 General network considerations

Keep in mind the following considerations about the use of IEEE 802.3ad Link Aggregation, 802.1Q VLAN tagging, and SEA:

- An 802.3ad Link Aggregation device has a maximum of eight active ports and eight standby ports.
- Set the speeds of each link in a 803.3ad Link Aggregation device to a common speed setting. For example, set all links to 1g/Full duplex.
- A virtual Ethernet adapter can support up to 20 VLANS (including the Port Virtual LAN ID (PVID)).
- A maximum of 16 virtual Ethernet adapters with 20 VLANS assigned to each adapter can be associated to an SEA.
- A maximum of 256 virtual Ethernet adapters can be assigned to a single virtual server, including the VIOS partitions.
- The IEEE 802.1Q standard supports a maximum of 4,096 VLANS.
- SEA failover is not supported in Integrated Virtualization Manager (IVM) because it only supports a single VIOS partition.

5.5 Advanced storage connectivity

You can physically cable additional Fibre Channel adapter ports in a dual VIOS setup by using either of the following approaches:

- Each VIOS partition can have its Fibre Channel adapter ports connected to the same SAN switch as illustrated in Figure 5-2.

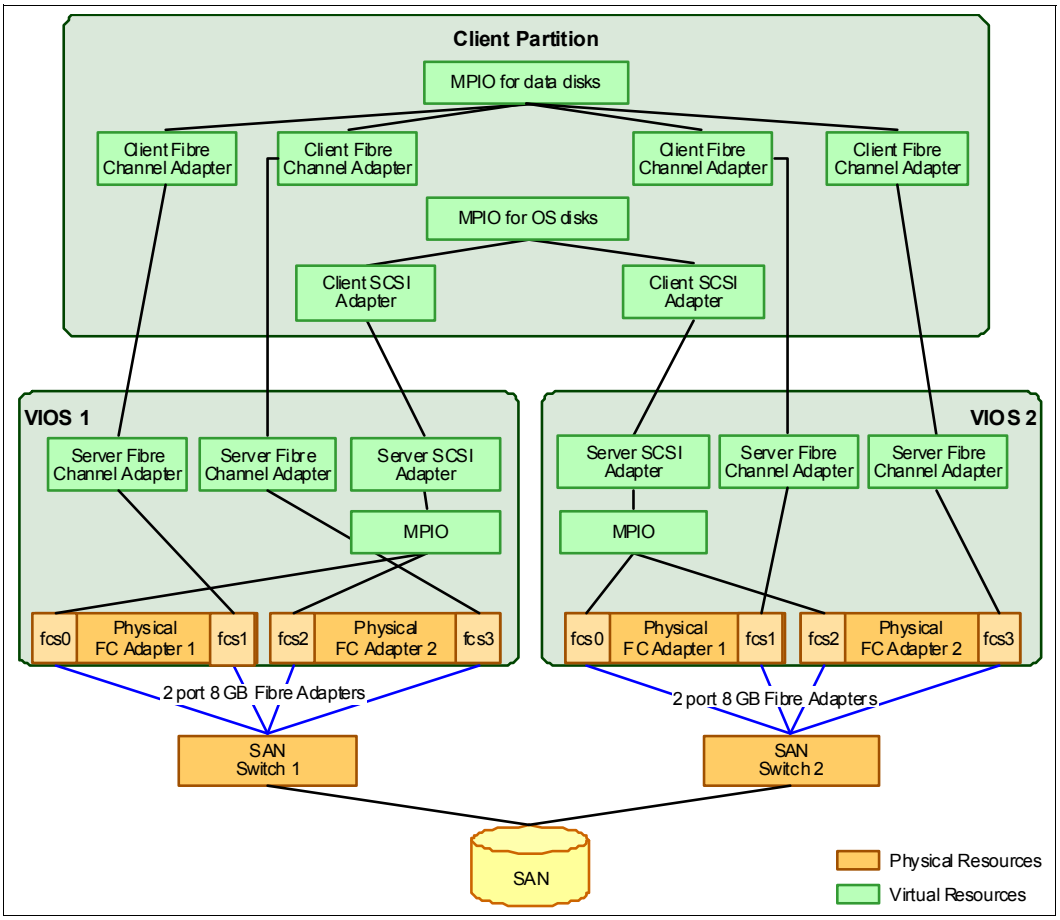

Figure 5-2 MPIO setup where each VIOS partition is connected to one SAN switch

- Each VIOS partition can have its Fibre Channel adapter ports connected to different SAN switches as illustrated in Figure 5-3.

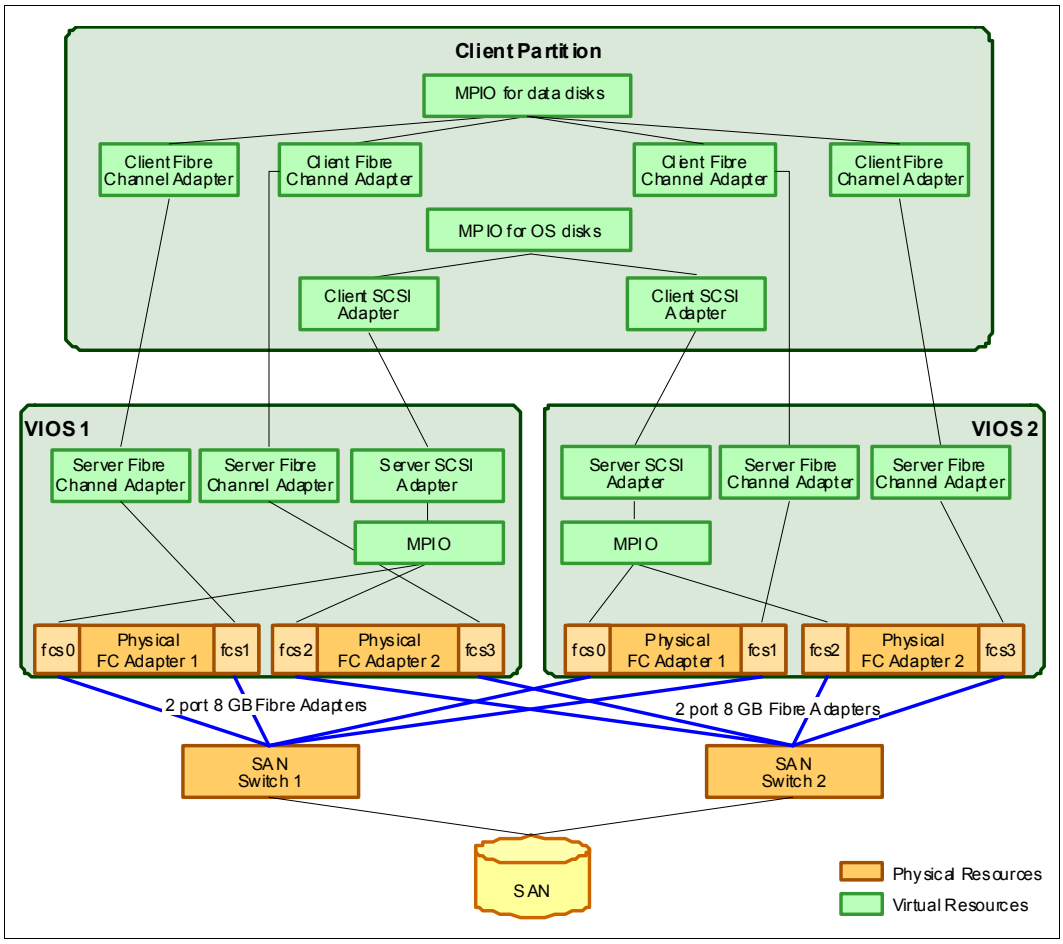

Figure 5-3 MPIO setup where the VIO partitions are connected to the 2 SAN switches

Table 5-2 highlights the benefits and drawbacks of the two approaches.

Table 5-2 Fibre Channel cabling scenarios

Scenario	VIOS partition connected to one SAN switch (Figure 5-2 on page 85)	VIOS Partition connected to two SAN switches (Figure 5-3)
SAN switch 1 is brought down for maintenance.	VIOS1 is unavailable for storage. The LUNS are accessible by using VIOS2.	Storage is available through both VIOS partitions.
SAN switch 1 is misconfigured.	VIOS1 is affected. VIOS2 is unaffected.	VIOS1 and VIOS2 are both impacted and might lose connectivity to the SAN.
Cabling issues	Easier to pinpoint cabling problems because all connections on VIOS1 are connected to SAN switch 1. For VIOS2, all connections are connected to SAN switch 2.	Harder to manage cable issues because VIOS1 and VIOS2 have connections to both SAN switch 1 and 2.

Consider the approach that best serves your environment.

5.6 Shared processor pools

Multiple shared-processor pools (MSPPs) refer to a capability supported on IBM POWER6 servers and later. With this capability, a system administrator can create a set of micropartitions with the purpose of controlling the processor capacity that can be consumed from the physical shared-processor pool.

For more information, see 6.1, “Multiple Shared Processor Pools” in *PowerVM Managing and Monitoring*, SG24-7590 or 2.3.3, “Shared-processor pools” in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940

5.7 Live Partition Mobility

PowerVM Live Partition Mobility allows for the movement of an active (running) or inactive (powered off) partition from one server to another with no application downtime. This method results in higher system utilization, improved application availability, and energy savings. With PowerVM Live Partition Mobility, planned application downtime due to regular server maintenance is no longer a concern.

PowerVM Live Partition Mobility requires systems with POWER6 or later processors running AIX or Linux operating systems and PowerVM Enterprise Edition. For more information about Live Partition Mobility, see *IBM PowerVM Live Partition Mobility*, SG24-7460.

5.8 Active Memory Sharing

IBM Active Memory™ Sharing enables the sharing of a pool of physical memory among AIX, IBM i, and Linux partitions on a single IBM server based on POWER6 technology or later, helping to increase memory utilization and drive down system costs.

The memory is dynamically allocated among the partitions as needed, to optimize the overall physical memory usage in the pool. A dedicated amount of physical memory is not assigned to each logical partition (LPAR) that uses shared memory (referred to as *shared memory partitions*). Instead, the IBM POWER Hypervisor constantly provides the physical memory from the shared memory pool to the shared memory partitions as needed. The POWER Hypervisor provides portions of the shared memory pool that are not being used by shared memory partitions to other shared memory partitions that need to use the memory.

For more information, see 2.4.1, “Active Memory Sharing” in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 or *IBM PowerVM Virtualization Active Memory Sharing*, REDP-4470.

5.9 Active Memory Deduplication

Active Memory Deduplication is a virtualization technology that allows memory pages with identical contents to be deduplicated in physical memory. This method frees physical memory positions so that more data can be held in memory at once.

Memory deduplication works in a shared memory environment. Therefore, it works with Active Memory Sharing. Active Memory Sharing allows multiple partitions on a system to share a pool of physical memory, sometimes creating an over commitment of this physical

memory. Active Memory Deduplication increases the performance of Active Memory Sharing because the savings can be used to either lower memory over commitment levels or to create room to increase the memory footprint of LPARs.

For more information about memory deduplication, see *Power Systems Memory Deduplication*, REDP-4827.

5.10 Shared storage pools

A *shared storage pool* is a pool of SAN storage devices that can span multiple VIOSs. It is based on a cluster of VIOSs and a distributed data object repository with a global namespace. Each VIOS that is part of a cluster represents a cluster node.

The distributed data object repository is using a cluster file system that was developed specifically for storage virtualization using the VIOS. It provides redirect-on-write capability and is highly scalable. The distributed object repository is the foundation for advanced storage virtualization features, such as thin provisioning. Additional features will be added in future releases that will provide significant benefits by facilitating key capabilities for emerging technologies, such as cloud computing.

When using shared storage pools, the VIOS provides storage through logical units that are assigned to client partitions. A *logical unit* is a file backed storage device that is in the cluster file system in the shared storage pool. It appears as a virtual SCSI disk in the client partition, in the same way as a virtual SCSI device, for example, backed by a physical disk or a logical volume.

Shared storage pools are best described and configured in 2.7.2, “Shared storage pools” in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860
- *Integrated Virtualization Manager on IBM System p5*, REDP-4061

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM i Information Center

  http://publib.boulder.ibm.com/infocenter/iseries/v6r1m0/index.jsp
- NIM installation and backup of the VIO server technote

  https://www.ibm.com/support/docview.wss?uid=isg3T1011386#4
- PowerVM QuickStart by William Favorite

  http://www.tablespace.net/quicksheet/powervm-quickstart.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM PowerVM Getting Started Guide

Step-by-step virtualization configuration to the first partition

Single and dual VIOS setups using three common management interfaces

Advanced configuration of a dual Virtual I/O Server setup

IBM PowerVM virtualization technology is a combination of hardware and software that supports and manages virtual environments on IBM POWER5, POWER5+, POWER6, and POWER7 processor-based systems. These systems are available on IBM Power Systems and IBM BladeCenter servers as optional editions, and are supported by the IBM AIX, IBM i, and Linux operating systems. With this set of comprehensive systems technologies and services, you can aggregate and manage resources with a consolidated, logical view.

By deploying PowerVM virtualization and IBM Power Systems, you can take advantage of the following benefits:

- Lower energy costs through server consolidation
- Reduced cost of your existing infrastructure
- Better management of the growth, complexity, and risk of your infrastructure

This IBM Redpaper publication is a quick start guide to help you install and configure a complete PowerVM virtualization solution on IBM Power Systems. It highlights how to use the following management console interfaces to configure PowerVM:

- Integrated Virtualization Manager (IVM)
- Hardware Management Console (HMC)
- Systems Director Management Console (SDMC)

This paper also highlights advanced configuration of a dual Virtual I/O Server setup.

This paper targets new customers who need assistance with quickly and easily installing, configuring, and starting a new PowerVM server in a virtualized environment.

REDP-4815-00

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks