

Integrated Virtualization Manager for IBM Power Systems Servers

**No dedicated Hardware Management
Console required**

**Powerful integration for
entry-level servers**

**Key administration tasks
explained**



**Marina Rodriguez Batalha
Raghavendra K Prasannakumar
Humberto Tadashi Tsubamoto**



International Technical Support Organization

**Integrated Virtualization Manager for IBM Power
Systems Servers**

December 2012

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

Third Edition (December 2012)

This edition applies to IBM Virtual I/O Server Version 2.2.1.4 that is part of every PowerVM edition on IBM Power Systems and BladeCenters. The firmware level that is used is AL730_087 for the IBM POWER7 technology-based server and EL350_132 for the IBM POWER6 technology-based server.

This document was created or updated on December 17, 2012.

© **Copyright International Business Machines Corporation 2012. All rights reserved.**

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

| | |
|---|------|
| Notices | vii |
| Trademarks | viii |
| Preface | ix |
| The team who wrote this paper | ix |
| Now you can become a published author, too! | xi |
| Comments welcome | xi |
| Stay connected to IBM Redbooks | xi |
| Summary of changes | xiii |
| December 2012, Third Edition | xiii |
| Chapter 1. Overview | 1 |
| 1.1 Hardware management | 2 |
| 1.1.1 Integrated Virtualization Manager (IVM) | 2 |
| 1.1.2 Hardware Management Console | 4 |
| 1.1.3 Advanced System Management Interface (ASMI) | 6 |
| 1.2 Integrated Virtualization Manager design | 7 |
| 1.2.1 Architecture | 8 |
| 1.2.2 Logical partition (LPAR) configuration | 9 |
| 1.3 Advanced PowerVM options | 16 |
| 1.3.1 Active Memory Sharing | 16 |
| 1.3.2 Live Partition Mobility | 16 |
| 1.3.3 N_Port ID Virtualization (NPIV) | 16 |
| 1.3.4 Shared Storage Pools | 17 |
| Chapter 2. Plan | 19 |
| 2.1 Considerations for partition setup | 20 |
| 2.2 Reset to Manufacturing Default Configuration | 21 |
| 2.3 Microcode update | 23 |
| 2.4 Advanced System Management Interface IP address setup | 26 |
| 2.4.1 Address setting using the ASMI | 27 |
| 2.4.2 Address setting using serial ports | 28 |
| 2.5 Virtualization feature activation | 29 |
| 2.6 Accessing the server using AMM in a blade environment | 36 |
| 2.7 Accessing the server by using serial link | 37 |
| Chapter 3. Install | 41 |
| 3.1 Virtual I/O Server image installation | 42 |
| 3.1.1 Virtual I/O Server image installation from the NIM server | 44 |
| 3.2 Initial configuration | 44 |
| 3.2.1 Set the date and time | 44 |
| 3.2.2 Initial network setup | 44 |
| 3.2.3 Changing the TCP/IP settings on the Virtual I/O Server | 45 |
| 3.3 Virtual I/O Server partition configuration | 46 |
| 3.4 Network management | 47 |
| 3.5 Virtual Storage management | 48 |
| Chapter 4. Set up | 49 |

| | | |
|------------------|---|------------|
| 4.1 | Configure and manage partitions | 50 |
| 4.2 | IVM graphical user interface | 50 |
| 4.2.1 | Connect to the IVM | 50 |
| 4.2.2 | Guided setup | 52 |
| 4.2.3 | Storage pool disk management | 52 |
| 4.3 | Active Memory Sharing | 61 |
| 4.3.1 | Creating the AMS pool | 62 |
| 4.3.2 | Changing the AMS pool configuration | 68 |
| 4.3.3 | Removing the AMS pool | 69 |
| 4.3.4 | Adding or removing specific paging devices to shared memory pool | 71 |
| 4.4 | Create logical partitions | 79 |
| 4.4.1 | Create an LPAR based on an existing partition | 89 |
| 4.4.2 | Shutting down logical partitions | 91 |
| 4.4.3 | Monitoring tasks | 92 |
| 4.4.4 | Hyperlinks for object properties | 94 |
| 4.5 | IVM command-line interface | 95 |
| 4.5.1 | Update the profile of the logical partition | 95 |
| 4.5.2 | Power on a logical partition | 96 |
| 4.5.3 | Console Terminal for IVM logical partition | 96 |
| 4.5.4 | Install an operating system on a logical partition | 98 |
| 4.6 | Optical device sharing | 99 |
| 4.7 | Logical partition configuration changes | 101 |
| 4.7.1 | Dynamic LPAR operations on an IVM partition | 101 |
| 4.7.2 | LPAR resources management | 106 |
| 4.7.3 | Adding a client LPAR to the partition workload group | 115 |
| 4.8 | Live Partition Mobility | 117 |
| 4.9 | Network management | 123 |
| 4.9.1 | Ethernet bridging | 125 |
| 4.9.2 | Ethernet link aggregation | 127 |
| 4.9.3 | Host Ethernet Adapter in IVM | 129 |
| 4.10 | Storage management | 133 |
| 4.10.1 | Virtual storage assignment to a partition | 133 |
| 4.10.2 | Virtual disk extension | 134 |
| 4.10.3 | IVM system disk mirroring | 135 |
| 4.11 | N_Port Identifier Virtualization (NPIV) | 138 |
| 4.11.1 | Adding a new virtual Fibre Channel adapter | 138 |
| 4.11.2 | Removing a virtual Fibre Channel adapter | 141 |
| 4.11.3 | Changing the physical port of an existing virtual Fibre Channel | 143 |
| 4.12 | Securing the Virtual I/O Server | 144 |
| 4.13 | Connecting to the Virtual I/O Server using OpenSSH | 147 |
| Chapter 5 | Maintain | 151 |
| 5.1 | IVM maintenance | 152 |
| 5.1.1 | Backup and restore of the logical partition definitions | 152 |
| 5.1.2 | Backup and restore of the IVM operating system | 155 |
| 5.1.3 | Backup and restore of the particular set of files in /home/padmin | 156 |
| 5.1.4 | Virtual I/O Server configuration backup/restore | 156 |
| 5.1.5 | IVM updates | 162 |
| 5.1.6 | The migration between HMC and IVM | 168 |
| 5.1.7 | Recovery after an improper HMC connection | 169 |
| 5.1.8 | Migration considerations | 170 |
| 5.1.9 | Migration from an HMC to an IVM environment | 174 |
| 5.1.10 | Migration from an IVM environment to an HMC | 180 |

| | |
|--|------------|
| 5.2 System maintenance | 185 |
| 5.2.1 Microcode update | 185 |
| 5.2.2 Capacity on demand operations | 187 |
| 5.2.3 Logical partition maintenance | 189 |
| 5.2.4 Backup of the operating system | 189 |
| 5.2.5 Restoration of the operating system | 190 |
| 5.3 Command logs | 190 |
| 5.4 Integration with IBM Director | 191 |
| 5.5 User Accounts | 195 |
| 5.6 IBM Electronic Service Agent | 198 |
| 5.7 Service Focal Point | 202 |
| 5.7.1 Manage serviceable events | 202 |
| 5.7.2 Service utilities | 203 |
| 5.8 Application logs | 204 |
| 5.9 Monitoring tasks | 204 |
| 5.10 Hardware inventory | 206 |
| Appendix A. IVM and HMC feature summary | 207 |
| Appendix B. System requirements | 211 |
| Related publications | 213 |
| IBM Redbooks related to PowerVM | 213 |
| IBM Redbooks related to hardware | 213 |
| Online resources | 214 |
| How to get IBM Redbooks | 214 |
| Help from IBM | 214 |

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 grant 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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](http://www.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™ | Global Technology Services® | PowerHA® |
| AIX 5L™ | IBM® | PowerVM® |
| AIX® | Micro-Partitioning® | POWER® |
| BladeCenter® | POWER Hypervisor™ | Redbooks® |
| developerWorks® | Power Systems™ | Redpaper™ |
| Electronic Service Agent™ | POWER6® | Redbooks (logo)  ® |
| Enterprise Workload Manager™ | POWER7 Systems™ | Tivoli® |
| Focal Point™ | POWER7® | |

The following terms are trademarks of other companies:

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

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

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

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

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

Preface

The *Virtual I/O Server (VIOS)* is part of the IBM PowerVM® feature on IBM® Power Systems™ and part of the IBM POWER® Hypervisor™. The VIOS is also supported on IBM BladeCenter®. The server is a single-function appliance that is in the *logical partition (LPAR)* of the Power Systems server. The VIOS facilitates the sharing of physical I/O resources between client partitions (IBM AIX®, IBM i, and Linux) within the server.

The Virtual I/O Server provides a virtual SCSI target, N_Port ID Virtualization (NPIV) target, and Shared Ethernet Adapter (SEA) virtual I/O function to client LPARs.

The Virtual I/O Server has the capability of a hardware management function, the *Integrated Virtualization Manager (IVM)*. IVM is a simplified hardware management solution that inherits most of the *Hardware Management Console (HMC)* features. The console manages a single server, avoiding the need of a dedicated personal computer. This device is designed to provide a solution that enables the administrator to reduce system setup time and to make hardware management easier, at a lower cost.

IVM provides a management model for a single system. Although it does not offer all of the HMC capabilities, it enables the exploitation of PowerVM technology. IVM targets the small and medium systems that are best suited for this product.

IVM provides the following functions:

- ▶ Shared Storage Pool
- ▶ IBM Active Memory™ Sharing (AMS)
- ▶ Live Partition Mobility (LPM)
- ▶ Task manager monitor for long-running tasks
- ▶ Security additions such as viosecure and firewall, and other improvements

There are many environments where there is the need for small partitioned systems, either for test reasons or for specific requirements, for which the HMC solution is not ideal. A sample situation is where there are small partitioned systems that cannot share a common HMC because they are in multiple locations. In these cases, IVM works.

Using IVM, companies can more cost-effectively consolidate multiple partitions onto a single server. With its intuitive, browser-based interface, the IVM is easy to use and significantly reduces the time and effort that is required to manage virtual devices and partitions.

This IBM Redpaper™ publication provides an introduction to IVM by describing its architecture and showing how to install and configure a partitioned server by using its capabilities. This document is intended for IT personnel who have a complete understanding of partitioning before reading this document.

The team who wrote this paper

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

Marina Rodriguez Batalha is a Client Technical Specialist (Field Technical Sales Specialist role) in IBM Brazil. She has six years of experience in IBM Power Systems servers. She holds a Technology degree in Data Processing from Centro Paula Souza (FATEC - Faculdade de Tecnologia). Her areas of expertise include operating systems such as IBM AIX, IBM i, Linux, and IBM PowerVM Virtualization implementation and administration.

Raghavendra K Prasannakumar is a Staff Software Engineer in IBM India. He has five years of experience in system testing domain since joining IBM. He holds a Bachelor's degree in Electronics and Communication Engineering from Visvesvaraya Technological University (VTU). His areas of expertise include system testing on IBM PowerVM, IBM AIX, Hardware Management Console, and storage area network (SAN). He wrote an IBM developerWorks® article on Live Partition Mobility. He is working as a test lead on the Integrated Software System test team.

Humberto Tadashi Tsubamoto is an IT Specialist in IBM Global Technology Services® in IBM Brazil. He has 13 years experience with UNIX and seven years of experience with storage systems. Humberto has worked at IBM for seven years. He holds a Bachelor's degree in Electrical Engineering from Escola de Engenharia Maua (EEM). His areas of expertise include implementation, support, and performance analysis of IBM Power Systems, IBM PowerHA®, IBM PowerVM, IBM AIX, Solaris, SAN, and network-attached storage (NAS).

The project that produced this paper was managed by:

Scott Vetter
IBM Austin

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

Richard M Conway, Ann Lund, and Linda Robinson
IBM International Technical Support Organization (ITSO)

Ivan Berrios and Cesar Diniz Maciel
IBM Atlanta

Bob G. Kovacs, Jim Partridge, Jim Parumi, Amartey S. Pearson, and Vani D. Ramagiri
IBM Austin

Rodrigo Mello Montagnana and Fernando Vermelho
IBM Brazil

Dennis Jurgensen
IBM Raleigh

Jaya Srikrishnan
IBM Poughkeepsie

Craig Wilcox
IBM Rochester

Peter Wuestefeld and Volker Haug
IBM Germany

Morten Vagmo
IBM Norway

Dai Williams and Nigel Griffiths
IBM U.K.

Thanks to the authors of the previous editions of this paper:

- ▶ Authors of the first edition, *Integrated Virtualization Manager for IBM Power Systems Servers*, published in October 2005, were:
 - Nicolas Guerin
 - Federico Vagnini
- ▶ Author of the second edition, *Integrated Virtualization Manager for IBM Power Systems Servers*, published in December 2006, was: Guido Somers.

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>

Summary of changes

The *summary of changes* section describes the technical changes that are made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for Integrated Virtualization Manager for IBM Power Systems Servers
as created or updated on December 17, 2012.

December 2012, Third Edition

This edition reflects the addition and modification of information as described in the following sections.

New information

The following topics highlight new features that were added since the last edition:

- ▶ This publication describes the configuration of newer virtualization features, such as N_Port ID Virtualization (NPIV), Active Memory Sharing (AMS), and Live Partition Mobility (LPM) from Integrated Virtualization Manager (IVM).
- ▶ This paper explains how to take backup and restore of the Virtual I/O Server configuration (**viosbr**).
- ▶ Use of the call home feature for hardware and software errors, by using IVM.
- ▶ Accessing Host Ethernet Adapter (HEA) by using IVM for Virtual I/O Server and VIO client LPAR.
- ▶ Accessing the server by using a serial port for IVM installation.

Changed information

The following topics highlight the major features that changed since the last edition:

- ▶ Dynamic Logical Partitioning (DLPAR) facility to physical adapters and HEA in Virtual I/O Server and VIO client LPARs.
- ▶ Using physical adapters to VIO client LPAR by using IVM.
- ▶ Taking backup and restore of important files of /home/padmin by using **File and Virtual Media Backup/Restore**.
- ▶ Revision of the previous content.



Overview

This chapter describes several available methods for hardware management and virtualization setup on IBM Power Systems and also introduces the *Integrated Virtualization Manager (IVM)*.

The IVM is a feature of the Virtual I/O Server, which is part of the IBM PowerVM technology. This device enables companies to consolidate multiple partitions onto a single server in a cost-effective way. With its intuitive and browser-based interface, the IVM is easy to use and significantly reduces the time and effort that is required to manage virtual devices and partitions.

1.1 Hardware management

With the use of virtualization techniques, hardware management has become more of an independent task. Operating systems have a less direct visibility and control over physical server hardware; therefore, system administrators must now focus on the management of resources that are assigned to them.

To be independent from operating system issues, hardware management requires a separate computing environment that can access, configure, control, monitor, and maintain the server hardware and firmware. This environment requires advanced platform management applications that provide the following functions:

- ▶ Server configuration before deployment of operating system
- ▶ Service when operating systems are unavailable
- ▶ Coordination of platform-related operations across multiple operating system images, within an independent security model
- ▶ Presentation of virtual operating system consoles

IBM developed several solutions for hardware management that target different environments which depend on the complexity of hardware setup.

1.1.1 Integrated Virtualization Manager (IVM)

The *Hardware Management Console (HMC)* is the comprehensive solution for hardware management that can be used either for a small configuration or for a multiserver environment. Complexity is kept low by design and many recent software revisions support this. However, the HMC solution might not fit in small and simple environments where only a few servers are deployed or not all HMC functions are required.

There are many environments where there is the need for small partitioned systems, either for test reasons or for specific requirements, for which the HMC solution is not ideal. A sample situation is where there are small partitioned systems that cannot share a common HMC because they are in multiple locations.

Integrated Virtualization Manager (IVM) is a simplified hardware management solution that inherits most of the HMC features. It manages a single server, avoiding the need of an independent personal computer. It is designed to provide a solution that enables the administrator to reduce system setup time and to make hardware management easier, at a lower cost.

IVM provides a management model for a single system. Although it does not offer all of the HMC capabilities, it enables the use of IBM PowerVM technology. IVM targets the small and medium systems that are best suited for this product. Table 1-1 lists the systems that are supported at the time of writing.

Table 1-1 Supported server models for IVM

| IBM Power Systems | Supported server models |
|----------------------|---|
| IBM POWER5 Systems | All server models, except p5-575, p5-590, and p5-595 |
| IBM POWER6@ Systems | All server models, except Power 575, Power 590, and Power 595 |
| IBM POWER7@ Systems™ | All server models, except Power 755, Power 770, Power 775, Power 780, and Power 795 |

| IBM Power Systems | Supported server models |
|-------------------|--|
| IBM BladeCenter | All POWER6 and POWER7 processor-based blades, and IBM BladeCenter JS21 |

IVM supported server models: For more information about IVM supported server models, see this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hch/iphchs/uppportedsystems.htm>

IVM is an enhancement of the *Virtual I/O Server*, the product that enables I/O virtualization in IBM Power Systems. It enables management of Virtual I/O Server functions and uses a web-based graphical interface that enables the administrator to remotely manage the server with a browser. The HTTPS protocol and server login with password authentication provide the security that is required by many enterprises.

Because one of the goals of IVM is simplification of management, some implicit rules apply to configuration and setup:

- ▶ When a system is designated to be managed by IVM, it must not be partitioned.
- ▶ The first operating system to be installed *must* be the Virtual I/O Server.

The Virtual I/O Server is automatically configured to own all of the I/O resources. The server can be configured to provide service to other *logical partitions (LPARs)* through its virtualization capabilities. However, all other LPARs can have either physical or virtual adapters for disk access, network, and optical devices.

Figure 1-1 shows a sample configuration that uses IVM. In this case, the Virtual I/O Server owns some of the physical adapters. LPAR1 is configured to use a virtual device and physical devices and LPAR2 is configured to use only virtual devices. The virtual devices are provided by the Virtual I/O Server. You can use either physical or virtual devices in the configuration of the LPAR. The administrator can use a browser to connect to the IVM to set up the system configuration.

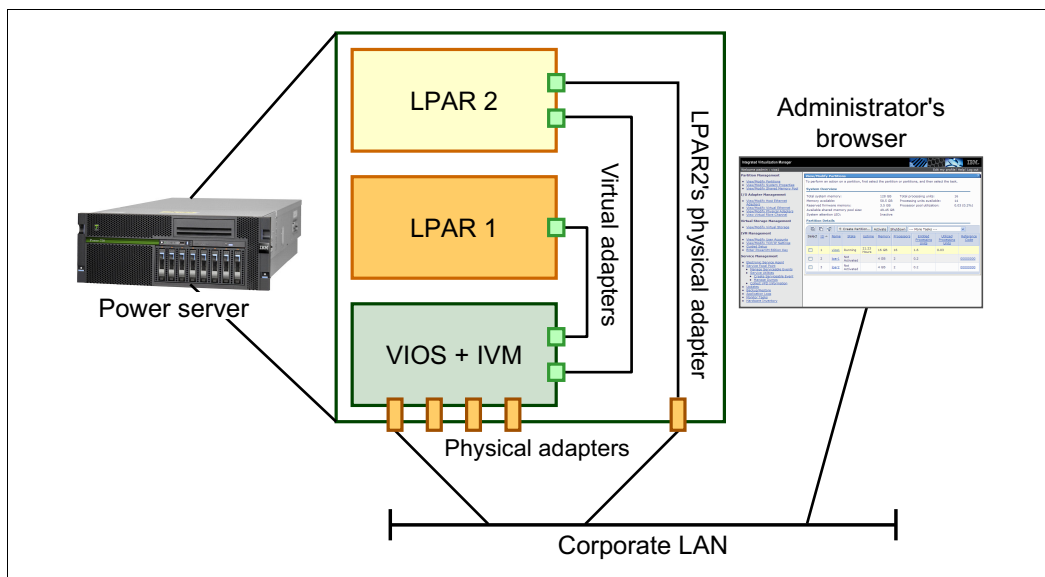


Figure 1-1 Integrated Virtualization Manager sample configuration

As Figure 1-1 on page 3 depicts, the system hypervisor is modified to enable the Virtual I/O Server to manage the partitioned system without an HMC. The software that is normally running on the HMC is rewritten to fit inside the Virtual I/O Server and to provide a simpler user interface. Because the IVM is running by using system resources, the design is developed to have a minimal affect on disk, memory, and processor resources.

The IVM does not interact with the service processor of the system. A specific device named *Virtual Management Channel (VMC)* was developed on the Virtual I/O Server to enable a direct hypervisor configuration without requiring more network connections. This device is activated, by default, when the Virtual I/O Server is installed as the first partition.

The VMC enables IVM to provide basic logical partitioning functions:

- ▶ Logical partitioning configuration
- ▶ Boot, start, and stop actions for individual partitions
- ▶ Display of partition status
- ▶ Management of virtual Ethernet
- ▶ Management of virtual storage
- ▶ Basic system management

Because IVM runs on an LPAR, there are limited service-based functions, and *Advanced System Management Interface (ASMI)* must be used. For example, power on the server by physically pushing the server power-on button or remotely accessing ASMI because IVM does not run while the server power is off. ASMI and IVM together provide a simple but effective solution for a single partitioned server.

LPAR management that uses IVM is through a common web interface that is developed for basic administration tasks. Being integrated within the Virtual I/O Server code, IVM also handles all virtualization tasks that normally require Virtual I/O Server commands to be run.

IVM interface with HMC: The IVM provides a unique setup and interface regarding the HMC for managing resources and partition configuration. An HMC expert needs to become familiar with the differences before using the IVM.

IVM has support for dynamic LPAR, starting with Version 1.3.0.0.

IVM and HMC are two unique management systems: The IVM is designed as an integrated solution designed to lower your cost of ownership, and the HMC is designed for flexibility and a comprehensive set of functions. This flexibility provides you the freedom to select the ideal solution for your production workload requirements.

Moving from IVM to HMC: The internal design of IVM requires that no HMC is connected to a working IVM system. If a client wants to migrate an environment from IVM to HMC, the configuration setup must be rebuilt manually. This setup includes systems that had previous software levels of Virtual I/O Server running on them because they were also managed by an HMC.

1.1.2 Hardware Management Console

The primary hardware management solution that is developed by IBM relies on an appliance server named the *Hardware Management Console (HMC)*, packaged as an external tower or rack-mounted personal computer. It is deployed on all Power Systems, each with its own specific set of management tools. Great effort is made to improve their functions and ease of use.

The HMC is a centralized point of hardware control. In a Power Systems environment, a single HMC can manage multiple POWER processor-based systems. Two HMCs can manage the same set of servers in a dual-active configuration that is designed for high availability.

Hardware management is performed by an HMC using a standard Ethernet connection to the service processor of each system. Interacting with the service processor, the HMC can modify the hardware configuration of the managed system, querying for changes, and managing service calls.

A hardware administrator can log in to the physical HMC and use the native graphical user interface (GUI). Or, the administrator can connect remotely to the HMC by using a web browser, which has the same presence of the native GUI.

Because it is a stand-alone personal computer, the HMC does not use any managed system resources and can be maintained without affecting system activity. Restarts and software maintenance on the HMC do not have any affect on the managed systems.

In the unlikely case that the HMC requires manual intervention, the systems continue to be operational and a new HMC can be plugged into the network and configured to download the current configuration from the managed systems. The new HMC becomes operationally identical to the replaced HMC.

The following list highlights the major HMC functions:

- ▶ Monitoring of system status
- ▶ Management of IBM capacity on demand (CoD)
- ▶ Creation of logical partitioning with dedicated processors
- ▶ Management of LPARs including power on, power off, and console
- ▶ Dynamic reconfiguration of partitions
- ▶ Management of virtual Ethernet among partitions
- ▶ Clustering
- ▶ Concurrent firmware updates
- ▶ Hot add and remove of I/O drawers

POWER processor-based systems are capable of IBM Micro-Partitioning®. The IBM POWER Hypervisor can support multiple LPARs, sharing the processors in the system and enabling I/O sharing. Power Systems servers require a PowerVM Edition feature.

On systems with Micro-Partitioning enabled, the HMC provides more functions:

- ▶ Creation of shared processor partitions
- ▶ Creation of the Virtual I/O Server partition for physical I/O virtualization
- ▶ Creation of virtual devices for Virtual I/O Server and client partitions

The HMC interacts with the hypervisor to create virtual devices among partitions, and the Virtual I/O Server partitions manage physical device sharing. Network, disk, and optical device access can be shared.

Partition configuration can be changed dynamically by issuing commands on the HMC or by using the HMC GUI. The allocation of resources, such as CPU, memory, and I/O, can be modified without making applications aware of the change.

To enable dynamic reconfiguration, an HMC requires TCP/IP communication with every involved LPAR besides the basic connection with the service processor. Using a Resource Monitoring and Control (RMC) protocol, the HMC can securely interact with the operating system to free and acquire resources and to coordinate these actions with hardware configuration changes.

The HMC also provides tools to ease problem determination and service support, such as the *Service Focal Point* feature, call-home, and error log notification through a modem or the Internet.

1.1.3 Advanced System Management Interface (ASMI)

Major hardware management activity is done by interacting with the service processor that is installed on all POWER processor-based systems. The HMC has access to the service processor through Ethernet and uses it to configure the system Hypervisor.

The service processor can be locally accessed through a serial connection by using system ports when the system is powered down. The processor can be remotely accessed in either power standby or powered-on modes. This access can occur by using an HTTPS session with a web browser that points to the IP address that is assigned to the Ethernet ports of the service processor.

The web GUI is called the *Advanced System Management Interface (ASMI)*, as shown in Figure 1-2.

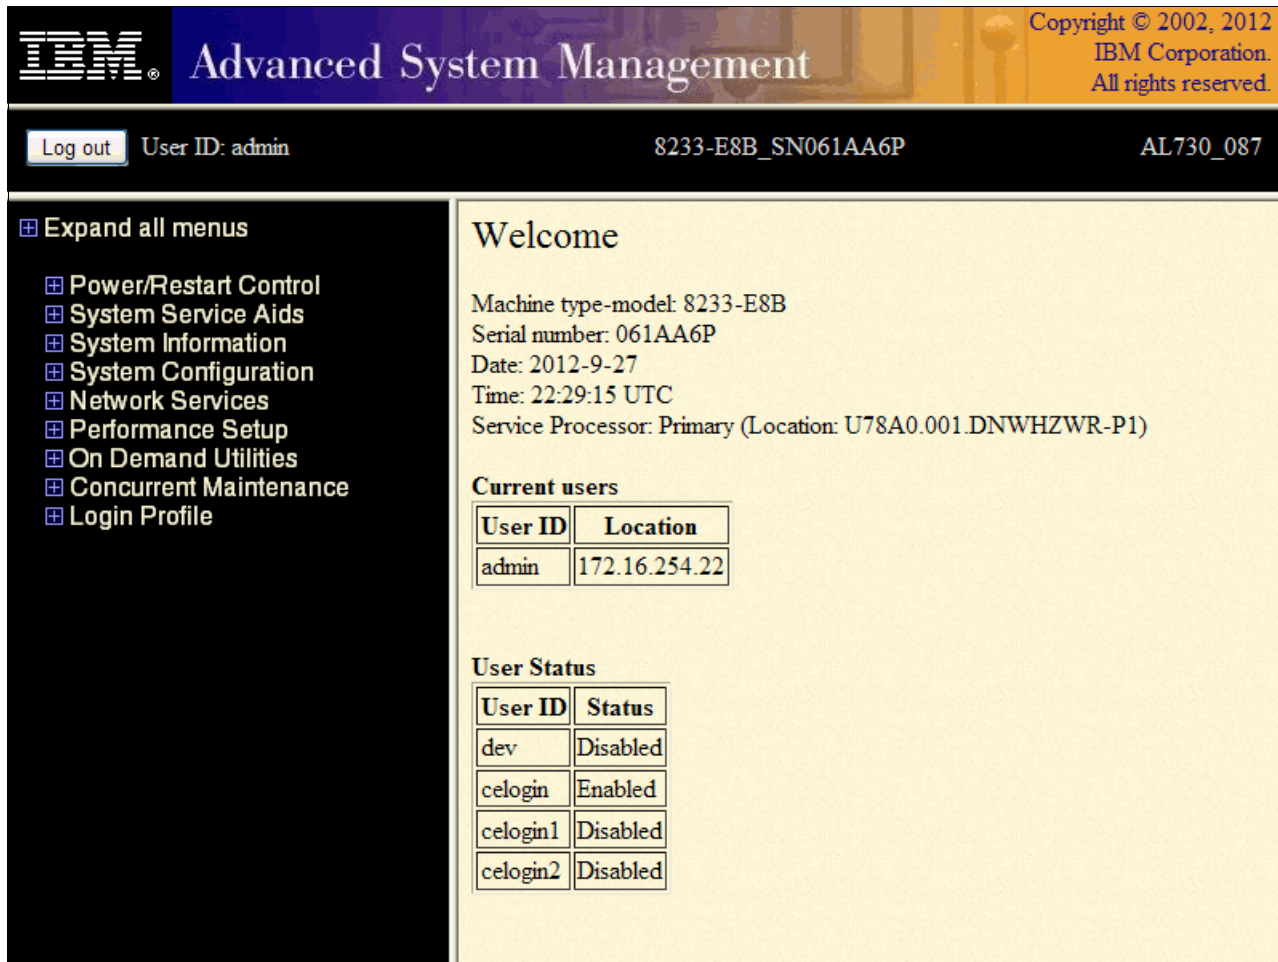


Figure 1-2 Advanced System Management Interface

ASMI is the major configuration tool for systems that are not managed by an HMC. The tool provides basic hardware setup features and is useful when the system is a stand-alone system. ASMI can be accessed and used when the HMC is connected to the system, but some of its features are disabled.

Using ASMI, the administrator can run the following basic operations:

- ▶ View system information
- ▶ Control system power
- ▶ Change the system configuration
- ▶ Set performance options
- ▶ Configure the network services of the service processor
- ▶ Use on-demand utilities
- ▶ Use concurrent maintenance utilities
- ▶ Run system service aids, such as accessing the error log of the service processor

The scope of every action is restricted to the same server. If there are multiple systems, the administrator must contact each of them independently, each in turn.

After the initial setup, typical ASMI usage is remote system power-on and power-off. The other functions are related to system configuration changes, for example, virtualization feature activation and troubleshooting, such as access to the logs of the service processor.

The ASMI does not allow LPARs to be managed. To deploy LPARs, a higher level of management is required, going beyond basic hardware configuration setup. This management can be done either with an HMC or by using the IVM.

1.2 Integrated Virtualization Manager design

Power Systems can be partitioned because they are all preinstalled with all of the necessary firmware support for a partitioned environment.

Because the partitioning schema is designed by the client, every system is set up by manufacturing in the same *Manufacturing Default Configuration* that can be changed or reset to when required.

Although configured by using the Manufacturing Default Configuration, the system has the following setup from a partitioning point of view:

- ▶ There is a single predefined partition.
- ▶ All hardware resources are assigned to the single partition.
- ▶ The partition has system *service authority*, so it can update the firmware.
- ▶ The partition is auto-started at power-on.
- ▶ Standard operating system installation methods apply for the partition (network or media-based).
- ▶ The physical control panel of the system is mapped to the partition, displaying its operating system messages and error codes.
- ▶ Base platform management functions, such as power control, are provided through integrated system control functions (for example, service processor and control panel).

The Manufacturing Default Configuration enables the system to be used immediately as a stand-alone server with all resources allocated to a single LPAR. If an HMC is attached to the service processor of a Power System, the system configuration can be changed to make the hypervisor manage multiple LPARs.

When an HMC is not available and the administrator wants to use virtualization features, the IVM can be used.

1.2.1 Architecture

The IVM was developed to provide a simple environment where a single control program has the ownership of the physical hardware and other LPARs use it to access resources.

The Virtual I/O Server has most of the required features because it can provide virtual disk and virtual networking capability. The current version of the Virtual I/O Server comes with several IVM features. The server includes features such as: dynamic LPAR-capability of the client LPARs, security improvements (firewall, viosecure), usability additions (TCP/IP GUI configuration, hyperlinks, simple LPAR creation, task monitor, to name a few), Active Memory Sharing (AMS), N_Port ID Virtualization (NPIV), and Live Partition Mobility (LPM).

To set up LPARs, the IVM requires management access to the hypervisor. It has no service processor connection that is used by the HMC and it relies on a new virtual I/O device type called *Virtual Management Channel (VMC)*. This device is activated only when the Virtual I/O Server installation detects that the environment must be managed by IVM.

VMC is present on Virtual I/O Server only when the following conditions are true:

- ▶ The virtualization feature is enabled.
- ▶ The system has not been managed by an HMC.
- ▶ The system is in Manufacturing Default Configuration.

To fulfill these requirements, an administrator must use the ASMI. By using the ASMI, they can enter the virtualization activation code, reset the system to the Manufacturing Default Configuration, and so on. A system reset removes any previous LPAR configuration and any existing HMC connection configuration.

On a Virtual I/O Server partition with IVM activated, a new `ibmvmc0` virtual device is present and a management web server is started. It listens to HTTP port 80 and to HTTPS port 443. The presence of the virtual device can be detected by using the `lsdev -virtual` command, as shown in Example 1-1.

Example 1-1 Virtual Management Channel device

```
$ lsdev -virtual | grep ibmvmc0
ibmvmc0          Available Virtual Management Channel
```

Because IVM relies on VMC to set up logical partitioning, it can manage only the system on which it is installed. For each IVM managed system, the administrator must open an independent web browser session.

Figure 1-3 on page 9 provides the schema of the IVM architecture. The primary user interface is a web browser that connects to port 80 of the Virtual I/O Server. The web server provides a simple GUI and runs commands by using the same command-line interface that can be used for logging in to the Virtual I/O Server. One set of commands provides LPAR management through the VMC, and a second set controls for Virtual I/O Server virtualization capabilities. It also enables Secure (encrypted) Shell (SSH) access.

Figure 1-3 shows the integration with IBM Director (Pegasus CIM Server).

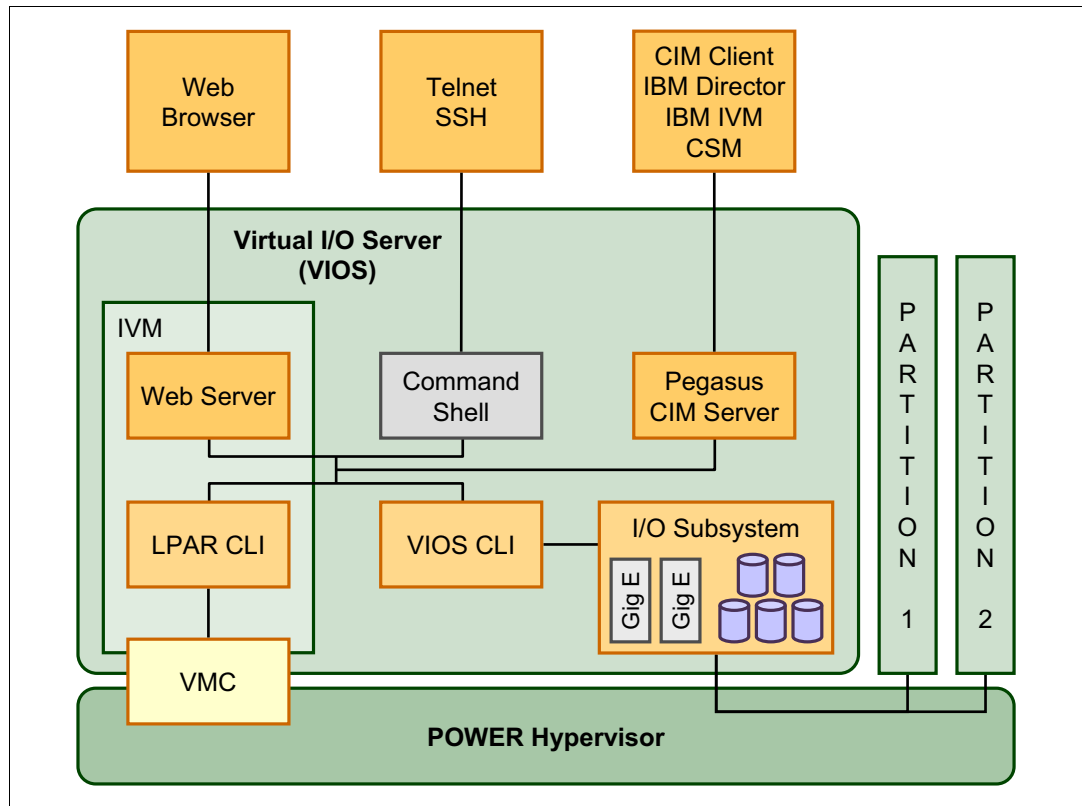


Figure 1-3 IVM high-level design

LPARs in an IVM managed system are isolated exactly as in an HMC managed system and they can talk using virtual and physical devices using normal TCP protocol. IVM is used to perform the following actions on the other LPARs such as:

- ▶ Activate and deactivate
- ▶ Send an emergency power-off warning (EPOW) signal to the operating system
- ▶ Create and delete
- ▶ View and change configuration
- ▶ Dynamic logical partitioning of processor, memory, physical, and virtual I/O
- ▶ To enable processor and memory sharing
- ▶ Migrate a partition to another IVM

1.2.2 Logical partition (LPAR) configuration

The simplification of the user interface of a single partitioned system is one of the primary goals of the IVM. LPAR management is designed to enable a quick deployment of partitions. Compared to HMC managed systems, configuration flexibility is reduced to provide a basic usage model. A new user with no HMC skills can easily manage the system in an effective way.

LPAR configuration is made by assigning processor, memory, and virtual I/O by using a web graphical user interface (GUI) wizard. At each step of the process, the administrator is asked simple questions, which provide the range of possible answers. Most of the parameters that are related to LPAR setup are hidden during creation time. This process is done to ease the setup and can be finely tuned, changing partition properties if needed after the initial setup.

Resources that are assigned to an LPAR are immediately allocated and are no longer available to other partitions, regardless of whether the LPAR is activated or powered down.

This behavior makes management more direct and it is a change when compared to HMC managed systems where resource over-commitment is allowed.

It is important to understand that any unused processor resources do become available to other partitions through the shared pool. This process happens when any LPAR is not using all of its processor entitlement.

System configuration is described in the GUI, as shown in Figure 1-4. In this example, an unbalanced system is manually prepared as a specific scenario. The system has 128 GB of global memory, 16 processing units, and four LPARs defined. In the Partition Details panel, the allocated resources are shown in terms of memory and processing units. Even if the LPAR2 and LPAR3 partitions are not activated, their resources are allocated and the available memory and processing units of the system are updated accordingly.

If a new LPAR is created, it cannot use the resources that belong to a powered-off partition. However, the partition can be defined by using the available free resources that are shown in the System Overview panel.

The processing units for the LPAR, named LPAR1 (ID 2), are changed from the default of 0.2, which was created by the wizard, to a processing unit of 0.1. LPAR1 can use up to one processor because it has one virtual processor and is guaranteed to use up to 0.1 processing units. Figure 1-4 depicts this system configuration status.

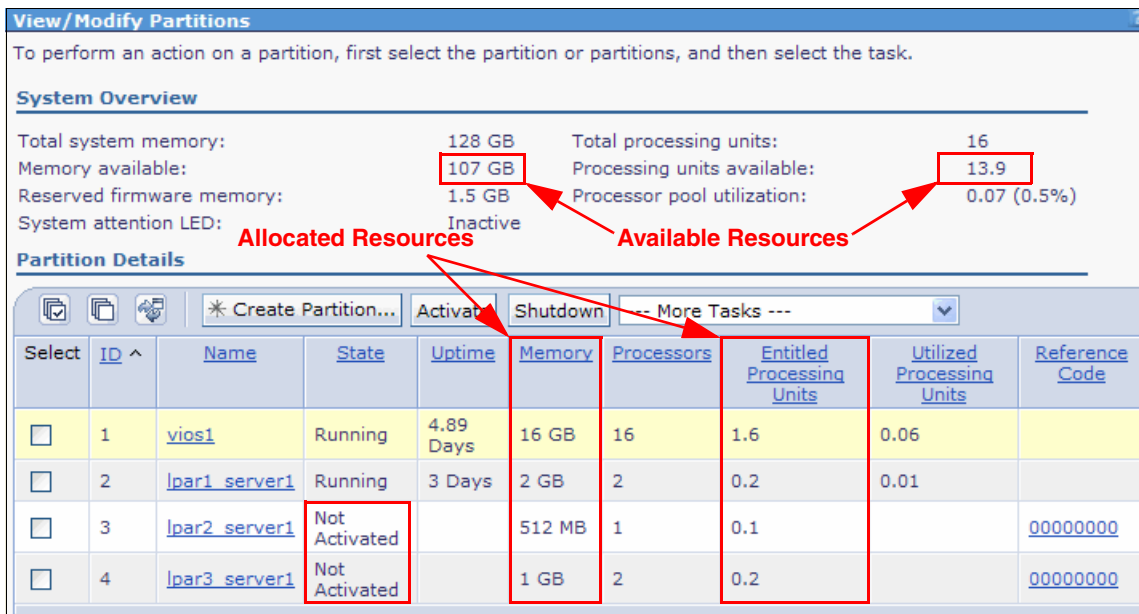


Figure 1-4 System configuration status

Memory

Memory is assigned to an LPAR by using available memory on the system with an allocation unit size that can vary from system to system, depending on its memory configuration. The wizard provides this information, as shown in Figure 1-5.

| | |
|-----------------------|---|
| Name | Memory |
| •• Memory | In dedicated mode, the partition uses assigned memory from total system memory. In shared mode, the partition uses the memory from the system shared memory pool. |
| Processors | |
| Ethernet | |
| Storage Type | You cannot create a partition that uses shared memory because there is no shared memory pool defined for the system. If you want to assign shared memory for the partition, use the View/Modify System Properties: Memory tab to exit the wizard and create a shared memory pool to enable shared memory on the system. |
| Storage | |
| Virtual Fibre Channel | |
| Optical/Tape | |
| Physical Adapters | If you want to assign dedicated memory for the partition, specify the amount of memory, in multiples of 64 MB, to assign for the partition. |
| Summary | Note: If you specify a number that is not a multiple of 64 MB, the wizard will round the number to the nearest multiple of 64 MB . |
| | Memory Mode |
| | <input checked="" type="radio"/> Dedicated |
| | <input type="radio"/> Shared |
| | Dedicated Mode |
| | Total system memory: 128 GB (131072 MB) |
| | Current memory available for partition usage: 107.19 GB (109760 MB) |
| | Assigned memory: <input type="text" value="512"/> MB <input type="button" value="v"/> |

Figure 1-5 Memory allocation to LPAR

The minimum allocation size of memory is related to the *logical memory block* (LMB) size of the system. The size is defined automatically at the start of the system, depending on the size of physical memory. However, it can be changed by using ASMI on the Performance Setup menu, as shown in Figure 1-6. The default automatic setting can be changed to the following values: 16 MB, 32 MB, 64 MB, 128 MB, or 256 MB.

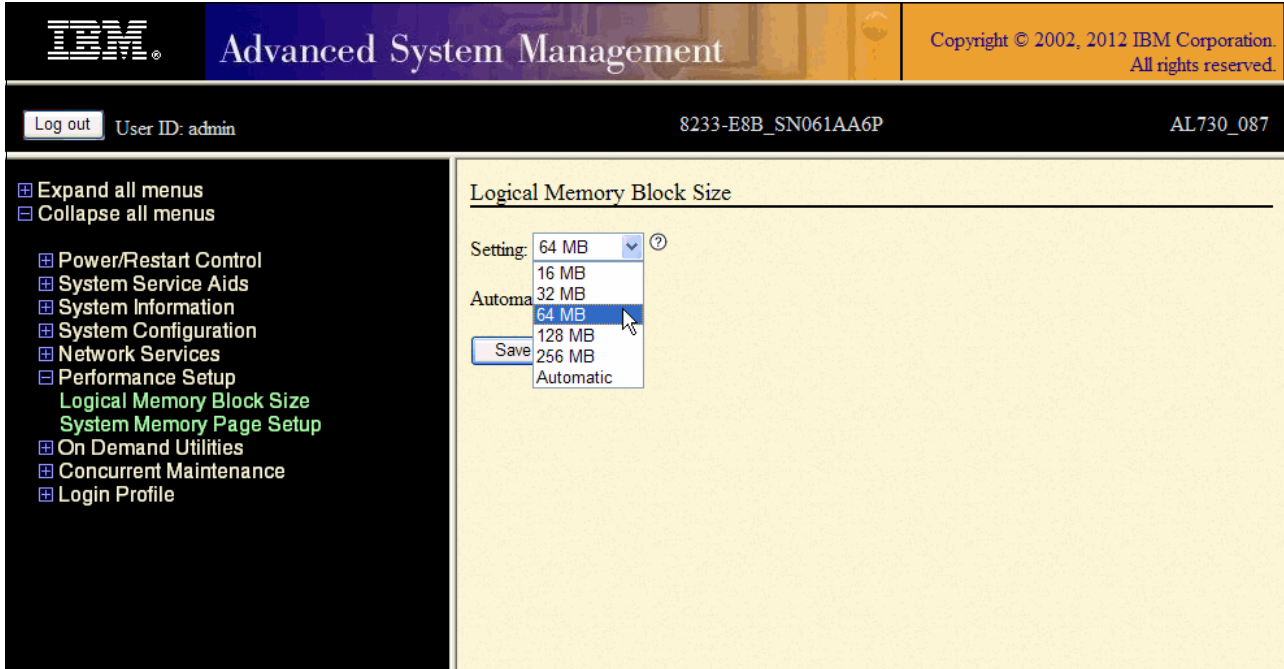


Figure 1-6 Logical Memory Block size setup

You can also change the value from IVM View/Modify **Shared Memory Pool** → **Memory** → **Memory Region Size**, as shown in Figure 1-7.

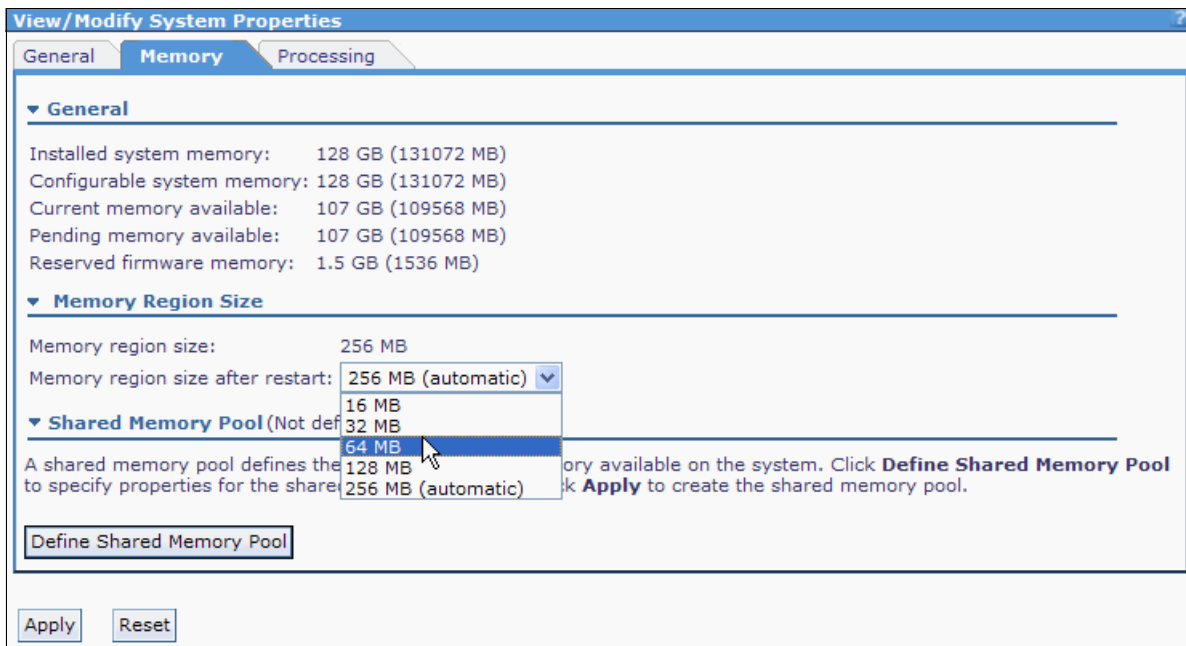


Figure 1-7 Logical Memory Block size setup from IVM

To change the LMB setting, the entire system must be shut down; you can do this from ASMI. If an existing partition has a memory size that does not fit in the new LMB size, the memory size is changed to the nearest value that can be allowed by the new LMB size. However, the value cannot exceed the original memory size.

A small LMB size provides a better granularity in memory assignment to the partitions, but requires higher memory allocation and deallocation times. This increase is because more operations are required for the same amount of memory. Larger LMB sizes can slightly increase the firmware reserved memory size. It is suggested to keep the default automatic setting.

Processors

An LPAR can be defined either with dedicated or with shared processors. The wizard provides available resources in both cases and asks which processor resource type to use.

When shared processors are selected for a partition, the wizard asks only the administrator to choose the number of virtual processors to be activated. The maximum value is equal to the number of system processors. For each virtual processor, 0.1 processing units are implicitly assigned and the LPAR is created in uncapped mode with a weight of 128.

Figure 1-8 explains virtual and real processor allocation to a freshly creating LPAR.

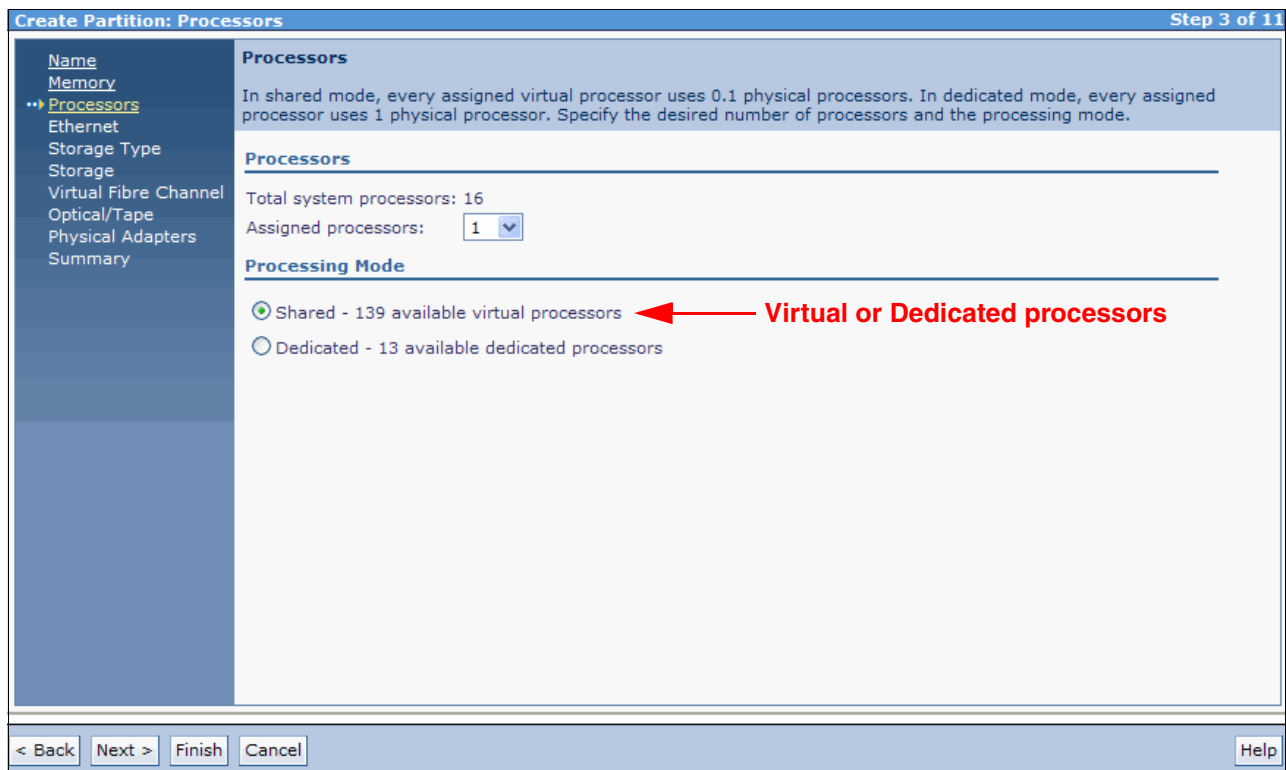


Figure 1-8 Processor allocation to LPAR

The LPAR configuration can be changed after the wizard finishes creating the partition. The following parameters are available for the partition:

- ▶ Processing unit value
- ▶ Virtual processor number
- ▶ Capped or uncapped property
- ▶ Uncapped weight

The default LPAR configuration that is provided by using the partition creation wizard is designed to keep the system balanced. Manual changes to the partition configuration can be made after careful planning of the resource distribution. The configuration that is described in Figure 1-4 on page 10 shows the manually changed processing units.

As a general suggestion, use these guidelines:

- ▶ For the LPAR configuration, select the appropriate virtual processors and keep the default processing units when possible.
- ▶ Leave some system processing units deallocated. They are available to all LPARs that require them.
- ▶ Do not underestimate the processing units that are assigned to the Virtual I/O Server. If the units are not needed, they remain available in the shared pool. However, on system peak utilization periods, the units can be important for the Virtual I/O Server to provide service to highly active partitions.

Virtual Ethernet

Every IVM managed system is configured with four predefined virtual Ethernet devices, each with a virtual Ethernet ID ranging 1 - 4. Every LPAR creation wizard by default shows two virtual Ethernet adapters that can be connected to any of the four virtual networks in the system. If needed, you can create more adapters.

Figure 1-9 shows the LPAR creation wizard. By default, adapter 1 is assigned to virtual Ethernet 1 and the second virtual Ethernet is unassigned.




Figure 1-9 Virtual Ethernet wizard of an LPAR creation

Each virtual Ethernet can be bridged by the Virtual I/O Server to a physical network by using only one physical adapter. If higher performance or redundancy is required, a physical adapter aggregation can be made on one of these bridges instead. The same physical adapter or physical adapter aggregation cannot bridge more than one virtual Ethernet. See 4.9, "Network management" on page 123 for more details.

Figure 1-10 shows a **View/Modify Virtual Ethernet** wizard panel. All four virtual networks are described with the corresponding bridging physical adapter, if configured. An administrator can decide how to configure the two available virtual adapters.

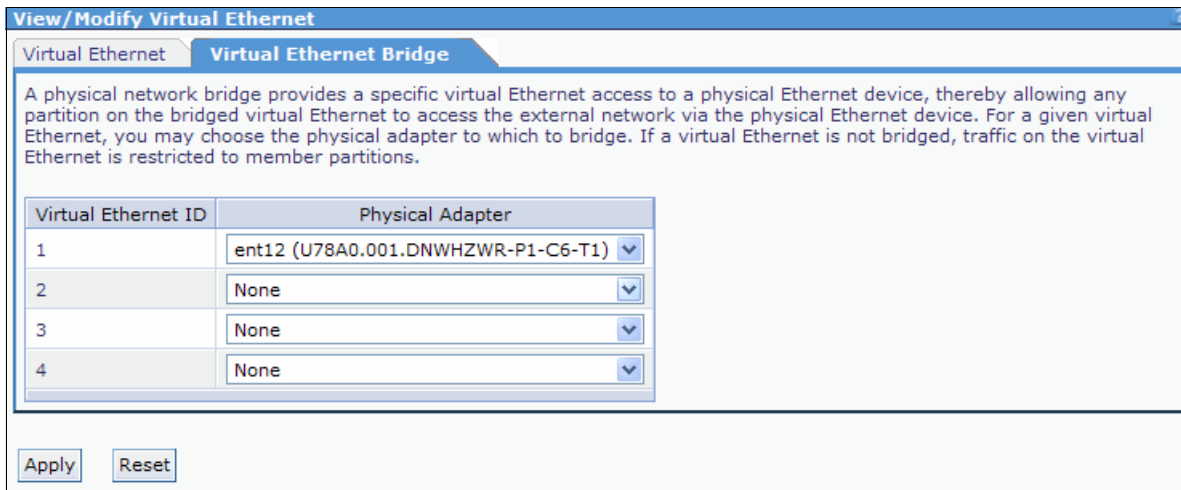


Figure 1-10 Virtual Ethernet allocation to an LPAR

The virtual Ethernet is a bootable device and can be used to install the operating system of the LPAR.

Virtual storage

Every LPAR can be equipped with one or more virtual devices by using a single virtual SCSI adapter. A virtual disk device has the following characteristics:

- ▶ The size is defined by the administrator.
- ▶ It is treated by the operating system as a normal SCSI disk.
- ▶ It is bootable.
- ▶ It is created by using the physical storage that is owned by the Virtual I/O Server partition. The storage can be either internal or external to the physical system (for example, on the storage area network (SAN)).
- ▶ It can be defined either by using an entire physical volume (SCSI disk or a logical unit number of an external storage server) or a portion of a physical volume.

Virtual disk device content is preserved if moved from one LPAR to another or increased in size. Before you make changes in the virtual disk device allocation or size, the owning partition unconfigures the device to prevent data loss.

A virtual disk device that does not require an entire physical volume can be defined by using disk space from a *storage pool*. This pool is created on the Virtual I/O Server, which is a set of physical volumes. Virtual disk devices can be created spanning multiple disks in a storage pool, and they can be extended if needed.

The IVM can manage multiple storage pools and change their configurations by adding or removing physical disks to them. To simplify management, one pool is defined to be the default storage pool and most virtual storage actions implicitly refer to it.

Virtual optical devices

Any optical device that is assigned to the Virtual I/O Server partition (either CD-ROM, DVD-ROM, or DVD-RAM) can be virtualized and assigned to any LPAR. This assignment is

done, one at a time, by using the same virtual SCSI adapter that is provided to the virtual disks. Virtual optical devices can be used to install the operating system and, when a DVD-RAM is available, to make backups.

Virtual TTY

To allow LPAR installation and management, the IVM provides a virtual terminal environment for LPAR console handling. When a new LPAR is defined, two matching virtual serial adapters are created for console access, one on the LPAR and one on the IVM. This configuration provides a connection from the IVM to the LPAR through the hypervisor. See more details in 4.5.3, “Console Terminal for IVM logical partition” on page 96.

1.3 Advanced PowerVM options

Through IVM, it is possible to configure many of the advanced virtualization functions that are available on the IBM Power Systems servers.

In this section, we explore the following topics:

- ▶ Active Memory Sharing
- ▶ Live Partition Mobility
- ▶ N_Port ID Virtualization (NPIV)
- ▶ Shared Storage Pools

In the following sections, different LPAR configurations are shown.

1.3.1 Active Memory Sharing

Active Memory Sharing (AMS) was announced with POWER6 processor-based servers and is available for servers that have the IBM PowerVM Enterprise Edition enabled. This feature helps to create a pool of memory to be shared among LPARs inside the server.

In this publication, we explore how to configure AMS by using the IVM GUI, in section 4.3, “Active Memory Sharing” on page 61.

1.3.2 Live Partition Mobility

Live Partition Mobility is an IBM PowerVM feature that is used to migrate an LPAR from one Power Systems server to another, if the pre-configuration requirements are met.

This function is available only with the PowerVM Enterprise Edition enabled for POWER6 processor-based and POWER7 processor-based servers.

In this publication, we show how to migrate an LPAR from a POWER6 to a POWER7 processor-based server through the IVM interface. For more information, see 4.8, “Live Partition Mobility” on page 117.

1.3.3 N_Port ID Virtualization (NPIV)

N_Port ID Virtualization (NPIV) is a virtualization standard feature. It has been supported since the introduction of the POWER6 processor-based servers with 8 Gb Fibre Channel adapters and the Virtual I/O Server Version 2.1 and later. Inside IVM, NPIV is represented by

a virtual Fibre Channel configuration which allow multiples LPARs to use the same Fibre Channel connection to access external storage.

In this publication, we describe how to configure NPIV, in section 4.11, “N_Port Identifier Virtualization (NPIV)” on page 138.

1.3.4 Shared Storage Pools

A *Shared Storage Pool (SSP)* is a new feature that is available in IVM from Version 2.2.1.2 and later. With this feature, you can combine multiple Virtual I/O Servers (beginning with Version 2.2.1.4 and later, is possible to configure four nodes) to use the SAN effectively.

With an SSP, you can achieve the following functions:

- ▶ Thin and thick provisioning
- ▶ Snapshot rollback
- ▶ One place to manage all types of disk storage devices
- ▶ Live Storage Mobility



Plan

Integrated Virtualization Manager (IVM) is shipped with the Virtual I/O Server media. The interface is activated during the Virtual I/O Server installation only, if all of the following conditions are true:

- ▶ The system is in the *Manufacturing Default Configuration*.
- ▶ The system has never been managed by a *Hardware Management Console (HMC)*.
- ▶ The virtualization feature is enabled.

A new system from manufacturing that was ordered with the virtualization feature is ready for the IVM. If the system was ever managed by an HMC, the administrator is required to reset it to the Manufacturing Default Configuration. If virtualization is not activated, the system cannot manage micropartitions. Contact an IBM marketing representative to order the activation code.

If a system supports the IVM, it can be ordered with the IVM preinstalled.

Ensure that you have the following equipment and IP addresses available before you install the IVM:

- ▶ A serial American Standard Code for Information Interchange (ASCII) console and cross-over cable (a physical ASCII terminal or a suitable terminal emulator) that is connected to one of the two system ports for initial setup
- ▶ An IP address for the IVM
- ▶ An optional, but recommended, IP address for the Advanced System Management Interface (ASMI)

This chapter describes the basic considerations, functions to set up in the ASMI, and how to apply the necessary virtualization code and microcodes to use the IVM interface.

2.1 Considerations for partition setup

When you use the IVM, it is easy to create and manage a partitioned system because most of the complexity of the logical partition (LPAR) setup is hidden. A new user can quickly learn an effective methodology to manage the system. However, it is important to understand how configurations are applied and can be changed.

The Virtual I/O Server is the only LPAR that can accomplish management interaction with the hypervisor and is able to react to hardware configuration changes. The configuration of the Virtual I/O Server can be changed dynamically while it is running. The other LPARs do not have access to the hypervisor and have no interaction with the IVM to be aware of possible system changes.

It is possible to change any resource allocation for the client LPARs through the IVM web interface. This action enables the user to change the processing unit configuration, memory allocation, and physical and virtual adapter setup while the LPAR is activated. This is possible through the Dynamic LPAR Manager (with a Resource Monitoring and Control (RMC) daemon). *Dynamic LPAR Manager* is a daemon that runs in an LPAR. The POWER Hypervisor communicates to the LPAR by using this daemon.

The IVM command-line interface (CLI) enables an experienced administrator to make modifications to a partition configuration. Changes that use the command line are shown in the web graphical user interface (GUI). A warning message is displayed to highlight that the resources of an affected LPAR are not yet synchronized.

Figure 2-1 shows a case where the memory was changed manually on the command line. To detect the actual values, the administrator must select the partition on the GUI and select **Properties** from the **More Tasks** drop-down menu or by just clicking the hyperlink for more details about the synchronization of the current and pending values.

The screenshot shows the IVM web interface. The main content area is titled 'View/Modify Partitions' and includes a 'System Overview' section with the following data:

| | | | |
|------------------------------------|----------|-----------------------------|-------------|
| Total system memory: | 128 GB | Total processing units: | 16 |
| Memory available: | 23.25 GB | Processing units available: | 14.1 |
| Reserved firmware memory: | 5.75 GB | Processor pool utilization: | 0.10 (0.6%) |
| Available shared memory pool size: | 64.53 GB | | |
| System attention LED: | Inactive | | |

Below the overview is a 'Partition Details' table:


| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
|--------------------------|------|---------------|---------------|--------------|---|------------|---------------------------|---------------------------|----------------|
| <input type="checkbox"/> | 1 | vios1 | Running | 2.04 Days | 20 GB | 13 | 1.3 | 0.09 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 55.6 Minutes | 4 GB | 2 | 0.3 | 0.01 | |
| <input type="checkbox"/> | 3 | lpar2_server1 | Not Activated | | 60 GB | 1 | 0.1 | | 00000000 |
| <input type="checkbox"/> | 4 | lpar3_server1 | Running | 1.98 Hours | 10 GB  [details] | 2 | 0.2 | 0.01 | |

Figure 2-1 Manual LPAR configuration

Figure 2-2 shows a generic LPAR schema from an I/O point of view. Every LPAR is created with one virtual serial and one virtual SCSI connection. There are four predefined virtual networks. The Virtual I/O Server already is equipped with one virtual adapter that is connected to each of them.

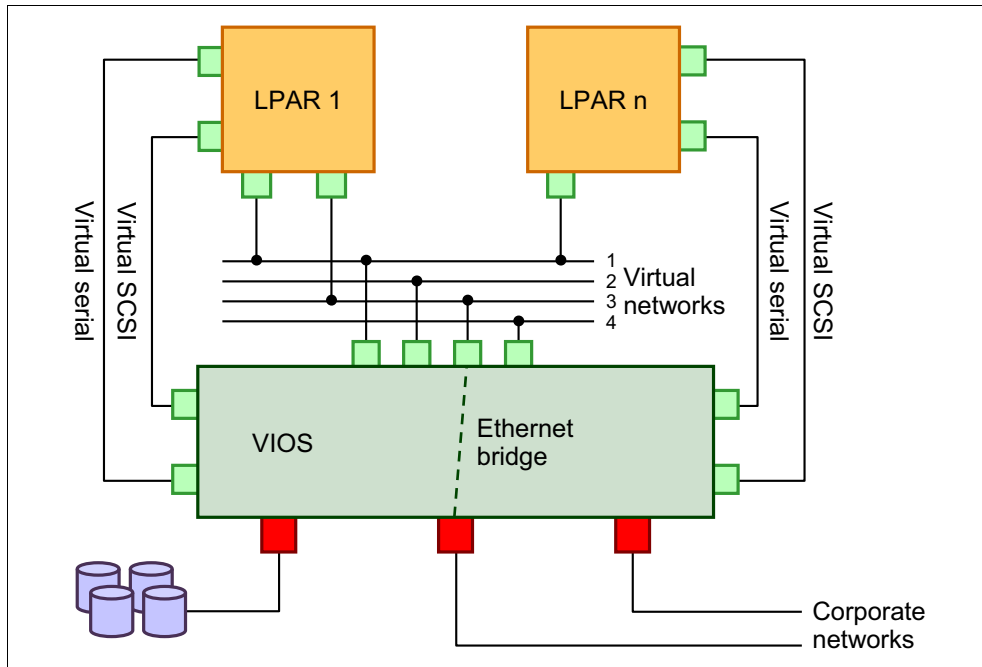


Figure 2-2 General I/O schema on an IVM managed system

There is only one virtual SCSI adapter for each LPAR. Therefore, the web GUI hides its presence and shows virtual disks and optical devices as assigned directly to the partition. When the CLI is used, the virtual SCSI adapter must be taken into account.

For a virtual Ethernet adapter configuration, the administrator needs to define only whether to create one or two virtual Ethernet adapters on each LPAR and on the virtual network to which it must be connected.

All remaining I/O configurations are done dynamically:

- ▶ An optical device can be assigned to any virtual SCSI channel.
- ▶ A virtual disk device can be created, deleted, or assigned to any virtual SCSI channel.
- ▶ A virtual Fibre Channel adapter can be created, deleted, or assigned to any LPAR.
- ▶ A physical adapter can be assigned to any LPAR.
- ▶ Ethernet bridging between a virtual network and a physical adapter can be created, deleted, or changed at any time.

2.2 Reset to Manufacturing Default Configuration

This operation is needed only if the system was previously managed by the HMC. It resets the system, removing all partition configurations and any personalization that was made to the service processor.

The following steps describe how to reset the system:

1. Power off the system.
2. Connect a serial ASCII console to a system port by using a null-modem (cross-over) cable. The following configuration is for the port settings:
 - 19,200 bits per second
 - Eight data bits
 - No parity
 - One stop bit
 - Xon/Xoff flow control
3. Press any key on the serial connection of the TTY to receive the service processor prompt.
4. Log in as the user `admin` and answer the questions about the number of lines and columns for the output. The default password is `admin`.
5. Type `2` to enter the **System Service Aids** menu and type `//` to select the **Factory Configuration** option. Type `1` to select **Reset all settings** from the **Factory Configuration** menu. A warning message similar to what is shown in Example 2-1 describes the effect of the reset and asks for confirmation. Type `1` and hit **Enter** to confirm.

Example 2-1 Factory configuration reset

Continuing will result in the loss of system settings that you may have set via

user interfaces. Before continuing with this operation make sure you have manually recorded all settings that need to be preserved.

This option will reset service processor settings such as the HMC access, ASMI passwords, time of day, network configuration, hardware deconfiguration policies, platform error logs, etc.

This option will reset the PCI bus configuration. It will also clear the firmware settings and all the partition information.

To complete this operation, the system will be powered on and then off.

Follow the instructions in the system service publications to configure the network interfaces after the reset.

Enter 1 to confirm or 2 to cancel: 1

Powering on the system server firmware and rebooting the service processor.
Attention: This operation may take some time, please be patient.

Factory configuration reset: After a factory configuration reset, the system activates the microcode version that is present in the permanent firmware image. Check the firmware levels in the permanent and temporary images before you reset the system.

Migration between HMC and IWM: More information about migration between the HMC and IWM can be found in section 5.1.6, “The migration between HMC and IWM” on page 168.

2.3 Microcode update

When you install the IVM, we suggest installing the latest available microcode. If the update is not needed, skip this section.

The active microcode level is provided by the service processor. If the system is powered off, connect to the system ports as described in section 2.2, “Reset to Manufacturing Default Configuration” on page 21, and log in as the admin user. The first menu shows the microcode level of the system in the Version line (Example 2-2).

Example 2-2 Current microcode level display using system port

System name: 8233-E8B_SN061AA6P

Version: AL730_087

User: admin

Copyright © 2002-2012 IBM Corporation. All rights reserved.

1. Power/Restart Control
2. System Service Aids
3. System Information
4. System Configuration
5. Network Services
6. Performance Setup
7. Concurrent Maintenance
8. Login Profile
99. Log out

S1>

If the IP address of the service processor is known, the same information is provided by using the ASMI in the upper panel of the web interface, as shown in Figure 2-3. For a description of the default IP configuration, see section 2.4, “Advanced System Management Interface IP address setup” on page 26.

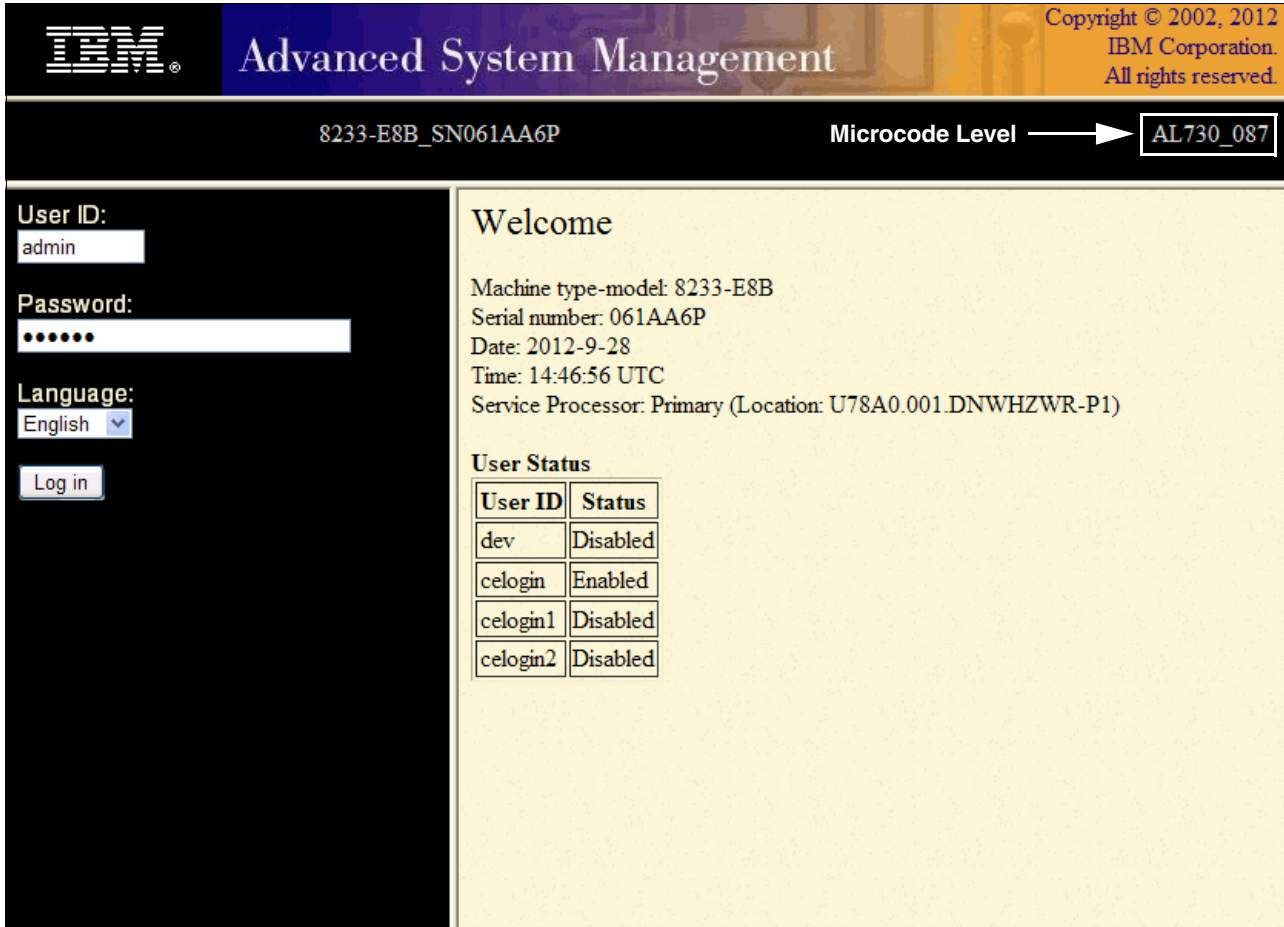


Figure 2-3 Current microcode level display by using the ASMI

If the system microcode must be updated, the code and installation instructions are available from the IBM Fix Central website:

<http://www.ibm.com/support/fixcentral>

Fix Central website: You are required to provide some information about your system in the Fix Central website to retrieve the correct code and instructions.

Microcode can be installed through one of the following methods:

- ▶ HMC
- ▶ Running operating system
- ▶ Running IVM
- ▶ Diagnostic CD

The HMC and running operating system methods require the system to be reset to the Manufacturing Default Configuration before you install the IVM. If the system is already running the IVM, refer to section 5.2.1, “Microcode update” on page 185 for instructions.

To use a diagnostic CD, a serial connection to the system port is required with the setup that is described in section 2.2, “Reset to Manufacturing Default Configuration” on page 21. You must download the latest **Microcode Update Files & Discovery Tool CD** set. You can download at the tool at this website:

<http://www-304.ibm.com/webapp/set2/firmware/gjsn?mode=10&page=cdrom.html>

The following steps describe how to update the microcode by using a diagnostic CD:

1. Insert the diagnostic CD in the system drive and start the system from it. Follow steps 1 - 8 as described in section 3.1, “Virtual I/O Server image installation” on page 42.
2. Follow the instructions on the window until the main menu window (Example 2-3) opens.

Example 2-3 Main diagnostic CD menu

FUNCTION SELECTION

- 1 Diagnostic Routines
This selection will test the machine hardware. Wrap plugs and other advanced functions will not be used.
- 2 Advanced Diagnostics Routines
This selection will test the machine hardware. Wrap plugs and other advanced functions will be used.
- 3 Task Selection (Diagnostics, Advanced Diagnostics, Service Aids, etc.)
This selection will list the tasks supported by these procedures. Once a task is selected, a resource menu may be presented showing all resources supported by the task.
- 4 Resource Selection
This selection will list the resources in the system that are supported by these procedures. Once a resource is selected, a task menu will be presented showing all tasks that can be run on the resource(s).
- 99 Exit Diagnostics

NOTE: The terminal is not properly initialized. You will be prompted to initialize the terminal after selecting one of the above options.

To make a selection, type the number and press Enter. []

3. Remove the diagnostic CD from the drive and insert the first *Microcode Update Files & Discovery Tool CD*.
4. Select **Task Selection (Diagnostics, Advanced Diagnostics, Service Aids, etc.)** → **Update and Manage System Flash** → **Validate and Update System Firmware**.
5. Select the CD drive from the menu.
6. Select the wanted flash image for your server by pressing the **F4** key (or use the **Esc-4** sequence if the console does not support it). This key provides a list of available flash images on the CD. If the wanted message does not show in the list, use the **F3** key (or **Esc-3** sequence) **two times** to return to the **Validate and Update System Firmware** menu. Then, change to the second Microcode Update Files & Discovery Tool CD, and try again.

If you select a flash update image that is not valid for your server, the window in Example 2-4 is displayed. If you get this message, use **F3** (or **Esc-3** sequence) to return to the previous menu option, and select a valid flash update image.

Example 2-4 Invalid firmware image file message

```
UPDATE AND MANAGE FLASH
802835
```

The file provided is not a valid firmware image. Verify that the file is a firmware image appropriate for this system's model and type. Verify that the firmware image is not corrupted. If neither check leads to the source of the problem, then contact your service support structure.

```
F3=Cancel          Esc+0=Exit          Enter
```

7. After you select the correct flash update image file, press the **F7** key to commit. If the console does not support it, use the **Esc-7** sequence.
8. On the final window, which is shown in Example 2-5, select **YES** and wait for the firmware update to be completed and for the subsequent system reboot to be executed.

Example 2-5 Confirmation panel for microcode update

```
UPDATE AND MANAGE FLASH                                802816
```

The image is valid and would update the temporary image to AL730_095.
The new firmware level for the permanent image would be AL730_087.

The current permanent system firmware image is AL730_087.
The current temporary system firmware image is AL730_087.

***** WARNING: Continuing will reboot the system! *****

Do you wish to continue?

Make selection, use 'Enter' to continue.

```
NO
YES
```

```
F3=Cancel          Esc+0=Exit
```

2.4 Advanced System Management Interface IP address setup

The service processor is equipped with two standard Ethernet ports, labeled HMC1 and HMC2, for network access. In an IVM environment, they are used to access Advanced System Management Interface (ASMI) menus by using a web browser. ASMI enables remote hardware administration and service agent setup and relies on the HTTPS protocol. Both Ethernet ports can be used if a valid network address is given.

By default, when the system is connected to a power source and the service processor starts, a *Dynamic Host Configuration Protocol (DHCP)* request is sent in the network through both HMC ports. If a DHCP server is available, it provides an IP address to the port; otherwise, Table 2-1 shows default IP addresses values that are used.

Table 2-1 Default Service Processor (SP) IP addresses

| Platform | SP | SP | Subnet Mask |
|------------------|---------------|---------------|---------------|
| | HMC1 | HMC2 | |
| POWER5 | 192.168.2.147 | 192.168.3.147 | 255.255.255.0 |
| POWER6 or POWER7 | 169.254.2.147 | 169.254.3.147 | 255.255.255.0 |
| Platform | Redundant SP | Redundant SP | Subnet Mask |
| | HMC1 | HMC2 | |
| POWER5 | 192.168.2.146 | 192.168.3.146 | 255.255.255.0 |
| POWER6 or POWER7 | 169.254.2.146 | 169.254.3.146 | 255.255.255.0 |

The DHCP-managed addresses are mainly intended to be used in an HMC environment. IVM can show the IP addresses of both HMC ports. However, when the system is powered off and IVM is not running, it might become difficult to contact ASMI because the addresses might change when the service processor reboots.

The IP configuration of the ports can be changed by using the ASMI menu or connecting to the system serial ports. ASMI can be reached only if the current IP configuration is known or if the default addresses are in use. Serial ports are available for service processor access only if the system is powered off.

2.4.1 Address setting using the ASMI

The following procedure uses the default address that is assigned to port HMC1. This address is in use if no other address has been manually configured and if no DHCP server gave an IP address to the port when the system was connected to a power source. If you are unsure about DHCP, disconnect the Ethernet cable from the HMC1 port, remove all power to the system, reconnect the power, and then wait for the service processor to start.

You need a system that is equipped with a web browser (Netscape 7.1, Microsoft Internet Explorer 6.0, Opera 7.23, or later versions) and configured with the following network configuration:

- ▶ IP 192.168.2.148 for POWER5, or 169.254.2.148 for POWER6 or POWER7
- ▶ Netmask 255.255.255.0

Use the following steps to set up the addressing:

1. Use an Ethernet cable to connect the HMC1 port with the Ethernet port of your system.
2. Connect the web browser by using the following URL:
<https://192.168.2.147> for POWER5; <https://169.254.2.147> for POWER6 or POWER7
3. Log in as the user admin with the password admin.

- Expand the **Network Services** menu and click **Network Configuration**, then select **IPv4**. After that, you can configure the IP. Figure 2-4 shows the corresponding menu.
- Complete the fields with the wanted network settings and click **Continue**. The Network interface eth0 corresponds to port HMC1; eth1 corresponds to HMC2.

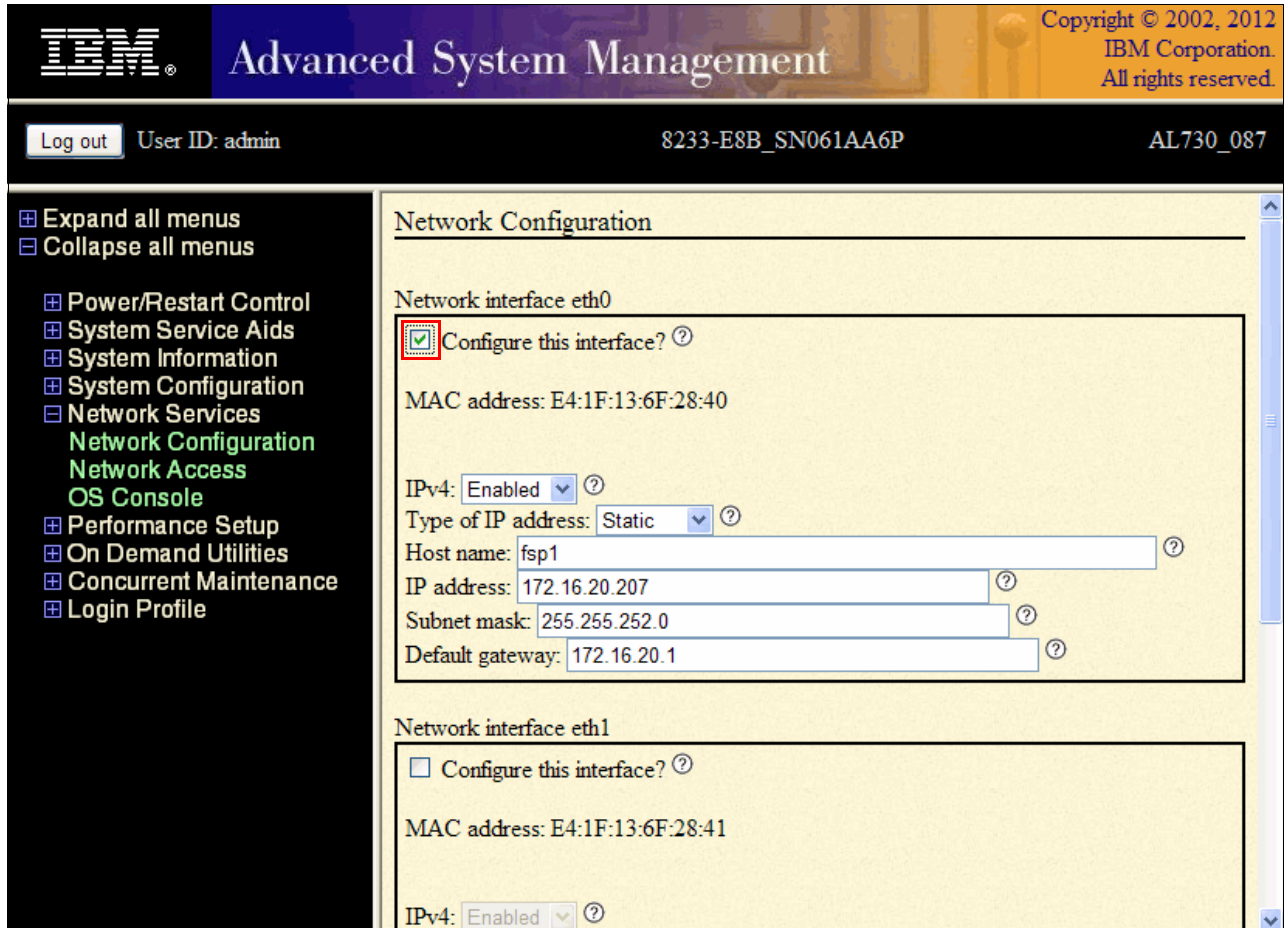


Figure 2-4 HMC1 port setup using the ASMI

- Review your configuration and click **Save settings** to apply the change.

2.4.2 Address setting using serial ports

When the IP addresses of the HMC port are not known and the ASMI cannot be used, it is possible to access the service processor by attaching an ASCII console to one of the system serial ports.

The following steps describe how to assign a fixed IP address to an HMC port:

- Power off the system.
- Connect to the system port as described in 2.2, “Reset to Manufacturing Default Configuration” on page 21.
- Select 5 for the **Network Services** menu and select 1 for **Network Configuration**.

Example 2-6 shows the steps to configure the port HMC1. The menu enables you to configure the interfaces Eth0 and Eth1 that correspond to system ports HMC1 and HMC2. To define a fixed IP address, provide the IP address, netmask, and, possibly, the default gateway.

Example 2-6 HMC1 port configuration

Network Configuration

1. Configure interface Eth0
2. Configure interface Eth1
98. Return to previous menu
99. Log out

S1> 1

Configure interface Eth0
MAC address: 00:02:55:2F:BD:E0
Type of IP address
Currently: Dynamic

1. Dynamic
 Currently: **169.254.2.147**
2. Static
98. Return to previous menu
99. Log out

S1> 2

Configure interface Eth0
MAC address: 00:02:55:2F:BD:E0
Type of IP address: Static

1. Host name
2. Domain name
3. IP address (Currently: **169.254.2.147**)
4. Subnet mask
5. Default gateway
6. IP address of first DNS server
7. IP address of second DNS server
8. IP address of third DNS server
9. Save settings and reset the service processor
98. Return to previous menu
99. Log out

S1>

2.5 Virtualization feature activation

This step is needed only if the system is not yet enabled with virtualization. Normally, new systems that are ordered with this feature arrive from manufacturing with the virtualization active.

Virtualization is enabled by using a specific code that is shipped with the system. Or, the code can be retrieved from the following Capacity on Demand website after you provide the system type and serial number:

<http://www-912.ibm.com/pod/pod>

The ASMI is used to activate the virtualization feature without an IVM or HMC.

Activation of virtualization using ASMI: For ASMI access, refer to the following website and look for your server Type/Model:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha8/hwicwe1come.htm>

As an example, here is the link for an 8202-E4B (Power 720) POWER7 processor-based server ASMI access:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hby/browser.htm>

This procedure works for any POWER7 processor-based server ASMI access.

To enter the activation feature, use the following steps:

1. Connect to the ASMI with a web browser by using the HTTPS protocol to the IP address of one of the HMC ports and log in as the user admin. The default password is admin.

2. Set the system in the standby state. Expand the **Power/Restart Control** menu and click **Power On/Off System**. Figure 2-5 shows the corresponding ASMI menu. In the boot to system server firmware field, select **Standby** and click **Save settings and power off**.

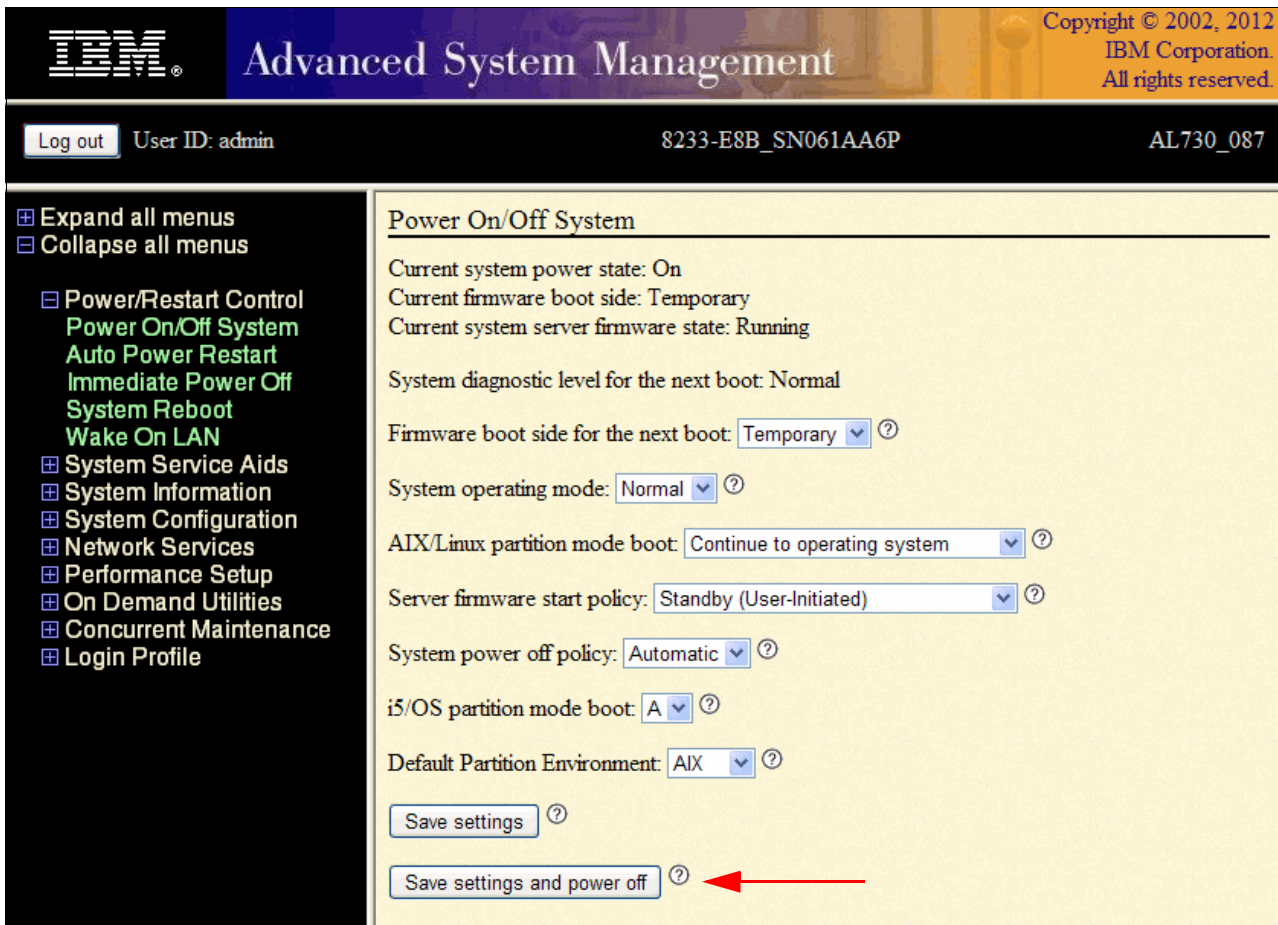


Figure 2-5 ASMI menu for Power On/Off System

3. You can look for on demand information about the system with the following ASMI panels, as shown in Figure 2-6.

The screenshot displays the IBM Advanced System Management (ASMI) web interface. At the top left is the IBM logo. The main header reads "Advanced System Management". On the top right, it says "Copyright © 2002, 2012 IBM Corporation. All rights reserved." Below the header, there is a navigation bar with a "Log out" button, the text "User ID: admin", the system identifier "8233-E8B_SN061AA6P", and the asset tag "AL730_087".

The main content area is split into two columns. The left column is a dark sidebar menu with the following items, each preceded by a plus sign in a square:

- Expand all menus
- Collapse all menus
- Power/Restart Control
- System Service Aids
- System Information
- System Configuration
- Network Services
- Performance Setup
- On Demand Utilities
 - CoD Order Information
 - CoD Activation
 - CoD Recovery
 - CoD Command
 - CoD Processor Information
 - CoD Memory Information
 - CoD VET Information
 - VET Capability Settings
- Concurrent Maintenance
- Login Profile

The right column displays the "CoD Order Information" panel, which contains the following text:

System type: 8233
System serial number: 06-1AA6P
Card type: 52B6
Card serial number: 07-808100E
Card ID: 0213204141420C5B

Figure 2-6 On Demand Utilities panel: CoD (capacity on demand) Order Information option

4. In the same On Demand Utilities panel, you can also find CoD VET (capacity on demand Virtualization Engine Technologies) information, as shown in Figure 2-7.

The screenshot displays the IBM Advanced System Management interface. At the top, the IBM logo is on the left, the title "Advanced System Management" is in the center, and the copyright notice "Copyright © 2002, 2012 IBM Corporation. All rights reserved." is on the right. Below the header, a navigation bar contains a "Log out" button, the user ID "admin", the system ID "8233-E8B_SN061AA6P", and the resource ID "AL730_087".

The main content area is divided into two sections. On the left is a dark sidebar menu with the following items:

- Expand all menus
- Collapse all menus
- Power/Restart Control
- System Service Aids
- System Information
- System Configuration
- Network Services
- Performance Setup
- On Demand Utilities
 - CoD Order Information
 - CoD Activation
 - CoD Recovery
 - CoD Command
 - CoD Processor Information
 - CoD Memory Information
 - CoD VET Information**
 - VET Capability Settings
- Concurrent Maintenance
- Login Profile

On the right is the "CoD VET Information" panel, which displays the following details:

- System type: 8233
- System serial number: 06-1AA6P
- Card type: 52B6
- Card serial number: 07-808100E
- Card ID: 0213204141420C5B
- Resource ID: CA1F
- Activated Resources: 0000
- Sequence number: 0040
- Entry check: C8

Figure 2-7 On Demand Utilities panel: CoD VET Information option

5. Enter the activation code as soon as the system finishes booting. Expand the **On Demand Utilities** menu and click **CoD Activation**. Figure 2-8 shows the corresponding menu. Enter the code that is provided to activate the feature in the specific system and click **Continue**. A confirmation message is displayed.

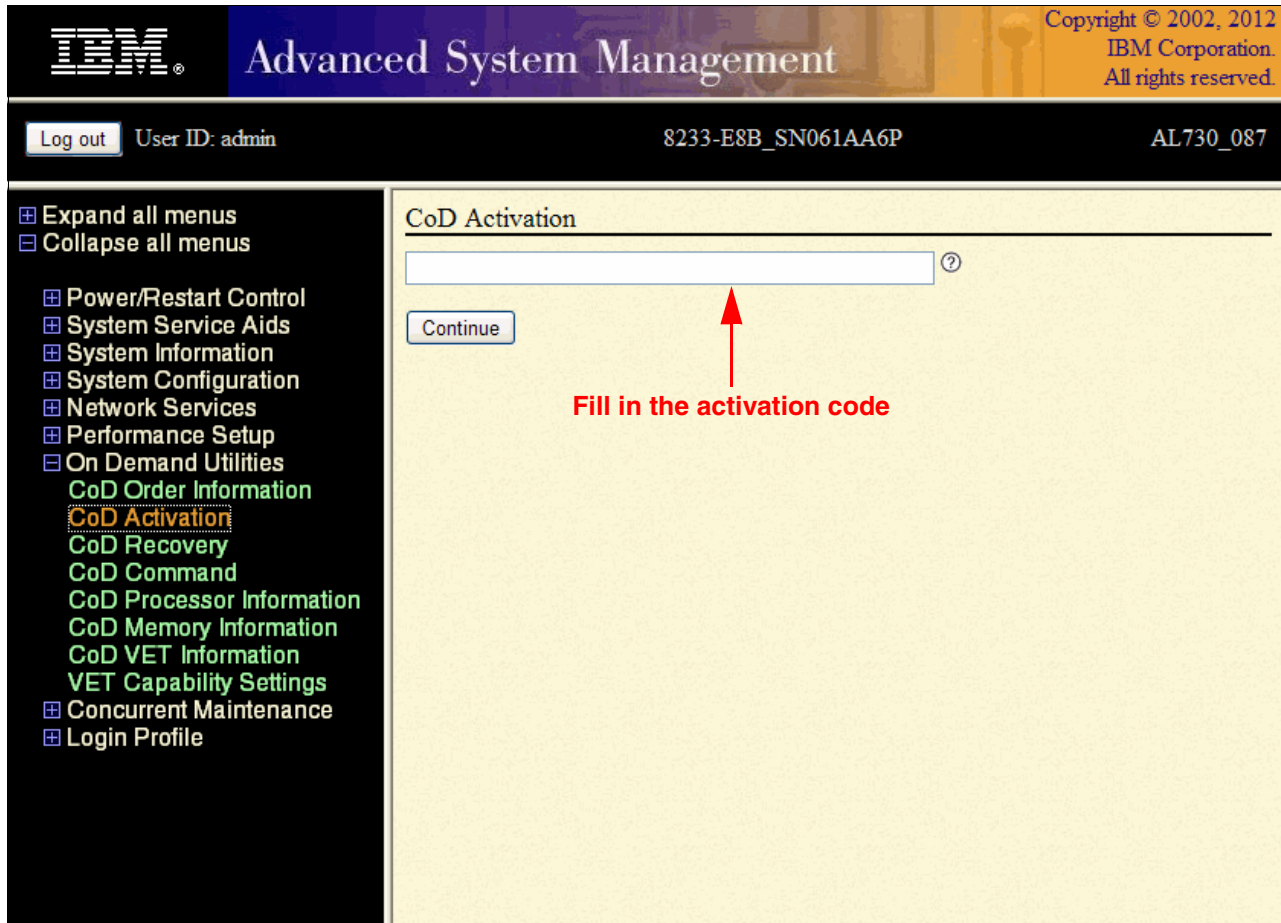


Figure 2-8 ASMI virtualization menu for code activation

- Set the system in running mode and shut it off. Again, select the **Power On/Off System** menu, select **Running** for the boot to system server firmware field, and click **Save settings and power off**, as shown in Figure 2-9.

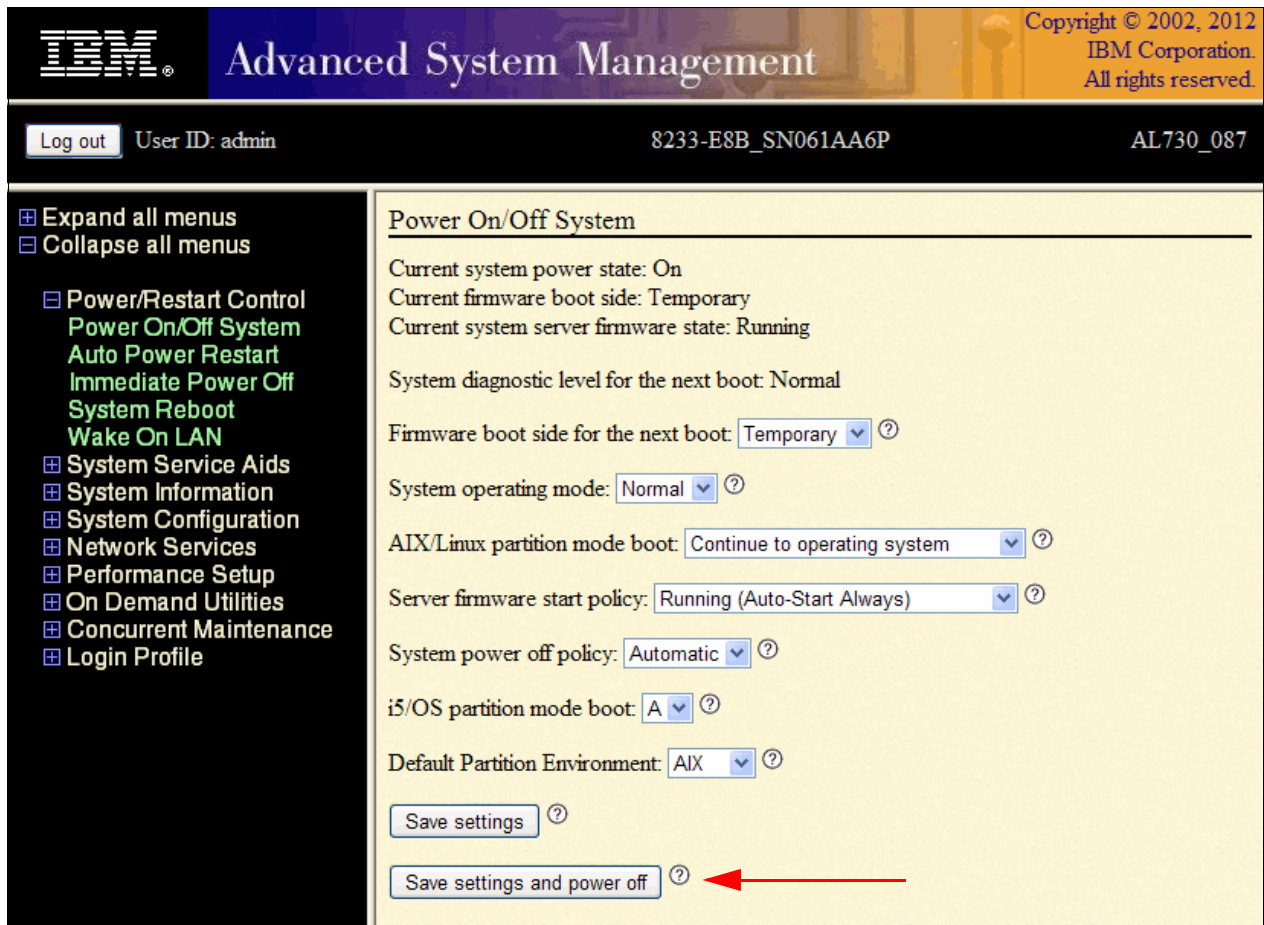


Figure 2-9 ASMI menu to bring system in running mode and power off

After you install the IVM on the server (for more information, see section 3.1, “Virtual I/O Server image installation” on page 42), you can also activate the virtualization feature through the IVM Management panel, as shown in Figure 2-10.

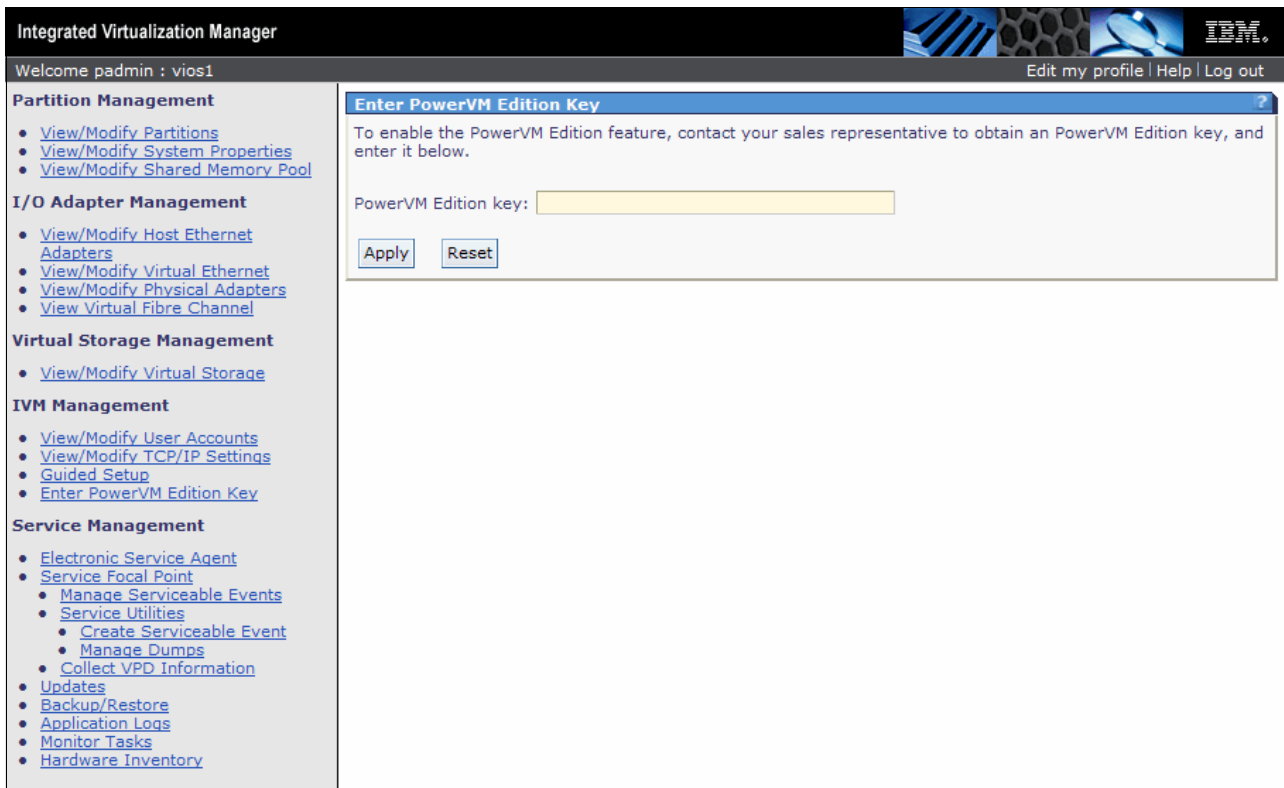


Figure 2-10 IVM menu for Enter PowerVM Edition Key

PowerVM activation: This PowerVM activation option can be helpful if you are doing a PowerVM Key upgrade. From PowerVM Express for PowerVM Standard, for example. That way, there is no need to access the ASMI.

2.6 Accessing the server using AMM in a blade environment

To use the PowerVM features in a POWER processor-based blade, it is mandatory to have IVM installed.

It is possible to do this installation by configuring Serial Over LAN (SoL) through the *advanced management module (AMM)* in the IBM BladeCenter chassis.

SoL in BladeCenter: For more details about how to configure SoL in a BladeCenter environment, access the information center website:

http://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp?topic=/com.ibm.bladecenter.advmgtmod.doc/bc_sol_guide.html

And for more information about enabling IVM on a BladeCenter with a POWER processor-based blade, see this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hchl/iphchinstallblade.htm>

After you access the **POWER processor-based blade** interface, follow installation procedures as described in section 3.1, “Virtual I/O Server image installation” on page 42.

2.7 Accessing the server by using serial link

For setup of serial link access to the server, place your server on standby state.

The serial links are used to connect to an ASCII terminal to manage a server that is running the AIX, Linux, or Virtual I/O Server operating systems. From the ASCII terminal, you can also access the ASMI to perform more installation tasks.

It is an alternative choice compared with ASMI web browser access, as mentioned in section 2.4, “Advanced System Management Interface IP address setup” on page 26.

The ASCII terminal is connected to the server through a serial link. The ASCII interface to the ASMI provides a subset of the web interface functions.

To access a Power server through its serial port, follow these steps:

1. By using a serial cable, connect the ASCII terminal to system connector 1 (P1-T1, which is the default) or 2 (P1-T2) on the rear of the server, as shown in Figure 2-11, Figure 2-12 on page 38, and Figure 2-13 on page 38.




Figure 2-11 Power 730/Power 740 rear view




Figure 2-12 Power 720/Power 740 rear view

Conversion cable: For Power 710, Power 720, Power 730, and Power 740 servers, if you are using a serial connection to the ASMI terminal, you must use a conversion cable. This cable (part number 46K5108; feature code #3930) is used to convert the ASCII terminal 9-pin Dshell connector to an RJ45 serial port connector on the system.




Figure 2-13 Power 750 rear view

2. Connect the power cord from the server to a power source.
3. Wait for the green light on the control panel to start flashing.

4. Ensure that your ASCII terminal is set to the general attributes, as shown in Table 2-2.

Table 2-2 ASCII general attributes

| General setup attributes | 3151 /11/31/41 settings | 3151 /51/61 settings | 3151 /51/61 settings | Description |
|--------------------------|-------------------------|----------------------|----------------------|---|
| Line speed | 19,200 | 19,200 | 19,200 | Uses the 19,200 (bits per second) line speed to communicate with the system unit. |
| Word length (bits) | 8 | 8 | 8 | Selects 8 bits as a data word length (byte). |
| Parity | No | No | No | Does not add a parity bit and is used together with the word length attribute to form the 8-bit data word (byte). |
| Stop bit | 1 | 1 | 1 | Places a bit after a data word (byte). |

5. Press a key on the ASCII terminal to allow the service processor to confirm the presence of the ASCII terminal.
6. When the login display is displayed for the ASMI, enter *admin* for the user ID and password.
7. Change the default password when you are prompted.
8. You completed the setup for an ASCII terminal and started the ASMI.

To see how to access ASMI through a web browser, refer to section 2.4.1, “Address setting using the ASMI” on page 27.

ASMI access: For more information about ASMI access, see the Systems Hardware information website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/topic/p7hby/ascii.htm>



Install

This chapter describes the different ways by which you can install the Virtual I/O Server. After that it explains the basic setup of Virtual I/O Server, network and storage management.

3.1 Virtual I/O Server image installation

The Virtual I/O Server is shipped as an optical media set that contains a bootable image of the software. It contains the IVM component. Installing it requires a serial connection to the system port with the setup described in 2.2, “Reset to Manufacturing Default Configuration” on page 21.

The following steps describe how to install the Virtual I/O Server:

1. Power on the system by using either the ASMI or pushing the power-on (white) button at the front of the system.
2. When you connect by using a TTY to the serial connection, you might be prompted to define it as an active console. If so, press the key that is indicated on the window.
3. If this is the first power-on of the system after a Manufacturing Default Configuration reset, you must accept the *License Agreement* to continue.
4. Wait for the *system management services (SMS)* menu selection, as shown in Example 3-1, and enter **1** after the word keyboard is displayed on the window.

Example 3-1 SMS menu selection

```
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
```

```
1 = SMS Menu                5 = Default Boot List
8 = Open Firmware Prompt    6 = Stored Boot List
```

```
memory      keyboard    network    scsi       speaker
```

5. Wait for *Language Selection Menu*. Select **2** to proceed to SMS main menu. When requested, provide the password for the admin user of the service processor. The default password is admin.
6. Insert the Virtual I/O Server installation media in the drive.
7. Use the SMS menus to select the CD or DVD device to boot. Select **Select Boot Options** → **Select Install/Boot Device** → **CD/DVD**. Select the CD/DVD media type of the system. Choose the correct device from a list, similar to the one shown in Example 3-2.

Example 3-2 Choose optical device from which to boot

```
PowerPC Firmware
Version EL350_132
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
```

```
-----
Select Device
Device Current Device
Number Position Name
1.      1      SATA CD-ROM
        ( loc=U78A0.0 1.DNWGPL0-P2-D2 )
```

```
-----
Navigation keys:
M = return to Main Menu
```

ESC key = return to previous screen X = eXit System Management Services

Type the number of the menu item and press Enter or select Navigation Key:1

8. Select **Normal Mode Boot** and exit from the SMS menu.
9. Select the console number and press **Enter**.
10. Select the preferred installation language from the menu.
11. Select the installation preferences. Choose the default settings, as shown in Example 3-3.

Example 3-3 Virtual I/O Server installation setup

```
Welcome to Base Operating System
                    Installation and Maintenance
```

Type the number of your choice and press Enter. Choice is indicated by >>>.

```
>>> 1 Start Install Now with Default Settings
    2 Change/Show Installation Settings and Install
    3 Start Maintenance Mode for System Recovery
    4 Configure Network Disks (iSCSI)
    5 Select Storage Adapters

    88 Help ?
    99 Previous Menu
```

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

12. Wait for the Virtual I/O Server restore. A progress status is shown, as in Example 3-4. At the end, Virtual I/O Server reboots.

Example 3-4 Virtual I/O Server installation progress status

```
Installing Base Operating System
```

```
Please wait...
```

| Approximate % tasks complete | Elapsed time (in minutes) | |
|---------------------------------|------------------------------|------------------------------|
| 28 | 7 | 29% of mksysb data restored. |

13. If your Virtual I/O Server media has more than one volume, the prompt that is shown in Example 3-5, is displayed. Remove volume 1 from the DVD drive and insert volume 2. Press **Enter** key to continue.

Example 3-5 Media change request

Please remove volume 1, insert volume 2, and press the ENTER key.

14. After you finish the Virtual I/O Server restore, the system automatically reboots. When the console prompt is displayed, log in to the Virtual I/O Server by using the user `padmin` and the default password `padmin`. When prompted, change the login password to something secure.

Accept the Virtual I/O Server license by issuing the `license -accept` command.

3.1.1 Virtual I/O Server image installation from the NIM server

It is possible to use a *Network Installation Management (NIM)* server to install the Virtual I/O Server. For more information, refer to the *NIM installation and backup of the VIO server* technote:

<http://www.ibm.com/support/docview.wss?uid=isg3T1011386>

3.2 Initial configuration

The new Virtual I/O Server requires a simple configuration setup using the command-line interface. Then, all management is done by using the web interface.

3.2.1 Set the date and time

Use the `chdate` command to set the Virtual I/O Server date and time, using the following syntax:

```
chdate [-year YYyy] [-month mm] [-day dd] [-hour HH] [-minute MM] [-timezone TZ]
chdate mmddHHMM[YYyy | yy] [-timezone TZ]
```

3.2.2 Initial network setup

The IVM web interface requires a valid network configuration to work. Configure the IP by choosing a physical network adapter and issuing the `mktcpip` command from the command line, using the following syntax:

```
mktcpip -hostname HostName -inetaddr Address -interface Interface
        [-start] [-netmask SubnetMask] [-cabletype CableType]
        [-gateway Gateway] [-nsrvaddr NameServerAddress]
        [-nsrvdomain Domain]
```

```
mktcpip -auto [-interface Interface]
```

```
mktcpip -hostname HostName -inetaddr Address -interface Interface
        [-start] [-cabletype CableType]
        [-gateway Gateway] [-nsrvaddr NameServerAddress]
        -nsrvdomain Domain] [-plen prefixLength]
```

Example 3-6 shows the setup of the host name, address, and IP address for the IVM.

Example 3-6 IVM network setup at the command line

```
$ mktcpip -hostname vios1 -inetaddr 172.16.20.200 -interface en0 -start -netmask  
255.255.255.000 -gateway  
172.16.20.1
```

Valid name resolution: The IVM, like a web server, requires a valid name resolution to work correctly. If DNS is involved, check that both the name and IP address resolution of the IVM host name are correct.

After the IVM web server has access to the network, it is possible to use the web graphical user interface (GUI) with the HTTP or the HTTPS protocol pointing to the IP address of the IVM server application. Authentication requires the use of the padmin user, unless other users were created.

3.2.3 Changing the TCP/IP settings on the Virtual I/O Server

The IVM enables you to change the TCP/IP settings on the Virtual I/O Server through the GUI.

Use any role other than the View Only role to do this task. Users with the View Only role can view the TCP/IP settings but cannot change them.

Before you can view or modify the TCP/IP settings, you must have an active network interface.

Important: Modifying your TCP/IP settings remotely might result in the loss of access to the current session. Ensure that you have physical console access to the IVM partition before you change the TCP/IP settings.

To view or modify the TCP/IP settings, do the following steps:

1. From the IVM Management menu, click **View/Modify TCP/IP Settings**. The View/Modify TCP/IP Settings panel opens, as shown in Figure 3-1.

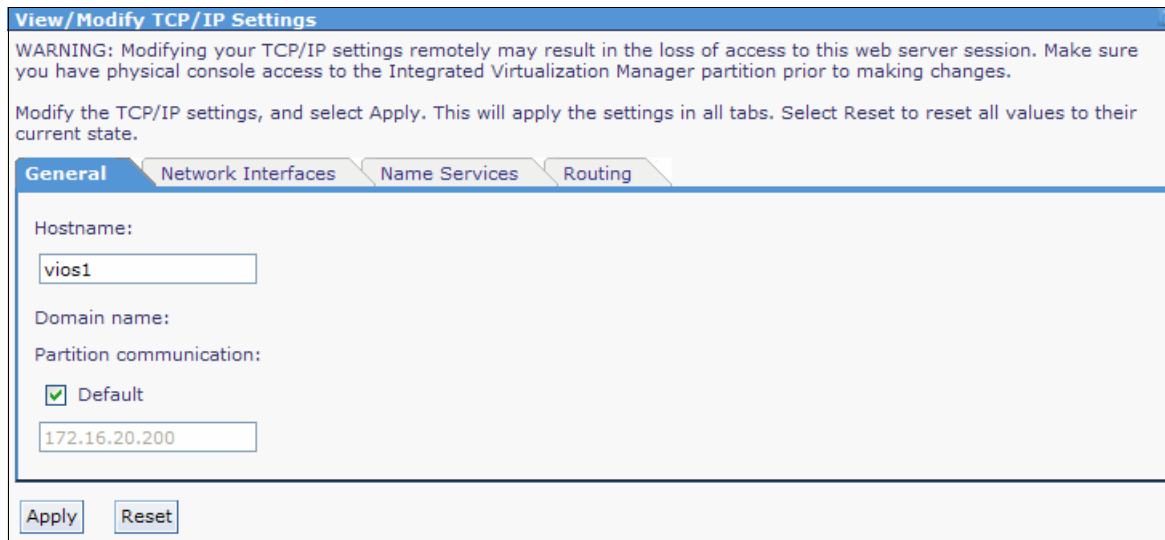


Figure 3-1 View/Modify TCP/IP settings

2. Depending on which setting you want to view or modify, click one of the following tabs:
 - **General** to view or modify the host name and the partition communication IP address
 - **Network Interfaces** to view or modify the network interface properties, such as the IP address, subnet mask, MAC Address, and the state of the network interface
 - **Name Services** to view or modify the domain name, name server search order, and domain server search order
 - **Routing** to view or modify the default gateway
3. Click **Apply** to activate the new settings.

3.3 Virtual I/O Server partition configuration

After you complete the network configuration of the Virtual I/O Server, the IVM interface is available and can be accessed by using a web browser. Connect using HTTP or HTTPS to the IP address assigned to Virtual I/O Server and log in as the user `padmin`.

The first panel that opens after the login process is the partition configuration, as shown in Figure 3-2 on page 47. After the initial installation of the IVM, there is only the Virtual I/O Server partition on the system with the following characteristics:

- ▶ The ID is 1.
- ▶ The name is equal to the serial number of the system.
- ▶ The state is Running.

- ▶ The allocated memory is the maximum value between 512 MB and one-eighth of the installed system memory. (In Figure 3-2, we are using a 32 GB server so it is showing 4 GB of allocated memory for the Virtual I/O Server partition).
- ▶ The number of virtual processors is equal to or greater than the number of processing units. The processing units are equal to at least 0.1 times the total number of virtual processors in the LPAR.

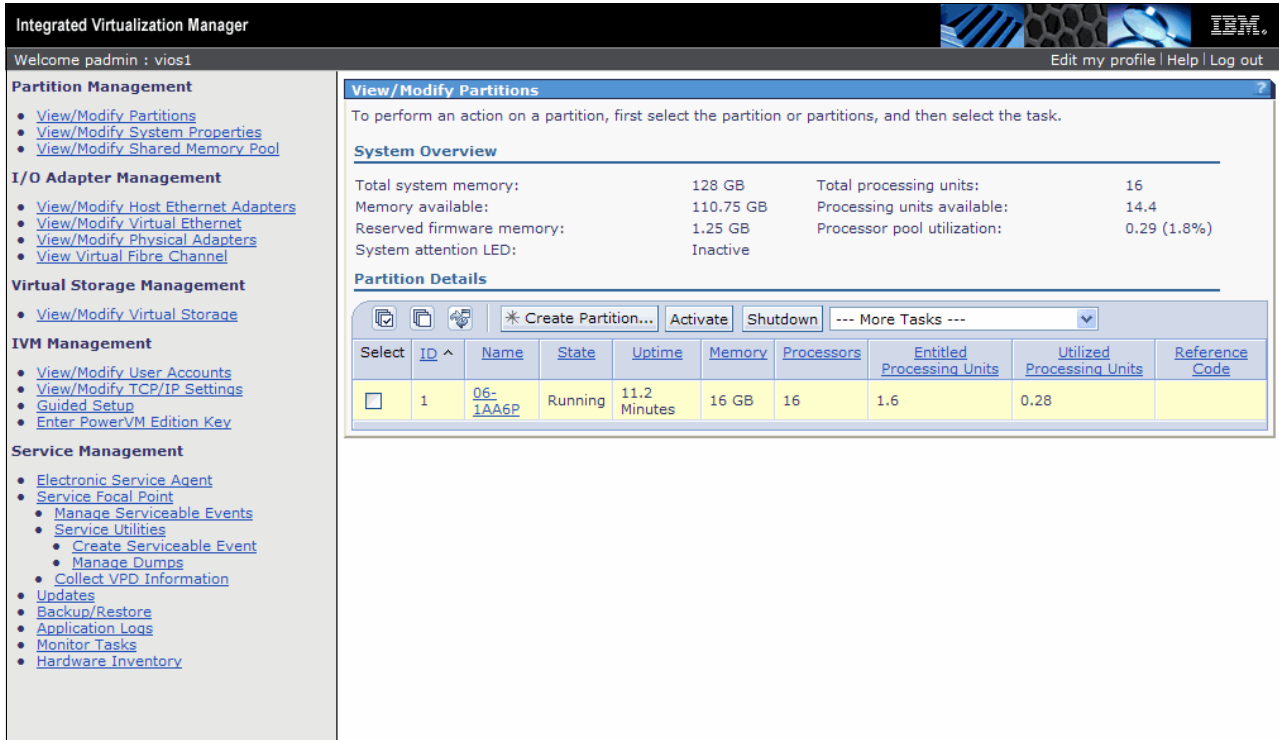


Figure 3-2 IVM (Virtual I/O Server) initial partition configuration

The default configuration for the partition is designed to be appropriate for most IVM installations. If the administrator wants to change memory or processing unit allocation of the Virtual I/O Server partition, a dynamic reconfiguration action can be made either using the web GUI or the command line, as described in 4.7, “Logical partition configuration changes” on page 101. Dynamic reconfiguration (Dynamic Logical Partitioning) of memory, processors, and I/O adapters of any client partitions is also supported.

3.4 Network management

When installed, the Virtual I/O Server configures one network device for each physical Ethernet that is present on the system. The server also creates four virtual Ethernet adapters, each belonging to a separate virtual network.

Any partition can be created with its own virtual adapters connected to any of the four available virtual networks. No bridging is provided with physical adapters at installation time.

The IVM enables any virtual network to be bridged by using any physical adapter, if the same physical adapter is not used to bridge more than one virtual network.

In 4.9, “Network management” on page 123, we describe the network bridging setup.

Since IVM 1.5, it is possible to configure dedicated I/O adapters for any client partitions.

Initially, all of the I/O resources are designated for the Virtual I/O Server LPARs. This configuration can be changed by the Dynamic Logical Partitioning process deallocating any I/O resources that you want to dedicate for client LPARs. You can use this process for Ethernet adapters also.

3.5 Virtual Storage management

The IVM uses the following virtual storage management concepts:

| | |
|------------------------|---|
| Physical volume | A physical disk or a logical unit number (LUN) on a storage area network (SAN). They are all owned by the IVM. A physical volume not belonging to any storage pool can be assigned whole to a single partition as a virtual device. |
| Storage pool | A set of physical volumes that is treated as a single entity. There can be multiple pools and they cannot share physical disks. One pool is defined as the <i>default storage pool</i> . |
| Virtual disk | Logical volume that the IVM assigns to a single partition as a virtual device. |

Both physical volumes and virtual disks can be assigned to an LPAR to provide disk space. Each of them is represented by the LPAR operating system as a single disk. For example, assigning a 146.8 GB physical disk and a 10 GB virtual disk to an LPAR running AIX 7.1 makes the operating system create two hdisk devices.

At installation time, there is only one storage pool, named *rootvg* (root volume group), which normally contains only one physical volume. All remaining physical volumes are available but not assigned to any pool.

The *rootvg* pool is used for IVM management, and we do not recommend using it to provide disk space to LPARs. Because it is the only pool that is available at installation time, it is also defined as the *default pool*. Create another pool and set it as the default before you create other partitions.

Important: Create at least one more storage pool so that the *rootvg* pool is not the default storage pool.

You can use *rootvg* as a storage pool on a system that is equipped with a SAS RAID adapter when all of the physical disks are configured as a single RAID array. In this case, the administrator must first boot the server using the Standalone Diagnostics CD-ROM that is provided with the system, and create the array. During the Virtual I/O Server image installation, only one disk is available, representing the array itself.

From any storage pool, virtual disks can be defined and configured. They can be created in several ways, depending on the IVM menu that is used:

- ▶ During LPAR creation. A virtual disk is created in the default storage pool and assigned to the partition.
- ▶ Using the View/Modify Virtual Storage menu on the Virtual Disk tab, using the **Create Virtual Disk** button. At this time, you can choose which *storage pool* belongs to the virtual disk and if you want to assign it for a specific partition or not (you can do it later, or in the *Partition Properties* window).

We describe basic storage management in 4.2.3, “Storage pool disk management” on page 52, and in 4.10, “Storage management” on page 133.



Set up

This chapter explains how to work with Integrated Virtualization Manager (IVM) both in a graphical user interface (GUI) and command-line interface (CLI).

It describes configuring and working on client logical partitions (LPARs), changing configurations both in storage and network, and also working on the console.

This chapter explains how to configure and work with the advanced virtualization features like:

- ▶ N_Port ID Virtualization
- ▶ Active Memory Sharing
- ▶ Live Partition Mobility

This chapter also describes the following advanced configurations on networking, storage management, and security:

- ▶ Virtual Ethernet bridging
- ▶ Ethernet link aggregation
- ▶ Disk space management
- ▶ Disk data protection
- ▶ Virtual I/O Server firewall
- ▶ Secure Shell (SSH) support

4.1 Configure and manage partitions

The IVM provides several ways to configure and manage LPARs:

- ▶ By use of a GUI, which is designed to be as simple and intuitive as possible, incorporating partition management, storage management, serviceability, and monitoring capabilities. See 4.2, “IVM graphical user interface” on page 50.
- ▶ By use of a CLI to enable scripting capabilities. See 4.5, “IVM command-line interface” on page 95.
- ▶ In the GUI of the IVM, there is also a so-called *simple partition creation* by using the option “Create Based On” in the task area. See 4.4.1, “Create an LPAR based on an existing partition” on page 89.

The following sections explain these methods.

4.2 IVM graphical user interface

The graphical user interface (GUI) of the IVM is an HTML-based interface. It enables you to create LPARs on a single managed system, manage system properties, manage physical adapters, manage the virtual storage and virtual Ethernet on the managed system, and view service information that is related to the managed system.

4.2.1 Connect to the IVM

Open a web browser window and connect by using the HTTP or HTTPS protocol to the IP address or host name that was assigned to the IVM during the installation process, as described in 3.2.2, “Initial network setup” on page 44. As a result, a Welcome window that contains the login and the password prompts opens, as shown in Figure 4-1. The default user ID is padmin, and the password you defined at IVM installation time.



Figure 4-1 IVM login page

After the authentication process, the default IVM console window opens, as shown in Figure 4-2. The IVM GUI is composed of two elements.

The following elements are the most important:

Navigation area The navigation area displays the tasks that you can access in the work area.

Work area The work area contains information that is related to the management tasks that you perform using the IVM and to the objects on which you can perform management tasks.

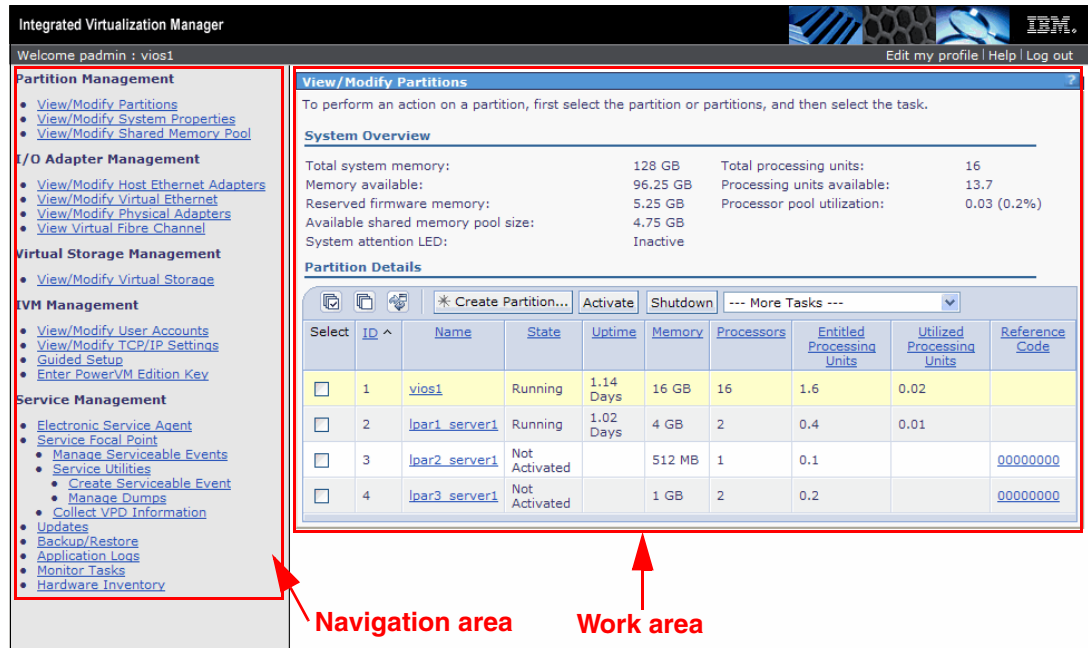


Figure 4-2 IVM console: View/Modify Partitions

4.2.2 Guided setup

The Guided Setup window provides quick access for some important configurations to do before you start creating partitions. This window also has a link for the IVM help page (if you click **More** in each topic). Figure 4-3 shows an example of the Guided Setup for the activity to Mirror the Integrated Virtualization Manager Partition.

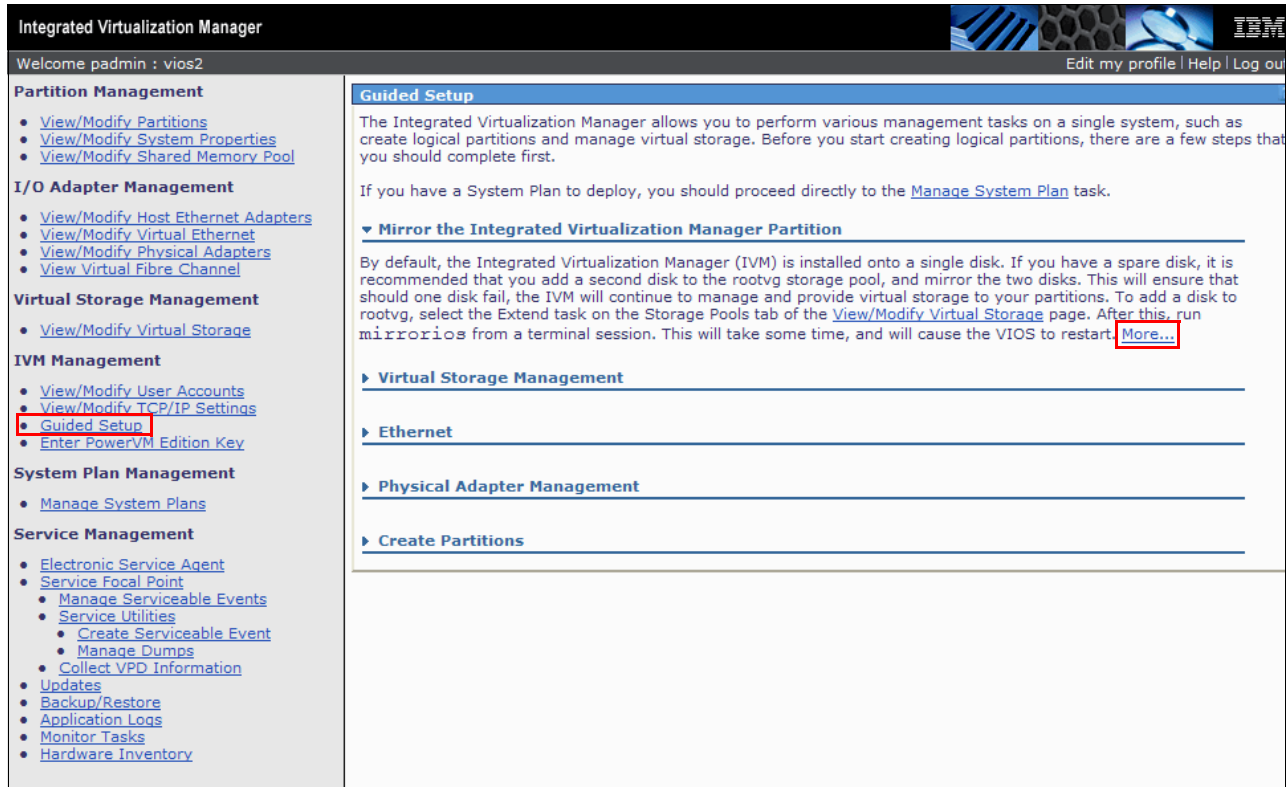


Figure 4-3 Guided Setup panel: Mirror the IVM Manager Partition option

After you click each link, you are redirected to other IVM menus. The configuration of these topics is addressed in this chapter.

4.2.3 Storage pool disk management

During the installation of the Virtual I/O Server, a *default storage pool* is created and named *rootvg*.

During the process of creating the LPAR, the IVM can automatically create virtual disks in the default storage pool. We recommend that you create a separate storage pool and add virtual disks to it for the LPARs. For advanced configuration of the storage pool, refer to 3.5, “Virtual Storage management” on page 48.

Storage pool creation

A *storage pool* consists of a set of physical disks that can be different types and sizes. You can create multiple storage pools; however, a disk can be a member of a single storage pool only.

Important: All data of a physical volume is erased when you add this volume to a storage pool.

The following steps describe how to create a storage pool:

1. Under the Virtual Storage Management menu in the navigation area, click the **View/Modify Virtual Storage** link.
2. Click **Storage Pools** tab in the work area, as shown in Figure 4-4.

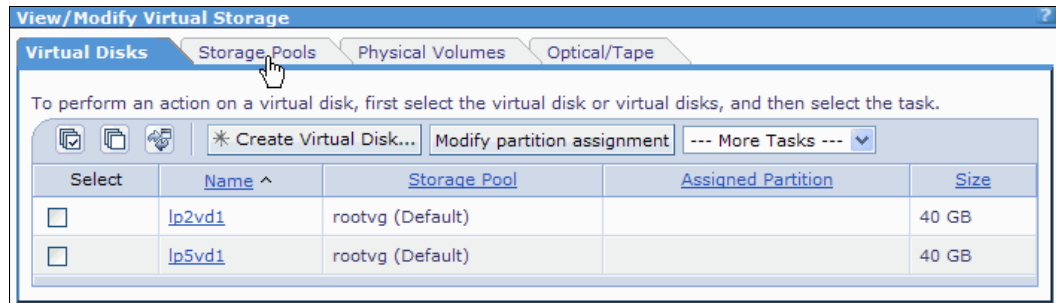


Figure 4-4 View/Modify Virtual Storage window

3. Click **Create Storage Pool** in the work area, as shown in Figure 4-5.




Figure 4-5 Create Virtual Storage

- Type a name in the *Storage pool name* field and select the wanted disks, as shown in Figure 4-6. To make this selection the default storage pool, select the **Assign as default storage pool** check box.

Create Storage Pool

To create a storage pool, enter the storage pool name and the type of storage pool to create. File based storage pools use files for virtual disks whereas logical volume based storage pools use logical volumes.

* Storage pool name:

Storage pool type:

Assign as default storage pool:

Logical volume based

A logical volume based storage pool is created using unassigned physical volumes. Select one or more physical volumes and select OK.

| Select | Physical Volume | Size | Physical Location Code |
|-------------------------------------|-----------------|-----------|---|
| <input checked="" type="checkbox"/> | hdisk1 | 136.73 GB | U5802.001.0086848-P3-D3 |
| <input checked="" type="checkbox"/> | hdisk2 | 136.73 GB | U5802.001.0086848-P3-D4 |
| <input type="checkbox"/> | hdisk3 | 65 GB | U5802.001.0086848-P1-C1-T1-LFF000-L0 |
| <input type="checkbox"/> | hdisk4 | 65 GB | U5802.001.0086848-P1-C1-T1-LFF0100-L0 |
| <input type="checkbox"/> | hdisk10 | 10 GB | U5802.001.0086848-P1-C3-T1-W201600A0B829AC12-L1000000000000 |
| <input type="checkbox"/> | hdisk11 | 10 GB | U5802.001.0086848-P1-C3-T1-W201600A0B829AC12-L2000000000000 |
| <input type="checkbox"/> | hdisk5 | 136.73 GB | U78A0.001.DNWHZWR-P2-D4 |
| <input type="checkbox"/> | hdisk6 | 136.73 GB | U78A0.001.DNWHZWR-P2-D5 |
| <input type="checkbox"/> | hdisk7 | 136.73 GB | U78A0.001.DNWHZWR-P2-D6 |

* Required field

Figure 4-6 Create Virtual Storage: Storage pool name

- Click **OK** to create the storage pool. A new storage pool called *ams_sp* with *hdisk1* and *hdisk2* is created.

Default storage pool

The default storage pool that is created during the IVM installation is *rootvg*. This is because *rootvg* is the only volume group that is created then.

Important: Create at least one more storage pool. The *rootvg* storage pool should not be the default storage pool. This would result in IVM and the user data being merged on the same storage devices.

Because the IVM is installed in rootvg, when IVM is reinstalled, the rootvg storage pool is overwritten. The default storage pool should also be changed to another one to avoid creating virtual disks within the rootvg by default, thus preventing the loss of user data during an IVM update.

The following steps describe how to change the default storage pool:

1. Under the Virtual Storage Management menu in the navigation area, click **View/Modify Virtual Storage**.
2. Click **Storage Pools** tab in the work area.
3. Select the storage pool that you want as the default, as shown in Figure 4-7.




Figure 4-7 View/Modify Virtual Storage: Storage Pools list

4. Select **Assign as default storage pool** from the **More Tasks** drop-down list, as shown in Figure 4-8.

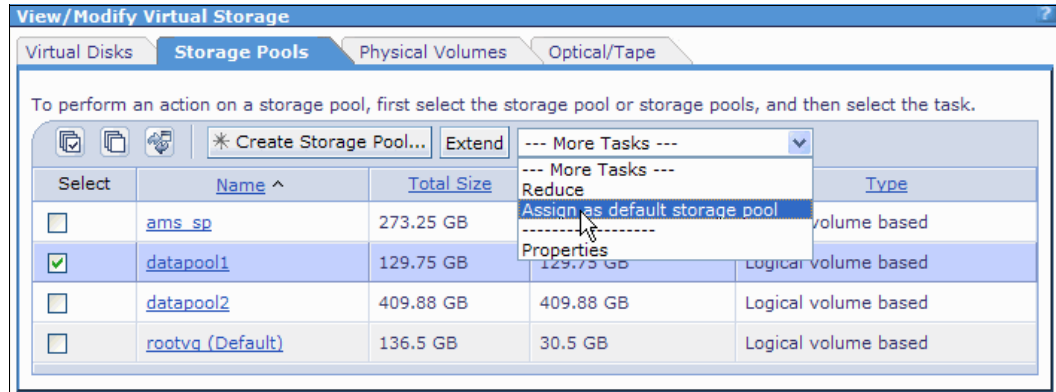


Figure 4-8 Assign as default storage pool option selection

5. A summary with the current and the next default storage pool opens, as shown in Figure 4-9. Click **OK** to validate the change. In this example, *datapool1* is the new default storage pool.



Figure 4-9 Assign as Default Storage Pool

Virtual disk/logical volume creation

Logical volumes belong to a storage pool and are also known as virtual disks. Logical volumes are used to provide disk space to LPARs but are not assigned to LPARs when you create them.

They can be created in several ways, depending on the menu that is in use:

- ▶ During LPAR creation: A logical volume is created in an existing storage pool and assigned to the partition.
- ▶ After or before LPAR creation: A virtual disk is not assigned to any partition and is created in an existing storage pool.

The following steps describe how to create a new logical volume:

1. Under the Virtual Storage Management menu in the navigation area, click **View/Modify Virtual Storage**.
2. Under the **Virtual Disks** tab in the work area, click **Create Virtual Disk** as shown in Figure 4-10.

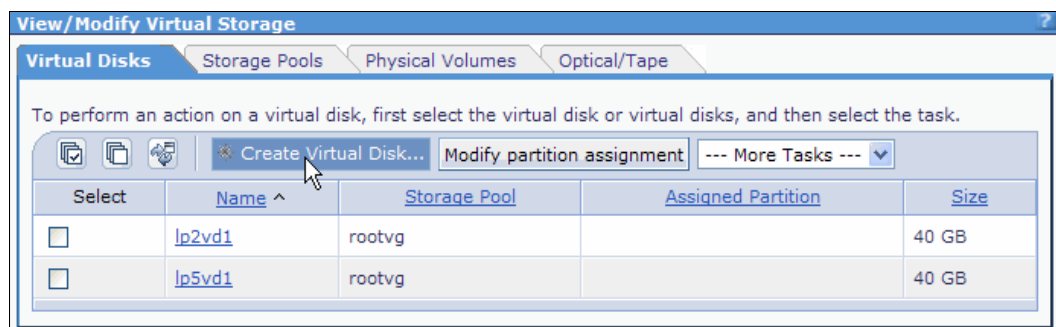


Figure 4-10 Create virtual storage

3. A new window is displayed. Enter a name for the virtual disk, select a storage pool name from the drop-down list, as shown in Figure 4-11.

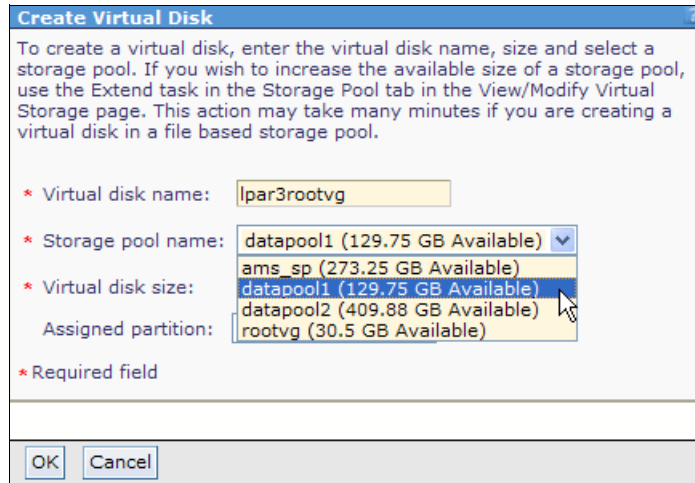


Figure 4-11 Create Virtual Disk: name and storage pool

4. Enter the size of the virtual disk and select an assigned partition for the new disk from the drop-down list. Select **None** if you do not want to assign the new disk to a partition now. Click **OK** to create the virtual disk. as depicted in Figure 4-12.

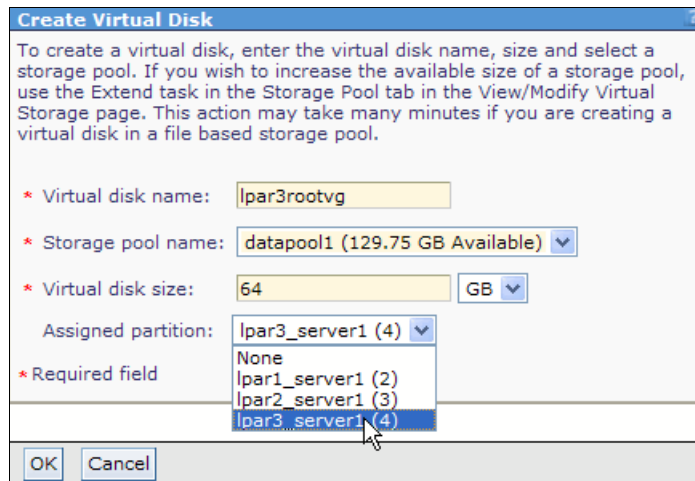


Figure 4-12 Create Virtual Disk: size and partition assignment

To view and use your new virtual disk/logical volume, select the **View/Modify Virtual Storage** link under the Virtual Storage Management menu in the navigation area. The list of available virtual disks is displayed in the work area, under the Virtual Disks tab. Figure 4-13 shows an example.

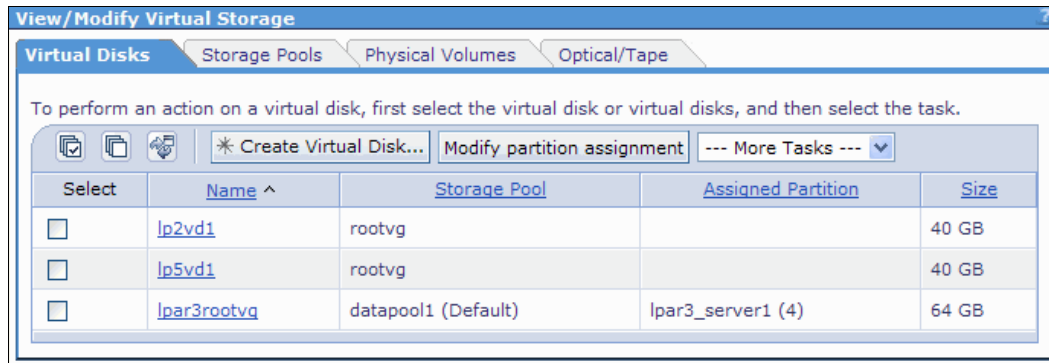


Figure 4-13 View existing virtual disks

Virtual disk/logical volume partition assignment

There are different ways to assign virtual disks to an LPAR. You can assign an existing virtual disk by using the View/Modify Virtual Storage menu in the navigation area, or by managing partition properties, where you can assign virtual disks.

The following steps describe how to assign virtual disks by using the View/Modify Virtual Storage menu:

1. Under the Virtual Storage Management menu in the navigation area, click **View/Modify Virtual Storage**.

- Under the Virtual Disks tab in the work area, select the virtual disk for which you want to change the assignment, as shown in Figure 4-14. You can select more than one virtual disk.

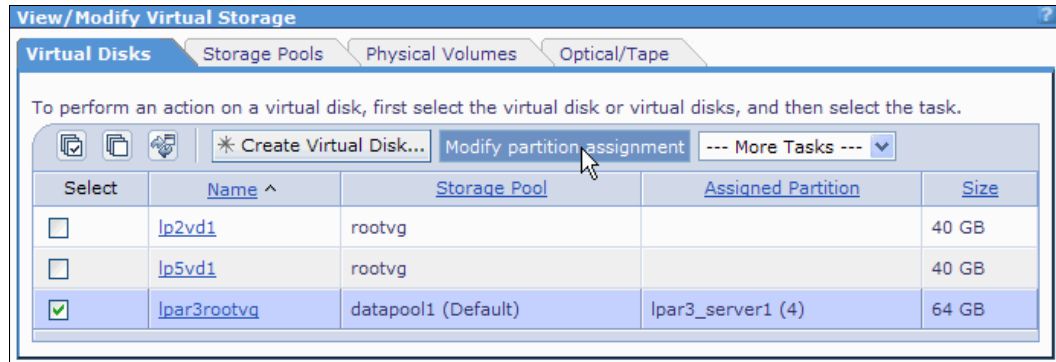


Figure 4-14 Modifying the virtual disk partition assignment

- Figure 4-15 shows the Modify Virtual Disk Partition Assignment window in the work area. A table shows the current assignment for the selected virtual disk and its current partition state. If the virtual disk is already assigned to another partition, ensure that this virtual disk is not being used.

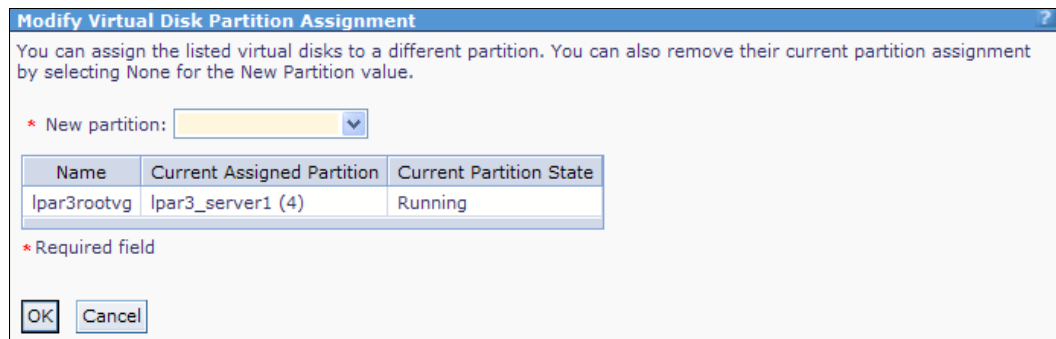


Figure 4-15 Current virtual disk assignment

- Select the destination partition for the selected virtual disk by using the drop-down list of the **New partition** field, as shown in Figure 4-16. If you want to unassign the virtual disk from an LPAR, select **None**. Click **OK** to assign the virtual disk to the destination partition.

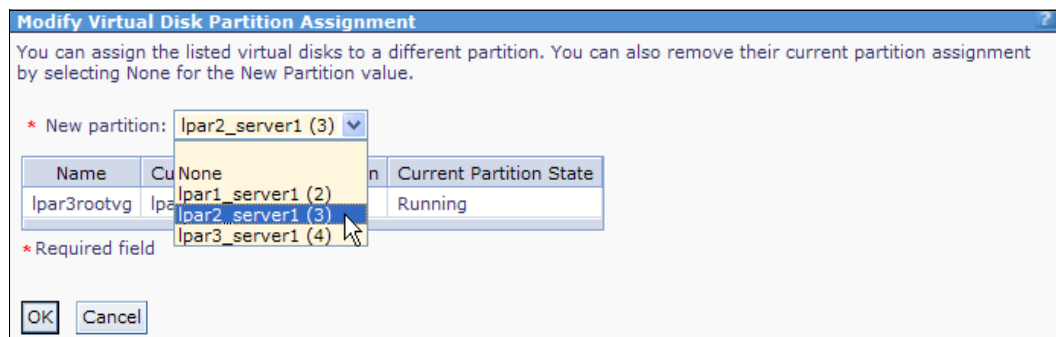


Figure 4-16 Destination partition selection

5. If the current assigned partition is in running state, a warning is issued, as depicted in Figure 4-17. Select **Force device removal from running partitions**. Ensure that the virtual disk is no longer being used by the current partition. Click OK to change the assignment of the virtual disk.

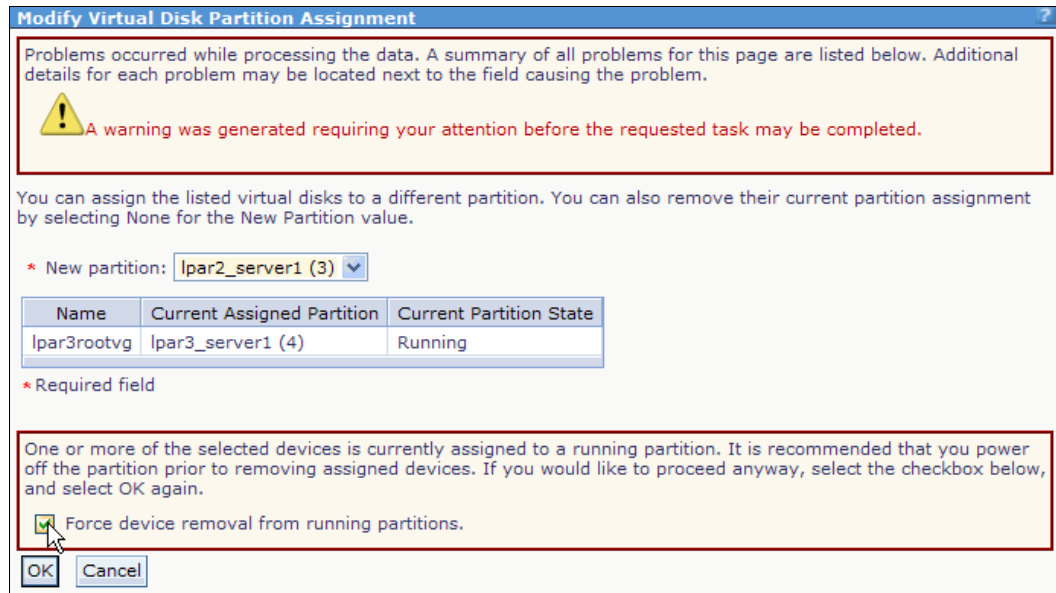


Figure 4-17 Virtual disk assignment warning

Removal of virtual disks: Removing virtual disks that are being used by running partitions can cause data loss. Use caution when you remove virtual disks from a running partition. Ensure that the partition is no longer using the virtual disk before you force the removal of the virtual disk.

Virtual disk/logical volume deletion

To remove a virtual disk, it cannot be assigned to a partition. The following steps describe how to delete a virtual disk:

1. Under the Virtual Storage Management menu in the navigation area, click **View/Modify Virtual Storage**. See Figure 4-18 on page 61.

- Under the Virtual Disks tab in the work area, select the virtual disk for which you want to change the assignment, as shown in Figure 4-18. You can select more than one virtual disk. Select **Delete** from the More Tasks drop-down list. Click **OK** to continue. See Figure 4-18.

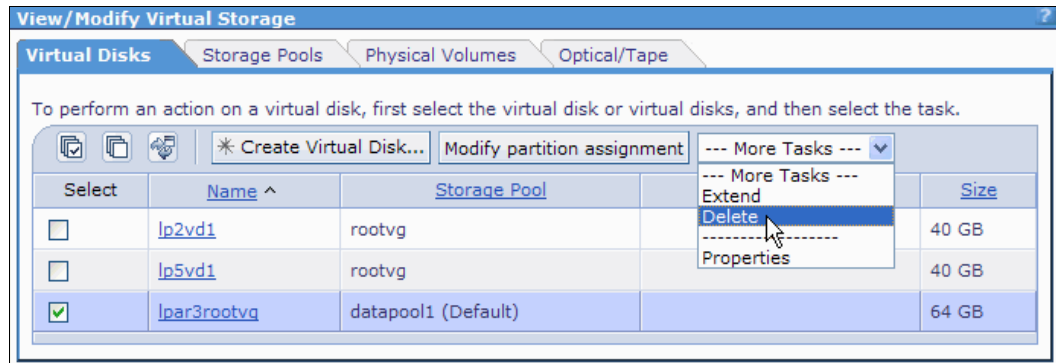


Figure 4-18 Virtual disk deletion

- Figure 4-19 shows the confirmation window that is displayed. Click **OK** to delete the virtual disk.




Figure 4-19 Virtual disk deletion confirmation

4.3 Active Memory Sharing

Active Memory Sharing (AMS) is an IBM PowerVM Enterprise Edition feature that enables advanced system memory virtualization capabilities for IBM Power Systems. With AMS, you can create a memory pool which can be shared between multiple partition, thus sharing physical memory.

To use the AMS feature, you need an existing storage pool with enough free space for AMS paging devices before you create the AMS pool. In IVM, the AMS paging devices are automatically managed. Although you can use any storage pool with enough free space, we suggest that you use a dedicated storage pool for the AMS paging devices, whenever possible.

After you create the AMS pool, you can manually add devices (physical or logical volumes), even if they do not belong to the paging storage pool. IVM automatically manages AMS paging devices. But you need to ensure that there are enough devices or free space in the paging storage pool to handle with the requirements of the LPAR.

See *IBM PowerVM Virtualization Active Memory Sharing*, REDP-4470 for more detailed information about AMS.

4.3.1 Creating the AMS pool

Before creating Active Memory Sharing (AMS) partitions, you must create a shared memory pool. This section describes the required steps to accomplish the creation of the AMS pool:

1. Connect to the management interface of the IVM by using a web browser. After you log in, the window in Figure 4-20 opens. Click **View/Modify Shared Memory Pool** in the navigation area.

The screenshot shows the IVM management interface. The left navigation pane includes sections for Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, and Service Management. The main content area is titled 'View/Modify Partitions' and displays system overview information and a table of partitions.

System Overview

| | | | |
|---------------------------|----------|-----------------------------|-------------|
| Total system memory: | 128 GB | Total processing units: | 16 |
| Memory available: | 108.5 GB | Processing units available: | 14.2 |
| Reserved firmware memory: | 1.5 GB | Processor pool utilization: | 0.19 (1.2%) |
| System attention LED: | Inactive | | |

Partition Details

| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
|--------------------------|------|---------------|---------|-----------|--------|------------|---------------------------|---------------------------|----------------|
| <input type="checkbox"/> | 1 | vios1 | Running | 4.84 Days | 16 GB | 16 | 1.6 | 0.18 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 2.95 Days | 2 GB | 2 | 0.2 | 0.01 | |

Figure 4-20 The IVM management interface

2. In the work area, the View/Modify System Properties window opens. Click **Define Shared Memory Pool**, as shown in Figure 4-21.

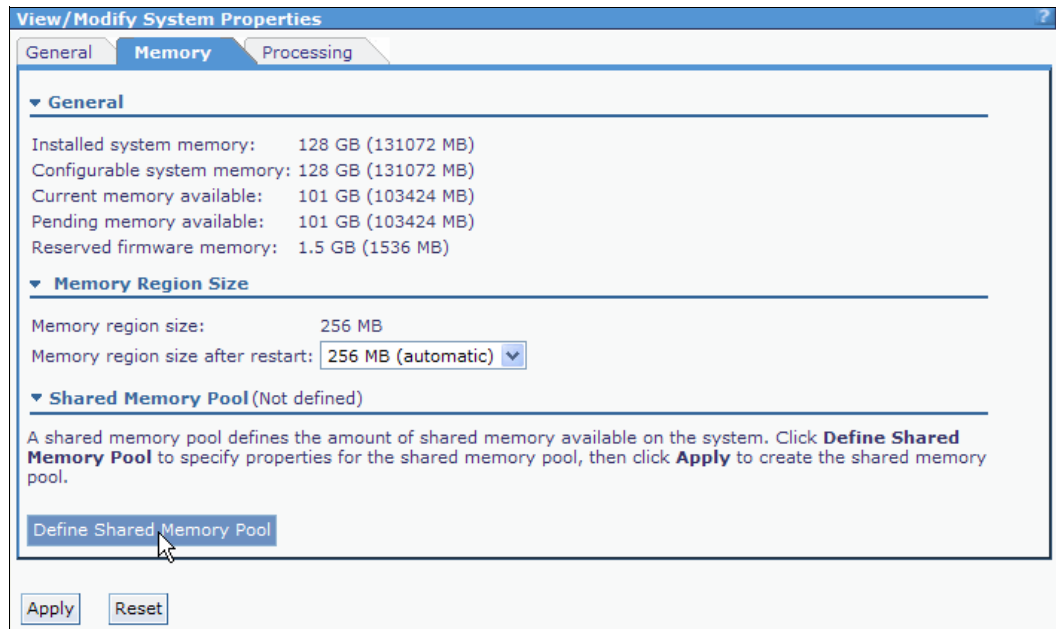


Figure 4-21 Defining Shared Memory Pool

3. The window in Figure 4-22 opens. Enter the wanted amount of memory for the shared memory pool by using the Assigned memory for the Shared Memory Pool field.



Figure 4-22 AMS pool size configuration

4. Choose the wanted storage pool for AMS paging devices by using the drop-down list of the Paging storage pool field, as shown in Figure 4-23. Remember that the wanted storage pool must be previously created before the AMS pool creation (see 4.2.3, “Storage pool disk management” on page 52 for more information about storage pools). Select GB from the drop-down list (MB is shown). Click **OK** to accept the options.

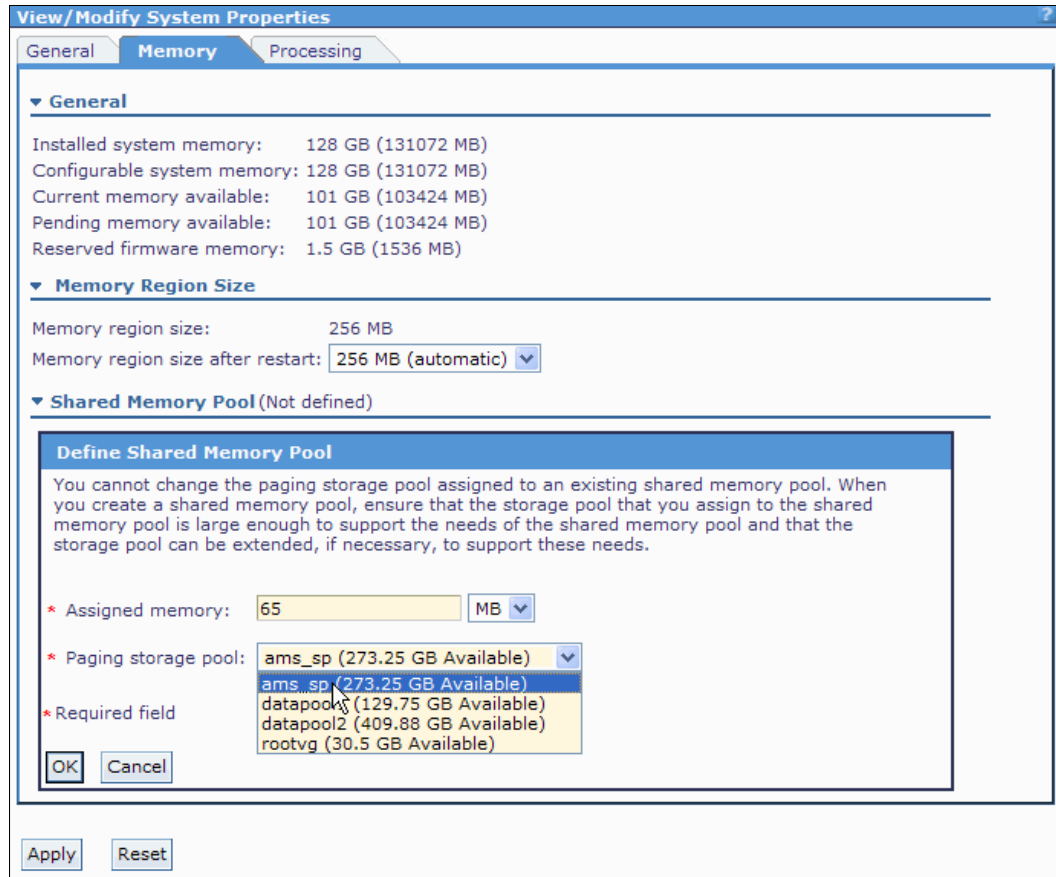


Figure 4-23 Storage pool selection for AMS paging devices

5. You will see exclamation marks on the window, as shown in Figure 4-24. This view shows that the AMS pool configuration is still pending. Click **Apply** to make the AMS pool configuration effective.

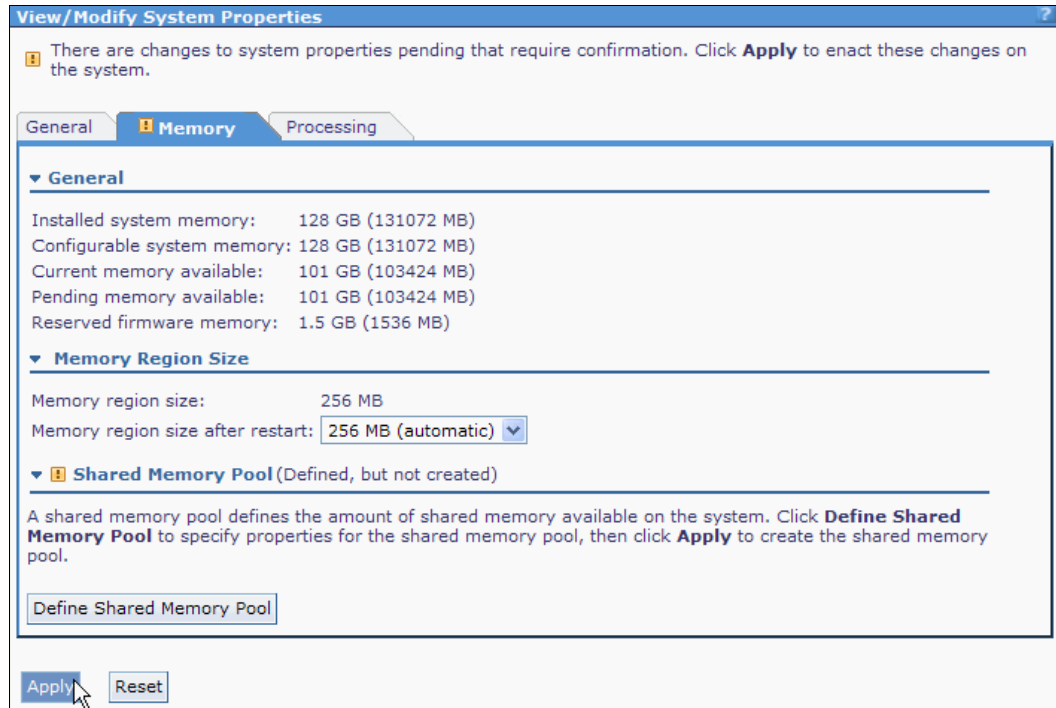


Figure 4-24 Applying AMS pool configuration

- After applying the AMS Pool configuration, the View/Modify System Properties window shows the new configuration, as depicted in Figure 4-25.

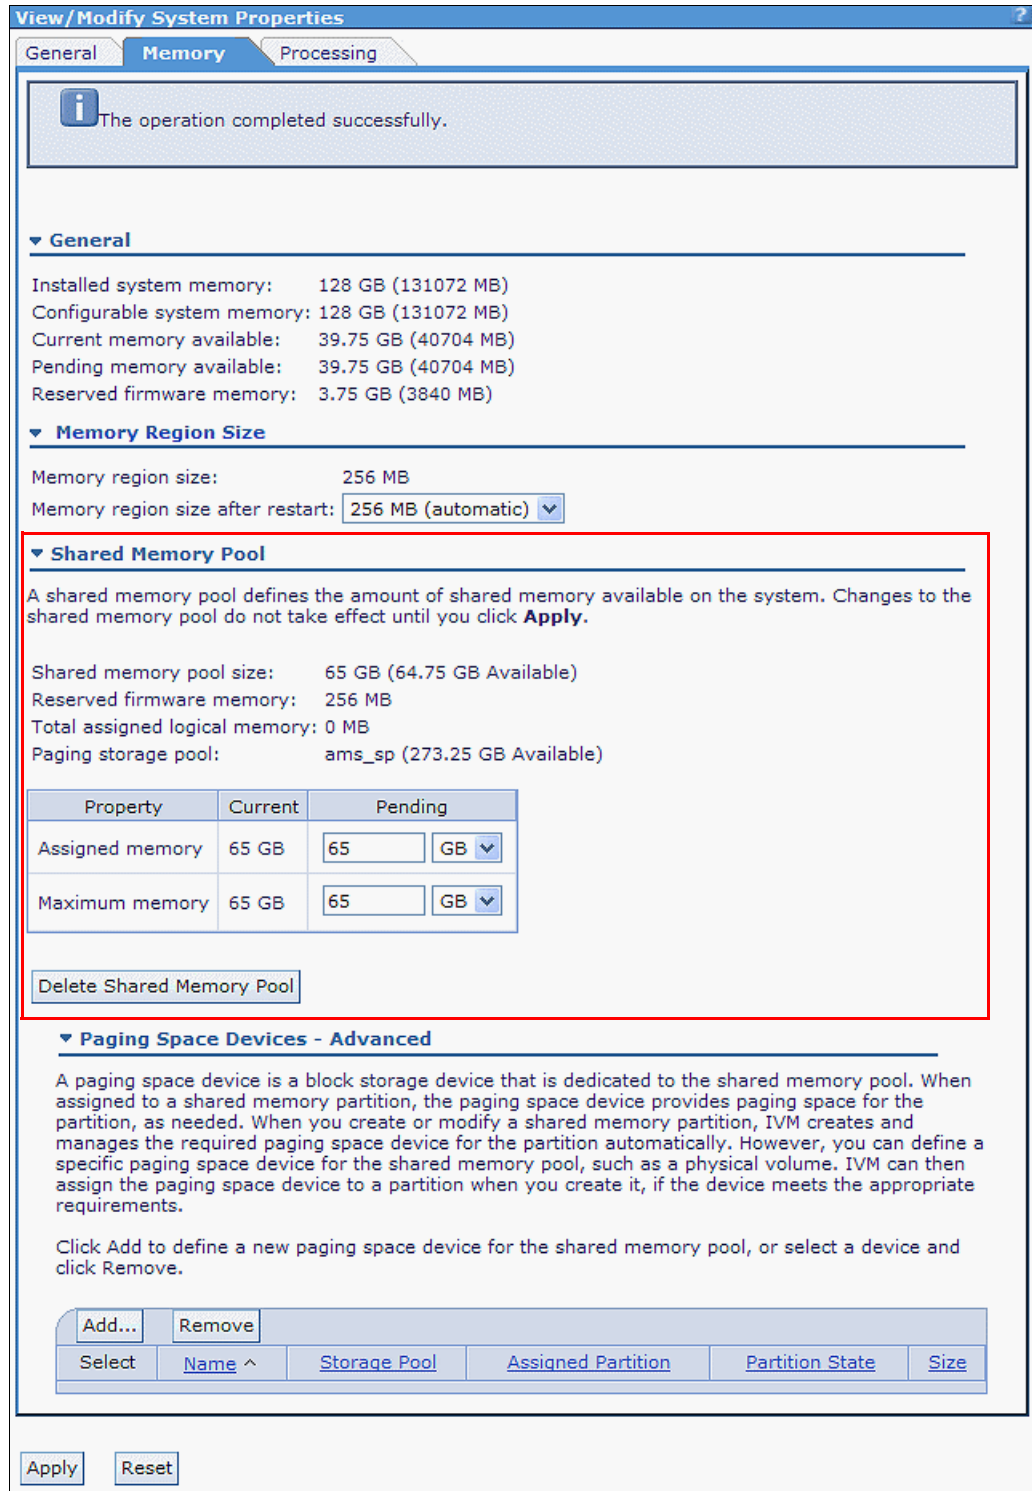


Figure 4-25 The IVM memory configuration after AMS pool creation

4.3.2 Changing the AMS pool configuration

In the IVM, you can change the assigned memory and the maximum memory of the Active Memory Sharing (AMS) pool. Use the following steps to change these parameters:

1. Click **View/Modify Shared Memory Pool** in the navigation area, under the Partition Management menu.
2. Change the Assigned memory field to the wanted amount. Change the Maximum memory field to at least the same amount of the Assigned memory field. The maximum memory for the shared pool must be equal or greater than the assigned memory. See Figure 4-26.




Figure 4-26 Changing the shared memory pool size

3. Click **Apply** to change the AMS pool size.

4.3.3 Removing the AMS pool

To remove the Active Memory Sharing (AMS) pool, you cannot have logical partitions using the shared memory pool. Ensure that you change all of your LPARs from the shared memory mode to the dedicated memory mode before you remove the AMS pool, otherwise the removal fails.

The following steps describe how to remove the AMS pool.

1. Click **View/Modify Shared Memory Pool** in the navigation area, under the Partition Management menu.
2. Click **Delete Shared Memory Pool** to remove the AMS pool, as shown in Figure 4-27.

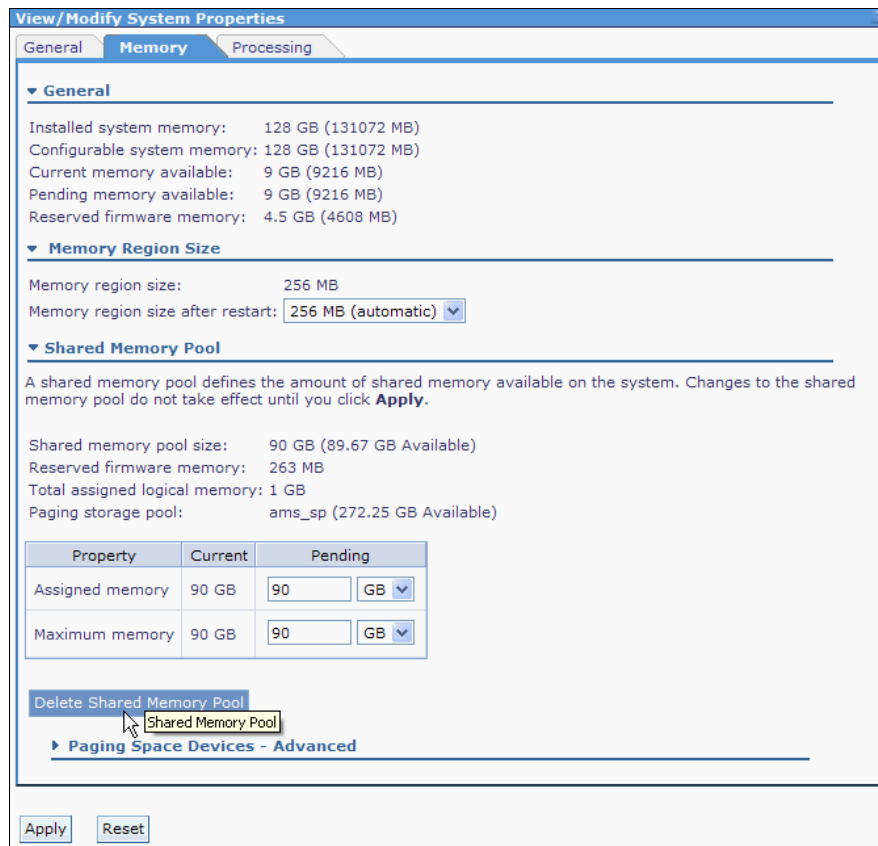


Figure 4-27 Shared memory pool deletion

3. A confirmation window opens, as depicted in Figure 4-28. Click **OK** to continue.

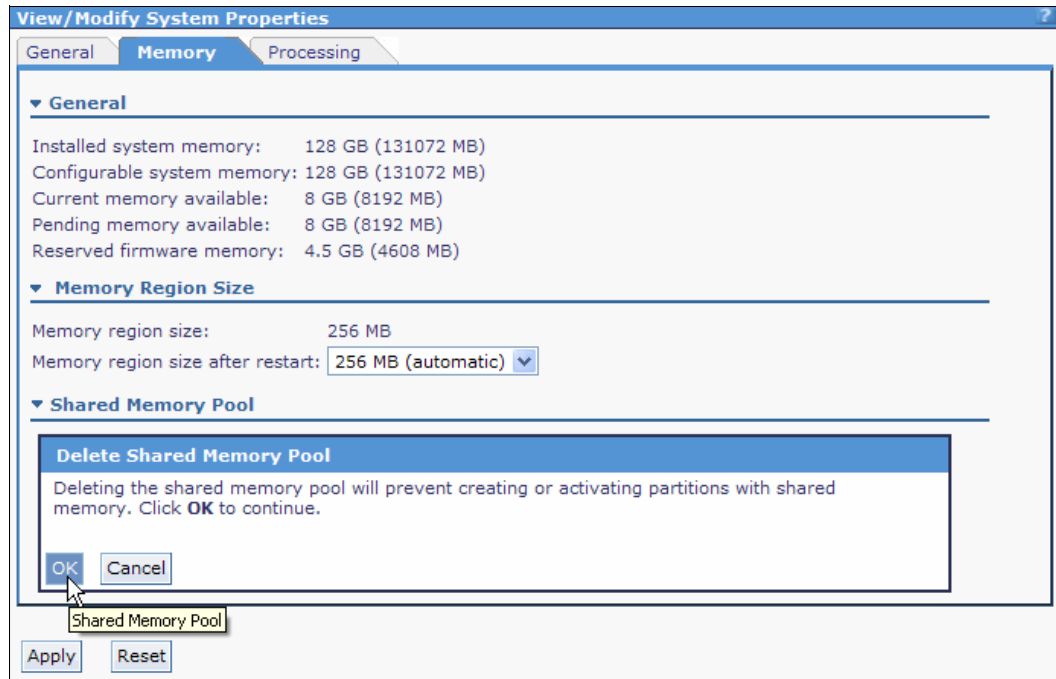


Figure 4-28 Shared memory pool deletion warning

4. The shared pool deletion is still pending, as shown in Figure 4-29. Click **Apply** to change and remove the AMS pool.

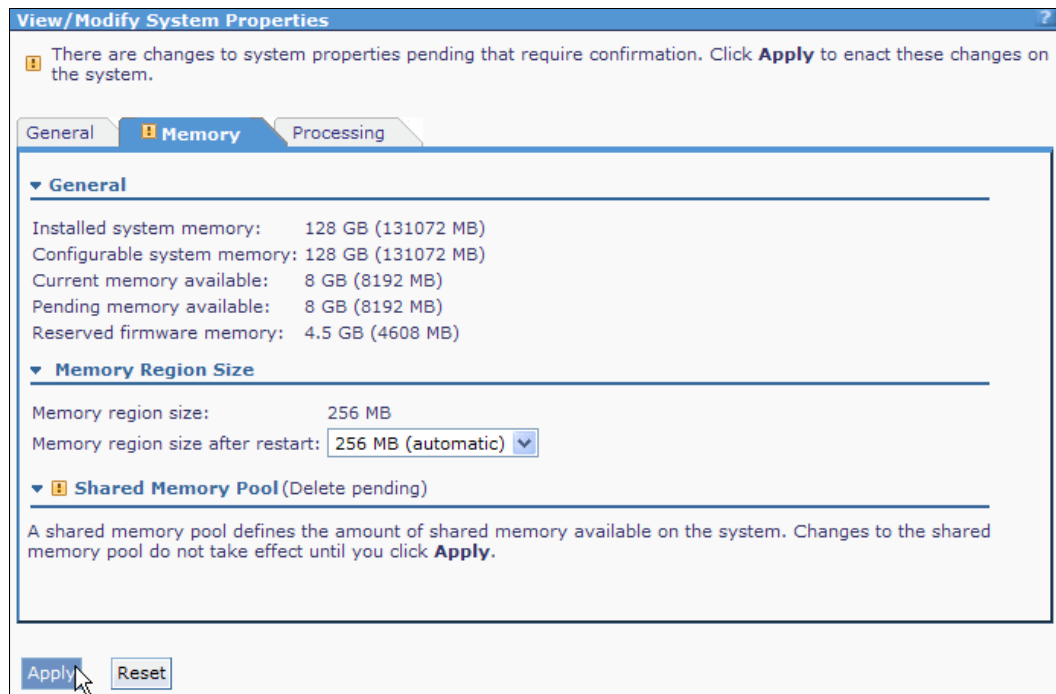


Figure 4-29 Applying shared memory pool deletion

4.3.4 Adding or removing specific paging devices to shared memory pool

The IVM manages shared memory paging devices automatically, by creating and extending paging space devices in the paging storage pool, as needed. But you can also add specific paging devices to the shared memory pool, even if they do not belong to the paging storage pool. You can even add physical volumes as paging space devices. IVM automatically assigns these paging space devices according to the requirements of the LPAR.

Adding specific paging devices

To add specific paging devices to the shared memory pool, follow these steps:

1. Click **View/Modify Shared Memory Pool** in the navigation area, under the Partition Management menu.
2. In the work area, expand the Paging Space Devices - Advanced menu. Click **Add** to assign more devices to the shared memory pool (Figure 4-30).

View/Modify System Properties

General **Memory** Processing

▼ **General**

Installed system memory: 128 GB (131072 MB)
 Configurable system memory: 128 GB (131072 MB)
 Current memory available: 34.75 GB (35584 MB)
 Pending memory available: 34.75 GB (35584 MB)
 Reserved firmware memory: 3.75 GB (3840 MB)

▼ **Memory Region Size**

Memory region size: 256 MB
 Memory region size after restart: 256 MB (automatic) ▼

▼ **Shared Memory Pool**

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 65 GB (64.67 GB Available)
 Reserved firmware memory: 263 MB
 Total assigned logical memory: 1 GB
 Paging storage pool: ams_sp (272.25 GB Available)

| Property | Current | Pending | |
|-----------------|---------|---------|------|
| Assigned memory | 65 GB | 65 | GB ▼ |
| Maximum memory | 65 GB | 65 | GB ▼ |

Delete Shared Memory Pool

▼ **Paging Space Devices - Advanced**

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click Add to define a new paging space device for the shared memory pool, or select a device and click Remove.

Add.. Remove

| Select | Name ^ | Storage Pool | Assigned Partition | Partition State | Size |
|-----------------------|--------|--------------|--------------------|-----------------|------|
| <input type="radio"/> | lv00 | ams_sp | lpar3_server1 (4) | Not Activated | 1 GB |

Apply Reset

Figure 4-30 Paging Space Devices - Advanced menu

3. A new window with a list of paging device candidates opens under the Shared Memory Pool menu in the work area, as shown in Figure 4-31. Select the wanted device to add as a paging space device and click **OK**.

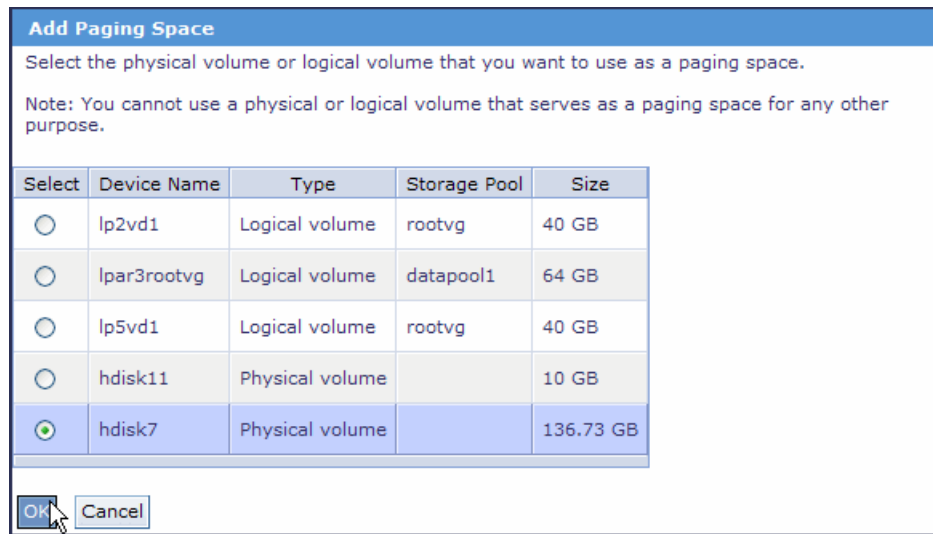


Figure 4-31 Paging space device selection

- You will see exclamation marks in the Paging Space Devices - Advanced menu. This view means that the configuration is still pending, as depicted in Figure 4-32. Click **Apply** to make the changes.

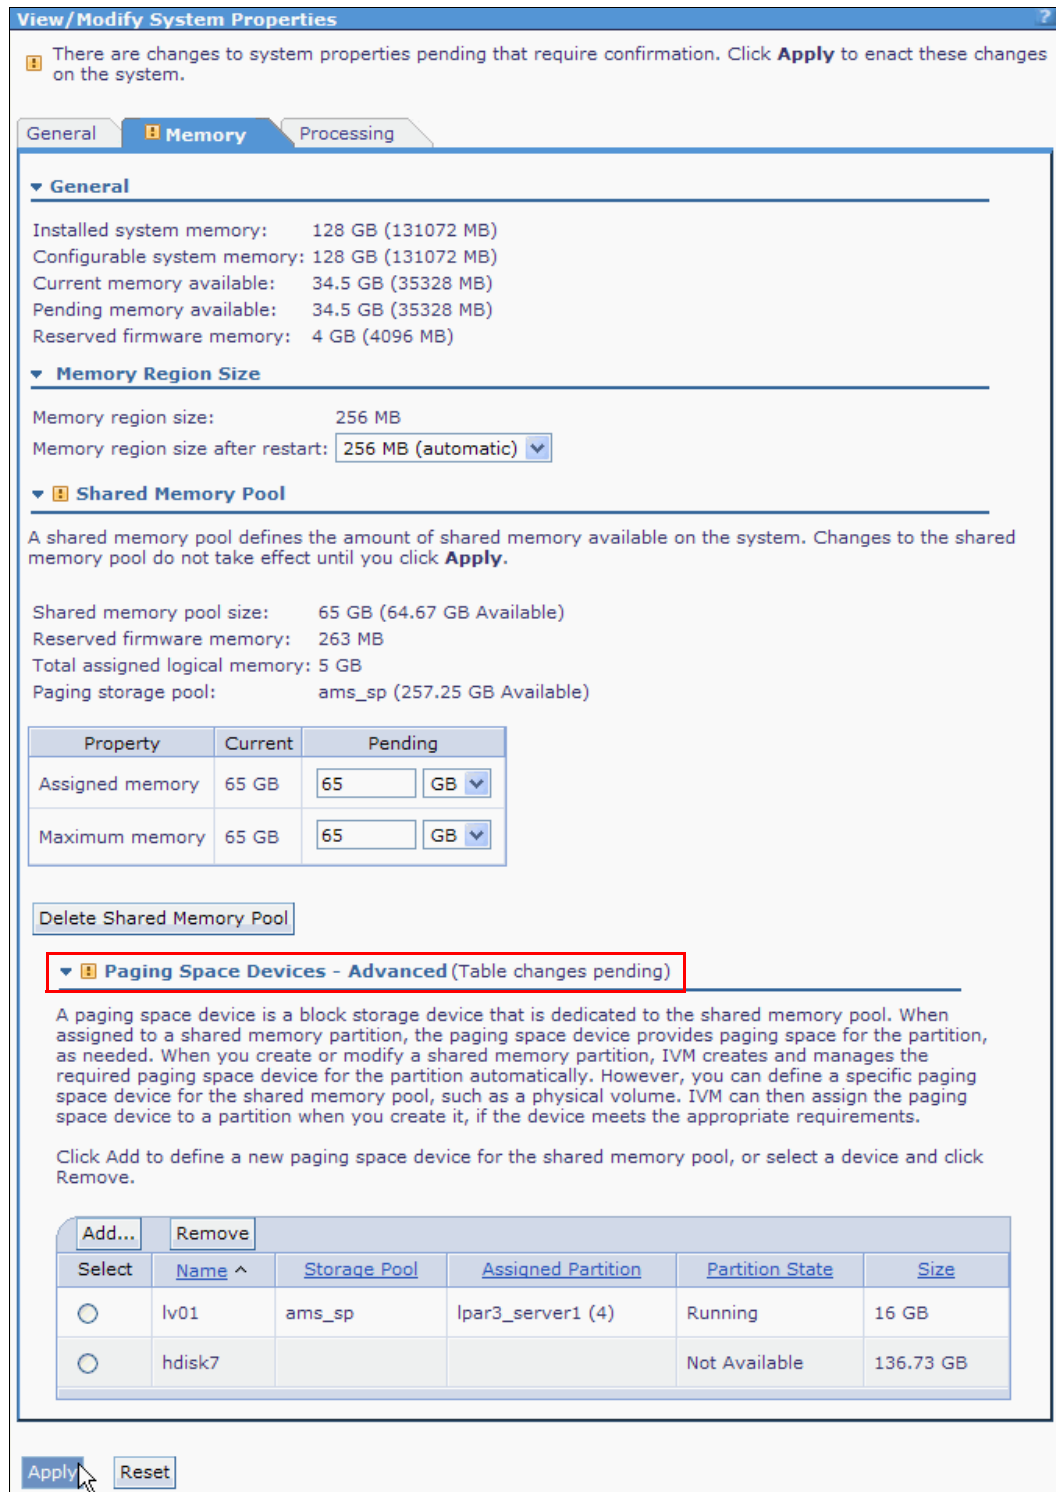


Figure 4-32 Adding paging space devices

Removing specific paging devices

To remove a specific paging device, it cannot be assigned to a logical partition. Ensure that the paging space device is not assigned to any LPAR before removing. Follow these steps to remove from the shared memory pool:

1. Click **View/Modify Shared Memory Pool** in the navigation area, under the Partition Management menu.

- In the work area, expand the Paging Space Devices - Advanced menu. Select the wanted paging device that you want to remove. Ensure that the device is not assigned to any partition. Click **Remove** to remove the device from the shared memory pool, as shown in Figure 4-33.

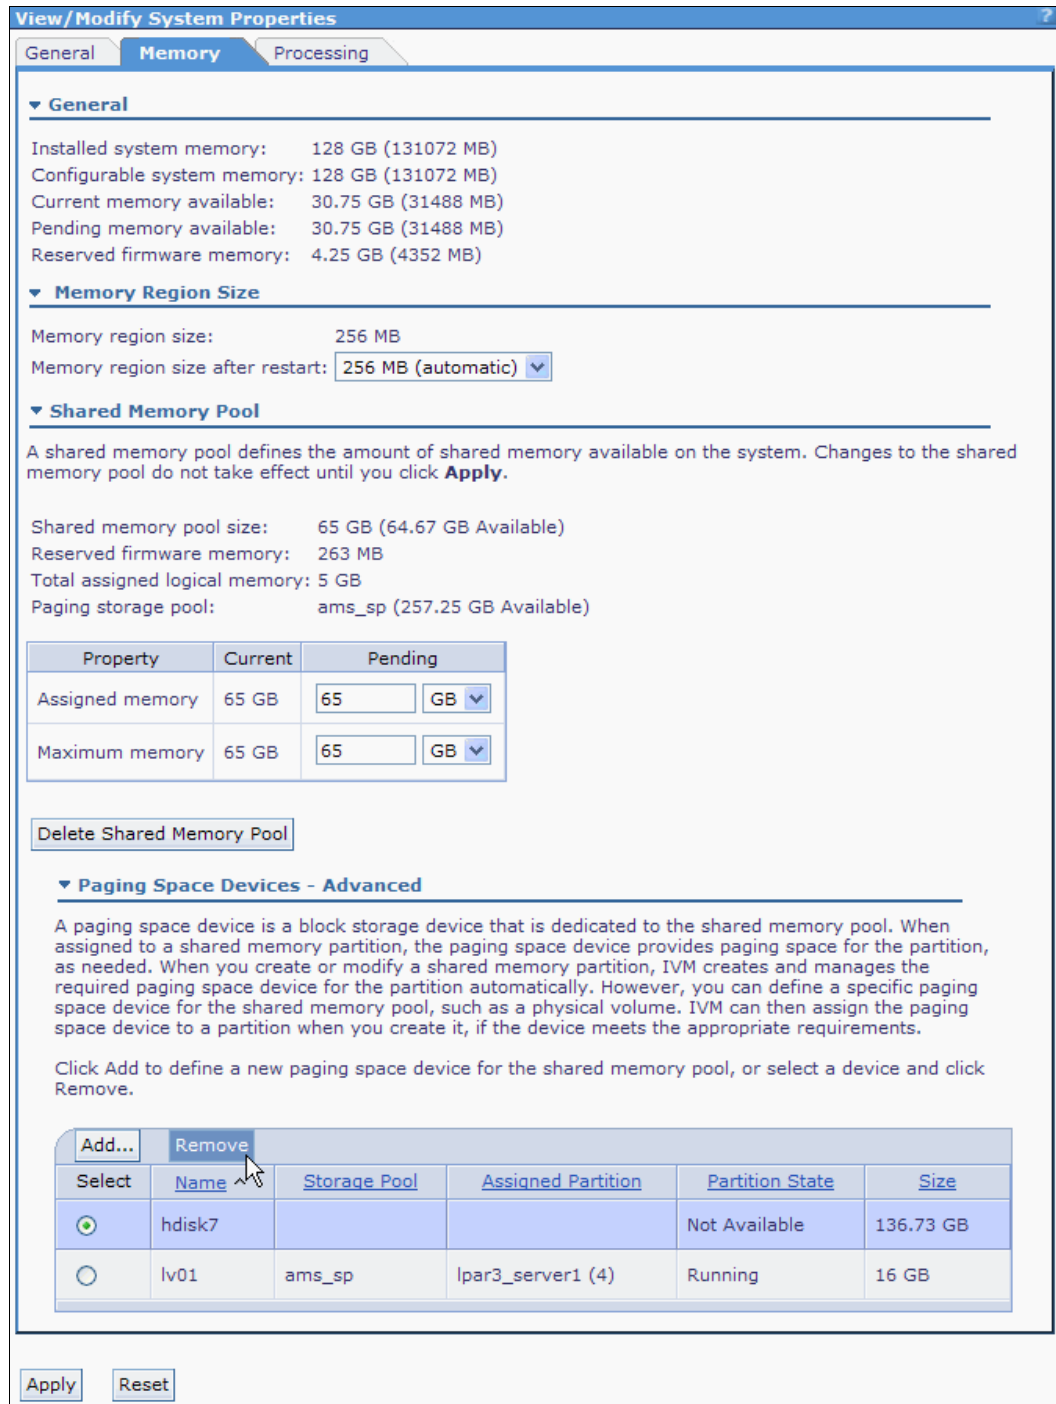


Figure 4-33 Removing paging space devices from the shared memory pool

3. A confirmation window is displayed, as shown in Figure 4-34. Click **Ok** to confirm.




Figure 4-34 *Paging space device removal confirmation*

- You will see exclamation marks in the Paging Space Devices - Advanced menu. This view means that the configuration is still pending, as depicted in Figure 4-35. Click **Apply** to make the changes.

View/Modify System Properties

There are changes to system properties pending that require confirmation. Click **Apply** to enact these changes on the system.

General **Memory** Processing

General

Installed system memory: 128 GB (131072 MB)
 Configurable system memory: 128 GB (131072 MB)
 Current memory available: 30.75 GB (31488 MB)
 Pending memory available: 30.75 GB (31488 MB)
 Reserved firmware memory: 4.25 GB (4352 MB)

Memory Region Size

Memory region size: 256 MB
 Memory region size after restart: 256 MB (automatic)

Shared Memory Pool

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 65 GB (64.67 GB Available)
 Reserved firmware memory: 263 MB
 Total assigned logical memory: 5 GB
 Paging storage pool: ams_sp (257.25 GB Available)

| Property | Current | Pending | |
|-----------------|---------|---------|----|
| Assigned memory | 65 GB | 65 | GB |
| Maximum memory | 65 GB | 65 | GB |

Delete Shared Memory Pool

Paging Space Devices - Advanced (Table changes pending)

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click Add to define a new paging space device for the shared memory pool, or select a device and click Remove.

Add... Remove

| Select | Name ^ | Storage Pool | Assigned Partition | Partition State | Size |
|-----------------------|--------|--------------|--------------------|-----------------|-------|
| <input type="radio"/> | lv01 | ams_sp | lpar3_server1 (4) | Running | 16 GB |

Apply Reset

Figure 4-35 Removing paging space devices from the shared memory pool

4.4 Create logical partitions

A *logical partition (LPAR)* is a set of resources: Processors, memory, and I/O devices. Each resource that is assigned to an LPAR is allocated, regardless of whether the LPAR is running or not.

The following steps describe how to create an LPAR:

1. Under the Partition Management menu in the navigation area, click **Create Partition** in the work area.
2. The Create Partition wizard window opens. Type a name for the new partition, as shown in Figure 4-36. Click **Next**.

The screenshot shows the 'Create Partition: Name' wizard window, Step 1 of 11. The window has a blue header bar with the title 'Create Partition: Name' and 'Step 1 of 11'. On the left side, there is a navigation pane with a tree view containing the following items: Name (selected), Memory, Processors, Ethernet, Storage Type, Storage, Virtual Fibre Channel, Optical/Tape, Physical Adapters, and Summary. The main area of the window is titled 'Name' and contains the text 'To create a partition complete the following information.' Below this text are four fields: 'System name: 8233-E8B_SN061AA6P', 'Partition ID: 5', '* Partition name: lpar4_server1' (with a red asterisk indicating a required field), and 'Environment: AIX or Linux' (with a dropdown arrow). At the bottom of the window, there are five buttons: '< Back', 'Next' (highlighted with a mouse cursor), 'Finish', 'Cancel', and 'Help'. A red asterisk and the text '* Required field' are located at the bottom left of the main area.

Figure 4-36 Create Partition: Name

3. Select the memory mode for the new partition. Figure 4-37 shows the window for the shared memory mode. Figure 4-38 shows the window for the dedicated memory mode. Enter the amount of memory needed. Click **Next**.

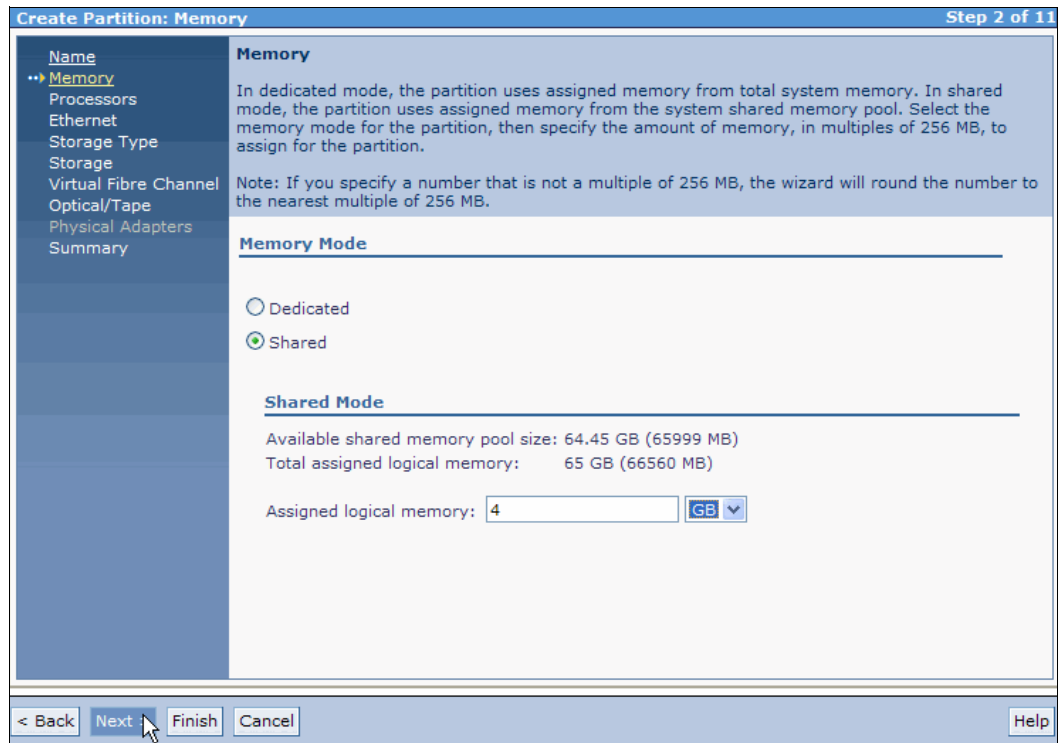


Figure 4-37 Create Partition: Assign memory for the shared memory mode

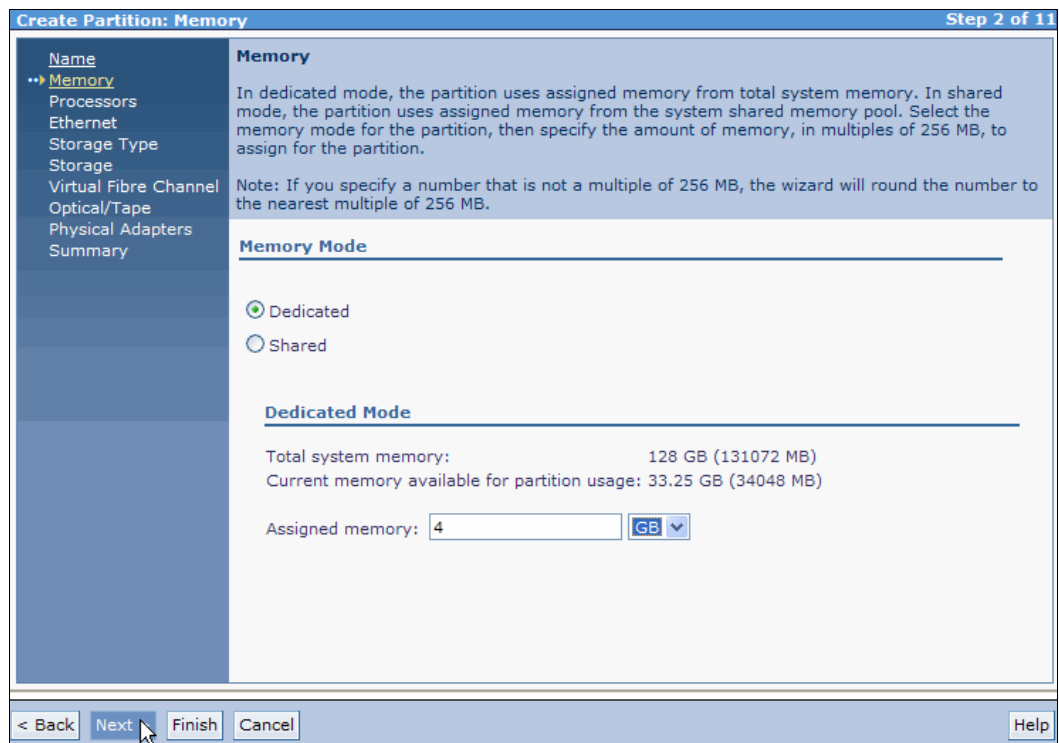


Figure 4-38 Create Partition: Assign memory for the dedicated memory mode

4. Select the number of processors that are needed and choose a processing mode, as shown in Figure 4-39. If you choose the shared memory mode, you cannot use dedicated processors. In the shared processor mode, each virtual processor uses 0.1 processing units at the time of LPAR creation. Click **Next**.

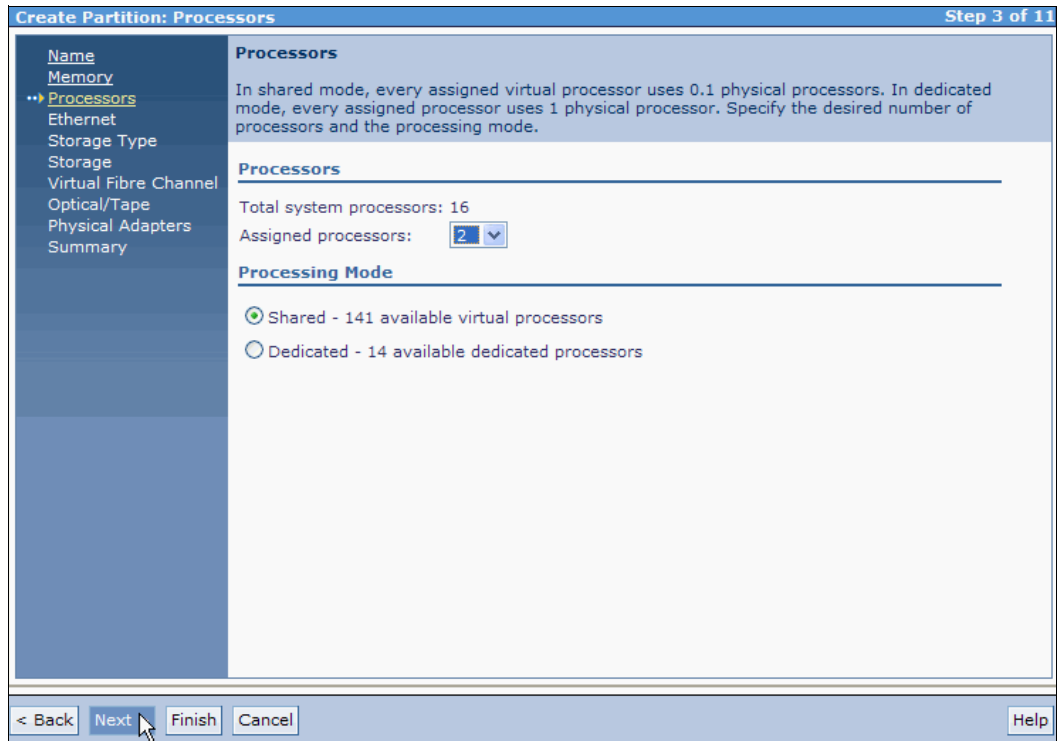


Figure 4-39 Create Partition: Processors (and Processing Mode)

- Each partition has two virtual Ethernet adapters that can be configured to one of the four available virtual Ethernet connections. In Figure 4-40, adapter 1 uses virtual Ethernet ID 1. If your system has a Host Ethernet Adapter (HEA), the HEA ports section opens. If you want to use a virtual port from an HEA, check the wanted HEA port check box.

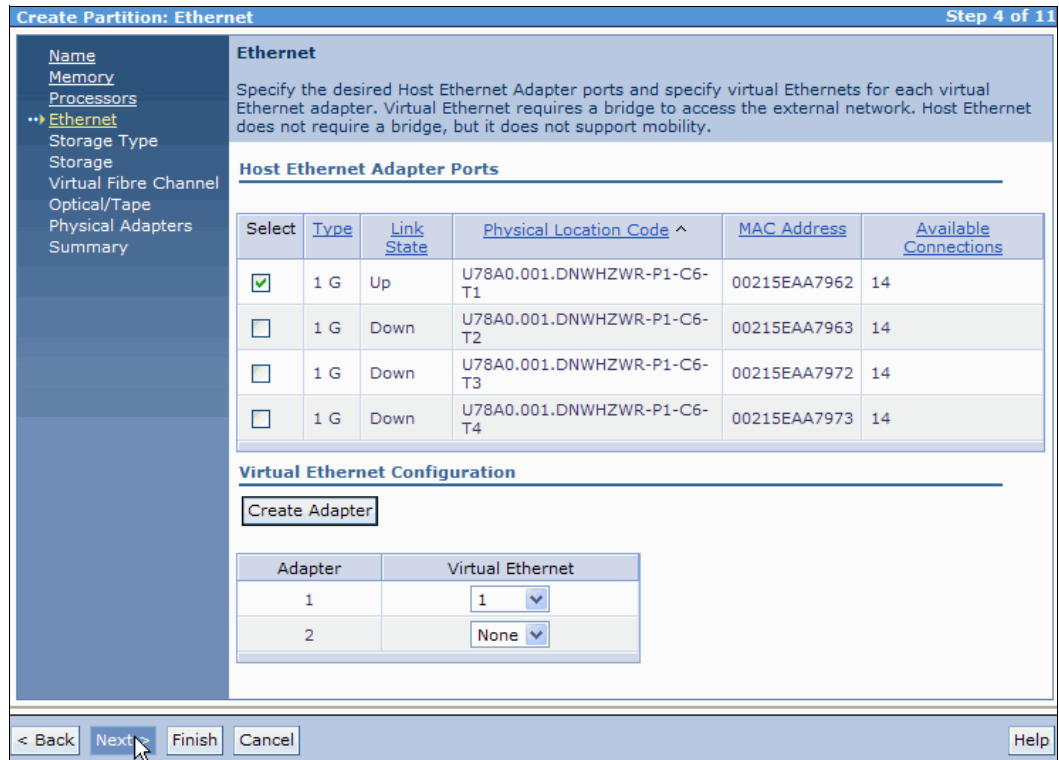


Figure 4-40 Create Partition: Virtual Ethernet and Host Ethernet Adapter

6. Select **Assign existing virtual disks and physical volumes**, as shown in Figure 4-41.

You can also allow the IVM to create a virtual disk for you by selecting **Create virtual disk**, but be aware that the virtual disk is created in the default storage pool. To create a storage pool and virtual disk, or to change the default storage pool, refer to 4.2.3, “Storage pool disk management” on page 52. Click **Next**.




Figure 4-41 Create Partition: Storage Type

- Select the needed virtual disks or physical volumes from the list, as shown in Figure 4-42. Click **Next**.

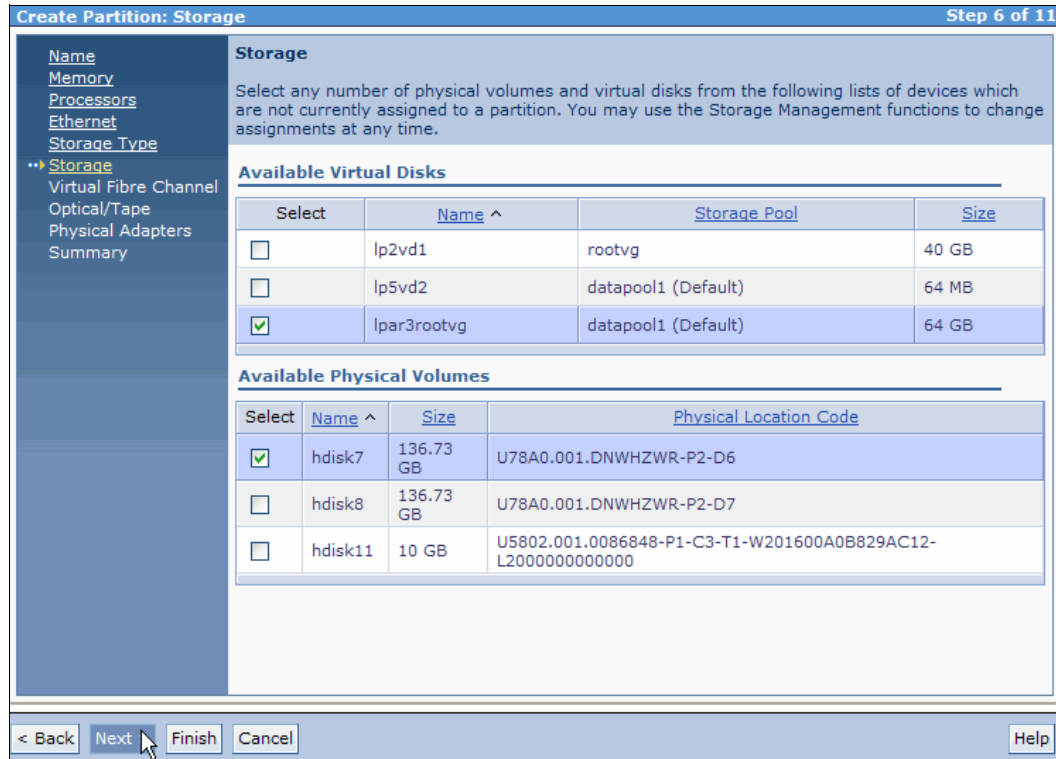


Figure 4-42 Create Partition: Storage

8. If you want to add a virtual Fibre Channel adapter to the new partition, click **Add**, as shown in Figure 4-43.

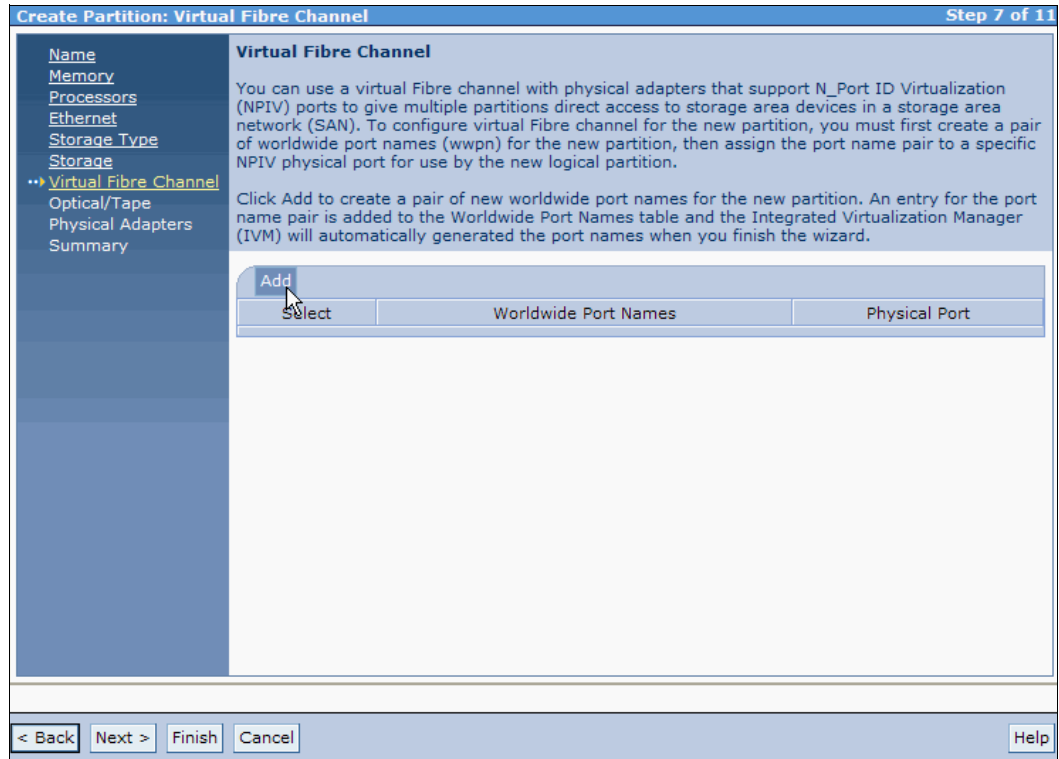


Figure 4-43 Create Partition: Virtual Fibre Channel

- Under the Physical Port column, select the wanted physical port for the virtual Fibre Channel adapter, as depicted in Figure 4-44. Click **Next**.

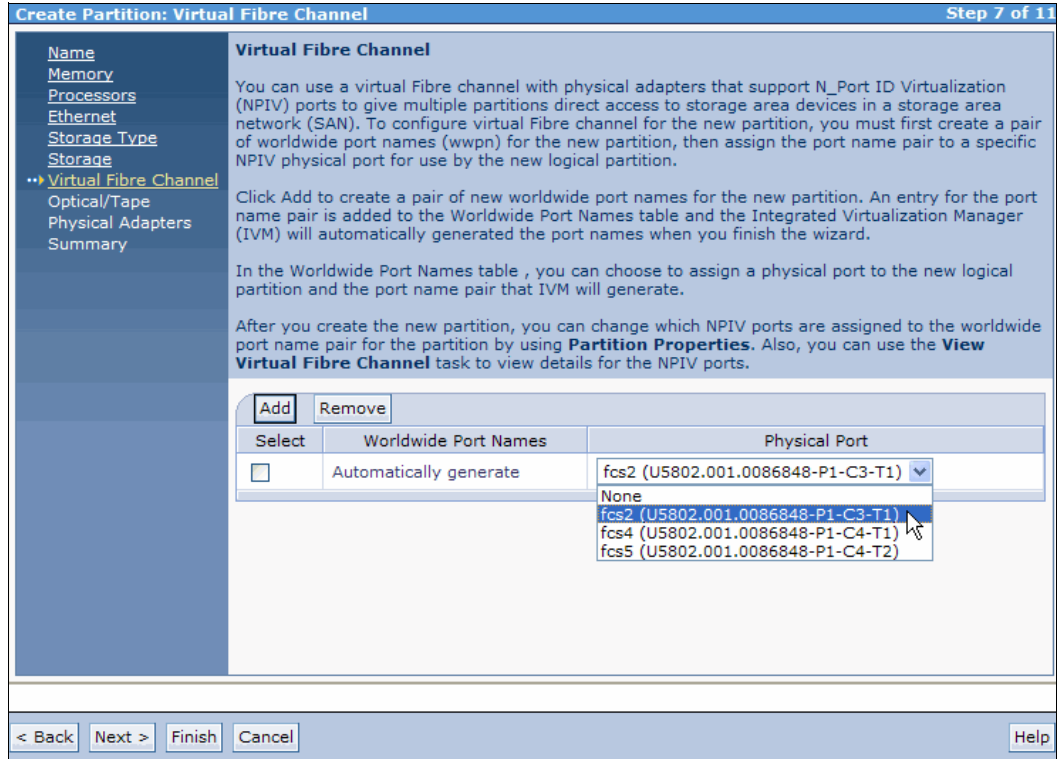


Figure 4-44 Virtual Fibre Channel adapter port selection

10. Select the needed virtual or physical optical and tape devices, as shown in Figure 4-45. Click **Next**.

Create Partition: Optical/Tape Step 8 of 11

Optical/Tape
Select optical or tape devices from the following list of devices which are not currently assigned to a partition.

Physical Optical Devices

Select one or more unassigned physical optical devices that you want to assign directly to the partition to use for storage.

| Select | Name ^ | Description | Physical Location Code |
|-------------------------------------|--------|--------------------|-------------------------|
| <input checked="" type="checkbox"/> | cd0 | SATA DVD-RAM Drive | U78A0.001.DNWHZWR-P2-D2 |

Virtual Optical Devices

You can use virtual optical devices to mount and unmount media files, such as an ISO image, that are in your media library for use by the partition. Select a virtual optical device in the table to assign it to the new partition. Clear the selection for a device if you do not want to assign it to the partition. Click Modify to change the mounted media for a specific optical device. Click Create Device to add a new optical device for the partition.

| Select | Name ^ | Current Media | Current Media Size | Mount Type |
|-------------------------------------|----------|-----------------------------|--------------------|------------|
| <input checked="" type="checkbox"/> | Unknown1 | None Modify | | |

Physical Tape Devices

Select one or more physical tape devices that you want to assign directly to the partition to use for storage.

| Select | Name ^ | Description | Physical Location Code |
|-------------------------------------|--------|-----------------------------------|-------------------------|
| <input checked="" type="checkbox"/> | rmt0 | SAS 1/2-inch Cartridge Tape Drive | U78A0.001.DNWHZWR-P2-D1 |

< Back **Next** Finish Cancel Help

Figure 4-45 Create Partition: Optical/Tape (Devices)

11. Select the needed physical adapters as shown in Figure 4-46. Click **Next**.

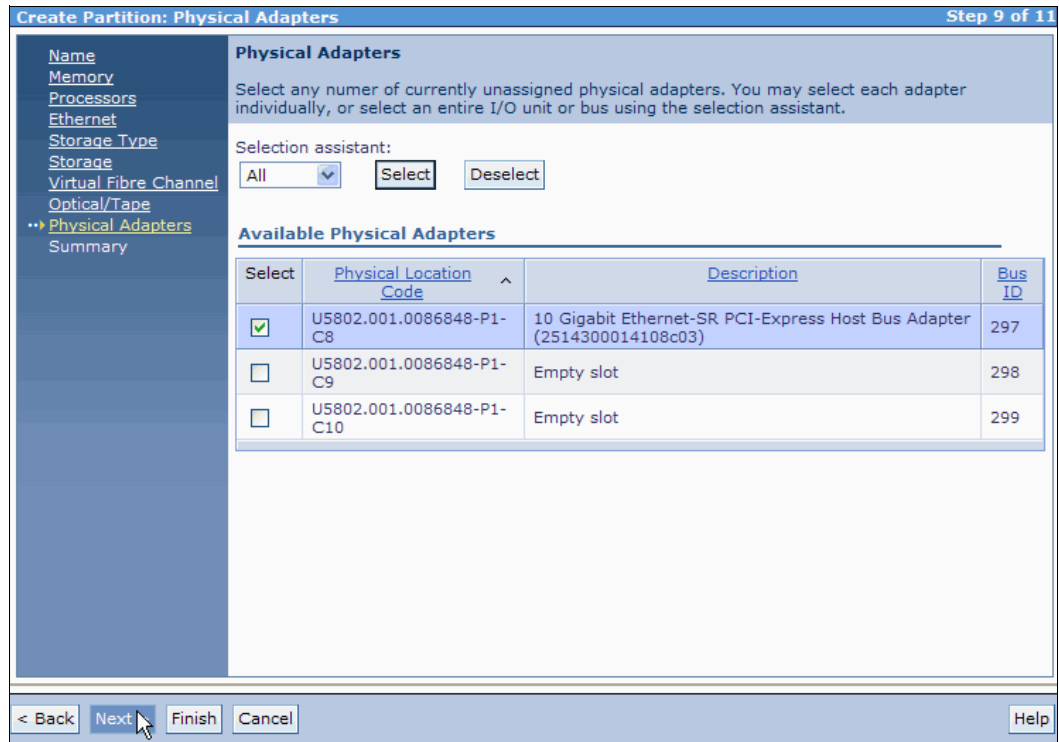


Figure 4-46 Create Partition: Physical Adapters

12. A summary of the partition to be created opens, as shown in Figure 4-47. Click **Finish** to create the LPAR.

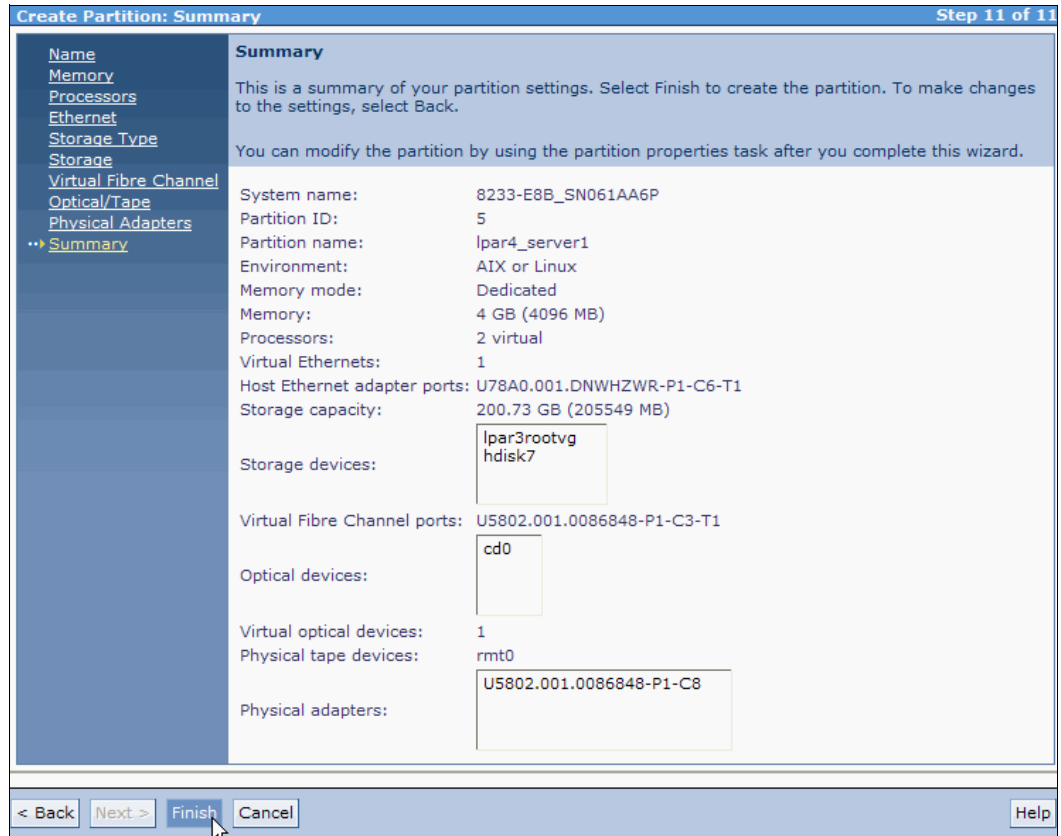


Figure 4-47 Create Partition: Summary

To view the new LPAR and use it, under the Partition Management menu in the navigation area, click the **View/Modify Partitions** link. A list opens in the work area.

4.4.1 Create an LPAR based on an existing partition

The IVM can be used to create an LPAR that is based on an existing partition on your managed system. Any role other than *View Only* can be used to do this task. This task enables you to create an LPAR with the same properties as the selected existing partition (except for ID, name, physical volumes, and optical devices).

To create an LPAR that is based on an existing partition, do the following steps:

1. Under Partition Management, click **View/Modify Partitions**. The View/Modify Partitions panel opens.
2. Select the LPAR that you want to use as a basis for the new partition.
3. In the **More Tasks** drop-down menu in the work area, select **Create based on** as shown in Figure 4-48.

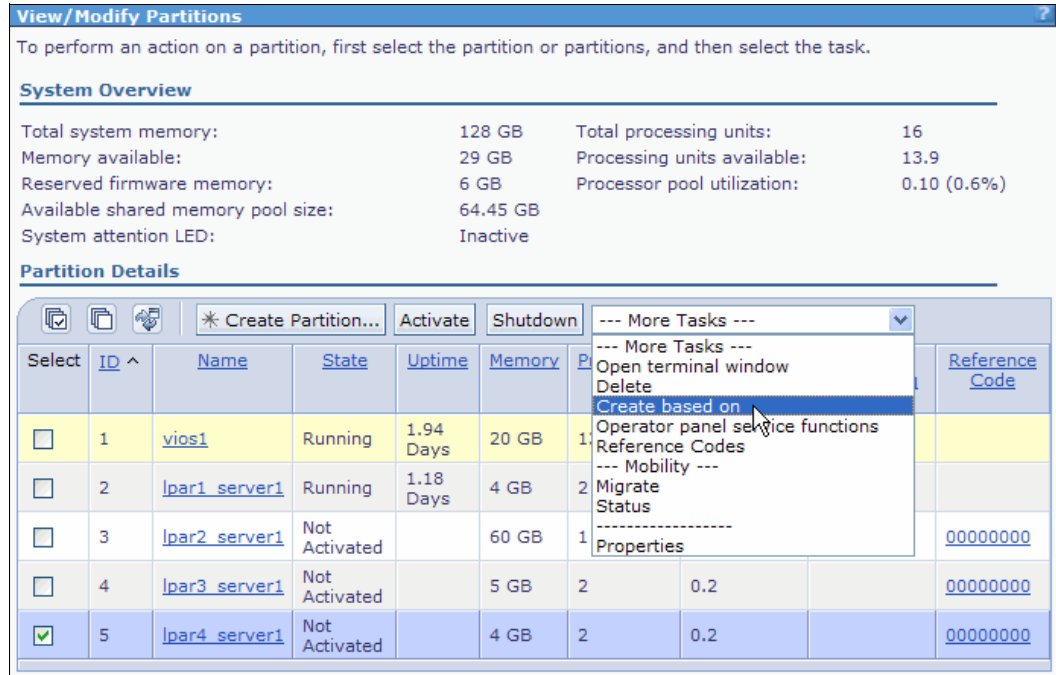


Figure 4-48 The Create based on selection from the More Tasks drop-down menu

4. The Create Based On panel opens (Figure 4-49). Enter the name of the new partition, and click **OK**.

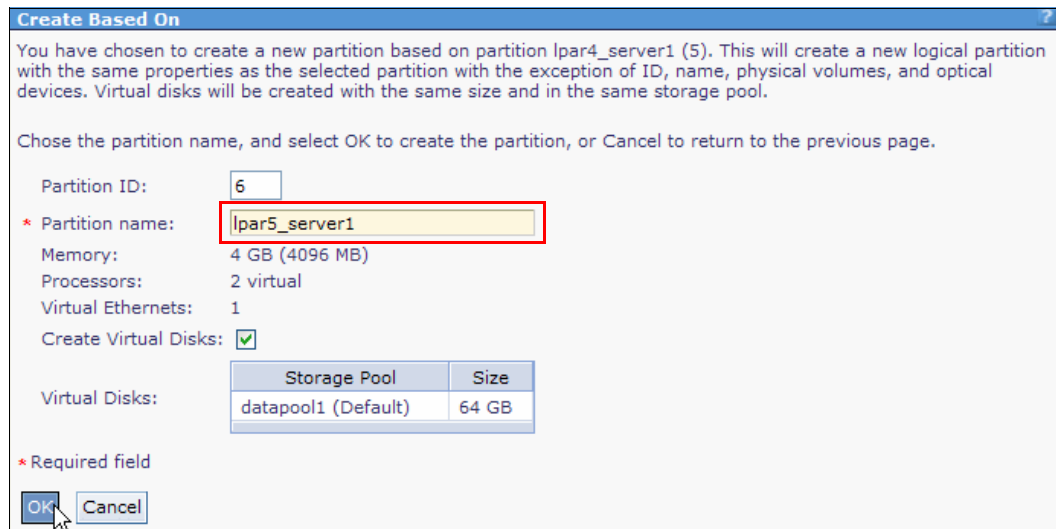


Figure 4-49 Create based on: name of the new LPAR

5. The View/Modify Partitions panel opens, showing the new partition (Figure 4-50).

The screenshot shows the 'View/Modify Partitions' window. At the top, it says 'To perform an action on a partition, first select the partition or partitions, and then select the task.' Below this is the 'System Overview' section with the following data:

| | | | |
|------------------------------------|----------|-----------------------------|-------------|
| Total system memory: | 128 GB | Total processing units: | 16 |
| Memory available: | 25 GB | Processing units available: | 13.7 |
| Reserved firmware memory: | 6 GB | Processor pool utilization: | 0.05 (0.3%) |
| Available shared memory pool size: | 64.45 GB | | |
| System attention LED: | Inactive | | |

Below the system overview is the 'Partition Details' section, which includes a toolbar with buttons for 'Create Partition...', 'Activate', 'Shutdown', and 'More Tasks'. Below the toolbar is a table of partitions:

| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
|--------------------------|------|-------------------------------|---------------|-----------|--------|------------|---------------------------|---------------------------|----------------|
| <input type="checkbox"/> | 1 | vios1 | Running | 1.95 Days | 20 GB | 13 | 1.3 | 0.04 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 1.19 Days | 4 GB | 2 | 0.3 | 0.01 | |
| <input type="checkbox"/> | 3 | lpar2_server1 | Not Activated | | 60 GB | 1 | 0.1 | | 00000000 |
| <input type="checkbox"/> | 4 | lpar3_server1 | Not Activated | | 5 GB | 2 | 0.2 | | 00000000 |
| <input type="checkbox"/> | 5 | lpar4_server1 | Not Activated | | 4 GB | 2 | 0.2 | | 00000000 |
| <input type="checkbox"/> | 6 | lpar5_server1 | Not Activated | | 4 GB | 2 | 0.2 | | 00000000 |

Figure 4-50 Create based on: New logical partition is created

The virtual disks that are created have the same size and are in the same storage pool as the selected partition. However, the data in these disks is not cloned.

4.4.2 Shutting down logical partitions

The IVM provides the following types of shutdown options for LPARs:

- ▶ Operating system (recommended)
- ▶ Delayed
- ▶ Immediate

The recommended shutdown method is to use the client operating system **shutdown** command. Use the immediate shutdown method as a last resort because it causes an abnormal shutdown that might result in data loss. (It is equivalent to pulling the power cord.)

If you choose not to use the operating system shutdown method, be aware of these considerations:

- ▶ Shutting down the LPARs is equivalent to pressing and holding the white control-panel power button on a server that is not partitioned.
- ▶ Use this procedure only if you cannot successfully shut down the LPARs through operating system commands. When you use this procedure to shut down the selected LPARs, the LPARs wait a predetermined amount of time to shut down. This procedure gives the LPARs time to end jobs and write data to disks. If the LPAR is unable to shut down within the predetermined amount of time, it ends abnormally, and the next restart might take a long time.

Take the following steps to shut down an LPAR:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions panel opens.
2. Select the LPAR that you want to shut down.
3. From the Tasks menu, click **Shutdown**. The Shutdown Partitions panel opens (Figure 4-51).

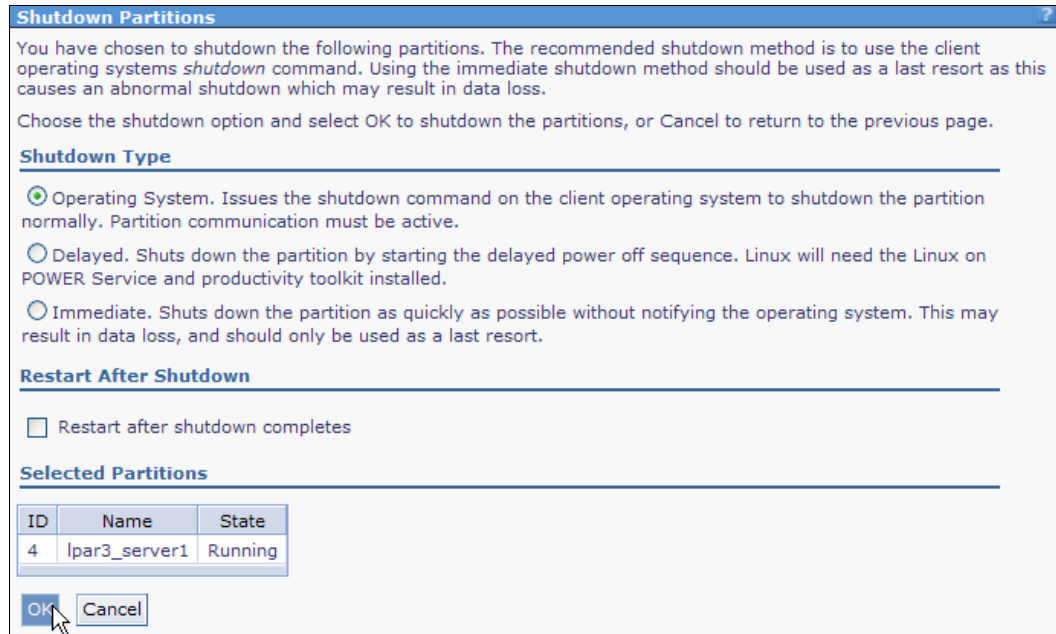


Figure 4-51 Shutdown Partitions: new options

4. Select **Shutdown Type**.
5. Optionally, select **Restart after shutdown completes** if you want the LPAR to start immediately after it shuts down.
6. Click **OK** to shut down the partition. The View/Modify Partitions panel is displayed, and the partition is shut down.

Shutdown of the IVM LPAR: If the LPAR does not have a Resource Monitoring and Control (RMC) connection, the operating system (OS) shutdown type is disabled. The *Delayed* type becomes the default selection. When the IVM logical partition is selected, the only available option is OS shutdown. In addition, a warning is placed at the top of the panel which indicates that shutting down the IVM LPAR affects other running LPARs.

4.4.3 Monitoring tasks

In IVM, you can view and monitor the most recent 40 tasks. All actions that a user can do in the GUI become *tasks*. All tasks are audited at the task level. Each task can have subtasks. The status of each subtask is managed. When you do a task, the user gets a *Busy* dialog, which indicates that the task is running. You can go away from the page and do other tasks.

To view the properties of the tasks, take the following steps:

1. In the Service Management menu, click **Monitor Tasks**. The Monitor Tasks panel opens.
2. Select the task for which you want to view the properties (Figure 4-52).




Figure 4-52 Monitor Tasks: the last 40 tasks

3. Click **Properties**. The Task Properties window opens, as shown in Figure 4-53.

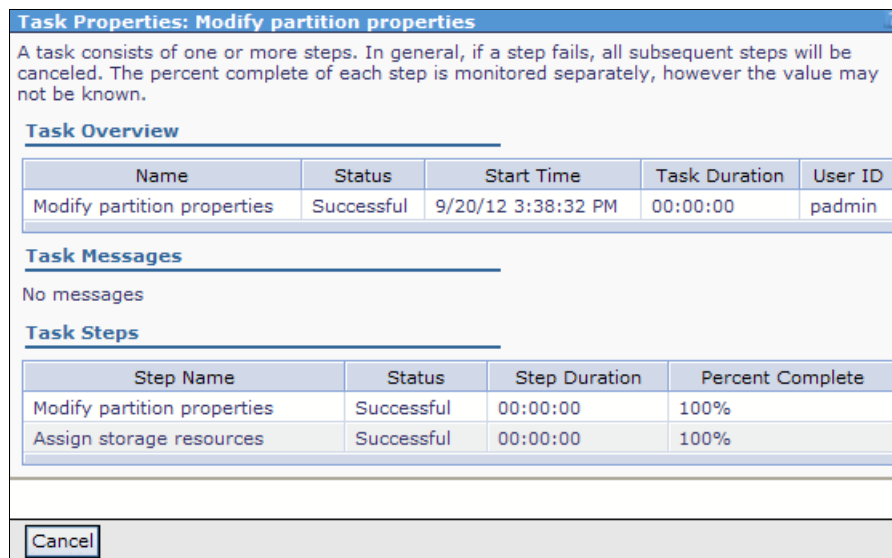


Figure 4-53 Task Properties

4. Click **Cancel** to close the Task Properties window. The Monitor Tasks panel opens.

You can also click the hyperlink of the task from which you want to view the properties. See more about hyperlinks in 4.4.4, “Hyperlinks for object properties” on page 94.

4.4.4 Hyperlinks for object properties

There are two ways to access the properties of an object. You can select the object and then select the properties task. Because this task is a frequent operation, there is another method to quickly access the properties of an object with one click: the hyperlinks. In a list view, the specified object has a hyperlink (typically on the Name) that, when selected, displays the properties sheet for the object. It behaves exactly as the Select → Properties method, but it requires only one click. Even if another object is selected, clicking the hyperlink of an object always brings it up, as shown in Figure 4-54.

The screenshot shows a window titled "View/Modify Partitions" with a toolbar and a table of partition details. The toolbar includes icons for selection, a "Create Partition..." button, and buttons for "Activate" and "Shutdown". A dropdown menu shows "More Tasks". The table below lists four partitions with hyperlinks in the Name column.

| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
|--------------------------|------|-------------------------------|---------------|-----------|--------|------------|---------------------------|---------------------------|----------------|
| <input type="checkbox"/> | 1 | vios1 | Running | 2 Days | 20 GB | 13 | 1.3 | 0.13 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 1.24 Days | 4 GB | 2 | 0.3 | 0.02 | |
| <input type="checkbox"/> | 3 | lpar2_server1 | Not Activated | | 60 GB | 1 | 0.1 | | 00000000 |
| <input type="checkbox"/> | 4 | lpar3_server1 | Running | 1 Hours | 5 GB | 2 | 0.2 | 0.01 | |

Figure 4-54 Hyperlinks

After you click the hyperlink, the properties of the objects open, as shown in Figure 4-55.

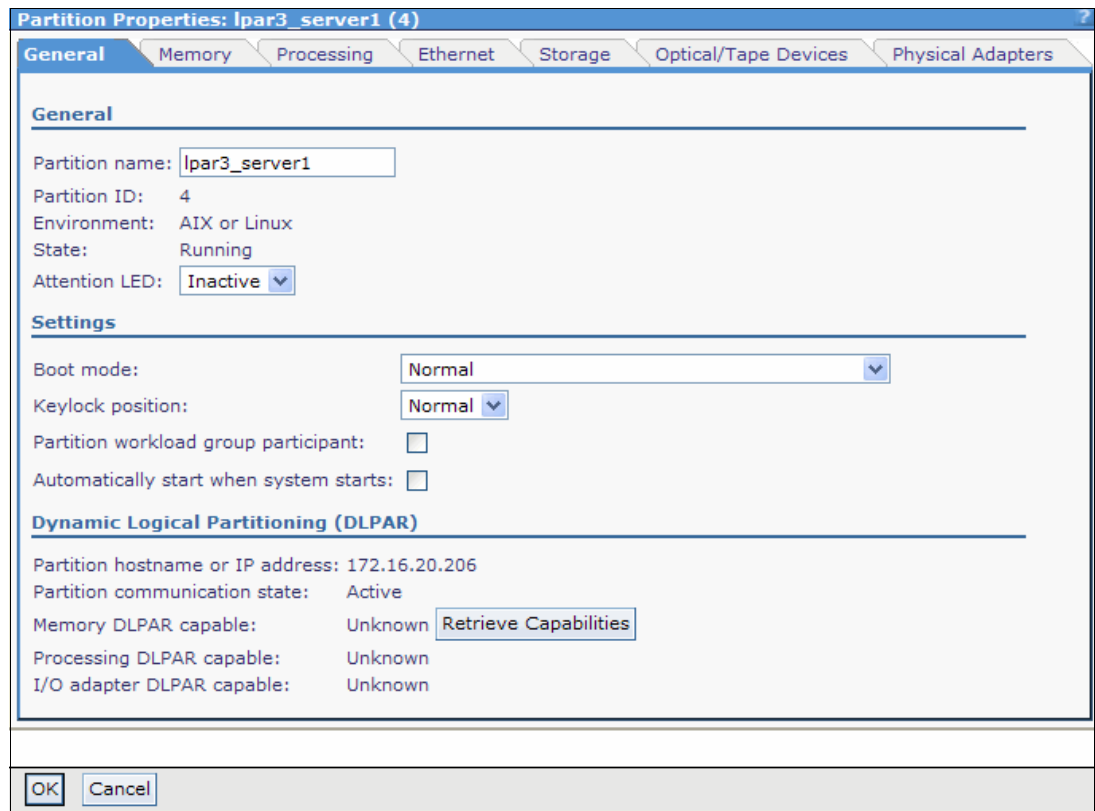


Figure 4-55 Properties of the object

4.5 IVM command-line interface

The text-based console with the command-line interface (CLI) is accessible through an ASCII terminal that is attached to a system port or through network connectivity by using the `telnet` command. The IP address is the same as the one used to connect to the GUI, and it was defined during the installation process. The CLI requires more experience to master than the GUI, but it offers more possibilities to tune the definitions of the partition and can be automated by using scripts.

4.5.1 Update the profile of the logical partition

Example 4-1 shows how to change the name of an LPAR with the `chsyscfg` command.

Example 4-1 Profile update

```
$ lssyscfg -r prof --filter "lpar_names=lpar1_server1" -F lpar_name  
lpar1_server1  
  
$ chsyscfg -r prof -i "lpar_name=lpar1_server1,new_name=LPAR1_new_name"  
  
$ lssyscfg -r prof --filter "lpar_names=LPAR1_new_name" -F lpar_name  
LPAR1_new_name
```

4.5.2 Power on a logical partition

Example 4-2 shows how to start an LPAR by using the `chsysstate` command. To follow the start process, use the `lsrefcode` command.

Example 4-2 Power on a partition

```
$ chsysstate -o on -r lpar -n lpar1_server1
$ lsrefcode -r lpar --filter "lpar_names=lpar1_server1" -F refcode
CA00E1F1
$ lsrefcode -r lpar --filter "lpar_names=lpar1_server1" -F refcode
CA00E14D
```

4.5.3 Console Terminal for IVM logical partition

There are two ways by which you can open a console for a VIO client LPAR:

- ▶ IVM GUI
- ▶ Virtual I/O Server command line

IVM GUI to open the console

Open the console for the Virtual I/O Server (the only way to access the console remotely for the Virtual I/O Server managed by IVM) and the Virtual I/O Server client by using this method.

To open console of a partition: This process requires a supported Java enabled browser.

Following are the steps to open the console of a partition:

1. Select the partition for which you want to open the console.
2. Then, select the **More Tasks** → **Open terminal window**.

Figure 4-56 shows the IVM menu to *Open terminal window*:

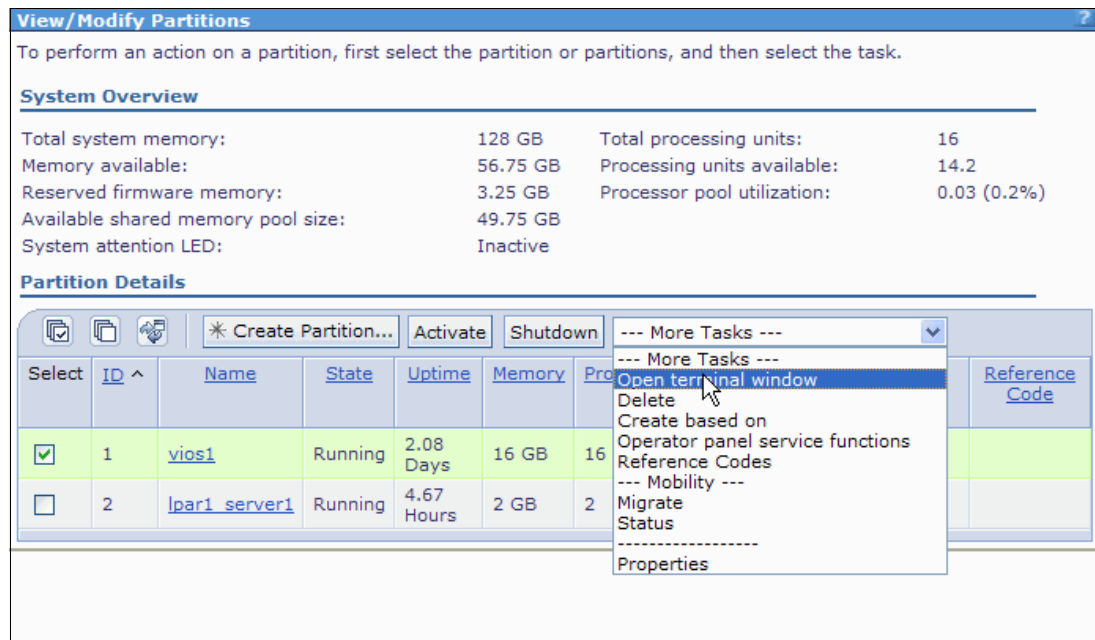


Figure 4-56 IVM menu to open the console

- When you click **Open terminal window**, the console window opens and prompts for the Virtual I/O Server (IVM) password (for both the Virtual I/O Server and virtual I/O client consoles). Then, enter the virtual I/O client password.

Figure 4-57 shows the console window that opens and waits for the user to enter their Virtual I/O Server password.

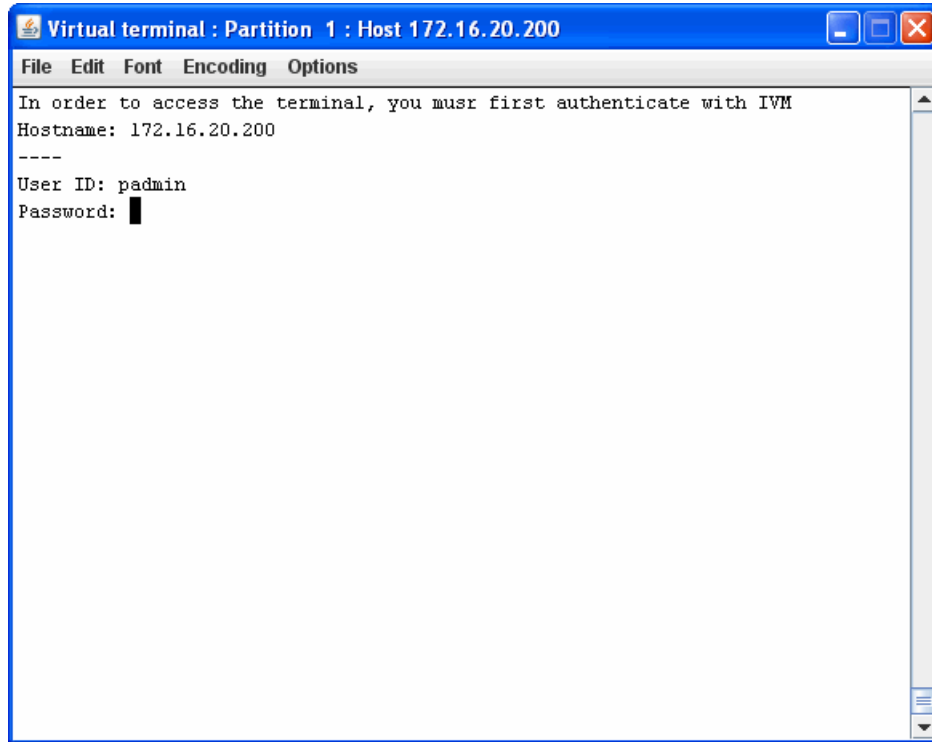


Figure 4-57 Console window of the IVM GUI

Using Virtual I/O Server command line to open the console

By using the command line, you can open the console only for VIO clients, not for the Virtual I/O Server.

By using the command line, take the following steps to open the console for VIO clients:

- Telnet or Secure Shell (SSH) to Virtual I/O Server.
- Use `mkvt -id <partition ID>` to open the console.

In Figure 4-58, the highlighted box in the IVM GUI shows which partition ID belongs to the corresponding partitions.

| Partition Details | | | | | | | | | |
|--------------------------|------|-------------------------------|---------|------------|--------|------------|---------------------------|---------------------------|----------------|
| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
| <input type="checkbox"/> | 1 | vios1 | Running | 2.11 Days | 16 GB | 16 | 1.6 | 0.05 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 5.34 Hours | 2 GB | 2 | 0.2 | 0.01 | |

Figure 4-58 IVM GUI for partition ID

Figure 4-59 shows how to open the console of a virtual I/O client through a Virtual I/O Server telnet session.

```
login: padmin
padmin's Password:
Last unsuccessful login: Fri Sep 14 11:23:18 CDT 2012 on /dev/pts/3 from 172.16.254.18
Last login: Fri Sep 14 15:39:12 CDT 2012 on /dev/pts/0 from 172.16.254.18

$ mkvt -id 2

AIX Version 7
Copyright IBM Corporation, 1982, 2012.
Console login:
```

Figure 4-59 Console window through Virtual I/O Server CLI

To close a console: To forcefully close a console that is already open, use the `rmvt -id <partition ID>` command.

4.5.4 Install an operating system on a logical partition

The operating system installation process is similar to the process for stand-alone systems. Use the following main steps to install an operating system on an LPAR:

1. Log in to the IVM partition.
2. Open a virtual terminal for the LPAR to be installed with the `mkvt` command in the Virtual I/O Server CLI, or with the IVM GUI. You must specify the ID of the LPAR when you use the CLI, as shown in Example 4-3.

Example 4-3 Open a virtual terminal

```
$ mkvt -id 2

AIX Version 7
Copyright IBM Corporation, 1982, 2012.
Console login:
```

3. Start the LPAR in the *system management services (SMS)* mode. You can change the boot mode in the properties of the partition profile before you start it. Or, you can enter 1 on the virtual terminal at the very beginning of the boot process, as shown in Example 4-4.

Example 4-4 Boot display

```

IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM

```

```

1 = SMS Menu                    5 = Default Boot List
8 = Open Firmware Prompt       6 = Stored Boot List

```

```

Memory      Keyboard      Network      SCSI      Speaker

```

4. Select a boot device, such as a virtual optical device, or a network for the *Network Installation Management (NIM)* installation.
5. Boot the LPAR.
6. Select the system console and the language.
7. Select the disks to be installed.

The installation of the operating system starts. Proceed as directed by your operating system installation instructions.

4.6 Optical device sharing

You can dynamically add, move, or remove optical devices from or to any LPAR, regardless of whether the LPAR is running.

The following steps describe how to change the assignment of an optical device:

1. Under the Virtual Storage Management menu in the navigation area, click **View/Modify Virtual Storage** and select the **Optical Devices** tab in the work area.

2. Select the optical device that you want to modify, as shown in Figure 4-60.
3. Click **Modify partition assignment** in the task area.



Figure 4-60 Optical Devices selection

4. Select the name of the LPAR to which you want to assign the optical device, as shown in Figure 4-61. You can also remove the optical device from the current LPAR by selecting **None**.

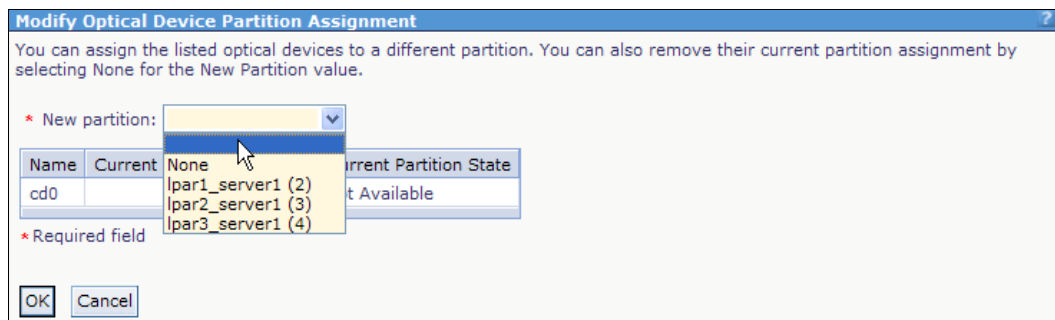


Figure 4-61 Optical Device Partition Assignment

5. Click **OK**.
6. If you move or remove an optical device from a running LPAR, you are prompted to confirm the forced removal before the device is removed. Because the device becomes unavailable, log in to the LPAR and remove the optical device before you go further. Press the Eject button. If the drawer opens, it is an indication that the device is not mounted.
7. Click **OK**.
8. The new list of optical devices is displayed with the changes that you made.

4.7 Logical partition configuration changes

As needed, you might want to modify the properties of the IVM or logical partitions (LPARs). With a dynamic logical partition (LPAR) operation, you can change the resources in run time.

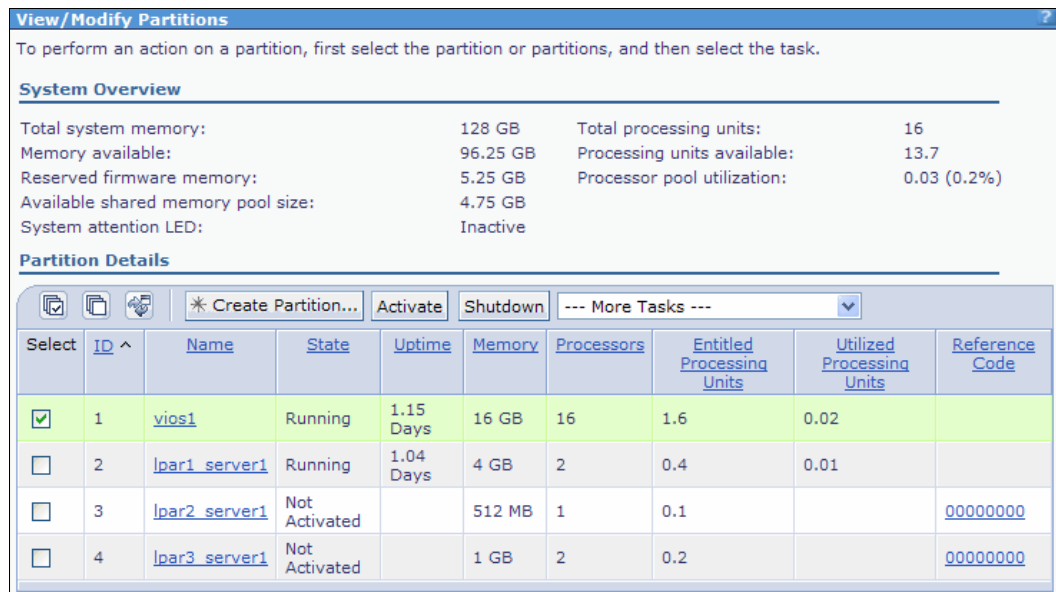
4.7.1 Dynamic LPAR operations on an IVM partition

Resources such as processors, memory, and physical/virtual adapters can be dynamically allocated or released on the IVM partition. You can run those operations either on the GUI or the CLI.

Dynamic LPAR on memory using the GUI

The following steps describe how to increase memory size dynamically for the IVM partition:

1. Click **View/Modify Partitions**.
2. Select the IVM partition, as shown in Figure 4-62.



The screenshot shows the 'View/Modify Partitions' window. At the top, it says 'To perform an action on a partition, first select the partition or partitions, and then select the task.' Below this is a 'System Overview' section with the following data:

| | | | |
|------------------------------------|----------|-----------------------------|-------------|
| Total system memory: | 128 GB | Total processing units: | 16 |
| Memory available: | 96.25 GB | Processing units available: | 13.7 |
| Reserved firmware memory: | 5.25 GB | Processor pool utilization: | 0.03 (0.2%) |
| Available shared memory pool size: | 4.75 GB | | |
| System attention LED: | Inactive | | |

Below the system overview is the 'Partition Details' section, which includes a toolbar with 'Create Partition...', 'Activate', and 'Shutdown' buttons, and a table of partitions:

| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
|-------------------------------------|------|-------------------------------|---------------|-----------|--------|------------|---------------------------|---------------------------|--------------------------|
| <input checked="" type="checkbox"/> | 1 | vios1 | Running | 1.15 Days | 16 GB | 16 | 1.6 | 0.02 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 1.04 Days | 4 GB | 2 | 0.4 | 0.01 | |
| <input type="checkbox"/> | 3 | lpar2_server1 | Not Activated | | 512 MB | 1 | 0.1 | | 00000000 |
| <input type="checkbox"/> | 4 | lpar3_server1 | Not Activated | | 1 GB | 2 | 0.2 | | 00000000 |

Figure 4-62 View/Modify Partitions: dynamic LPAR memory operation

3. Select **Properties** from the **More Tasks** drop-down menu (or use the one-click hyperlink method explained in 4.4.4, “Hyperlinks for object properties” on page 94).
4. Modify the pending values as needed. In Figure 4-63, the assigned memory is increased by 4 GB. Click **OK**.

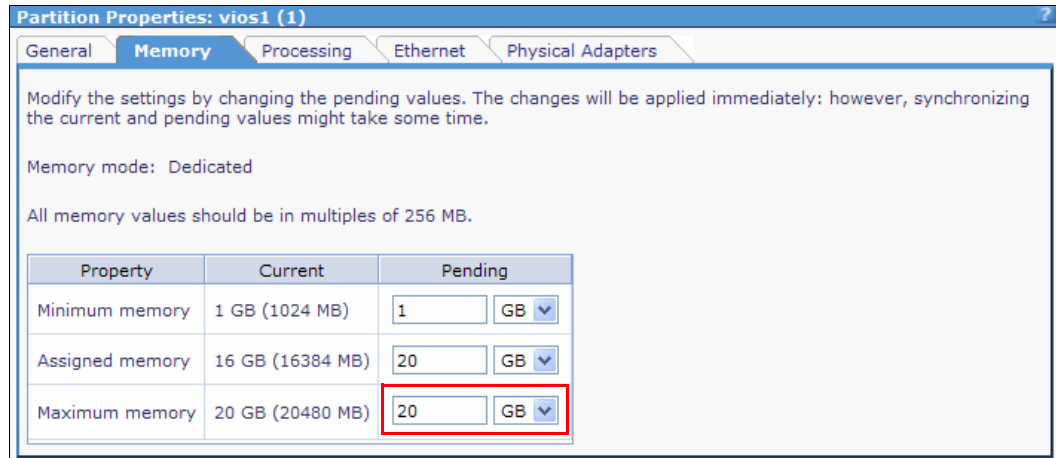


Figure 4-63 Partition Properties: Virtual I/O Server; Memory tab; increase memory size

Memory is not added or removed in a single operation. You can monitor the status by looking at the partition properties.

Dynamic LPAR operation on virtual processors using the command line

Log in to the IVM by using the CLI and run your dynamic LPAR operation. Example 4-5 shows how to add a 0.1 processing unit dynamically to the IVM by using the **chsyscfg** command.

Example 4-5 Dynamic LPAR virtual processor operation

```
$ lshwres -r proc --level lpar --filter "lpar_names=vios1" -F curr_proc_units
1.60
$ chsyscfg -r prof -i lpar_name=vios1,desired_proc_units=1.3,desired_procs=13
$ lshwres -r proc --level lpar --filter "lpar_names=vios1" -F curr_proc_units
1.30
```

Dynamic LPAR operation on physical adapters using graphical interface

In the partition properties of the IVM, click the **Physical Adapters** tab. Select or clear the adapters on which you want to perform a dynamic LPAR operation, and click **OK**.

Figure 4-64 shows the dynamic LPAR removal of two empty slots from the Virtual I/O Server.

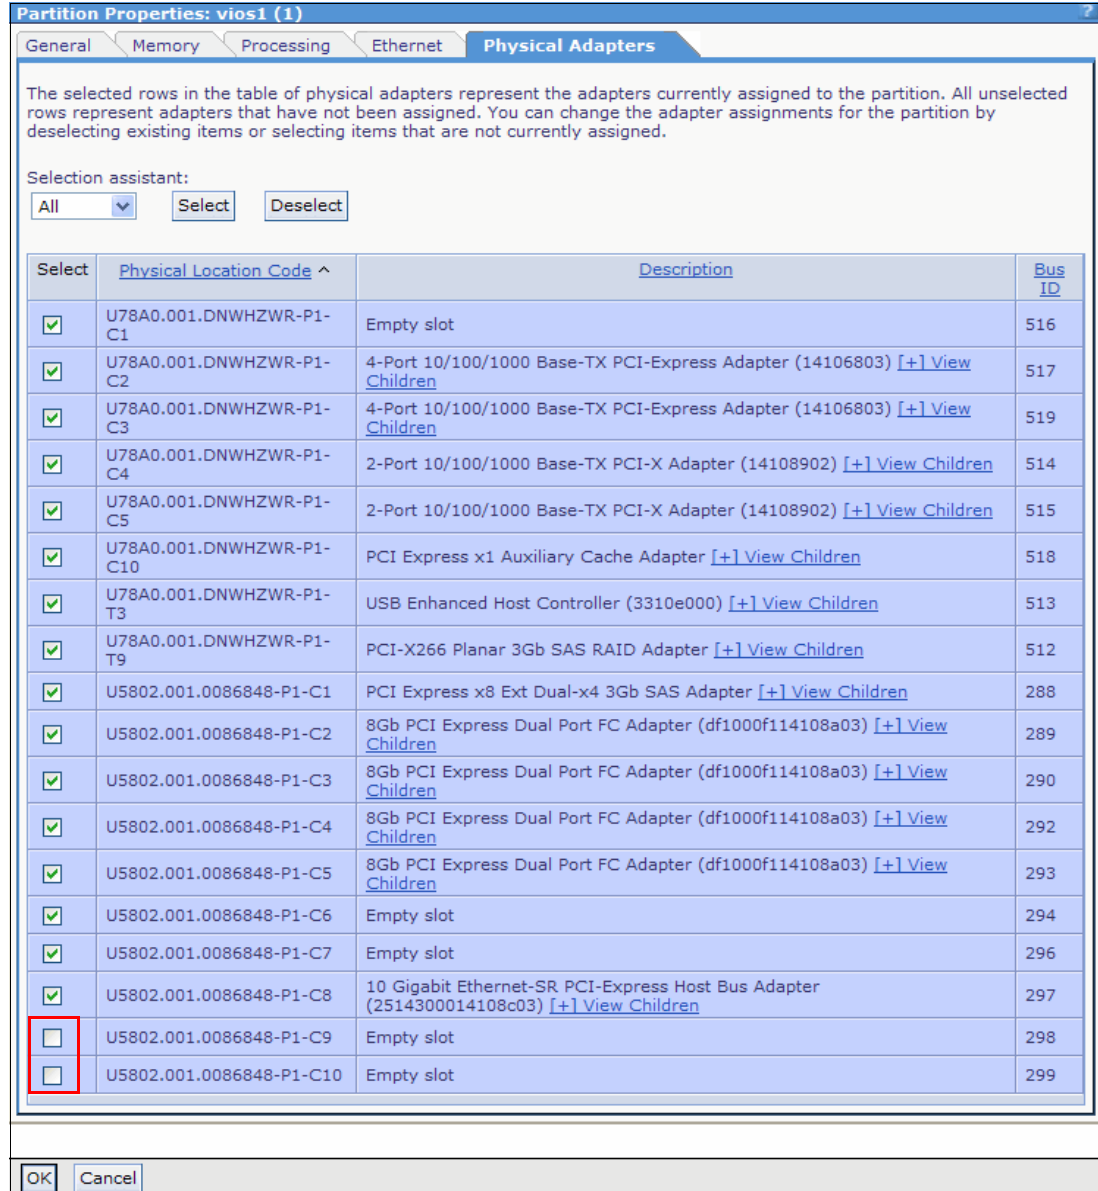


Figure 4-64 Dynamic LPAR of physical adapters

Dynamic LPAR operation on HEA adapters using graphical interface

In the partition properties of the IVM, click the **Ethernet** tab. Select or clear the HEA adapters on which you want to perform a dynamic LPAR, and click **OK**.

Figure 4-65 shows the dynamic LPAR removal of one HEA port.

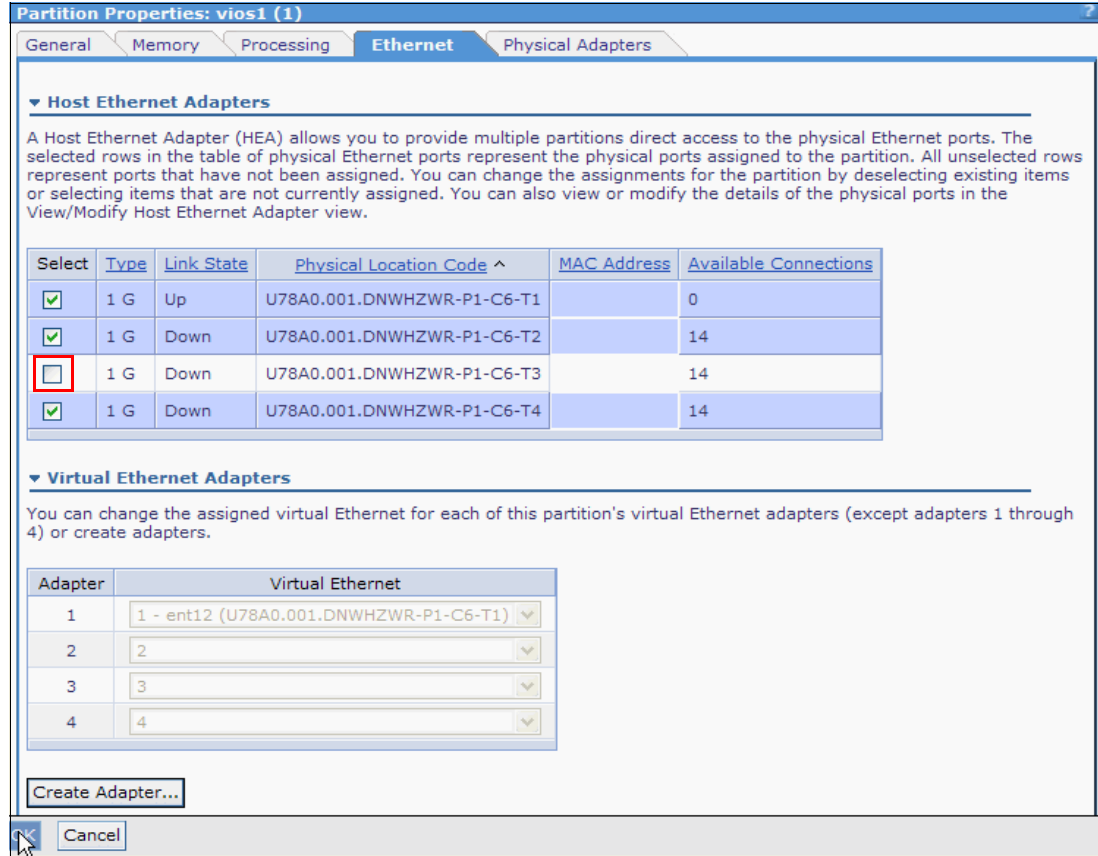


Figure 4-65 Dynamic LPAR of an HEA

Dynamic LPAR operation on virtual adapters using graphical interface

You are unable to remove virtual Ethernet adapters, but you can add them by using dynamic LPAR.

1. In the partition properties of the IVM, click the **Ethernet** tab. Click **Create Adapter**, as shown in Figure 4-66.

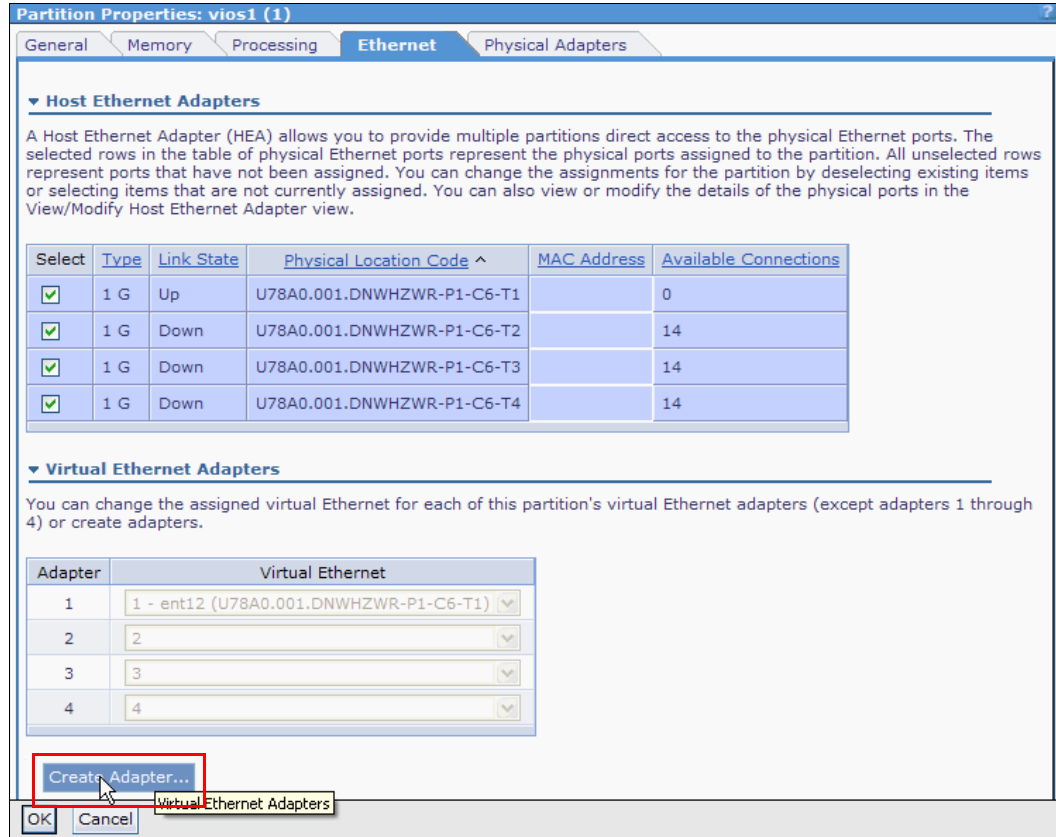


Figure 4-66 Dynamic LPAR for a virtual adapter

2. In the Enter Virtual Ethernet ID pop-up window, enter the **Virtual Ethernet ID** and click **OK**.

Figure 4-67 shows how to add a wanted VLAN at the time of dynamic LPAR.

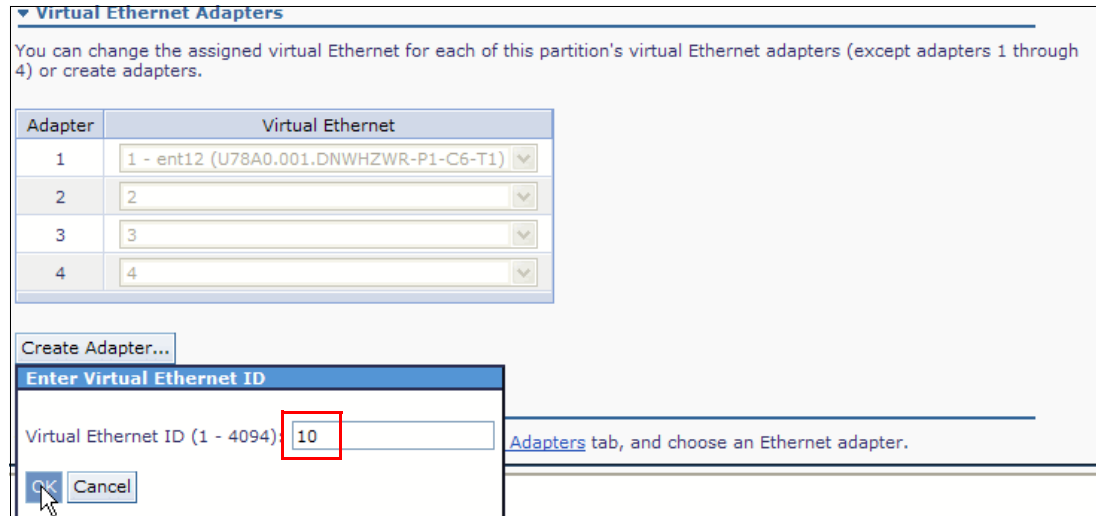


Figure 4-67 Dynamic LPAR for a virtual adapter by adding a VLAN

3. Click **OK** in the **View/Modify Partitions** wizard.

4.7.2 LPAR resources management

The Virtual I/O Server also allows dynamic operations on resources such as the processor, memory, and virtual/physical I/O on a client partition. To accomplish this goal, the concept of a dynamic LPAR manager is introduced. This concept is a daemon task that runs in the Virtual I/O Server LPAR and monitors the pending and runtime values for processing and memory values. It also drives the runtime and pending values to be in sync.

To perform a dynamic LPAR operation, the user must modify the pending value. The dynamic LPAR manager does the appropriate operations on the LPAR to complete the runtime change. If the runtime and pending values are in sync, then the dynamic LPAR manager blocks until another configuration change is made. It does not poll when in this state. The dynamic LPAR manager (*dlparmgr*) manages the Virtual I/O Server LPAR directly, and the client LPARs are managed through the RMC daemon. Dynamic operations on the disks, optical devices, partition name, and boot mode were already allowed.

Changes to pending value of LPAR: If a change is made to a pending value of an LPAR in a workload management group with another LPAR, the workload management software must be aware of this change and dynamically adapt to it. Otherwise, manual intervention is required. This process only applies to processors and memory.

When the *dlparmgr* encounters an error, it is written to the *dlparmgr* status log, which can be read with the `lssvcevents -t dlpar` command. This log contains the last *drmgr* command that is run for each object type, for each LPAR. It includes the responses from the *drmgr* command. The user is not notified directly of these errors; their indication is that the pending values are out of sync. The GUI enables you to see the state and gives you more information about the result of the operation. See Figure 4-72 on page 111.

The dynamic LPAR capabilities for each logical partition are returned as an attribute on the `lssyscfg -r lpar` command. This process allows the GUI to selectively enable/disable a dynamic LPAR based on the current capabilities of the logical partition.

Setup of dynamic LPAR

To enable dynamic LPAR operations, you must enable RMC communication between the IVM LPAR and the client LPARs:

1. Transmission Control Protocol/Internet Protocol (TCP/IP) must be configured in the IVM LPAR, or you cannot use the IVM browser interface.
2. For the IVM LPAR, select which physical interface will be used to provide the bridge between the internal VLAN and the real external network, in this case *VLAN-1* and *en0*. The client LPAR can have a TCP/IP configuration on that same subnet or on another one, depending on their external network or switch.
3. If multiple interfaces are configured in the Virtual I/O Server (and thus, multiple TCP/IP addresses), the LPARs must be on the same network. Therefore, they must be able to ping each other.
4. The client LPAR must have TCP/IP configured (on the same subnet that you selected for IVM-to-Client communication). You then must wait 2 or 3 minutes while the RMC subsystem completes the *handshake* between the client LPAR and the IVM LPAR. Viewing the client LPAR properties then shows the IP address, and the communication status equal to Active. Clicking **Retrieve Capabilities** verifies that dynamic LPAR operations are possible.

RMC and RSCT requirements: The following website contains the RMC and RequestSecurityTokenCollection (RSCT) requirements for dynamic LPAR, including the additional file sets that must be installed on Linux clients:

<https://www14.software.ibm.com/webapp/set2/sas/f/1opdiags/home.html>

Graphical user interface changes for dynamic LPAR

The GUI for performing dynamic LPAR operations is the same as the interface for performing static operations. Specifically, the user changes the pending value. Because the `chsyscfg` command can be used for dynamic LPAR, the GUI can use the same commands for both static and dynamic reconfiguration. The GUI disables changes for LPARs that do not support dynamic LPAR. For LPARs that do support certain dynamic LPAR operations (for example, processing, but not memory), the operation that is not supported is disabled appropriately. An inline message is then displayed indicating why they might not change the value.

The GUI also displays a *details* link next to the warning exclamation point when resources are out of sync. This view yields a pop-up window with the last run status of the `dlparmgr` for the specified LPAR and resource.

Partition properties changes for dynamic LPAR

A new section was added to the *General* panel of the properties sheet. This section indicates the RMC connection state, the IP address of the LPAR, and the dynamic LPAR capabilities of the logical (client) partition.

Because retrieving these dynamic LPAR capabilities can be a time-consuming operation (generally less than one second, but can be up to 90 seconds with a failed network connection), the initial capabilities are displayed as *Unknown* (if the Partition communication state is Active). Click **Retrieve Capabilities** to retrieve them and update the fields.

Table 4-1 provides more information about the different fields that are shown in Figure 4-68.

Table 4-1 Field name values

| Field name | Values |
|-----------------------------------|---|
| Partition host name or IP address | The IP address (or DNS host name) of the LPAR. This can be blank if RMC is not configured. |
| Partition communication state | Inactive, Active, Not Configured. Note: The name of this field matches the “Partition communication” field in the TCP/IP settings. |
| Memory dynamic LPAR capable | Yes, No, Unknown. Note: This will always default to “No” if you have not successfully been able to query the partition (RMC state is not active). No Linux LPARs are currently memory dynamic LPAR capable. <i>Unknown</i> will be the default state if communication is Active but the user has not selected Retrieve Capabilities. |
| Processing dynamic LPAR capable | Yes, No, Unknown. Note: If RMC is active, this will nearly always be Yes, but there is no guarantee. <i>Unknown</i> will be the default state if communication is Active but the user has not selected Retrieve Capabilities. |
| Retrieve Capabilities | Button that is visible if the Partition Communication State is active and the user has not previously selected the button on this properties sheet. |

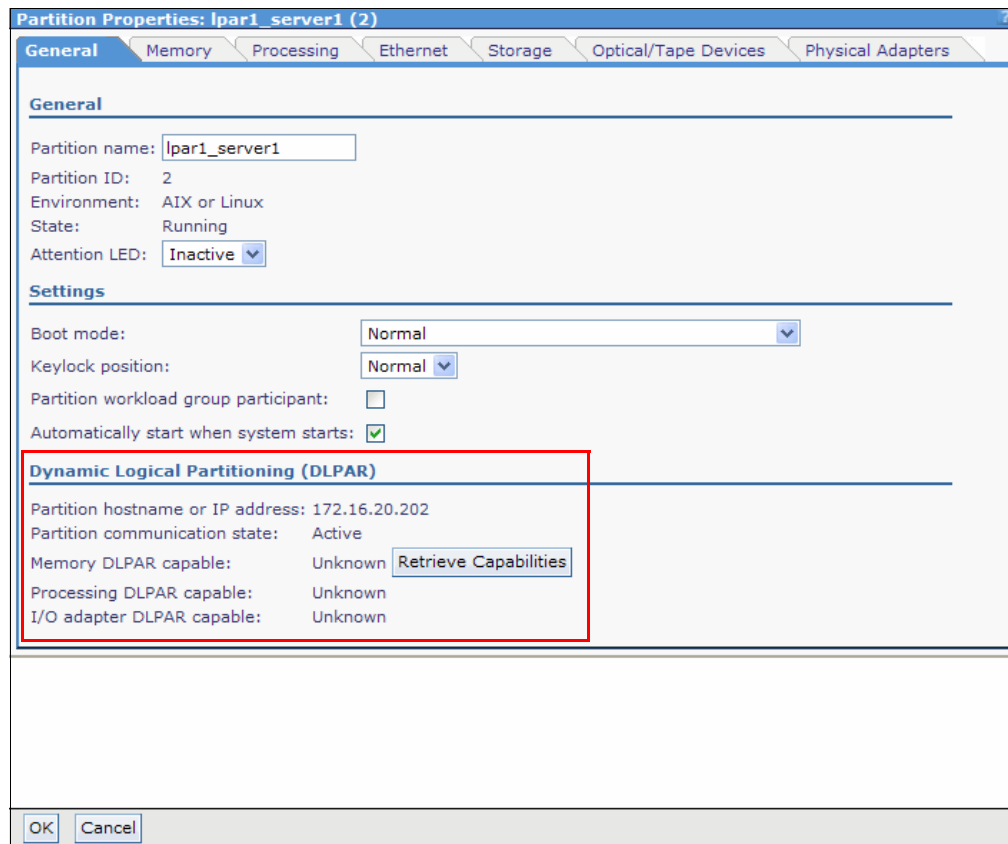


Figure 4-68 Dynamic LPAR properties

Memory tab

If the LPAR is powered on and the memory is dynamic-LPAR capable (see capabilities on the General tab), then the Pending assigned value is enabled. (Minimum memory and Maximum memory are still disabled). The user can change this value and select OK. The change takes effect immediately for the pending value. The `dlparmgr` daemon then works to bring the pending and current (runtime) values into sync. If these values are not in sync, the user sees the warning icon, as shown in Figure 4-71 on page 110. Figure 4-69 and Figure 4-70 show a change.

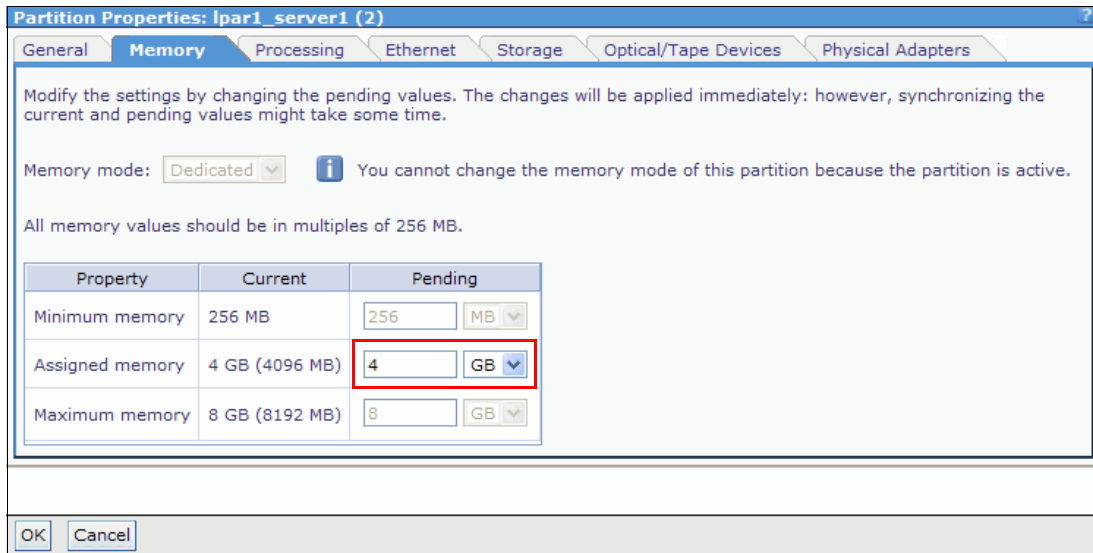


Figure 4-69 Partition Properties: Memory tab

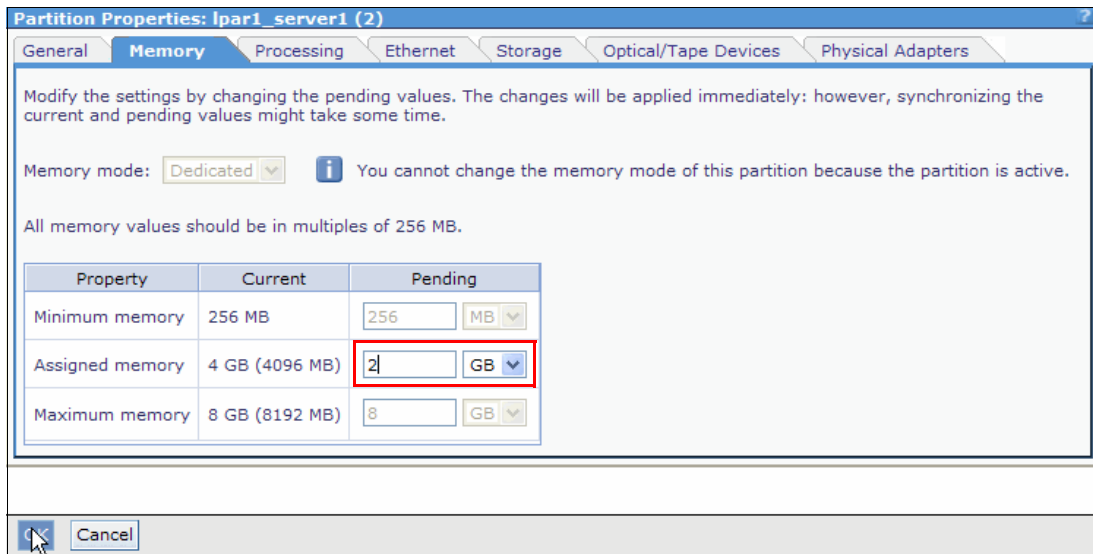


Figure 4-70 Dynamic LPAR of memory: removal of 2 GB of memory

Figure 4-71 depicts a warning in the work area because pending and current values are not in sync.

View/Modify Partitions

To perform an action on a partition, first select the partition or partitions, and then select the task.

System Overview

| | | | |
|------------------------------------|----------|-----------------------------|-------------|
| Total system memory: | 128 GB | Total processing units: | 16 |
| Memory available: | 101 GB | Processing units available: | 14 |
| Reserved firmware memory: | 1.75 GB | Processor pool utilization: | 0.06 (0.4%) |
| Available shared memory pool size: | 0 MB | | |
| System attention LED: | Inactive | | |

Partition Details


| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Reference Code |
|--------------------------|------|-------------------------------|---------------|-------------|--|------------|---------------------------|---------------------------|--------------------------|
| <input type="checkbox"/> | 1 | vios1 | Running | 17.23 Hours | 20 GB | 13 | 1.3 | 0.03 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 17.16 Hours | 2 GB  [details] | 2 | 0.4 | 0.01 | |
| <input type="checkbox"/> | 3 | lpar2_server1 | Open Firmware | 17.16 Hours | 512 MB | 1 | 0.1 | 0.01 | AA00E158 |
| <input type="checkbox"/> | 4 | lpar3_server1 | Open Firmware | 17.16 Hours | 1 GB | 2 | 0.2 | 0.01 | AA00E158 |

Figure 4-71 Warning in work area because pending and current values are not in sync

For more information about the resource synchronization, click the details hyperlink, as shown in Figure 4-72.

Resource Synchronization Details: lpar1_server1 (2)

Changing resource allocations while a partition is active may result in pending and current resource values not being synchronized. The most common reason for this is that certain resource changes may take some time to synchronize - particularly memory changes. Synchronizing these values requires that the partition communication state be active.

The resource types are listed below, along with their current state. If the resource is not synchronized, then details about the latest synchronization commands run will be displayed.

Memory

Resource synchronized: No
Reason: Synchronization in progress (code 1)
Latest commands run on partition:

| Reason | Time | Return Code | Command | Output |
|--------------------------------------|--------------------|-------------|---------------------------|--------|
| Synchronization in progress (code 1) | 9/19/12 9:45:00 AM | 0 | drmgr -c mem -q 1 -r -w 5 | |
| Synchronization successful (code 0) | 9/19/12 9:44:42 AM | 0 | | |

Processing Units

Resource synchronized: Yes

Processors

Resource synchronized: Yes

Uncapped Weight

Resource synchronized: Yes

Figure 4-72 Resource synchronization details

Memory values: The minimum and maximum memory values are enabled for the Virtual I/O Server/IVM LPAR at all times.

Table 4-2 provides possible memory field values.

Table 4-2 Possible field modifications: memory

| Capability setting | Enabled fields | Introduction text |
|--------------------|-----------------|---|
| Yes | Assigned Memory | Modify the settings by changing the pending values. Changes will be applied immediately, but synchronizing the current and pending values might take some time. |
| No | None | Modify the settings by changing the pending values. This LPAR does not currently support modifying these values while running, so pending values can be edited only when the LPAR is powered off. |
| Unknown | Assigned Memory | Modify the settings by changing the pending values. This LPAR does not currently support modifying these values while running, so pending values can be edited only when the LPAR is powered off. |

Processing tab

If the LPAR is powered on and is processor dynamic-LPAR capable (see capabilities on the General tab), then the Pending assigned values are enabled. The minimum and maximum (processing units and virtual processors) values are still disabled. The user can change these values and select OK. The change takes effect immediately for the pending value. The `dlparmgr` daemon then works to bring the pending and current (runtime) values into sync. If these values are not in sync, the user sees the warning icon as in the memory panel. As with the memory panel, the same rules apply regarding the enabled fields and introductory text for the various capability options.

Dynamic logical partition status

When a resource is not in sync, the warning icon is displayed on the View/Partition modify page with a *details* link. Selecting the details link yields the synchronization details pop-up window, as shown in Figure 4-72 on page 111.) In the first IVM release, when a dynamic resource is not in sync, a warning icon is displayed in the main partition list view next to the Memory or Processors resource. This icon also is displayed with text in the partition properties sheet.

A details link is now added next to the icon. This yields a pop-up window which shows the current dynamic LPAR status of the logical partition. All resources are shown in this window. If a resource is out of sync, the reason is provided. In addition, details about the previous two `drmgr` commands that were run against the LPAR in an attempt to synchronize the pending and current (runtime) values, are shown.

The Reason field generally matches the latest command run reason. But if the user modifies the minimum and maximum values without changing the assigned value, the `dlparmgr` considers the assigned values in sync, but the warning icon is still present.

Dynamic LPAR operation on virtual disks by using the GUI

The following steps describe how to assign *Virtual Disks* and *Physical Volumes* that are present in the Virtual I/O Server's Virtual Fibre Channel dynamically to a partition using the GUI:

1. To add or remove virtual disks or physical volumes, click the **Storage** tab of Partition Properties, **Select** the disks that you want to add and clear the disks that you want to remove and click **OK**, as shown in Figure 4-73.

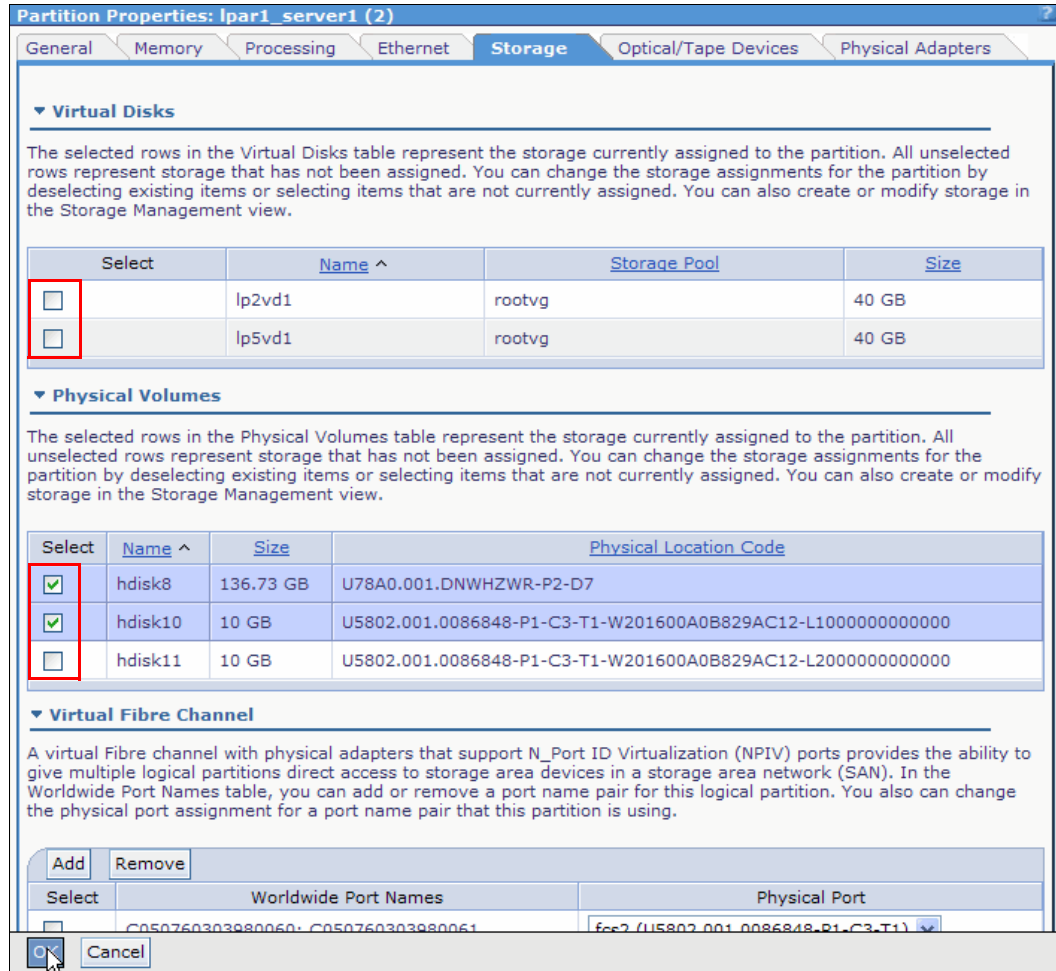


Figure 4-73 View/Modify Virtual Storage: Virtual Disks selection

2. Log in to the related LPAR and discover the new disk by running the `cfgmgr` command. Example 4-6 shows how the partition discovers one new virtual disk.

Example 4-6 Virtual disk discovery

```
# lsdev -Cc disk
hdisk0 Available          Virtual SCSI Disk Drive
hdisk1 Available C5-T1-01 MPIO DS4800 Disk

# cfgmgr

# lsdev -Cc disk
hdisk0 Available          Virtual SCSI Disk Drive
hdisk1 Available C5-T1-01 MPIO DS4800 Disk
hdisk2 Available          Virtual SCSI Disk Drive
```

Use the same procedure to add or remove virtual Fibre Channel adapters.

The same dynamic LPAR operation of disks can be done by using the View/Modify Virtual Storage menu in the navigation area. See more details in the section “Virtual disk/logical volume partition assignment” on page 58.

Operation on partition definitions by using the command-line interface

The command-line interface (CLI) for performing dynamic LPAR operations is the same as on the HMC. The dynamic LPAR Manager keys off differences between runtime and pending operations. The `chsyscfg` command is used for dynamic configuration changes because it updates the pending values.

You can perform the same operations with the CLI as with the GUI. Example 4-7 shows how to decrease processing units for an LPAR by using the `chsyscfg` command.

Example 4-7 Decreasing the processing units of lpar1_server1

```
$ lssyscfg -r prof --filter "lpar_names=lpar1_server1" -F desired_proc_units
0.40

$ chsyscfg -r prof -i "lpar_name=lpar1_server1,desired_proc_units=0.3"

$ lssyscfg -r prof --filter "lpar_names=lpar1_server1" -F desired_proc_units
0.30
```

A warning icon with an exclamation point is displayed in the View/Modify Partitions window if the current and pending values are not synchronized.

Example 4-8 shows an increase of memory operation.

Example 4-8 Increasing the memory of lpar1_server1 with 2048 MB

```
$ lssyscfg -r prof --filter "lpar_names=lpar1_server1" -F desired_mem
2048
$ chsyscfg -r prof -i "lpar_name=lpar1_server1,desired_mem+=2048"

$ lssyscfg -r prof --filter "lpar_names=lpar1_server1" -F desired_mem
4096
```

4.7.3 Adding a client LPAR to the partition workload group

If you want to manage logical partition resources by using a workload management tool, you must add the client LPAR to the partition workload group.

A *partition workload group* identifies a set of LPARs that are on the same physical system. Some workload management tools require that additional software is installed on the LPARs to monitor its workload, manage its resources, or both. Workload management tools use partition workload groups to identify which LPARs they can manage. For example, *IBM Enterprise Workload Manager™ (EWLM)* can dynamically and automatically redistribute processing capacity within a partition workload group to satisfy workload performance goals. EWLM adjusts processing capacity that is based on calculations that compare actual performance of work that is processed by the partition workload group to the business goals defined for the work.

Workload management tools use dynamic LPAR to make resource adjustments that are based on performance goals. Therefore, each LPAR in the partition workload group must support dynamic LPAR. Verify that the LPAR that you want to add to the partition workload group supports dynamic LPAR for the resource type that your workload management tool adjusts, as shown in Table 4-3.

Partition workload groups: Systems that are managed by the IVM can have only one partition workload group for each physical server.

It is not required that all LPARs on a system participate in a partition workload group. Workload management tools manage the resources of only those LPARs that are assigned to a partition workload group. Workload management tools can monitor the work of an LPAR that is not assigned to a partition workload group, but they cannot manage the resources of the LPAR.

Table 4-3 Dynamic LPAR support

| Logical partition type | Supports processor dynamic LPAR | Supports memory dynamic LPAR |
|------------------------|---------------------------------|---|
| AIX | Yes | Yes |
| Linux | Yes | Yes/no (SLES less than version 11 and RHEL less than 6.1 support memory add, but not memory removal). |
| IBM i | Yes | Yes |

For example, the partition management function of EWLM adjusts processor resources that are based on workload performance goals. Thus, EWLM can adjust the processing capacity for AIX and Linux LPARs.

The following recommendations apply workload management:

- ▶ Do not add the management partition to the partition workload group. To manage LPAR resources, workload management tools often require that you install some type of management or agent software on the LPARs. To avoid creating an unsupported environment, do not install more software on the management partition.
- ▶ The dynamic LPAR support that is listed in Table 4-3 is not the same as the dynamic LPAR capabilities that are in the partition properties for an LPAR. The dynamic LPAR support that is listed in Table 4-3 reflects what each operating system supports regarding dynamic

LPAR functions. The dynamic LPAR capabilities that are shown in the partition properties for an LPAR reflect the following combinations:

- An RMC connection between the management partition and the client LPAR
- The operating system's support of dynamic LPAR (see Table 4-3 on page 115)

For example, an AIX client LPAR does not have an RMC connection to the management partition, but AIX supports both processor and memory dynamic LPAR. In this situation, the dynamic LPAR capabilities that are shown in the partition properties for the AIX LPAR indicate that the AIX LPAR is not capable of processor or memory dynamic LPAR. However, because AIX supports processor and memory dynamic LPAR, a workload management tool can dynamically manage its processor and memory resources. Workload management tools are not dependent on RMC connections to dynamically manage LPAR resources.

- ▶ If an LPAR is part of the partition workload group, you cannot dynamically manage its resources from the IVM because the workload management tool is in control of dynamic resource management. Not all workload management tools dynamically manage both processor and memory resources. When you implement a workload management tool that manages only one resource type, you limit your ability to dynamically manage the other resource type. For example, EWLM dynamically manages processor resources, but not memory. AIX supports both processor and memory dynamic LPAR. EWLM controls dynamic resource management of both processor resources and memory for the AIX LPAR, but EWLM does not dynamically manage memory. Because EWLM has control of dynamic resource management, you cannot dynamically manage memory for the AIX LPAR from the IVM.

To add an LPAR to the partition workload group, complete the following steps:

1. Select the logical partition that you want to include in the partition workload group and click **Properties**. The Partition Properties window opens (Figure 4-74).
2. In the Settings section, select **Partition workload group participant**. Click **OK**.

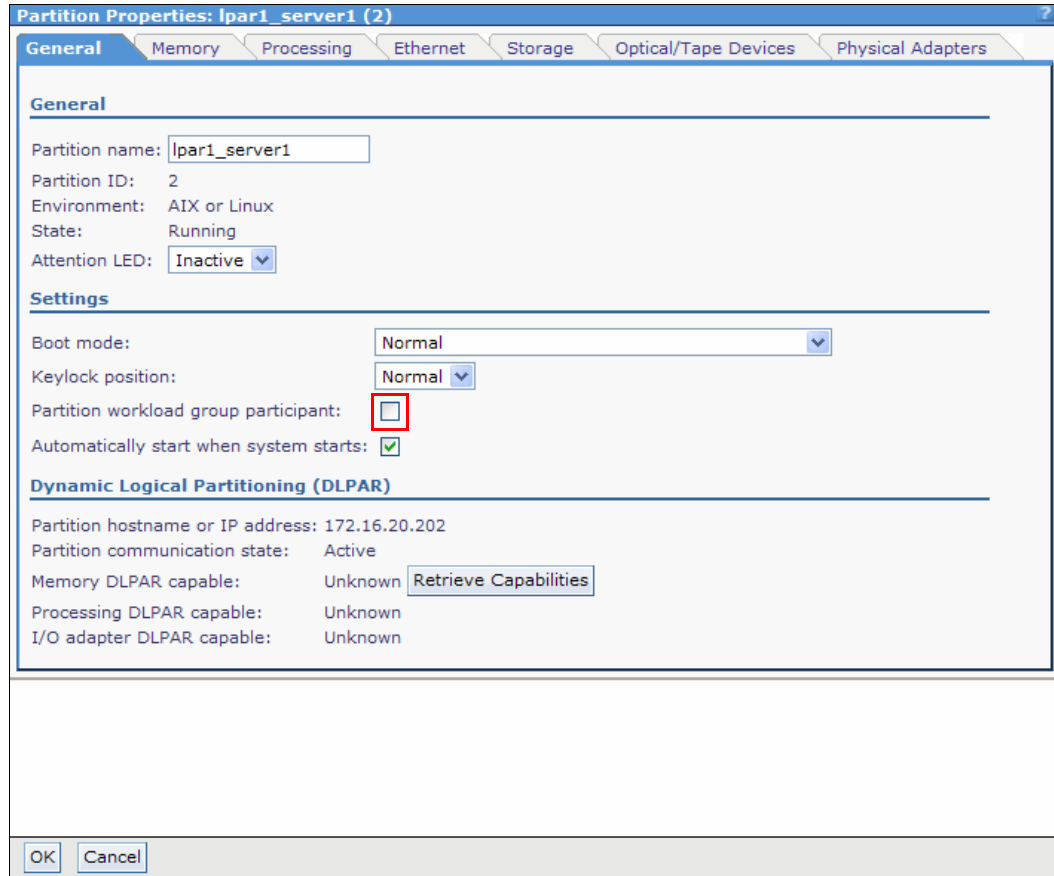


Figure 4-74 Partition Properties: selecting the partition workload group

Workload management: For more information about partition workload management, refer to the following resources:

For AIX 7.1, refer to this website:

<http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.baseadmndmndita%2Fwlm.htm>

For AIX 6.1, refer to this website:

http://pic.dhe.ibm.com/infocenter/aix/v6r1/index.jsp?topic=%2Fcom.ibm.aix.baseadmndmndita%2Fht_cnfgwlm.htm

4.8 Live Partition Mobility

Live Partition Mobility (LPM) is a feature which enables an active (or inactive) LPAR migration from one server to another. In this publication, we show how to do this migration through IVM between POWER6 and POWER7 processor-based servers.

Follow the requirements for LPM on an IVM environment:

- ▶ The source and destination system must be an IBM Power Systems POWER6 or POWER7 technology-based model.
- ▶ Both servers must have the PowerVM Enterprise feature activated.
- ▶ Both source and destination systems must be at firmware level 01Ex320 or later, where x is an S for BladeCenter, an L for entry servers (such as the Power 710, Power 720, Power 740, and Power 750).
- ▶ Virtual I/O Server must be installed at release level 1.5 or higher both on the source and destination systems.
- ▶ At the time of writing, the following operating systems are supported by LPM:
 - AIX 5.3 Level 5300-07; AIX 6.1, or later; and AIX 7.1, or later
 - Red Hat Enterprise Linux version 5 Update1, or later; and Red Hat Enterprise Linux version 6 Update1, or later
 - SUSE Linux Enterprise Server 10 Service Pack 1, or later; and SUSE Linux Enterprise Server 11 Service Pack 1
 - IBM i 7.1 TR4 PTF group - SF9907 level 4, or later

Although we suggest that you always use operating systems in the latest version and service pack, there are other considerations:

- ▶ All of the resources from the LPAR that you want to move must be virtualized and must use external storage, including the boot disks.
- ▶ Although there is a minimum required firmware level that we suggest, have the most current system firmware installed that is available.

To know the procedure for IVM updates, refer to 5.1.5, “IVM updates” on page 162.

If you are migrating between two different POWER processor technologies (as in our example) and you want to migrate from a POWER7 to a POWER6 processor-based server, ensure that the partition that you want to migrate is in the **POWER6** compatibility mode. The processor compatibility mode is shown in Figure 4-75.

This mode can be changed after the partition creation at the Processing tab of the Partition Properties window.

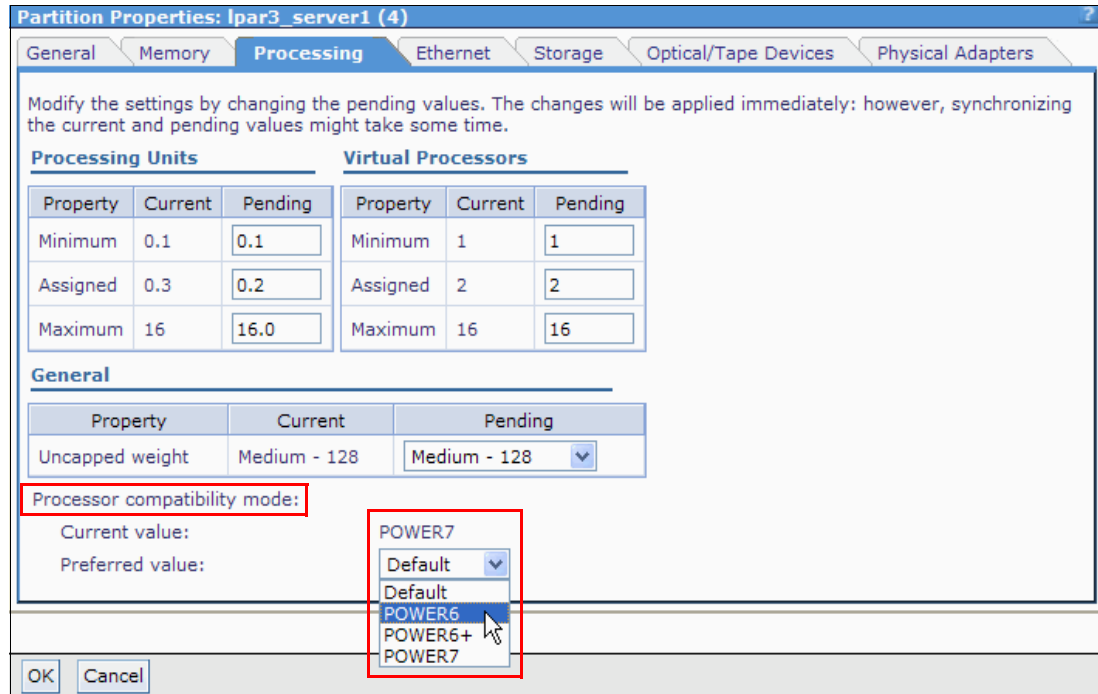


Figure 4-75 Partition Properties: Setting the POWER6 compatibility mode on the Processing tab

Follow these steps to perform the LPM process in an IVM environment:

1. **Select** the partition that you want to migrate on the **View/Modify Partitions** panel, as shown in Figure 4-76.




Figure 4-76 Selecting the LPAR on the View/Modify Partitions panel

2. On the **More Tasks** drop-down menu, select **Migrate** (Figure 4-77).

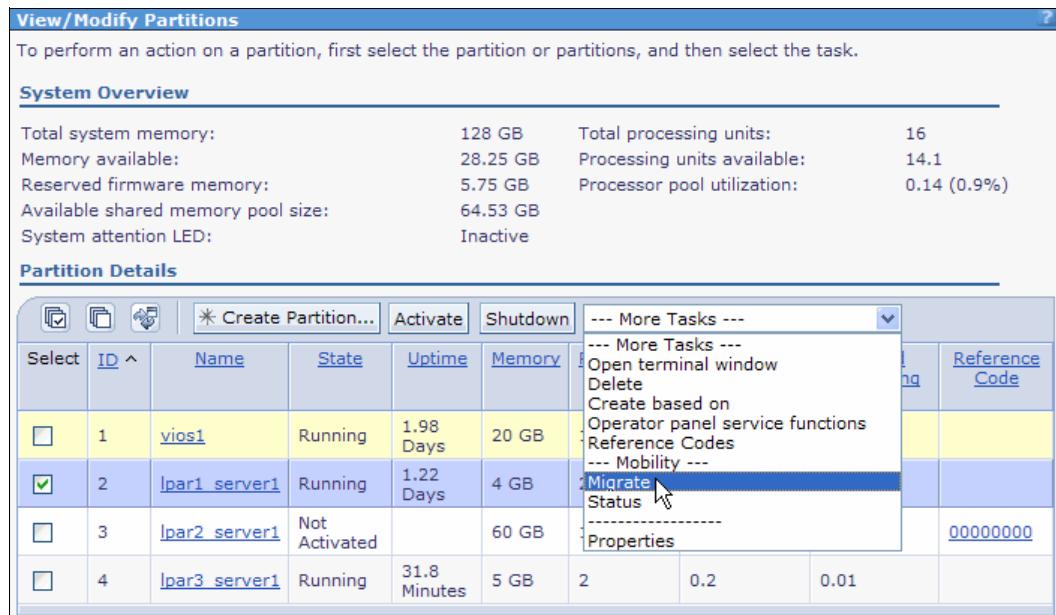


Figure 4-77 Migrate option for selected LPAR on More Tasks in the View/Modify Partitions panel

3. In the next window, provide the **Remote IVM** IP address or the host name from the server to where you want to migrate your LPAR. Also, enter the padmin user's **Password**. Then, click **Validate**, as shown in Figure 4-78.

Migrate Partition: lpar1_server1 (2)

It might be possible to migrate this partition to run on another managed system. In order to migrate this partition, it must meet certain conditions. For details, consult your documentation. Specify the hostname or IP address of the remote Integrated Virtualization Manager (IVM) that controls the target managed system, and select Validate or Migrate.

* Remote IVM:

Remote user ID:

Password:


* Required field

Figure 4-78 Migrate Partition panel with information fields

4. If there is an error which is not allowing the LPM process to occur, a message is shown in the Migrate Partition window. We assigned a physical adapter to the client partition to simulate an error, as shown in Figure 4-79.

Migrate Partition: lpar1_server1 (2)

Problems occurred while processing the data. A summary of all problems for this page are listed below. Additional details for each problem may be located next to the field causing the problem.

 [VIOSE01042028-0415] The partition cannot be migrated because it has physical I/O assignments.

It might be possible to migrate this partition to run on another managed system. In order to migrate this partition, it must meet certain conditions. For details, consult your documentation. Specify the hostname or IP address of the remote Integrated Virtualization Manager (IVM) that controls the target managed system, and select Validate or Migrate.

* Remote IVM:

Remote user ID:

Password:

* Required field

Figure 4-79 Migrate Partition panel with an error message

- After the physical adapter was unassigned for that partition, the operation was successfully completed, as shown in Figure 4-80. Now select **Migrate**.

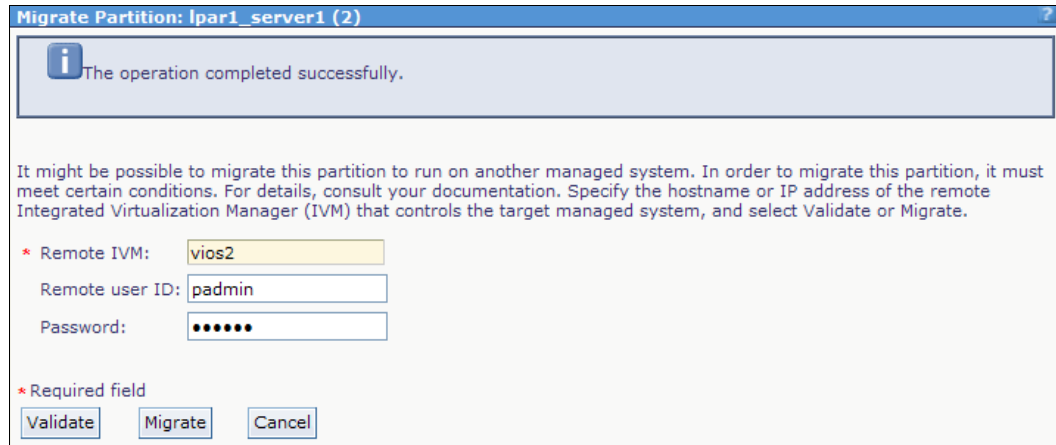


Figure 4-80 Migrate Partition window after validation

- While the task is running, it is possible to see the migration status. **Select** the LPAR which is migrating on the **View/Modify Partitions** window. Then, click the **More Tasks** drop-down menu and select **Status**, as shown in Figure 4-81.

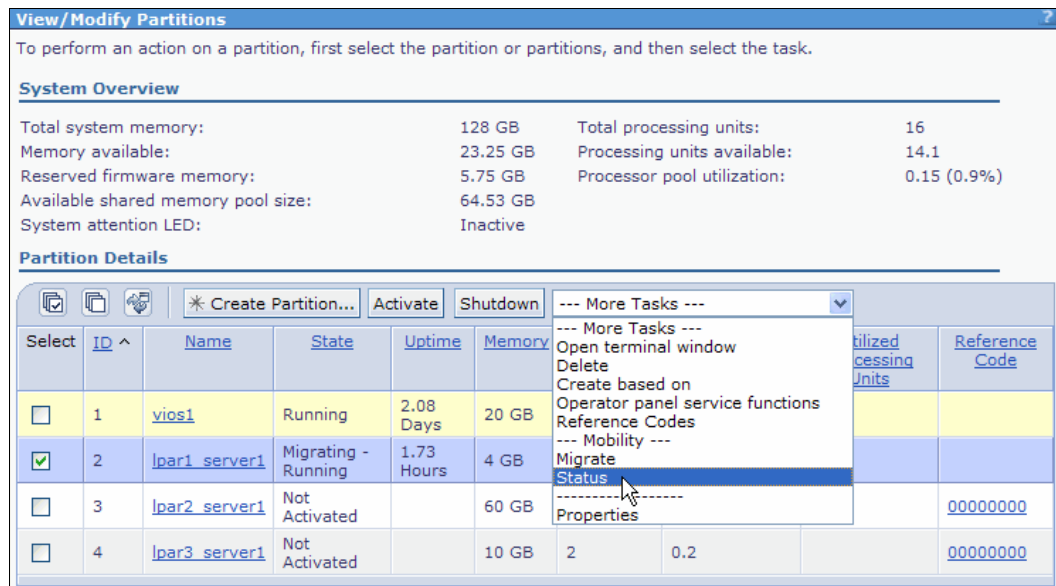


Figure 4-81 View/Modify Partitions panel: select Status for the LPAR in the migration process

- ▶ According to our example, **lpar1_server1** is no more in the POWER7 processor-based server (Serial Number 061AA6P). Now it is running in the POWER6 processor-based server (Serial Number 1059C72), as shown in the AIX 7.1 command line on Example 4-9, and on Example 4-10.

Example 4-9 AIX (lpar1) serial number before LPM

```
# hostname
lpar1
# lsconf | grep -i Serial
Machine Serial Number: 061AA6P
* vsa0                U8233.E8B.061AA6P-V2-C0
LPAR Virtual Serial Adapter
```

Example 4-10 AIX (lpar1) serial number after LPM

```
# hostname
lpar1
# lsconf | grep -i Serial
Machine Serial Number: 1059C72
* vsa0                U9409.M50.1059C72-V3-C0
LPAR Virtual Serial Adapter
```

LPAR migration: The migrated LPAR keeps the same IP address when moved and does not cause any service interruption into the system. There is no need to migrate the LPAR after the process. LPM can also be done when the LPAR is inactive.

4.9 Network management

When IVM is installed, all physical Ethernet adapters that are installed in the system are managed by it. Logical partitions can have at most two virtual Ethernet adapters, each connected to one of the four virtual networks that are present in the system.

You can also configure physical Ethernet adapters for the client LPARs. The way to accomplish this is to deallocate the selected Ethernet adapter from the Virtual I/O Server partition, as shown in Figure 4-82 on page 124, and allocate it in the target LPAR. This process can be done during the LPAR creation or it can be changed later in the client LPAR profile, as shown in Figure 4-83 on page 124. There is no need to reboot the partition because of this change.

To allow partitions to access any external corporate network, every virtual network can be bridged to a physical adapter. For each network, a separate adapter is required. IVM provides a web interface to configure bridging.

When higher throughput and better link availability is required, Ethernet link aggregation is also available by using the Virtual I/O Server capabilities.

Figure 4-82 depicts how to deallocate the selected Ethernet adapter from the Virtual I/O Server partition.



Figure 4-82 Virtual I/O Server Partition Properties: Physical Adapters tab

Figure 4-83 shows how to allocate an adapter in the Physical Adapters tab of the Partition Properties panel.

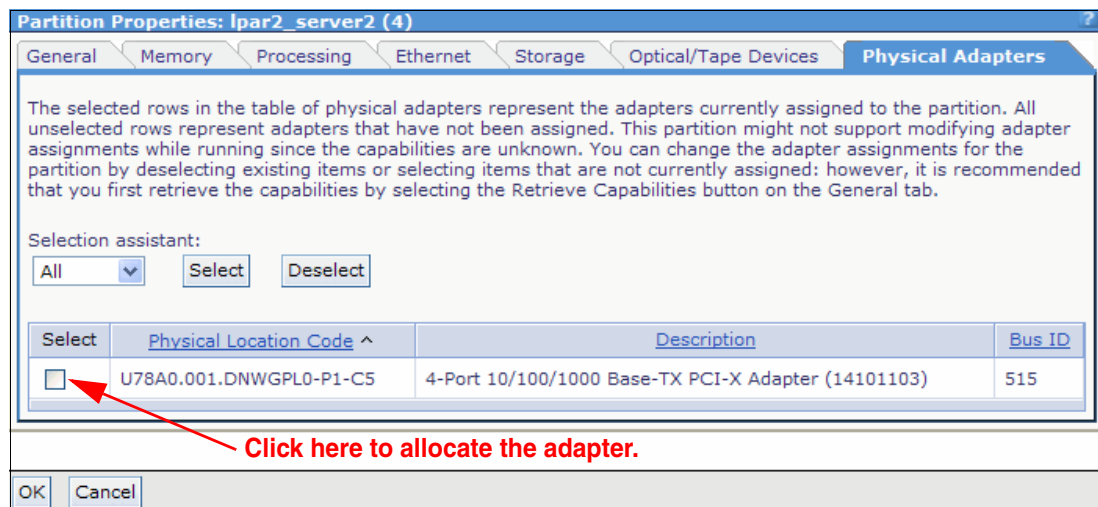


Figure 4-83 Client Partitions Properties: Physical Adapters tab

4.9.1 Ethernet bridging

Under **I/O Adapter Management** in the navigation area, which is shown in Figure 4-84, click **View/Modify Virtual Ethernet**. In the work area, the Virtual Ethernet panel shows what partitions are connected to the four available networks. Go to the **Virtual Ethernet Bridge** tab to configure bridging, as shown in Figure 4-84. For each virtual Ethernet, you can select one physical device. Use the drop-down menu to select the **Physical Adapter** and click **Apply** to create the bridging device.

IP address of a physical adapter: If the physical adapter that is selected for bridging is already configured with an IP address by using the CLI, all connections to that address are reset.

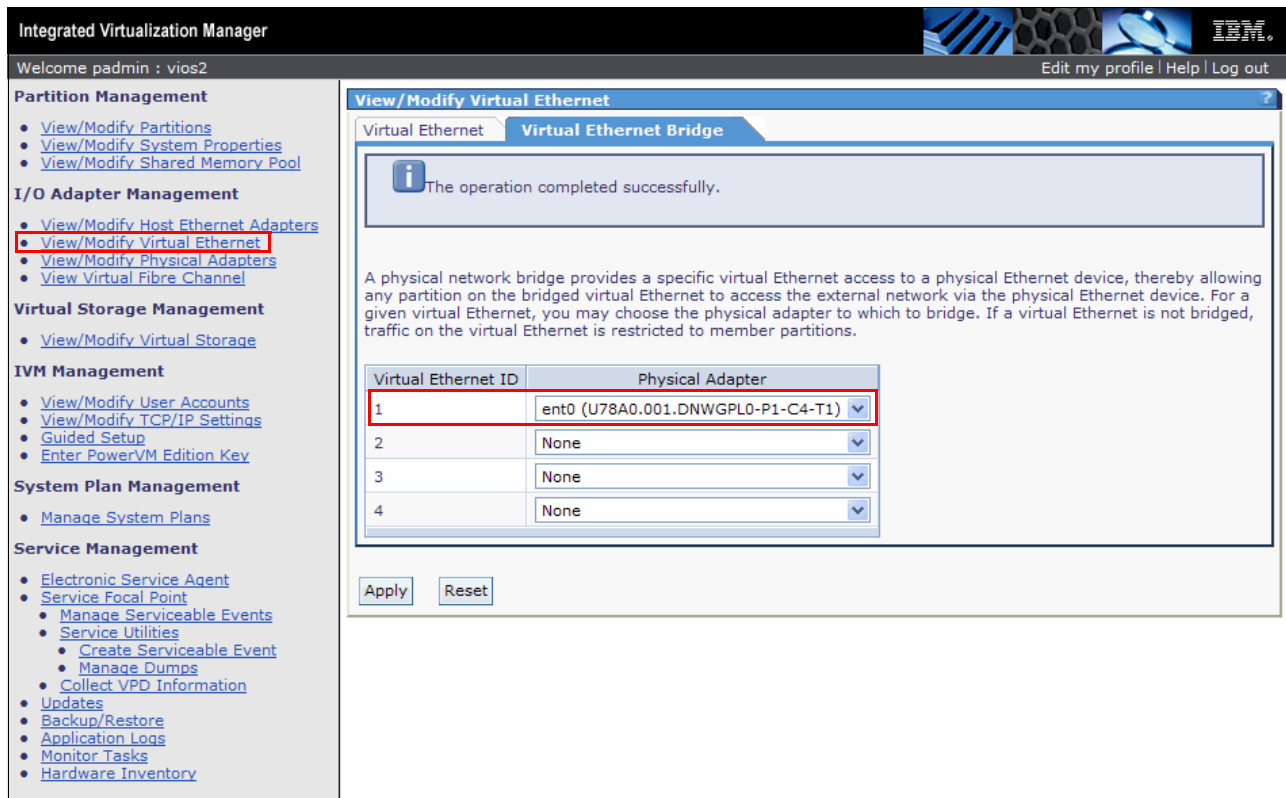


Figure 4-84 View/Modify Virtual Ethernet: Virtual Ethernet Bridge creation

The web graphical interface hides the details of the network configuration. Example 4-11 describes the Virtual I/O Server configuration before the creation of the bridge. For each physical and virtual network adapter, an Ethernet device is configured. The IVM is connected to a physical network and four virtual network adapters are available.

Example 4-11 Virtual I/O Server Ethernet adapters with no bridging

```
$ lsdev | grep ^en
en0          Available  Standard Ethernet Network Interface
en1          Defined   Standard Ethernet Network Interface
en2          Defined   Standard Ethernet Network Interface
en3          Defined   Standard Ethernet Network Interface
en4          Defined   Standard Ethernet Network Interface
en5          Defined   Standard Ethernet Network Interface
```

```

en6          Defined      Standard Ethernet Network Interface
en7          Defined      Standard Ethernet Network Interface
en8          Defined      Standard Ethernet Network Interface
en9          Defined      Standard Ethernet Network Interface
en10         Defined      Standard Ethernet Network Interface
en11         Defined      Standard Ethernet Network Interface
en12         Defined      Standard Ethernet Network Interface
en13         Defined      Standard Ethernet Network Interface
ent0         Available    2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent1         Available    2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent2         Available    4-Port 10/100/1000 Base-TX PCI-X Adapter (14101103)
ent3         Available    4-Port 10/100/1000 Base-TX PCI-X Adapter (14101103)
ent4         Available    4-Port 10/100/1000 Base-TX PCI-X Adapter (14101103)
ent5         Available    4-Port 10/100/1000 Base-TX PCI-X Adapter (14101103)
ent6         Available    Logical Host Ethernet Port (1p-hea)
ent7         Available    Logical Host Ethernet Port (1p-hea)
ent8         Available    Logical Host Ethernet Port (1p-hea)
ent9         Available    Logical Host Ethernet Port (1p-hea)
ent10        Available    Virtual I/O Ethernet Adapter (1-lan)
ent11        Available    Virtual I/O Ethernet Adapter (1-lan)
ent12        Available    Virtual I/O Ethernet Adapter (1-lan)
ent13        Available    Virtual I/O Ethernet Adapter (1-lan)

```

```
$ !stcpip
```

| Name | Mtu | Network | Address | Ipkts | Ierrs | Opkts | Oerrs | Coll |
|------|-------|-----------|-----------------|----------|-------|---------|-------|------|
| en0 | 1500 | link#2 | 0.9.6b.6e.12.4e | 11287173 | 0 | 3147434 | 3 | 0 |
| en0 | 1500 | 172.16.20 | vios2 | 11287173 | 0 | 3147434 | 3 | 0 |
| lo0 | 16896 | link#1 | | 741903 | 0 | 741903 | 0 | 0 |
| lo0 | 16896 | 127 | loopback | 741903 | 0 | 741903 | 0 | 0 |
| lo0 | 16896 | ::1%1 | | 741903 | 0 | 741903 | 0 | 0 |

When a virtual Ethernet bridge is created, a new *Shared Ethernet adapter* (SEA) is defined, binding the physical device with the virtual device. If a network interface was configured on the physical adapter, the IP address is migrated to the new SEA.

Example 4-12 shows the result of bridging *virtual network 1* with the physical adapter *ent0* when the IVM is using the network interface *en0*. A new *ent2* SEA device is created, and the IP address of the IVM is migrated on the *en2* interface. Because of the migration, all active network connections on *en0* are reset.

Example 4-12 Shared Ethernet adapter configuration

```

$ !sdev | grep ^en
en0          Available    Standard Ethernet Network Interface
en1          Available    Standard Ethernet Network Interface
en2          Available    Standard Ethernet Network Interface
en3          Defined      Standard Ethernet Network Interface
en4          Defined      Standard Ethernet Network Interface
en5          Defined      Standard Ethernet Network Interface
en6          Defined      Standard Ethernet Network Interface
en7          Defined      Standard Ethernet Network Interface
en8          Defined      Standard Ethernet Network Interface
en9          Defined      Standard Ethernet Network Interface
en10         Defined      Standard Ethernet Network Interface
en11         Defined      Standard Ethernet Network Interface

```

```

en12          Defined      Standard Ethernet Network Interface
en13          Defined      Standard Ethernet Network Interface
ent0          Available   2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent1          Available   2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent2         Available Shared Ethernet Adapter
ent6          Available   Logical Host Ethernet Port (lp-hea)
ent7          Available   Logical Host Ethernet Port (lp-hea)
ent8          Available   Logical Host Ethernet Port (lp-hea)
ent9          Available   Logical Host Ethernet Port (lp-hea)
ent10         Available   Virtual I/O Ethernet Adapter (l-lan)
ent11         Available   Virtual I/O Ethernet Adapter (l-lan)
ent12         Available   Virtual I/O Ethernet Adapter (l-lan)
ent13         Available   Virtual I/O Ethernet Adapter (l-lan)
$ lstcpip

```

| Name | Mtu | Network | Address | Ipkts | Ierrs | Opkts | Oerrs | Coll |
|------|-------|-----------|-----------------|--------|-------|--------|-------|------|
| en2 | 1500 | link#3 | 0.9.6b.6e.12.4e | 9205 | 0 | 227 | 0 | 0 |
| en2 | 1500 | 172.16.20 | vios2 | 9205 | 0 | 227 | 0 | 0 |
| lo0 | 16896 | link#1 | | 742305 | 0 | 742305 | 0 | 0 |
| lo0 | 16896 | 127 | loopback | 742305 | 0 | 742305 | 0 | 0 |
| lo0 | 16896 | ::1% | | 742305 | 0 | 742305 | 0 | 0 |

4.9.2 Ethernet link aggregation

Link aggregation is a network technology that enables several Ethernet adapters to be joined to form a single virtual Ethernet device. This solution can be used to overcome the bandwidth limitation of a single network adapter and to avoid bottlenecks when sharing one network adapter among many client partitions.

The aggregated device also provides high-availability capabilities. If a physical adapter fails, the packets are automatically sent on the other available adapters without disruption to existing user connections. The adapter is automatically returned to service on the link aggregation when it recovers.

Link aggregation is an expert-level configuration and is not managed by the IVM GUI. The configuration is defined by using the Virtual I/O Server functions with the command line. However, the IVM can use the link aggregation for network configuration after it is defined.

To create the link aggregation, use the `mkvdev` command with the following syntax:

```
mkvdev -lnagg TargetAdapter ... [-attr Attribute=Value ...]
```

In the environment that is shown in Example 4-13, it is possible to aggregate the two physical Ethernet adapters `ent6` and `ent7`. A new virtual adapter `ent3` is created, as described in Example 4-13.

Example 4-13 Ethernet aggregation creation

```

$ lsdev | grep ent
ent0          Available   2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent1          Available   2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent2          Available   Shared Ethernet Adapter
ent6          Available   Logical Host Ethernet Port (lp-hea)
ent7          Available   Logical Host Ethernet Port (lp-hea)
ent8          Available   Logical Host Ethernet Port (lp-hea)
ent9          Available   Logical Host Ethernet Port (lp-hea)

```

```
ent10      Available  Virtual I/O Ethernet Adapter (1-1an)
ent11      Available  Virtual I/O Ethernet Adapter (1-1an)
ent12      Available  Virtual I/O Ethernet Adapter (1-1an)
ent13      Available  Virtual I/O Ethernet Adapter (1-1an)
ibmvmc0    Available  Virtual Management Channel
```

```
$ mkvdev -lnagg ent6 ent7
ent3 Available
```

```
$ lsdev -dev ent3
name      status      description
ent3      Available  EtherChannel / IEEE 802.3ad Link Aggregation
```

```
$ lsdev -dev en3
name      status      description
en3       Defined    Standard Ethernet Network Interface
```

Aggregated devices can be used to define an SEA. The SEA must be created by using the **mkvdev** command with the following syntax:

```
mkvdev -sea TargetDevice -vadapter VirtualEthernetAdapter ...
       -default DefaultVirtualEthernetAdapter
       -defaultid SEADefaultPVID [-attr Attributes=Value ...]
       [-migrate [-auto]]
```

Example 4-14 shows the bridging of *virtual network 4* with *SEA ent3*. The **mkvdev** command requires the identification of the virtual Ethernet adapter that is connected to virtual network 4.

The **lssyscfg** command with the parameter **lpar_names** that is set to the Virtual I/O Server partition's name provides the list of virtual adapters that are defined for the Virtual I/O Server. The adapters are separated by commas, and their parameters are separated by slashes. The third parameter is the network number (4 in the example) and the first is the slot identifier (6 in the example).

The **lsdev** command with the **-vpd** flag provides the physical location of a virtual Ethernet adapter that contains the letter C followed by its slot number. In the example, ent13 is the virtual Ethernet adapter that is connected to network 4.

The created *ent10* adapter is the new SEA. Refer to Example 4-14.

Example 4-14 Creating an SEA adapter with link aggregation

```
$ lsdev -virtual | grep 1-1an
ent10      Available  Virtual I/O Ethernet Adapter (1-1an)
ent11      Available  Virtual I/O Ethernet Adapter (1-1an)
ent12      Available  Virtual I/O Ethernet Adapter (1-1an)
ent13      Available  Virtual I/O Ethernet Adapter (1-1an)

$ lssyscfg -r prof --filter lpar_names=10-59C72 -F virtual_eth_adapters
"3/0/1//1/0,4/0/2//1/0,5/0/3//1/0,6/0/4//1/0"

$ lsdev -dev ent10 -vpd | grep Location
      Hardware Location Code.....U9409.M50.1059C72-V1-C3-T1
      Physical Location: U9409.M50.1059C72-V1-C3-T1
$ lsdev -dev ent11 -vpd | grep Location
      Hardware Location Code.....U9409.M50.1059C72-V1-C4-T1
```



```

Physical Location: U9409.M50.1059C72-V1-C4-T1
$ lsdev -dev ent12 -vpd | grep Location
  Hardware Location Code.....U9409.M50.1059C72-V1-C5-T1
  Physical Location: U9409.M50.1059C72-V1-C5-T1
$ lsdev -dev ent13 -vpd | grep Location
  Hardware Location Code.....U9409.M50.1059C72-V1-C6-T1
  Physical Location: U9409.M50.1059C72-V1-C6-T1

$ mkvdev -sea ent3 -vadapter ent13 -default ent13 -defaultid 4
ent4 Available

```

After the SEA is created by using the command line, it is available from the IVM panels. It is displayed as a device with no location codes inside the parenthesis because it uses a virtual device.

Figure 4-85 shows how IVM represents an SEA that was created by using an Ethernet link aggregation.

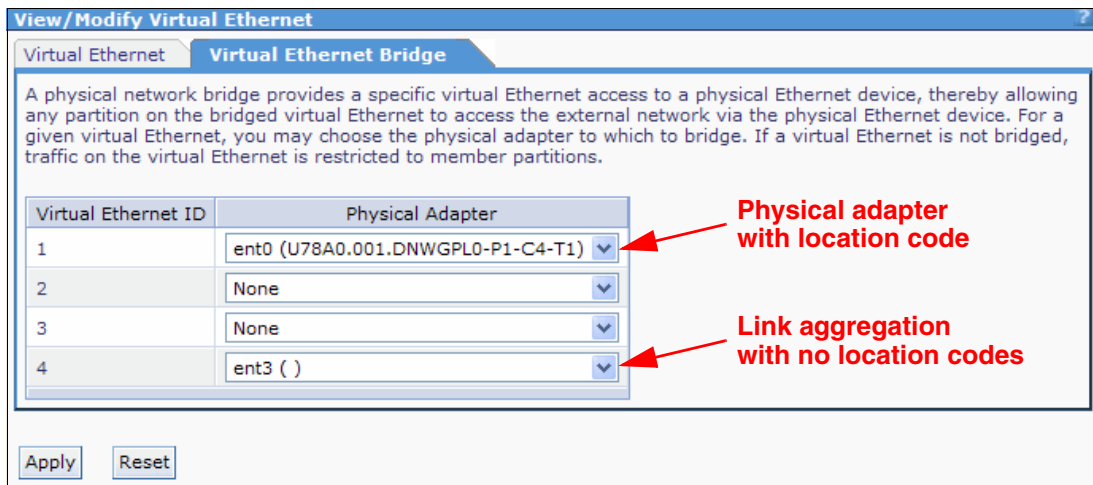


Figure 4-85 Virtual Ethernet bridge with a link aggregation device

As shown in Figure 4-85, the SEA can be removed by using the IVM and selecting **None** as the Physical Adapter for the virtual network. When you click **Apply**, the IVM removes all devices that are related to the SEA, but the link aggregation remains active.

4.9.3 Host Ethernet Adapter in IVM

You can use a *Host Ethernet Adapter (HEA)* as a network interface card for both the IVM partition and its clients.

To see whether an HEA is connected, which model it is (a two- or four-port HEA), and to see the status of all the ports in the IVM, complete the following process.

In the navigation panel, under the **I/O Adapter Management** section, **View/Modify Host Ethernet Adapters** provides the details of the HEA on that IVM, as shown in Figure 4-86.

The screenshot shows the IVM interface with the following navigation menu items:

- Partition Management**
 - [View/Modify Partitions](#)
 - [View/Modify System Properties](#)
 - [View/Modify Shared Memory Pool](#)
- I/O Adapter Management**
 - View/Modify Host Ethernet Adapters** (highlighted)
 - [View/Modify Virtual Ethernet](#)
 - [View/Modify Physical Adapters](#)
 - [View Virtual Fibre Channel](#)
- Virtual Storage Management**
 - [View/Modify Virtual Storage](#)
- IVM Management**
 - [View/Modify User Accounts](#)
 - [View/Modify TCP/IP Settings](#)
 - [Guided Setup](#)
 - [Enter PowerVM Edition Key](#)
- Service Management**
 - [Electronic Service Agent](#)
 - [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)

The main content area is titled "View/Modify Host Ethernet Adapters" and includes a description: "A Host Ethernet Adapter (HEA) allows you to provide multiple partitions direct access to the physical Ethernet ports. To perform an action on a physical port, first select the port, and then select the task." Below this is a "Properties" tab and a table:

| Select | Type | Link State | Physical Location Code ^ | Connected Partitions | Available Connections |
|-----------------------|------|------------|--|----------------------|-----------------------|
| <input type="radio"/> | 1 G | Up | U78A0.001.DNWHZWR-P1-C6-T1 | 2 | 13 |
| <input type="radio"/> | 1 G | Down | U78A0.001.DNWHZWR-P1-C6-T2 | 1 | 13 |
| <input type="radio"/> | 1 G | Down | U78A0.001.DNWHZWR-P1-C6-T3 | 1 | 14 |
| <input type="radio"/> | 1 G | Down | U78A0.001.DNWHZWR-P1-C6-T4 | 1 | 14 |

Figure 4-86 Host Ethernet Adapter details

Assigning a Host Ethernet Adapter port to newly created LPAR

You can assign the HEA port to an LPAR at the creation time in the **Ethernet** window, by selecting the HEA port, as shown in Figure 4-87.

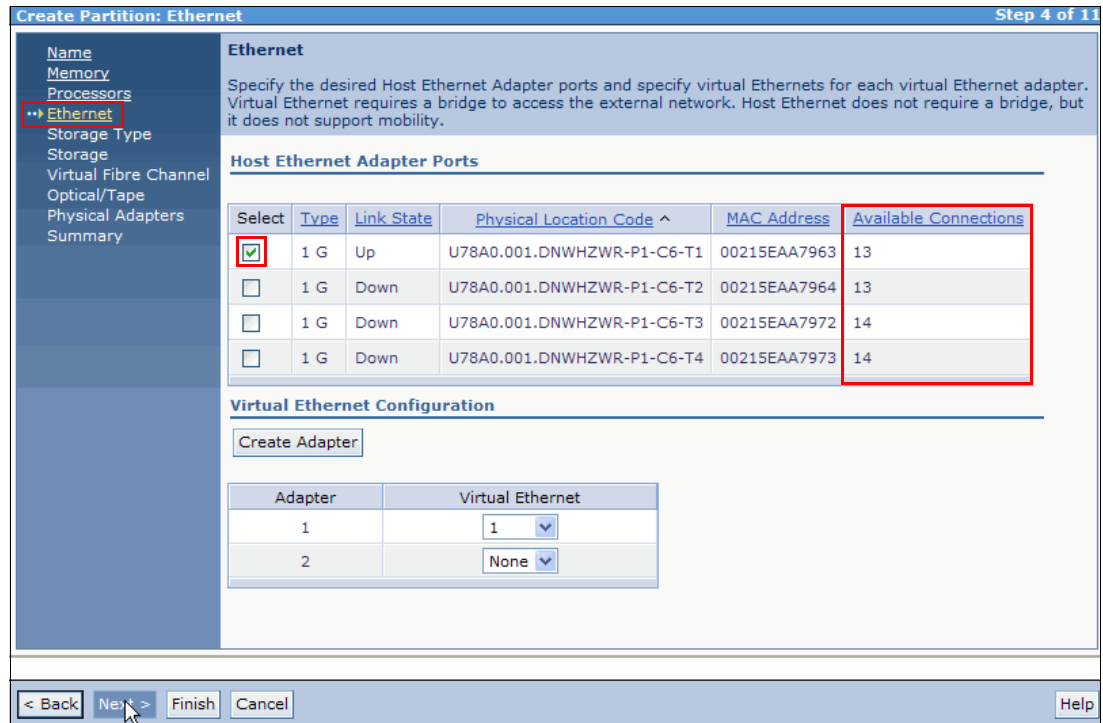


Figure 4-87 Assigning an HEA port to a new LPAR

You can see how many HEA ports are still available for each physical port in the **Available Connections** column, as shown in Figure 4-87.

Properties of the Host Ethernet Adapter physical port

You can see details of the *Per HEA physical port* for Physical Location Code, Speed, and Maximum transmission unit (MTU) by selecting the physical port in the **View/Modify Host Ethernet Adapters** panel and clicking **Properties**, as shown in Figure 4-88 and Figure 4-89 on page 132.

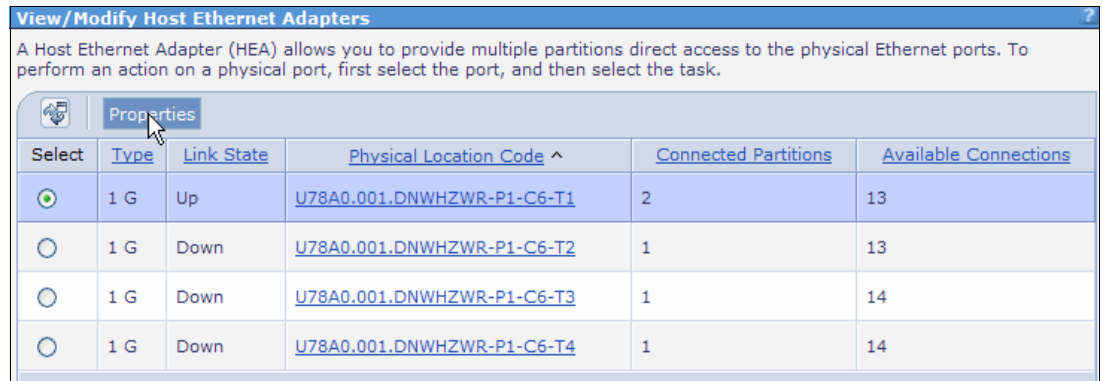


Figure 4-88 To check properties of the physical adapter

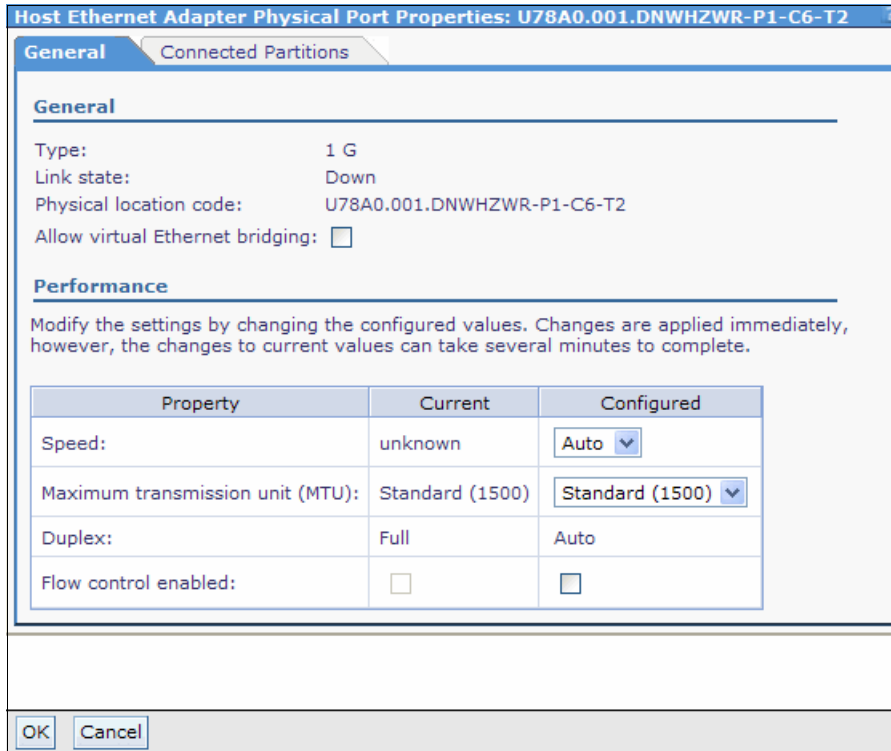


Figure 4-89 HEA physical adapter properties

You can see the partitions which are sharing the *Per HEA physical port* and its MAC address in the **Connected Partitions** tab of the HEA properties page. Also, you can assign the available free ports to other partitions, as shown in Figure 4-90.

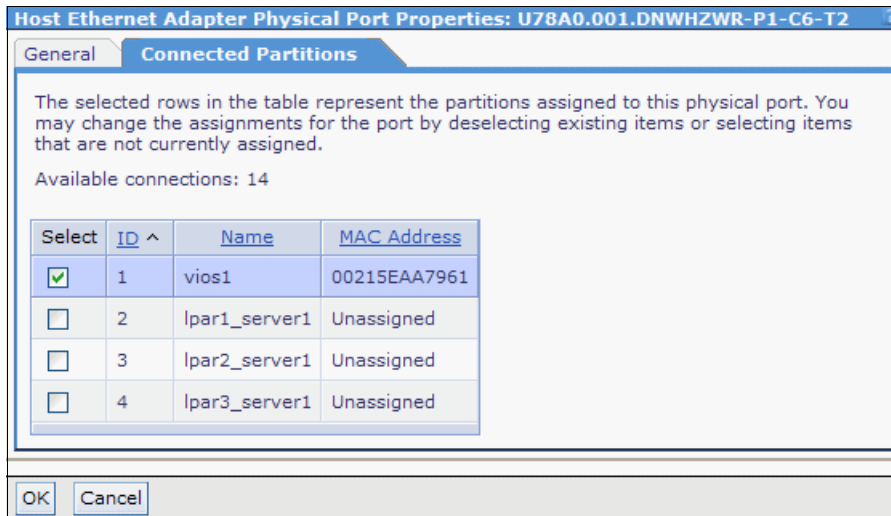


Figure 4-90 Partitions served by an HEA port

Using an HEA port for Ethernet bridging or Shared Ethernet Adapter

To use an HEA port for bridging client partitions with virtual Ethernet, you must assign a complete HEA physical port to the Virtual I/O Server; it cannot be shared by any other LPAR. To do this step, click **Allow virtual Ethernet bridging**, as shown in Figure 4-91.

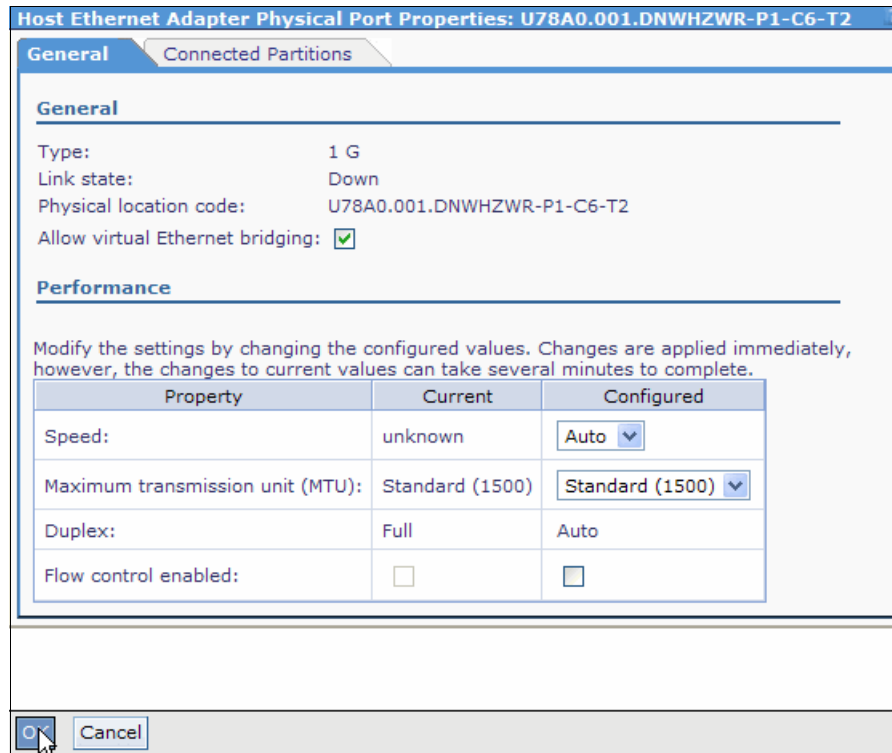


Figure 4-91 Ethernet bridging using HEA

You can assign the HEA to the Virtual I/O Server and LPARs by using dynamic LPAR at run time. See details in the dynamic LPAR on HEA section, “Dynamic LPAR operation on HEA adapters using graphical interface” on page 104.

4.10 Storage management

Virtual disks and physical volumes can be assigned to any LPAR, one at a time. Storage allocation can be changed over time, and the content of the virtual storage is kept. When a virtual disk is created by using a logical volume, its size can also be increased.

Data protection against single disk failure is available by using *software mirroring*:

- ▶ Use on IVM to protect it, but not the managed systems' data.
- ▶ Use two virtual disks for each of the managed systems' LPARs to protect its data.

4.10.1 Virtual storage assignment to a partition

Unassigned virtual disks and physical volumes can be associated to a running partition. After the operation completes, the operating system of the LPAR must issue its device discovery procedure to detect the newly added disk.

Before you remove a physical disk or a virtual disk from a running partition, the operating system removes the corresponding disk device because it becomes unavailable.

On the web graphical interface, it is possible to remove a virtual disk or a physical volume from a running LPAR by using the **View/Modify Virtual Storage** menu. Click **Modify Virtual Disk Partition Assignment**. A warning sign opens which requires an additional confirmation. Figure 4-92 shows an example of this message.

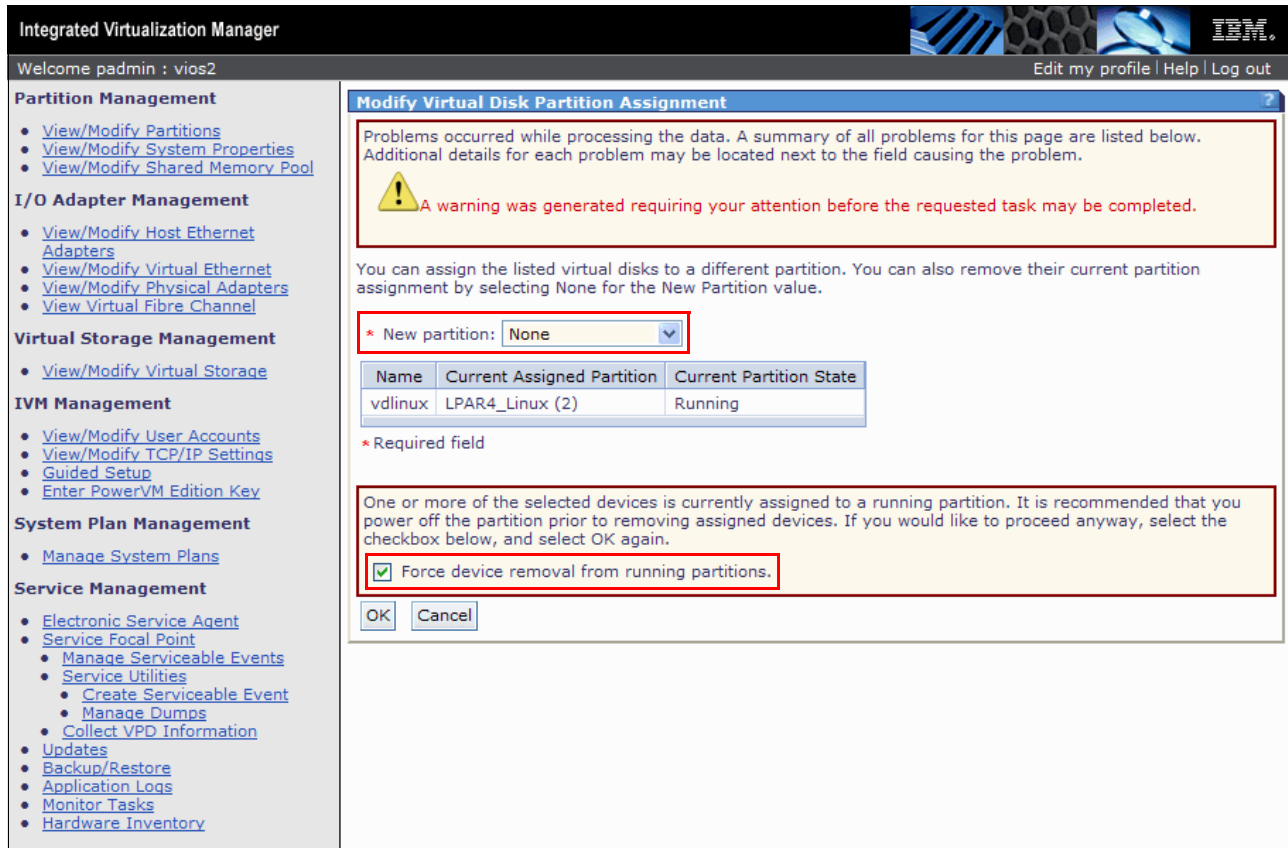


Figure 4-92 Forced removal of a virtual disk

4.10.2 Virtual disk extension

Several options are available to provide more disk space to an LPAR. The primary solution is to create a new virtual disk or select an entire physical disk and dynamically assign it to a partition. Because these operations can be done when the partition is running, it is preferred. After the partition's operating system issues its own device reconfiguration process, a new virtual disk is available for use. This disk can be used to extend existing data structures when using Linux with a logical volume manager or AIX 6.1 and 7.1.

When disk space is provided to a partition using a virtual disk, a secondary solution is to extend it using the IVM.

This operation can be executed when the partition is running, but the virtual disk must be taken offline to activate the change. Schedule the disk outages carefully so that they do not affect overall application availability. Consider the use of this solution when an existing operating system's volume must be increased in size and a new virtual SCSI disk cannot be added for this purpose; that is, when using Linux without a logical volume manager.

The following steps describe how to extend a virtual disk:

1. On the operating system, halt any activity on the disk to be extended. If stopping is not possible, shut down the partition.

- From the Virtual Storage Management menu in the IVM navigation area, click **View/Modify Virtual Storage**. From the work area, select the **Virtual Disks** tab. Then, **Select** the virtual disk that you want to extend, and choose **Extend** in the **More Tasks** drop-down menu. See Figure 4-93.

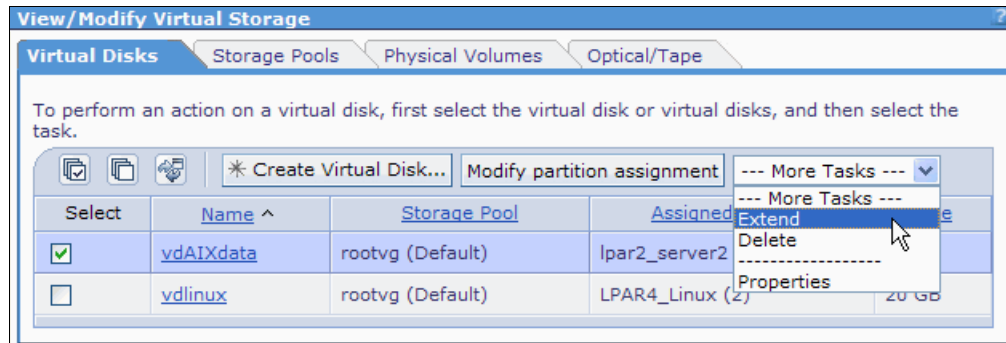


Figure 4-93 Extending virtual disks

- Enter the disk space to be added and click **OK**. If the virtual disk is owned by a running partition, a warning message opens, as shown in Figure 4-94. You must then select the check box to force the expansion. The additional disk space is allocated to the virtual disk and it is available for the LPAR.

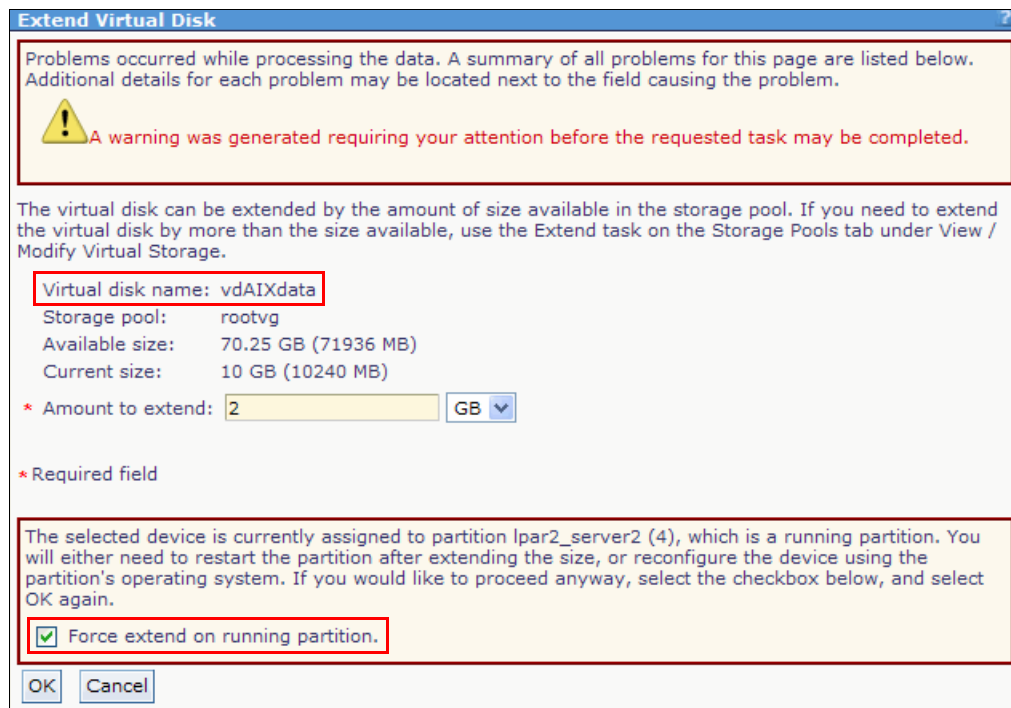


Figure 4-94 Forced expansion of a virtual disk

- On the operating system, issue the appropriate procedure to recognize the new disk size.

4.10.3 IVM system disk mirroring

To prevent an IVM outage because of a system disk failure, make the rootvg storage pool of the Virtual I/O Server redundant. The default installation of IVM uses only one physical disk.

Disk mirroring on the IVM is an advanced feature that, at the time of writing, is not available on the web GUI. It can be configured by using the Virtual I/O Server capabilities on the CLI. Only system logical volumes can be mirrored.

Mirrored logical volumes: Mirrored logical volumes are supported as virtual disks. This procedure mirrors all logical volumes that are defined in the rootvg and must not run if rootvg contains virtual disks.

The following steps describe how to provide a mirrored configuration for the rootvg storage pool:

1. Use the IVM to add a second disk of a similar size to rootvg. Under Virtual Storage Management in the navigation area, click **View/Modify Virtual Storage**, then go to the **Physical Volumes** tab. **Select** a disk of a similar size that is not assigned to any storage pool. From the **More Tasks** drop-down menu, click **Add to storage pool**, as shown in Figure 4-95.

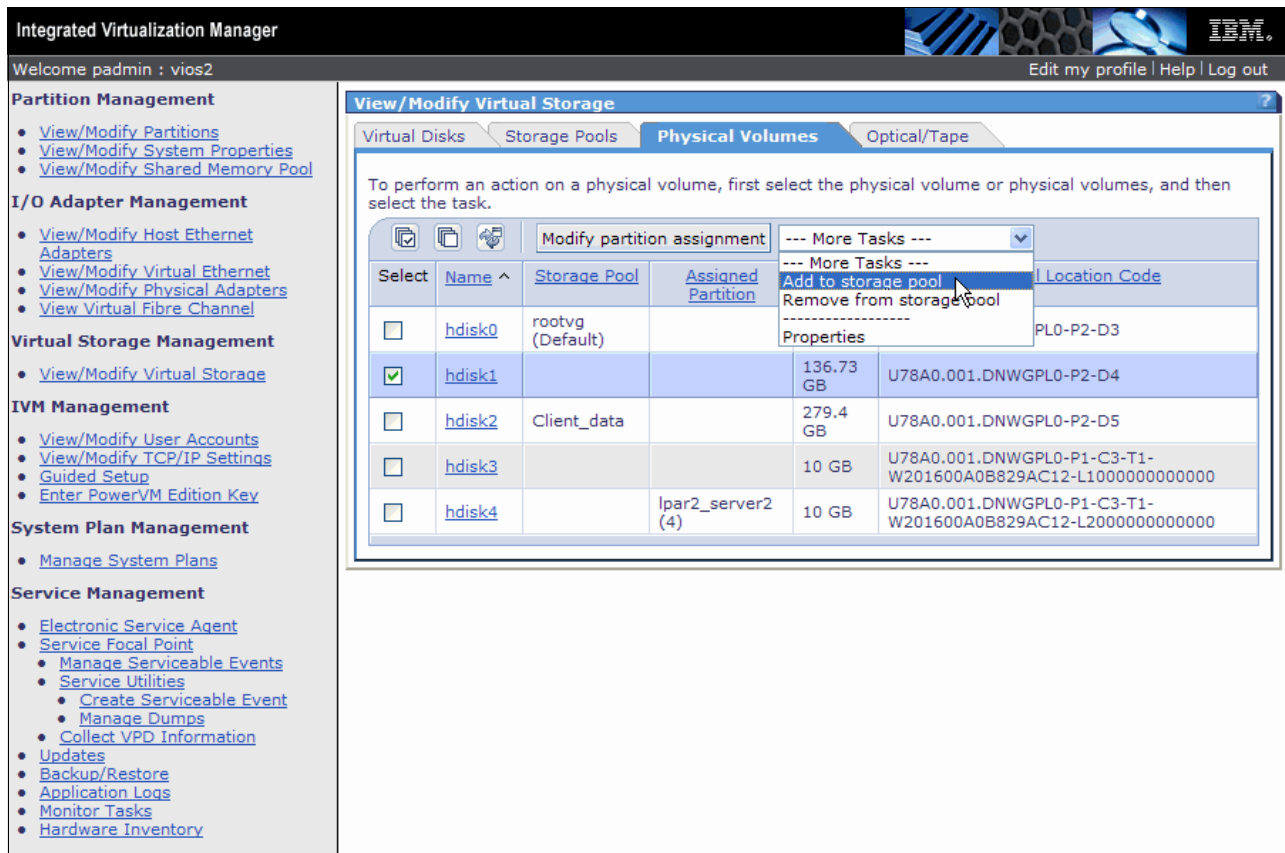


Figure 4-95 Adding a second disk in a storage pool for rootvg

2. In the Storage pool field, select **rootvg** and click **OK**.

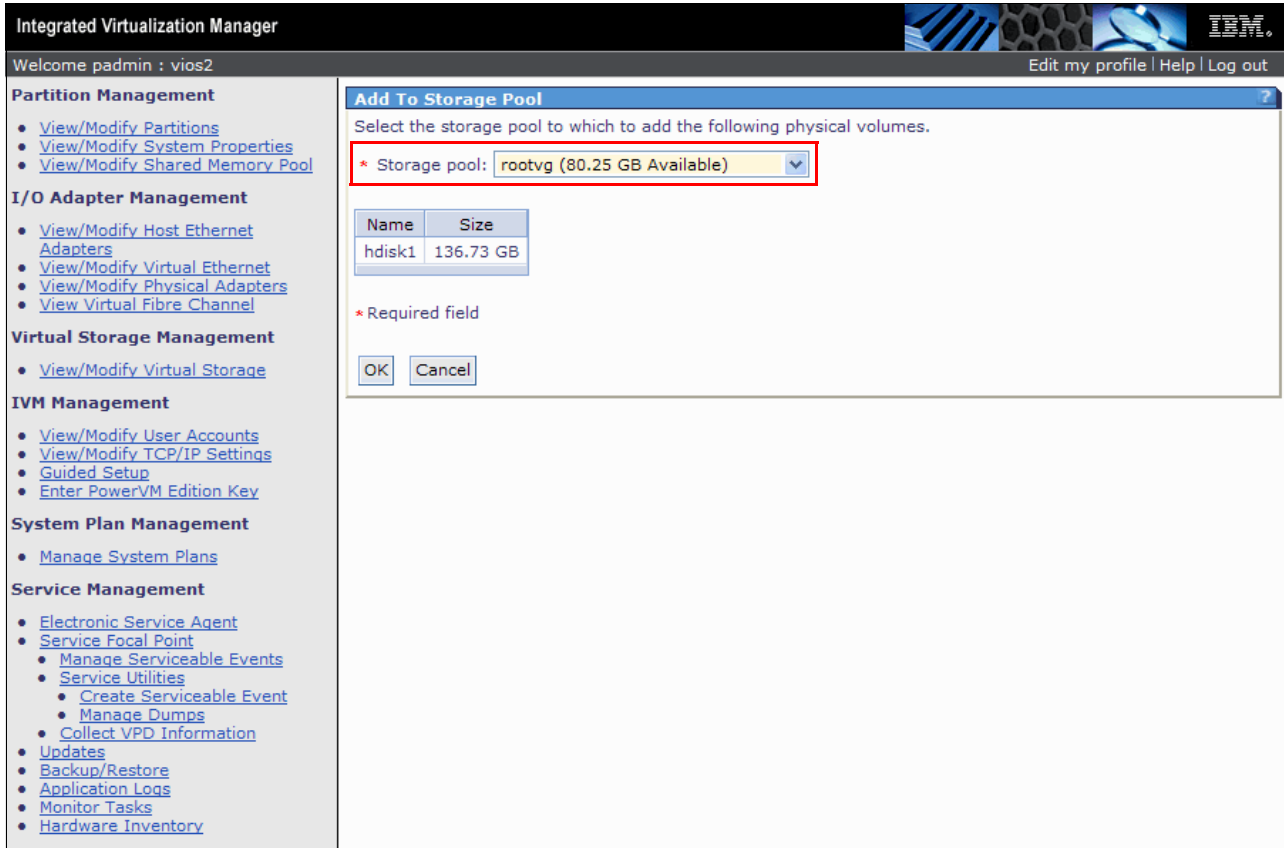


Figure 4-96 Specify addition to storage pool

3. The actual mirroring is done by using the Virtual I/O Server command line. Log in with the padmin user ID and issue the **mirrorios** command, as shown in Example 4-15. The command asks for confirmation and causes a Virtual I/O Server reboot to activate the configuration after performing data mirroring.

Example 4-15 rootvg mirroring at command line

```

$ lspv
NAME              PVID              VG              STATUS
hdisk0            000fe401e81abecd rootvg          active
hdisk1            000fe401e81abfd7 rootvg          active
hdisk2            00c1f170e937b535 Client_data     active
hdisk3            00f61aa6c6cbd642 None
hdisk4            00f61aa6df2ffe71 None

$ mirrorios
This command causes a reboot. Continue [y|n]?
y

```

4.11 N_Port Identifier Virtualization (NPIV)

IVM supports virtual Fibre Channel with physical adapters that support NPIV. To create a virtual Fibre Channel adapter in IVM, a client LPAR is needed. You can either define during LPAR creation (see 4.4, “Create logical partitions” on page 79) or by adding the virtual Fibre Channel adapter to an existing LPAR. This section describes how to add and remove a virtual Fibre Channel adapter to an existing LPAR.

Changes to a virtual Fibre Channel: To make changes (such as add or remove operation) to a virtual Fibre Channel on an active logical partition, you must have a Dynamic Logical Partitioning-capable partition. Otherwise, you must shut down the LPAR to finish the changes.

4.11.1 Adding a new virtual Fibre Channel adapter

You can add a new virtual Fibre Channel adapter by using the IVM management interface. The following steps describe how to add a new virtual Fibre Channel adapter:

1. Go to **View/Modify Partitions** in the IVM management interface. **Select** the wanted LPAR link, as shown in Figure 4-97.

The screenshot displays the IVM management interface. On the left is a navigation menu with categories like Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, and Service Management. The main area is titled 'View/Modify Partitions' and includes a 'System Overview' section with metrics such as Total system memory (128 GB), Memory available (39.75 GB), and Total processing units (16). Below this is a 'Partition Details' table with columns for Select, ID, Name, State, Uptime, Memory, Processors, Entitled Processing Units, and Utilized Processing Units. The table lists four partitions: 'vios1' (Running, 16 GB), 'lpar1_server1' (Running, 2 GB), 'lpar2_server1' (Not Activated, 512 MB), and 'lpar3_server1' (Not Activated, 1 GB). A mouse cursor is pointing at the 'lpar2_server1' row.

| Select | ID ^ | Name | State | Uptime | Memory | Processors | Entitled Processing Units | Utilized Processing Units | Ref C |
|--------------------------|------|---------------|---------------|------------|--------|------------|---------------------------|---------------------------|-------|
| <input type="checkbox"/> | 1 | vios1 | Running | 3.86 Hours | 16 GB | 16 | 1.6 | 0.03 | |
| <input type="checkbox"/> | 2 | lpar1_server1 | Running | 1.08 Hours | 2 GB | 2 | 0.2 | 0.01 | |
| <input type="checkbox"/> | 3 | lpar2_server1 | Not Activated | | 512 MB | 1 | 0.1 | | 0000 |
| <input type="checkbox"/> | 4 | lpar3_server1 | Not Activated | | 1 GB | 2 | 0.2 | | 0000 |

Figure 4-97 View/Modify Partitions panel

2. The Partition Properties window is started. Select the **Storage** tab, as depicted in Figure 4-98.

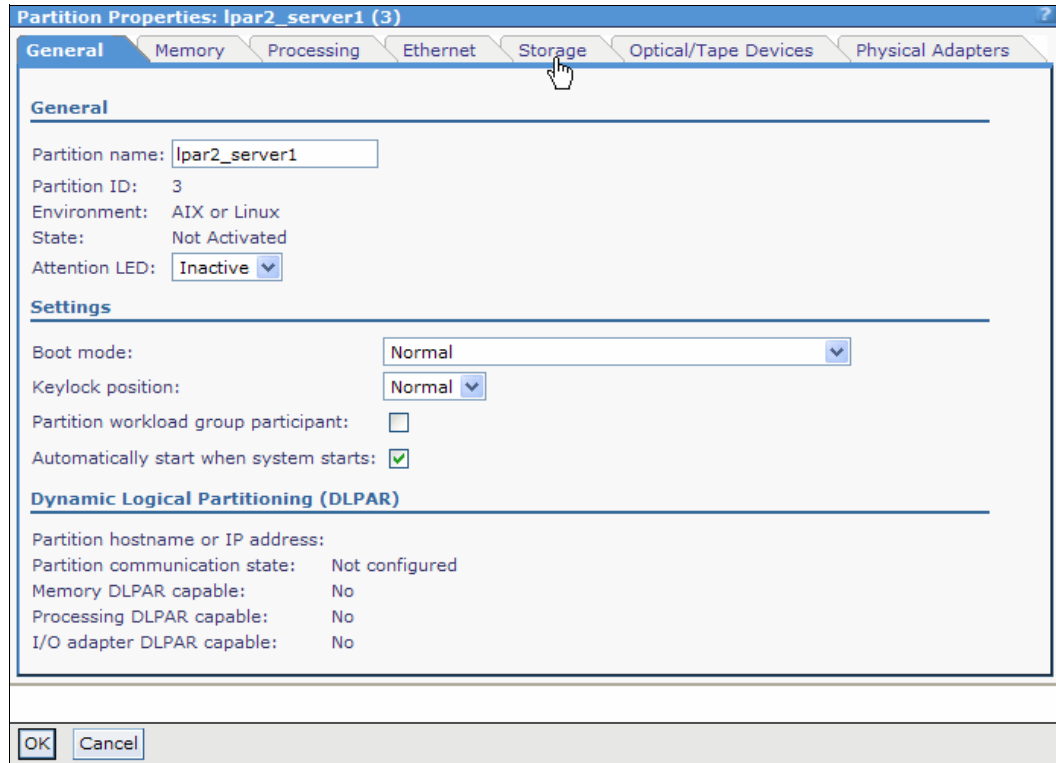


Figure 4-98 Partition Properties window

3. Figure 4-99 shows the Storage tab. Expand the **Virtual Fibre Channel** menu and click **Add** to create a new virtual Fibre Channel adapter.

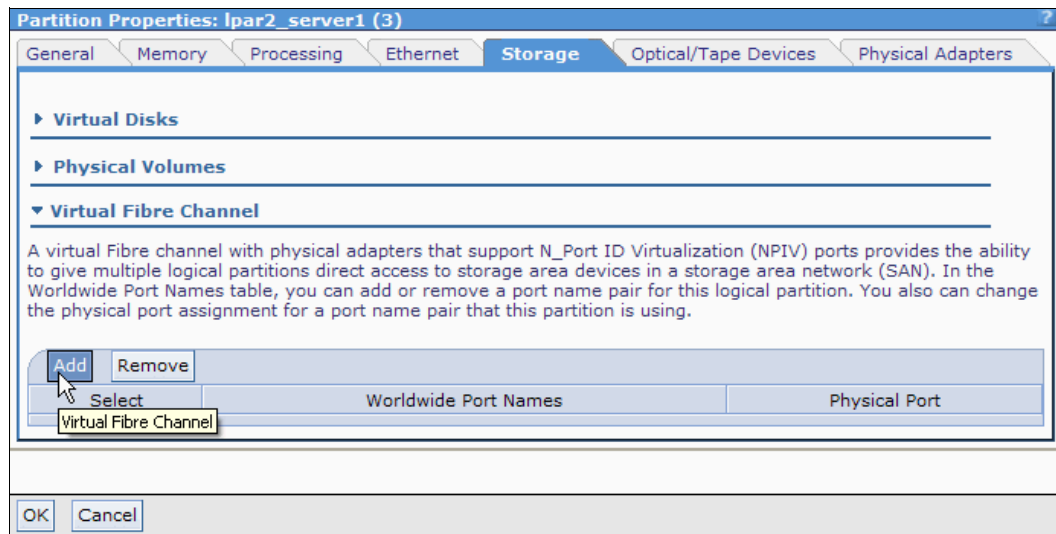


Figure 4-99 Storage tab

4. Ensure that the check box for the new virtual Fibre Channel is selected. Using the **Physical Port** pull-down menu, choose the physical port that you want to associate with the new virtual Fibre Channel adapter, as shown in Figure 4-100. Click **OK** to make the configuration effective.

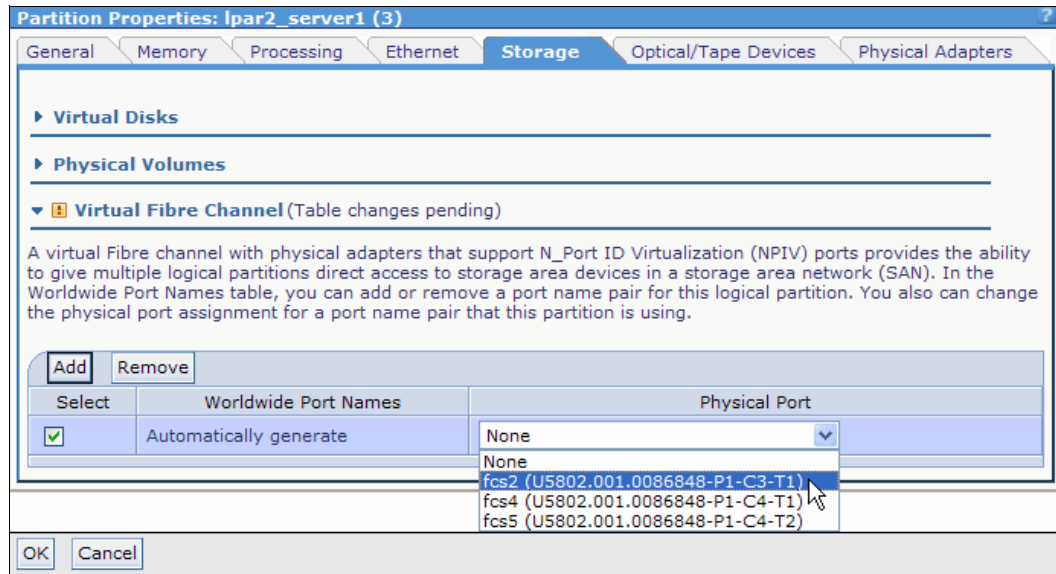


Figure 4-100 Adding a new virtual Fibre Channel adapter

5. The window in Figure 4-101 opens. It is closed when the configuration change is finished.

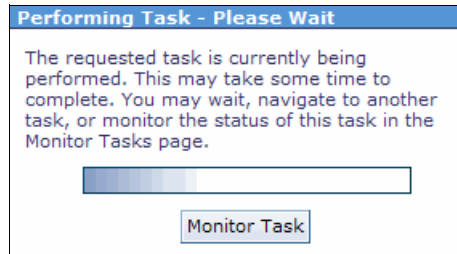


Figure 4-101 Performing Task panel

4.11.2 Removing a virtual Fibre Channel adapter

The following steps describe how to remove a virtual Fibre Channel adapter:

1. Click **View/Modify Partitions** in the Partition Management menu in the navigation area. In the work area, click the wanted LPAR link.
2. Select the **Storage** tab in the Partition Properties window, as shown in Figure 4-102.
3. Expand the **Virtual Fibre Channel** menu in the Storage tab. A list with available virtual Fibre Channel adapters is displayed, as shown in Figure 4-102. Select the check box of the virtual Fibre Channel adapter that you want to remove. Click **Remove**.

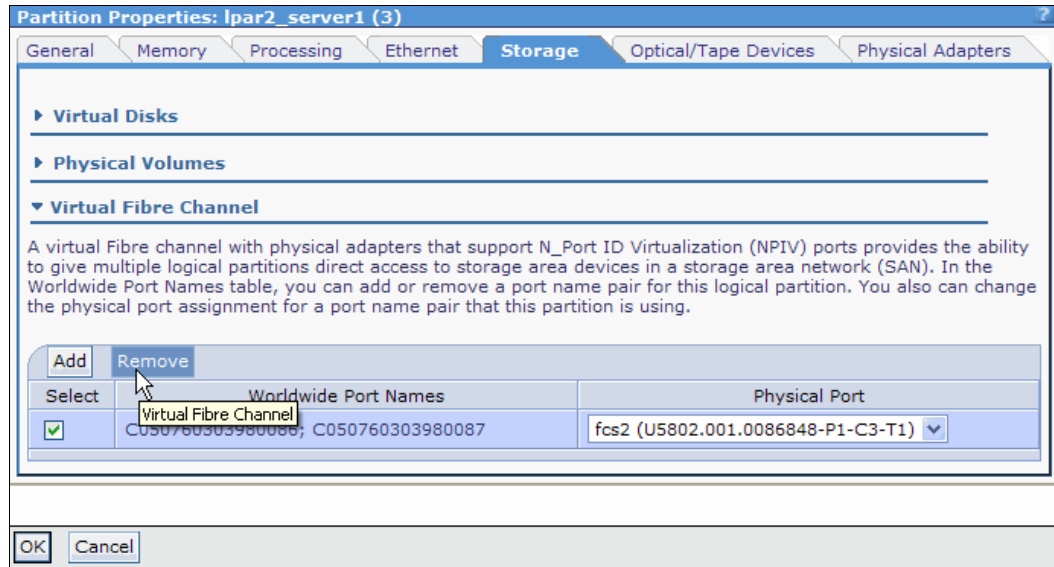


Figure 4-102 Removing a virtual Fibre Channel adapter

4. A confirmation panel is displayed, as shown in Figure 4-103. Click **OK** to continue.

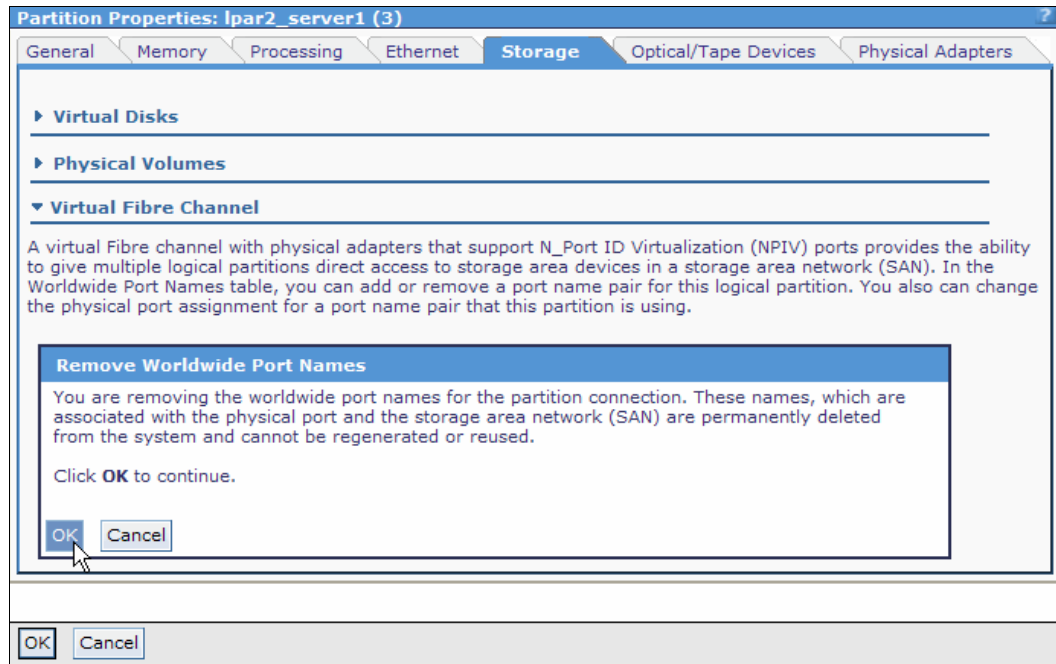


Figure 4-103 Confirmation window for removing a virtual Fibre Channel adapter

5. The window in Figure 4-104 is displayed. There is an exclamation mark near the Virtual Fibre Channel menu; this means that the configuration change is pending. Click **OK** to apply the changes.

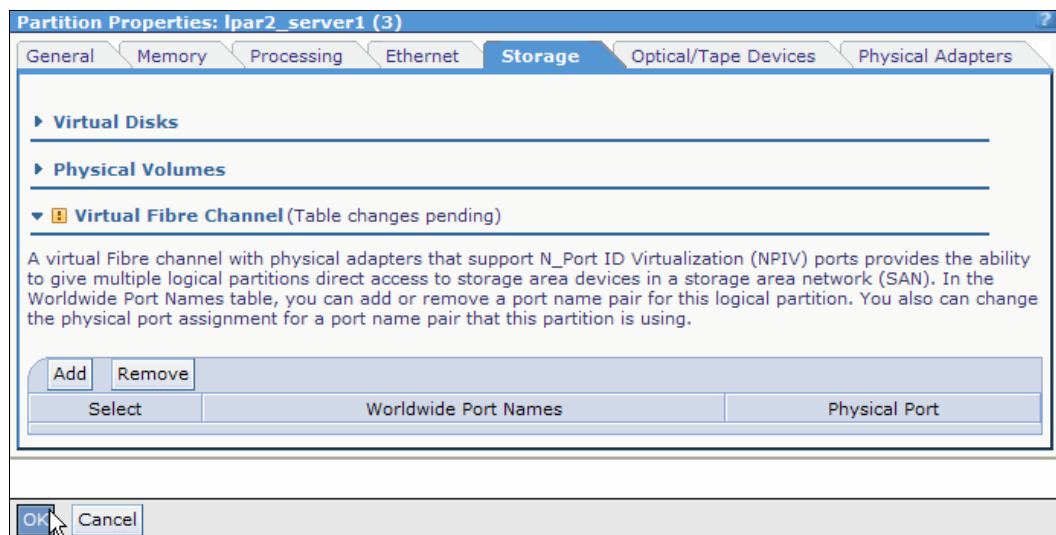


Figure 4-104 Applying changes to the LPAR configuration

6. A Performing Task window is shown, as depicted in Figure 4-101 on page 140. The window closes when the new configuration is applied.

4.11.3 Changing the physical port of an existing virtual Fibre Channel

You can change a physical port of an existing virtual Fibre Channel adapter without removing it. The following steps describe how to modify an existing virtual Fibre Channel adapter physical port's assignment:

1. Go to **View/Modify Partitions** in the IVM management interface. Click the wanted LPAR link.
2. Select the **Storage** tab in the Partition Properties window.
3. Expand the **Virtual Fibre Channel** menu in the Storage tab. A list with available virtual Fibre Channel adapters is displayed. Using the pull-down menu under the **Physical Port** column, choose the new physical port for the virtual Fibre Channel adapter that you want to change. See Figure 4-105.

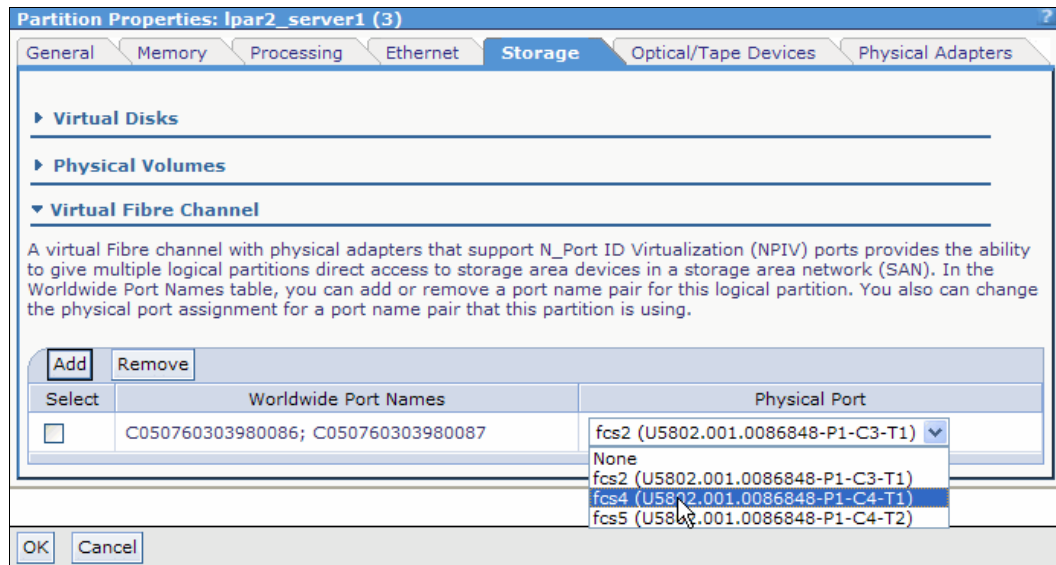


Figure 4-105 Changing the physical port for a virtual Fibre Channel adapter

- There is an exclamation mark near the Virtual Fibre Channel menu in Figure 4-106. It means that the new configuration is pending. Click **OK** to apply the changes.

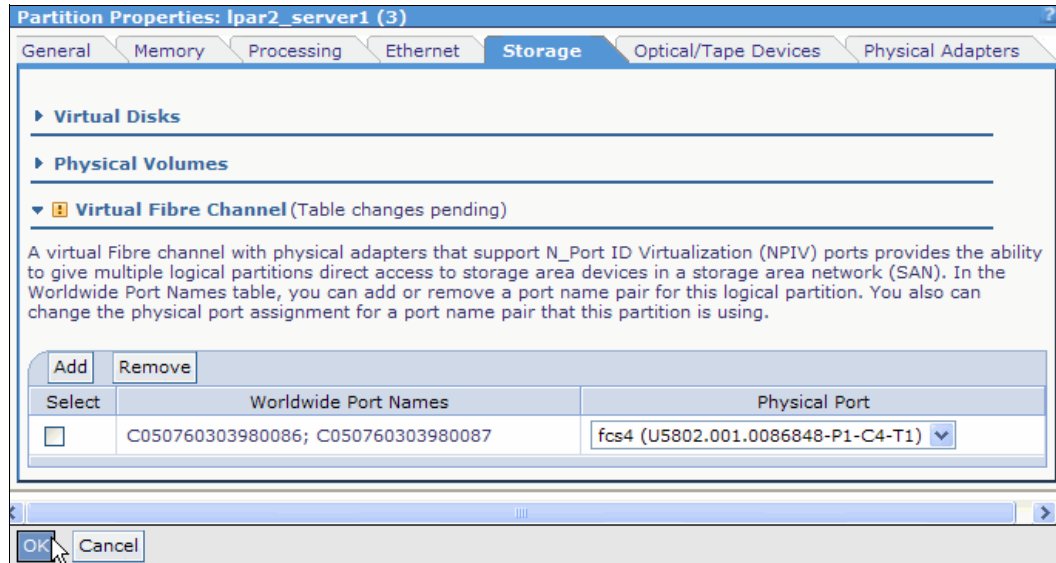


Figure 4-106 Applying a new physical port for a virtual Fibre Channel adapter

- A Performing Task window is shown, as depicted in Figure 4-101 on page 140. The window closes when the new configuration is applied.

4.12 Securing the Virtual I/O Server

The *Virtual I/O Server* provides extra security features that enable you to control access to the virtual environment and ensure the security of your system. The following topics describe the available security features and provide tips for ensuring a secure environment for your Virtual I/O Server setup.

Introduction to Virtual I/O Server security

You can set security options that provide tighter security controls over your Virtual I/O Server environment. These options enable you to select a level of system security hardening and specify settings that are allowable within that level. The Virtual I/O Server security feature also enables you to control network traffic by enabling the Virtual I/O Server firewall. You can configure these options by using the **viopsecure** command.

The **viopsecure** command activates, deactivates, and displays security hardening rules. By default, none of the security hardening features are activated after installation. Upon running the **viopsecure** command, the command guides the user through the proper security settings, which range from High to Medium to Low. After this initial selection, a menu is displayed which itemizes the security configuration options that are associated with the selected security level in sets of 10. These options can be accepted in whole, individually toggled off or on, or ignored. After any changes, **viopsecure** continues to apply the security settings to the computer system.

The **viopsecure** command also configures, unconfigures, and displays network firewall settings. Using the **viopsecure** command, you can activate and deactivate specific ports and specify the interface and IP address from which connections are allowed.

For more information about this command, see the **viosecure** command in the Virtual I/O Server Commands Reference. The following sections provide an overview of the Virtual I/O Server features.

System security hardening

The *system security hardening* feature protects all elements of a system by tightening security or implementing a higher level of security. Although hundreds of security configurations are possible with the Virtual I/O Server security settings, you can easily implement security controls by specifying a high, medium, or low security level. Configuration of the Virtual I/O Server system security hardening is described in one of the next sections.

The system security hardening features provided by the Virtual I/O Server enable you to specify the following values:

- ▶ Password policy settings
- ▶ The `usrck`, `pwdck`, `grpck`, and `sysck` actions
- ▶ Default file creation settings
- ▶ System crontab settings

Configuring a system at too high of a security level might deny services that are needed. For example, the `telnet` and `rlogin` commands are disabled for high-level security because the login password is sent over the network unencrypted. If a system is configured at a too low security level, the system might be vulnerable to security threats. Because each enterprise has its own unique set of security requirements, the predefined High, Medium, and Low security configuration settings are best suited as a starting point for security configuration rather than an exact match for security requirements. As you become more familiar with the security settings, you can make adjustments by choosing the hardening rules that you want to apply. You can get information about the hardening rules by running the `man` command.

Virtual I/O Server firewall

The Virtual I/O Server firewall enables you to enforce limitations on IP activity in your virtual environment. With this feature, you can specify which ports and network services are allowed access to the Virtual I/O Server system. For example, if you need to restrict login activity from an unauthorized port, you can specify the port name or number and specify *deny* to remove it from the Allow list. You can also restrict a specific IP address.

Before configuring firewall settings, you must first enable the Virtual I/O Server firewall. The following topic describes this action.

Configuring firewall settings

Enable the Virtual I/O Server (VIOS) firewall to control IP activity.

The Virtual I/O Server firewall is not enabled by default. To enable the firewall, you must turn it on by using the **viosecure** command with the `-firewall` option. When you enable it, the default setting is activated, which allows access for the following IP services:

- ▶ ftp
- ▶ ftp-data
- ▶ dns
- ▶ ssh
- ▶ web
- ▶ https
- ▶ rmc
- ▶ wbem-rmi
- ▶ wbem-http
- ▶ wbem-https

telnet command: The `telnet` command is disabled when the firewall is turned on. Therefore, if you are using Telnet to set the security settings, you lose your connection or session.

Firewall settings: The firewall settings are in the `viosecure.ct1` file in the `/home/ios/security` directory. You can use the `-force` option to enable the standard firewall default ports. For more about the force option, see the `viosecure` command description.

You can use the default setting or configure the firewall settings to meet the needs of your environment by specifying which ports or port services to allow. You can also turn off the firewall to deactivate the settings.

Use the following tasks at the Virtual I/O Server command line to configure the Virtual I/O Server firewall settings:

1. Enable the Virtual I/O Server firewall by issuing the following command:

```
viosecure -firewall on -reload
```

2. Specify the ports to allow or deny, by using the following command:

```
viosecure -firewall allow | deny -port number
```

3. View the current firewall settings by issuing the following command:

```
viosecure -firewall view
```

4. If you want to disable the firewall configuration, issue the following command:

```
viosecure -firewall off
```

For more about any `viosecure` command option, see the `viosecure` command description.

Configuring Virtual I/O Server system security hardening

Set the security level to specify the security hardening rules for your Virtual I/O Server (VIOS) system.

To implement system security hardening rules, you can use the `viosecure` command to specify a security level of High, Medium, or Low. A default set of rules is defined for each level. You can also set a level of default, which returns the system to the system standard settings and removes any level settings that are applied.

The low-level security settings are a subset of the medium-level security settings, which are a subset of the high-level security settings. Therefore, the High level is the most restrictive and provides the greatest level of control. You can apply all of the rules for a specified level or select which rules to activate for your environment. By default, no Virtual I/O Server security levels are set; you must run the `viosecure` command to enable the settings.

Use the following tasks to configure the system security settings of the Virtual I/O Server.

Setting a security level

To set a Virtual I/O Server security level of High, Medium, or Low, use the **viosecure -level** command, as in the following example:

```
viosecure -level low -apply
```

Changing the settings in a security level

To set a Virtual I/O Server security level in which you specify which hardening rules to apply for the setting, run the **viosecure** command interactively, as in the following example:

1. At the Virtual I/O Server command line, type **viosecure -level high**. All security level options (hardening rules) at that level are displayed, 10 at a time. (Pressing Enter displays the next set in the sequence.)
2. Review the displayed options and make your selection by entering the numbers that you want to apply, separated by a comma. Type ALL to apply all of the options, or type NONE to apply none of the options.
3. Press Enter to display the next set of options, and continue entering your selections.

To exit the command without making any changes, enter q.

Viewing the current security setting

To display the current Virtual I/O Server security level setting, use the **viosecure** command with the -view flag, as in the following example:

```
viosecure -view
```

Removing security level settings

To unset any previously set system security levels and return the system to the standard system settings, issue the following command:

```
viosecure -level default
```

For more information about using the **viosecure** command, see the **viosecure** command description.

4.13 Connecting to the Virtual I/O Server using OpenSSH

This topic describes how to set up remote connections to the Virtual I/O Server by using Open Secure Shell (OpenSSH) secure connections.

Setting up Secure Shell authorization for a non-prompted connection

Use the following steps to set up Secure Shell (SSH) authorization for a non-prompted connection:

1. If the `id_dsa` files do not exist on your workstation, create them using the **ssh-keygen** command (press Enter for passphrases), as shown in Example 4-16.

Example 4-16 Create the id_dsa files on your workstation

```
nimres2:/ # ssh-keygen -t dsa
Generating public/private dsa key pair.
Enter file in which to save the key (//.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
```

```

Enter same passphrase again:
Your identification has been saved in //.ssh/id_dsa.
Your public key has been saved in //.ssh/id_dsa.pub.
The key fingerprint is:
24:96:65:93:72:04:37:76:f2:09:98:90:ff:84:15:d4 root@nimres2
The key's randomart image is:
+--[ DSA 1024]-----+
|      .o.*&+.      |
|      . +*+BE.     |
|      .+=.  o      |
|      .oo.         |
|      oS          |
|      .           |
+-----+
nimres2:/ #

```

- Verify that the keys are generated on your workstation (Example 4-17).

Example 4-17 Verify successful creation of id_dsa files

```

nimres2:/.ssh # ls -l
total 16
-rw----- 1 root    system    668 Sep 21 11:03 id_dsa
-rw-r--r-- 1 root    system    602 Sep 21 11:03 id_dsa.pub
nimres2:/.ssh #

```

- Now, log in to the IVM through SSH. There is not yet a *known_hosts* file that is created, which is done during the first SSH login (Example 4-18).

Example 4-18 First SSH login toward IVM - known_hosts file creation

```

nimres2:/.ssh # ssh padmin@vios2
The authenticity of host 'vios2 (172.16.20.201)' can't be established.
RSA key fingerprint is 30:aa:f2:c2:30:4f:ad:6c:05:a8:91:35:b1:9b:6d:0a.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'vios2' (RSA) to the list of known hosts.
Last unsuccessful login: Thu Sep 20 16:47:41 CDT 2012 on ssh from 172.16.254.22
Last login: Fri Sep 21 09:10:56 CDT 2012 on /dev/pts/2 from 172.16.254.14
$
Connection to vios2 closed.
nimres2:/.ssh # ls -l
total 24
-rw----- 1 root    system    668 Sep 21 11:03 id_dsa
-rw-r--r-- 1 root    system    602 Sep 21 11:03 id_dsa.pub
-rw-r--r-- 1 root    system    401 Sep 21 11:06 known_hosts
nimres2:/.ssh #

```

The *known_hosts* file is created.

- The next step is to copy the `authorized_keys2` file using the `scp` command. (Example 4-19).

Example 4-19 Transfer of authorized_keys2 file

```
nimres2:/.ssh # scp padmin@vios2:/home/padmin/.ssh/authorized_keys2 ~/.ssh
padmin@vios2's password:
authorized_keys2                100% 394    0.4KB/s  00:00
nimres2:/.ssh #
```

- Add the contents of your local SSH public key (`id_dsa.pub`) to the `authorized_keys2` file (Example 4-20).

Example 4-20 Add contents of local SSH public key to authorized_keys2 file

```
nimres2:/.ssh # cat id_dsa.pub >> authorized_keys2
nimres2:/.ssh #
```

- Verify the successful addition of the public key by comparing the size of the `authorized_keys2` file to the `id_dsa.pub` file (Example 4-21).

Example 4-21 Compare addition of public key

```
nimres2:/.ssh # ls -l
total 32
-rw-r--r--  1 root    system    996 Sep 21 11:13 authorized_keys2
-rw-----  1 root    system    668 Sep 21 11:03 id_dsa
-rw-r--r--  1 root    system    602 Sep 21 11:03 id_dsa.pub
-rw-r--r--  1 root    system    401 Sep 21 11:06 known_hosts
```

- Transfer the `authorized_keys2` file back to the IVM into the directory `/home/padmin/.ssh` (Example 4-22).

Example 4-22 File Transfer Protocol (FTP) of the authorized key back to IVM

```
nimres2:/.ssh # scp ~/.ssh/authorized_keys2 padmin@vios2:/home/padmin/.ssh
padmin@vios2's password:
authorized_keys2                100% 996    1.0KB/s  00:00
nimres2:/.ssh #
```

- Verify that the key can be read by the SSH daemon on the IVM and test the connection by typing the `ioslevel` command (Example 4-23).

Example 4-23 Test the configuration

```
nimres2:/.ssh # ssh padmin@vios2
Last unsuccessful login: Thu Sep 20 16:47:41 CDT 2012 on ssh from 172.16.254.22
Last login: Fri Sep 21 10:15:39 CDT 2012 on ssh from 172.16.20.41
$ ioslevel
2.2.1.4
$
```

After you establish these secure remote connections, you can run several commands. The following examples show some of the commands that you can run:

► **ssh padmin@vios2**

This command provides an interactive login (IP address is also possible).

► **ssh -t padmin@vios2 ioscli mkvt -id 2**

This command enables us to get a console directly to a client LPAR with id 2.

► **ssh padmin@vios2 lssyscfg -r sys**

Example 4-24 shows the output of the **lssyscfg** command.

Example 4-24 Output of the lssyscfg command

```
nimres2:/ # ssh padmin@vios2 lssyscfg -r sys
name=M50_1_9409-M50-1059C72,type_model=9409-M50,serial_num=1059C72,"ipaddr=172.16.20.209
,fe80:0000:0000:0000:0214:5eff:fe0:28d6",state=Operating,sys_time=06/08/72
20:03:09,power_off_policy=0,active_lpar_mobility_capable=1,inactive_lpar_mobility_capabl
e=1,cod_mem_capable=0,cod_proc_capable=1,vet_activation_capable=1,os400_capable=1,active
_lpar_share_idle_procs_capable=1,micro_lpar_capable=1,dlpar_mem_capable=1,assign_phys_io
_capable=1,lhea_capable=1,active_mem_sharing_capable=1,max_lpars=40,max_power_ctrl_lpars
=1,service_lpar_id=1,service_lpar_name=10-59C72,mfg_default_config=0,curr_configured_max
_lpars=96,pend_configured_max_lpars=96,config_version=0100030000000000,pend_lpar_config
state=enabled,lpar_avail_priority_capable=1,lpar_proc_compat_mode_capable=1,"lpar_proc_c
ompat_modes=default,POWER6_enhanced,POWER6",max_vios_lpar_id=1,virtual_io_server_capable
=0,virtual_fc_capable=1
nimres2:/ #
```



Maintain

This chapter provides information about maintenance operations on the IBM Integrated Virtualization Manager (IVM).

This chapter describes the following topics:

- ▶ IVM backup and restore
- ▶ Logical partition backup and restore
- ▶ IVM upgrade
- ▶ Managed system firmware update
- ▶ IVM migration
- ▶ Command logging
- ▶ Integration with IBM Director
- ▶ Reporting hardware and software errors to IBM automatically

5.1 IVM maintenance

You can use the IVM to perform operations such as backup, restore, or upgrade. Some operations are available by using the *graphical user interface (GUI)*, the *command-line interface (CLI)*, or the *Advanced System Management Interface (ASMI)* menus.

5.1.1 Backup and restore of the logical partition definitions

Logical partition (LPAR) configuration information can be backed up to a file. This file is used to restore information if required and can also be exported to another system.

The following steps describe how to back up the LPAR configuration:

1. Under the Service Management menu in the navigation area, click **Backup/Restore**.
2. Select **Generate Backup** in the work area, as shown in Figure 5-1.

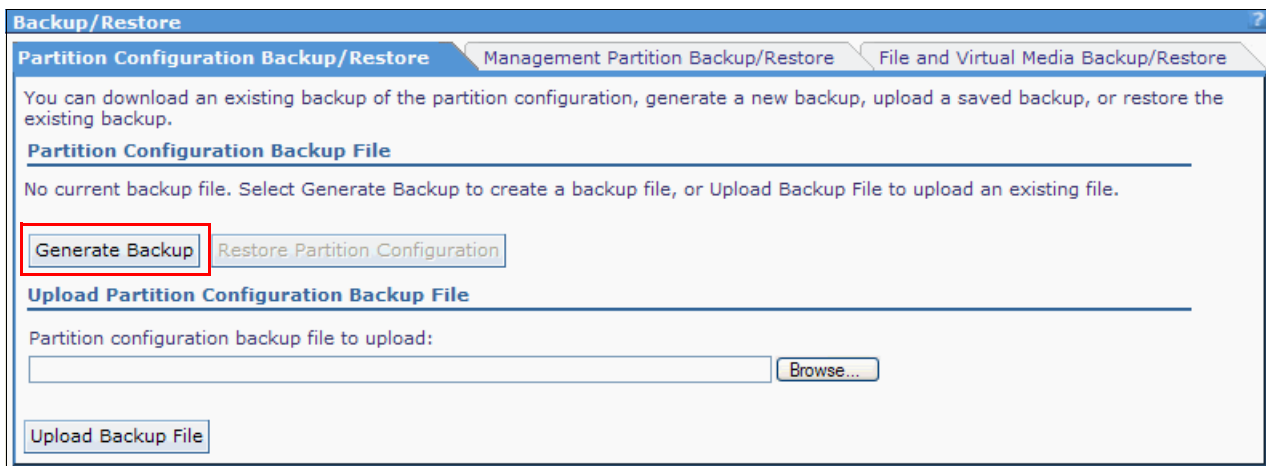


Figure 5-1 Partition Configuration Backup/Restore

A file named *profile.bak* is generated and stored under the user's home directory. In the work area, you can select this file name and save it to a disk. There is only one unique backup file at a time, and a new backup file replaces an existing one, as shown in Figure 5-2.

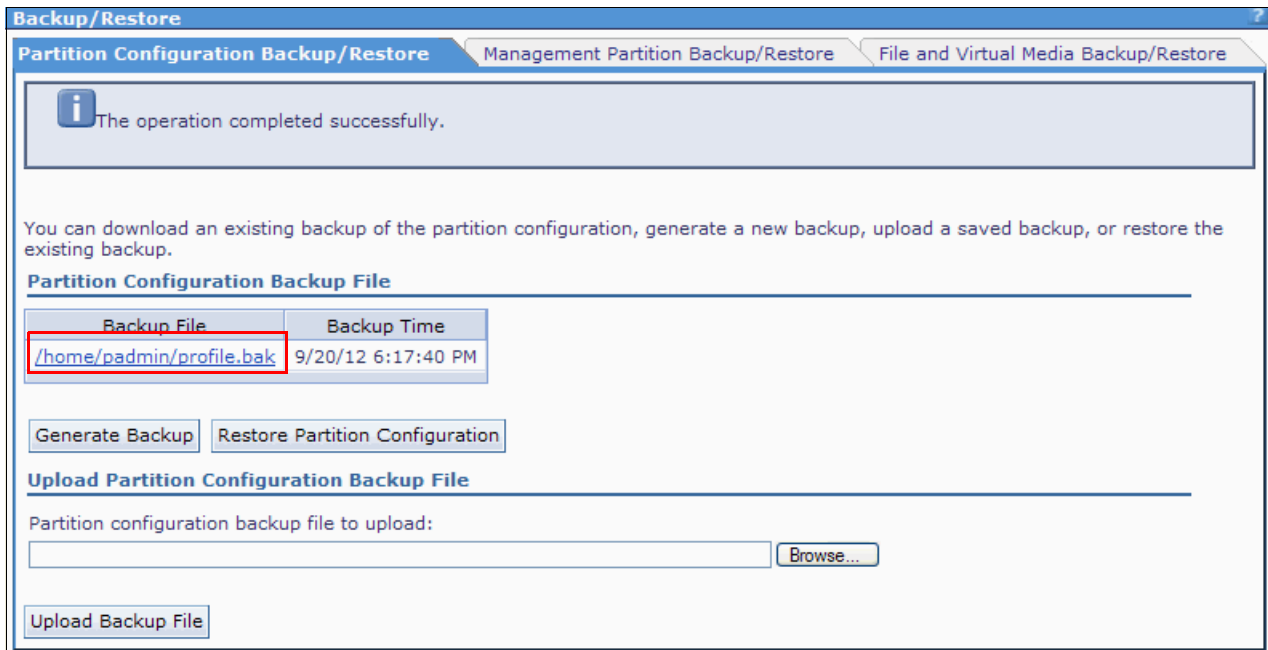


Figure 5-2 Profile backup stored in the /home directory

The backup file contains the LPAR's configuration, such as processors, memory, network, physical adapters, Host Ethernet Adapter (HEA), virtual SCSI, and virtual Fibre Channel. Information about virtual disks or physical disks is not included in the backup file; for this data, you must take a **viosbr** backup. Refer to section 5.1.4, "Virtual I/O Server configuration backup/restore" on page 156.

To perform a restore operation, the system must not have any LPAR configuration defined. Click **Restore Partition Configuration** to restore the last backed-up file. If you want to restore a backup file that is stored on your disk, follow these steps:

1. Click **Browse** and select the file.
2. Click **Upload Backup File**. The uploaded file replaces the existing backup file.
3. Click **Restore Partition Configuration** to restore the uploaded backup file. See details in Figure 5-3.

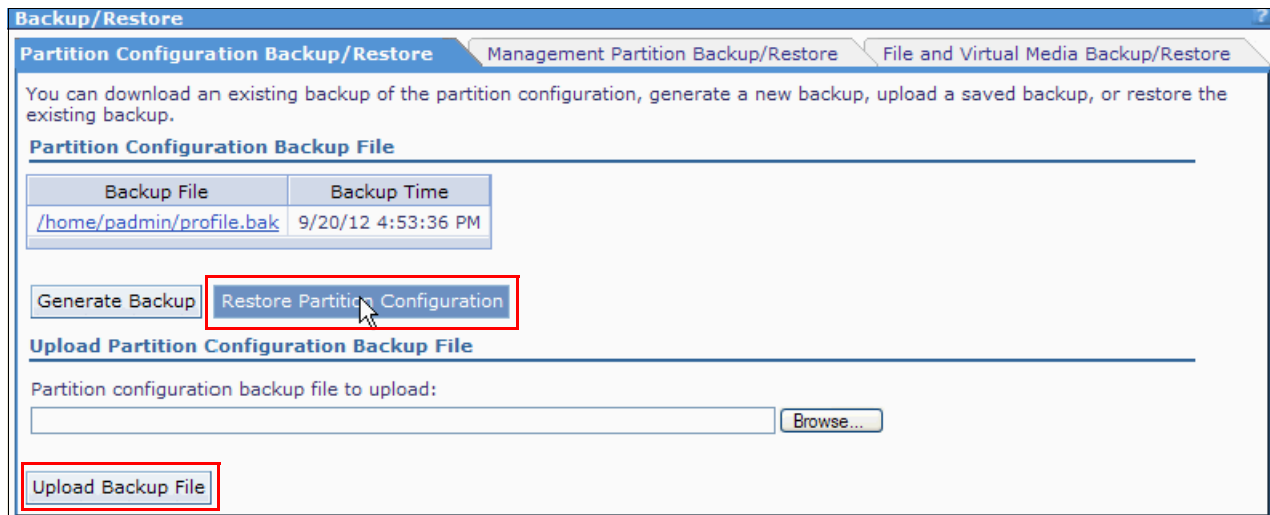


Figure 5-3 Restoring the profile

You can also back up and restore LPAR configuration information from the CLI. Use the **bkprofdata** command to back up the configuration information and the **rstprofdata** command to restore it. See the Virtual I/O Server and Integrated Virtualization Manager command descriptions on the Systems Hardware information web page:

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp?topic=/iphb1/iphb1_vios_commandslist.htm

5.1.2 Backup and restore of the IVM operating system

The only way to back up the IVM operating system is with the **backupios** command. No operating system backup operation is available within the GUI. This command creates a bootable image that includes the IVM partition's root volume group (rootvg). It can also contain the storage pool structure, depending on the flags used.

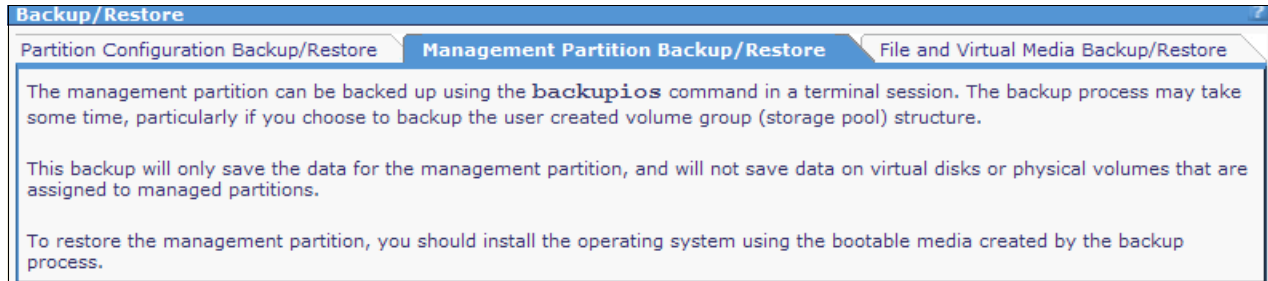


Figure 5-4 Backup/Restore of the Management Partition

LPAR backup operation: The backup operation does not save the data that is contained in virtual disks or physical volumes that are assigned to the LPARs.

The backup can use one of the following media types:

- ▶ File
- ▶ Tape
- ▶ CD-R
- ▶ DVD-RAM

To restore the management partition, install the operating system by using the bootable media that is created by the backup process.

5.1.3 Backup and restore of the particular set of files in /home/padmin

You can take backup of a particular set of files to a tape that is present in `/home/padmin` or the complete directory and restore them; see Figure 5-5.

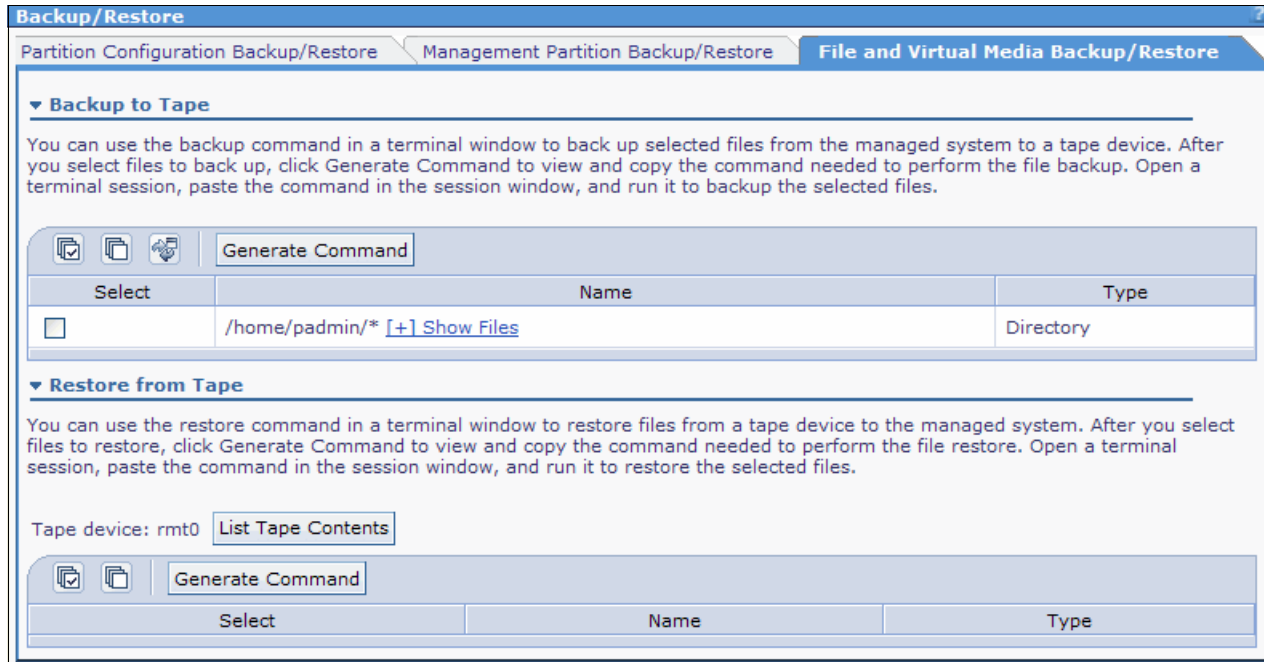


Figure 5-5 Backup of /home/padmin to tape

5.1.4 Virtual I/O Server configuration backup/restore

With `viosbr`, you can take backup of a Virtual I/O Server configuration. Examples of this type of setup include AMS, SEA, virtual SCSI mappings for virtual disk, NPIV, and SSP cluster (across all the nodes/Virtual I/O Server). You can restore these configurations completely if the VIO is unavailable or upgraded to a new level. Or, when there is no need for any of these features, just delete them and when needed, you can restore them using `viosbr`.

You can also specify `viosbr` to take backups at regular intervals, such as daily, weekly, and monthly. Then, the `viosbr` automatically takes the backup of the Virtual I/O Server with the current configuration. Therefore, there is no need to be concerned about the changes that occurred on the server.

Example 5-1 on page 157 shows the usage of `viosbr`.

Example 5-1 *viosbr usage*

```
viosbr -h
Usage: viosbr -backup -file FileName [-frequency daily|weekly|monthly [-numfiles fileCount]]
       viosbr -backup -clustername clusterName -file FileName [-frequency daily|weekly|monthly
[-numfiles fileCount]]
       viosbr -nobackup
       viosbr -view -file FileName [-type devType] [-detail | -mapping]
       viosbr -view -file FileName -clustername clusterName [-type devType][ -detail | -mapping]
       viosbr -view -list [UserDir]
       viosbr -restore -file FileName [-validate | -inter] [-type devType]
       viosbr -restore -file FileName [-type devType] [-force]
       viosbr -restore -file FileName -skipcluster
       viosbr -restore -clustername clusterName -file FileName -subfile NodeFile [-validate | -inter |
-force][ -type devType][ -skipcluster]
       viosbr -restore -clustername clusterName -file FileName -repopvs list_of_disks [-validate |
-inter | -force][ -type devType][ -currentdb]
       viosbr -restore -clustername clusterName -file FileName -subfile NodeFile -xmlvtds
       viosbr -recoverdb -clustername clusterName [ -file FileName ]
       viosbr -migrate -file FileName
```

Backup virtual & logical configurations, lists configurations and restores configurations.

| | |
|--------------|---|
| -backup | Takes backup of VIOS configuration. A compressed file with extension .tar.gz will be created. |
| -clustername | Specifies name of cluster to generate, display, or restore configuration file. |
| -currentdb | Restores the cluster without restoring database from the backup. |
| -detail | Displays all the devices from the backup file with all their attribute values. |
| -file | Name of the backup configuration file. Either absolute path or relative path can be given. |
| -force | In non interactive mode, restoration of a device that has not been successfully validated, will be tried. This option cannot be used in combination with -inter or -validate. |
| -frequency | Frequency of the backup to run automatically. |
| -inter | Interactively deploy each device with user confirmation. |
| -list | Lists the backup files either from default directory or from user given directory. |
| -mapping | Displays mapping information for SPs, FBSP, repositories, etherchannel, SEA, vscsi adapters, vfc adapters and paging devices. |
| -migrate | Migrates earlier cluster version of backup file to current version. New file will be created with _MIGRATED string appended to the given filename |
| -nobackup | Stops auto backups. |
| -numfiles | When backup runs automatically, this number indicates the maximum amount of backup files that can be saved. Oldest file is deleted for next cycle of backup. |
| -recoverdb | Recovers from SSP database corruption, either from backup file or from solidDB backup. |
| -repopvs | Space separated list of physical volumes to be used as repository disks, The given physical volume(s) cannot contain repository signature. |

| | |
|--------------|--|
| -restore | Takes backup file as input and brings the VIOS partition to the same state when the backup was taken. |
| -skipcluster | Restores all local devices, except cluster0. |
| -subfile | Node configuration file to be restored. This option needs to be used when the valid cluster repository exists on the disks. It cannot be used with '-repopvs' option. This option is ignored if the backup file is not a cluster backup. |
| -type | Only devices of given type are used for list or restore operation. |
| -validate | Validates the devices on the VIO Server against the devices listed in the backup file. |
| -view | Display the information of all virtual and logical devices from backup file. |
| -xmlvtds | Restores SSP mappings which are not in database, but are in backup xml file. This option is valid only when restoring a node using cluster backup file |

Example 5-2 shows how to take backup of a Virtual I/O Server.

Example 5-2 Virtual I/O Server Config backup

```
$ viosbr -backup -file vios1_backup
Backup of this node (vios1) successful
$
```

Example 5-3 shows the method to check the contents of a **viosbr** backup.

Example 5-3 Virtual I/O Server (VIOS) Config details

```
$ viosbr -view -file vios1_backup.tar.gz
```

Details in: vios1_backup

=====

Controllers:

=====

| Name | Phys Loc |
|---------|--------------------------------|
| ---- | ----- |
| iscsi0 | |
| sissas0 | U5802.001.0086848-P1-C1-T1 |
| sissas1 | U78A0.001.DNWHZWR-P1-T9 |
| sissas2 | U78A0.001.DNWHZWR-P1-C10-T1 |
| ent10 | U78A0.001.DNWHZWR-P1-C3-T3 |
| ent11 | U78A0.001.DNWHZWR-P1-C3-T4 |
| ent12 | U78A0.001.DNWHZWR-P1-C6-T1 |
| pager0 | U8233.E8B.061AA6P-V1-C16-L0-L0 |
| ent13 | U78A0.001.DNWHZWR-P1-C6-T2 |
| pager1 | U8233.E8B.061AA6P-V1-C17-L0-L0 |
| ent14 | U78A0.001.DNWHZWR-P1-C6-T3 |
| vasi0 | U8233.E8B.061AA6P-V1-C16 |
| pager2 | U8233.E8B.061AA6P-V1-C19-L0-L0 |
| ent15 | U78A0.001.DNWHZWR-P1-C6-T4 |
| vasi1 | U8233.E8B.061AA6P-V1-C17 |
| pager3 | U8233.E8B.061AA6P-V1-C20-L0-L0 |

| | |
|----------|----------------------------------|
| pager4 | U8233.E8B.061AA6P-V1-C7-L0-L0 |
| usbhc0 | U78A0.001.DNWHZWR-P1 |
| vasi2 | U8233.E8B.061AA6P-V1-C19 |
| ent17 | U8233.E8B.061AA6P-V1-C3-T1 |
| usbhc1 | U78A0.001.DNWHZWR-P1 |
| vasi3 | U8233.E8B.061AA6P-V1-C20 |
| ent18 | U8233.E8B.061AA6P-V1-C4-T1 |
| vasi4 | U8233.E8B.061AA6P-V1-C7 |
| usbhc2 | U78A0.001.DNWHZWR-P1 |
| vbsd0 | U8233.E8B.061AA6P-V1-C16-L0 |
| ent19 | U8233.E8B.061AA6P-V1-C5-T1 |
| vbsd1 | U8233.E8B.061AA6P-V1-C17-L0 |
| vbsd2 | U8233.E8B.061AA6P-V1-C19-L0 |
| sata0 | U5802.001.0086848-P1-C1-T1 |
| vbsd3 | U8233.E8B.061AA6P-V1-C20-L0 |
| sata1 | U78A0.001.DNWHZWR-P1-T9 |
| vbsd4 | U8233.E8B.061AA6P-V1-C7-L0 |
| lhea0 | U78A0.001.DNWHZWR-P1 |
| ibmvmc0 | U8233.E8B.061AA6P-V1-C2 |
| fcs0 | U5802.001.0086848-P1-C2-T1 |
| fcs1 | U5802.001.0086848-P1-C2-T2 |
| fcs2 | U5802.001.0086848-P1-C3-T1 |
| vt1 | U8233.E8B.061AA6P-V1-C10 |
| fcs3 | U5802.001.0086848-P1-C3-T2 |
| vt2 | U8233.E8B.061AA6P-V1-C12 |
| fcs4 | U5802.001.0086848-P1-C4-T1 |
| fcs5 | U5802.001.0086848-P1-C4-T2 |
| fcs6 | U5802.001.0086848-P1-C5-T1 |
| sfwcomm0 | U5802.001.0086848-P1-C2-T1-W0-L0 |
| fcs7 | U5802.001.0086848-P1-C5-T2 |
| sfwcomm1 | U5802.001.0086848-P1-C2-T2-W0-L0 |
| sfwcomm2 | U5802.001.0086848-P1-C3-T1-W0-L0 |
| sfwcomm3 | U5802.001.0086848-P1-C3-T2-W0-L0 |
| fscsi0 | U5802.001.0086848-P1-C2-T1 |
| sfwcomm4 | U5802.001.0086848-P1-C4-T1-W0-L0 |
| ent0 | U78A0.001.DNWHZWR-P1-C4-T1 |
| fscsi1 | U5802.001.0086848-P1-C2-T2 |
| sfwcomm5 | U5802.001.0086848-P1-C4-T2-W0-L0 |
| ent1 | U78A0.001.DNWHZWR-P1-C4-T2 |
| fscsi2 | U5802.001.0086848-P1-C3-T1 |
| sfwcomm6 | U5802.001.0086848-P1-C5-T1-W0-L0 |
| ent2 | U78A0.001.DNWHZWR-P1-C5-T1 |
| fscsi3 | U5802.001.0086848-P1-C3-T2 |
| sfwcomm7 | U5802.001.0086848-P1-C5-T2-W0-L0 |
| ent3 | U78A0.001.DNWHZWR-P1-C5-T2 |
| fscsi4 | U5802.001.0086848-P1-C4-T1 |
| sfwcomm8 | SAS |
| ent4 | U78A0.001.DNWHZWR-P1-C2-T1 |
| ent20 | U8233.E8B.061AA6P-V1-C6-T1 |
| fscsi5 | U5802.001.0086848-P1-C4-T2 |
| sfwcomm9 | SAS |
| ent5 | U78A0.001.DNWHZWR-P1-C2-T2 |
| fscsi6 | U5802.001.0086848-P1-C5-T1 |
| ent22 | U8233.E8B.061AA6P-V1-C22-T1 |
| ent6 | U78A0.001.DNWHZWR-P1-C2-T3 |

```

sas0          U5802.001.0086848-P1-C1-T1
fscsi7       U5802.001.0086848-P1-C5-T2
ent7         U78A0.001.DNWHZWR-P1-C2-T4
sfw0
sas1         U78A0.001.DNWHZWR-P1-T9
ent8         U78A0.001.DNWHZWR-P1-C3-T1
ent9         U78A0.001.DNWHZWR-P1-C3-T2
fcnet0       U5802.001.0086848-P1-C2-T1
fcnet1       U5802.001.0086848-P1-C2-T2
fcnet2       U5802.001.0086848-P1-C3-T1
fcnet3       U5802.001.0086848-P1-C3-T2
fcnet4       U5802.001.0086848-P1-C4-T1
fcnet5       U5802.001.0086848-P1-C4-T2
fcnet6       U5802.001.0086848-P1-C5-T1
fcnet7       U5802.001.0086848-P1-C5-T2

```

Physical Volumes:

=====

| Name | Phys Loc |
|---------|--|
| ---- | ----- |
| hdisk0 | U5802.001.0086848-P3-D2 |
| hdisk1 | U5802.001.0086848-P3-D3 |
| hdisk2 | U5802.001.0086848-P3-D4 |
| hdisk3 | U5802.001.0086848-P1-C1-T1-LFF0000-L0 |
| hdisk4 | U5802.001.0086848-P1-C1-T1-LFF0100-L0 |
| hdisk5 | U78A0.001.DNWHZWR-P2-D4 |
| hdisk6 | U78A0.001.DNWHZWR-P2-D5 |
| hdisk7 | U78A0.001.DNWHZWR-P2-D6 |
| hdisk8 | U78A0.001.DNWHZWR-P2-D7 |
| hdisk9 | U78A0.001.DNWHZWR-P1-T9-LFF0000-L0 |
| hdisk10 | U5802.001.0086848-P1-C3-T1-W201600A0B829AC12-L1000000000000 |
| hdisk11 | U5802.001.0086848-P1-C3-T1-W201600A0B829AC12-L20000000000000 |

Optical Devices:

=====

```

cd0          U78A0.001.DNWHZWR-P2-D2

```

Tape Devices:

=====

```

rmt0        U78A0.001.DNWHZWR-P2-D1

```

Ethernet Interfaces:

=====

```

Name
----
en10
en11
en12
en13
en14
en15
en16
en17
en18
en19

```


en0
 en20
 en1
 en22
 en2
 en3
 en4
 en5
 en6
 en7
 en8
 en9

Storage Pools:
 =====

| SP Name | PV Name |
|-----------|------------------|
| ----- | ----- |
| rootvg | hdisk0 |
| ams_sp | hdisk1 hdisk2 |
| datapool1 | hdisk3 hdisk4 |
| datapool2 | hdisk5 hdisk6 |

File Backed Storage Pools:
 =====

Optical Repository:
 =====

Ether Channel Adapters:
 =====

Shared Ethernet Adapters:
 =====

| Name | Physical Adapter | Default Adapter | Virtual Adapters |
|-------|------------------|-----------------|------------------|
| ---- | ----- | ----- | ----- |
| ent16 | ent12 | ent17 | ent17 |

Virtual Server Adapters:
 =====

| SVSA | Phys Loc | VTD |
|--------|--------------------------|---------|
| ---- | ----- | --- |
| vhost1 | U8233.E8B.061AA6P-V1-C11 | |
| vhost2 | U8233.E8B.061AA6P-V1-C13 | vtscsi2 |

Virtual Suspended Adapters:
 =====

Virtual Suspended Devices:
 =====

| Suspend Device | Phys Loc |
|----------------|----------|
|----------------|----------|

```

-----
suspended0

SVFC Adapters:
=====
Name                FC Adapter          Phys Loc
----                -
vfchost1            fcs2                U8233.E8B.061AA6P-V1-C21

VBSD Pools:
=====
Name
----
pool0

VRM(AMS) Pages:
=====
Name                StreamID
----                -
vrmpage0            0x100017d536f5c

RSCT:
=====
ct_node_id
-----
471df119a2752e29

Cluster:
=====
Cluster             State
-----             -

Virtual Log Repositories:
=====
Virtual Log Repository  State
-----             -

vlogrepo0          AVAILABLE
$

```

viosbr backups: By default, **viosbr** backups are stored in `/home/padmin/cfgbackups`. You can use a different location if you want. You can even specify Network File Systems (NFSs) as a path.

5.1.5 IVM updates

IBM periodically provides updates (or fix packs) for the Virtual I/O Server (VIOS). These fix packs can be downloaded from the Virtual I/O Server support web page:

<http://www14.software.ibm.com/support/customer/sas/f/vios/home.html>

Updates are necessary whenever new functions or fixes are introduced.

Fix packs can cause IVM to restart: Applying a fix pack can cause the restart of the IVM. That means that all LPARs must be stopped during this reboot.

Determining the current Virtual I/O Server level

By running the `ioslevel` command from the Virtual I/O Server command line, the padmin user can determine the actual installed level of the Virtual I/O Server software (Example 5-4).

Example 5-4 Using the ioslevel command

```
$ ioslevel  
2.2.1.4
```

In the example, the level of the Virtual I/O Server software is 2.2.1.4.

You can refer to the Fix Central website (<http://www.ibm.com/support/fixcentral>) to confirm the latest available version of the Virtual I/O Server software level. You must select some options in the website to check for the latest fix level that is available.

In the Fix Central website, select **Virtualization software** in the **Product Group** drop-down menu. Then, select **PowerVM Virtual I/O Server** in the **Select from Virtualization software** drop-down menu. In the **Installed Version** drop-down menu, select your Virtual I/O Server version. Click **Continue**. Figure 5-6 shows an example.

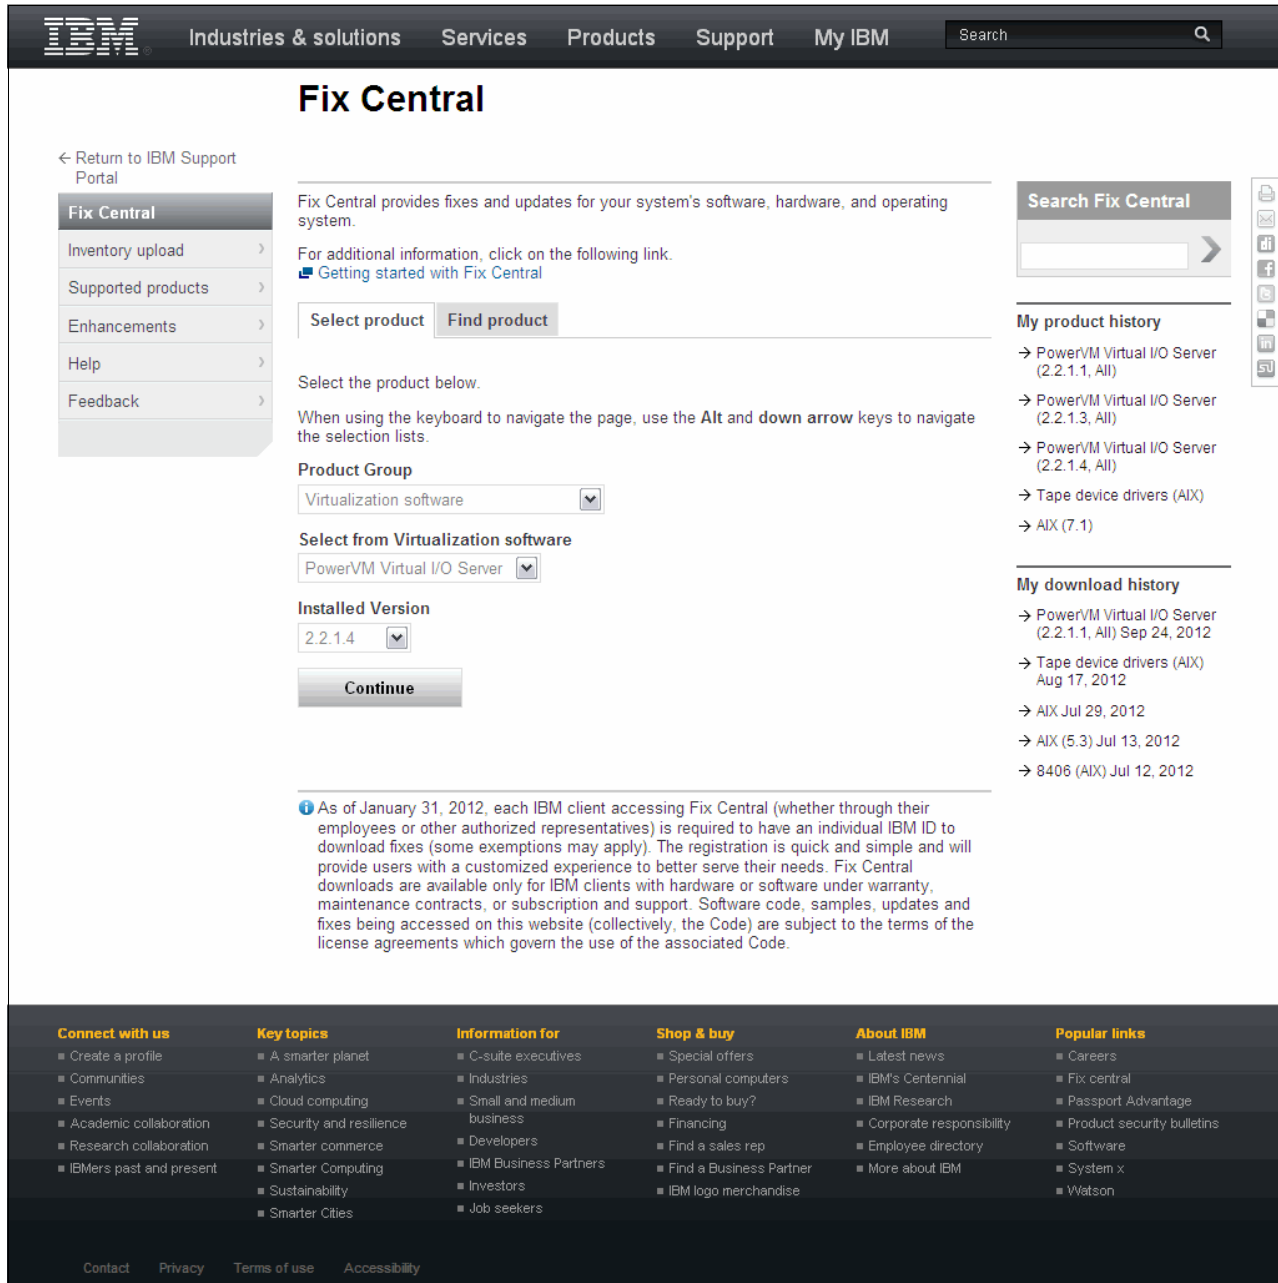


Figure 5-6 Fix Central website

The **Select fixes** page opens, and shows the latest available fixes (fix packs, service packs, and interim fixes) that are available for the selected version. Figure 5-7 shows an example.

The screenshot shows the IBM Fix Central 'Select fixes' page for 'Virtualization software, PowerVM Virtual I/O Server (2.2.1.4, All platforms)'. The page is divided into several sections:

- Navigation:** Includes 'Return to IBM Support Portal', 'Fix Central' menu, and 'Change your selection' section with 'Product selector' (PowerVM Virtual I/O Server) and 'Installed Version' (2.2.1.4).
- Download options:** Shows 'Download method: HTTP' and 'Include requisites: Yes'.
- Select fixes:** A list of two interim fixes:
 - Interim fix: VIOS_2.2.1.4_IV20593s01.120511** (dated Jun 8, 2012). Includes a notice: 'Interim Fix IV20593s01 is required to be installed if both conditions below are true: 1) IBMi or Linux Client partitions are using virtualized USB tape devices [Feature Codes EU16 and 5673] 2) VIOS partitions were updated to VIOS release 2.2.1.4 via the VIOS DVD Media. Possible data error when a USB tape device is exported as a virtual device from a VIOS to either an IBMi or Linux client partition. AIX partitions are not affected. This intermittent problem can cause incorrect data on the tape device, either on the write path (incorrect data is written to the tape) or the read path (the data is correct on the tape, but incorrect data is returned to the client partition).' A 'Readme' link is provided.
 - Interim fix: VIOS_2.2.1.4_IC83004_IZ94423_Fix** (dated May 24, 2012). A 'Readme' link is provided.
- Filter your content:** Includes filters for Component (5765G3400), And Platform (VIOS), And Fix type (Interim fix), and And Applies to (2.2.1.4).
- Actions:** 'Continue', 'Clear selections', and 'Back' buttons are present at the bottom of each results section.
- Footer:** Contains navigation links like 'Connect with us', 'Key topics', 'Information for', 'Shop & buy', 'About IBM', and 'Popular links'.

Figure 5-7 Fixes list from the Fix Central website

All Virtual I/O Server fix packs are cumulative and contain all fixes from previous fix packs. For service packs, you must have previously installed the related fix pack.

We suggest that you keep your server always updated with the latest available system firmware. See 2.3, “Microcode update” on page 23 for more information about microcode updates.

It is also a good practice to look for your server in the Fix Level Recommendation Tool website: <http://www14.software.ibm.com/webapp/set2/flrt/home>

Important: Ensure that you have the correct level of firmware before you update the IVM.

All interim fixes that are applied to the Virtual I/O Server must be manually removed before applying a fix pack or service pack. Virtual I/O Server clients who applied interim fixes to the Virtual I/O Server can use the following procedure to remove them before applying fix or service packs. Example 5-5 shows how to list fixes from the VIOS command line.

Example 5-5 Listing fixes

```
$ oem_setup_env
# emgr -P /*gives a list of the installed efix's (by label)
There is no efix data on this system.
# exit
$
```

Message for no fixes installed: Because there are no fix packs installed in our server, the output from the `emgr -P` command shows this message, which is shown in Example 5-5.

Downloading the fix packs

You can download the fix packs by using four different options:

- ▶ Download Director
- ▶ FTP
- ▶ HTTP
- ▶ ISO image

It is possible to choose the download options in the Fix Central website. After selecting the fixes that you want to download in the Select fixes page, a list of options is displayed, similar to what is shown in Figure 5-8. **Select your preferred download option.**

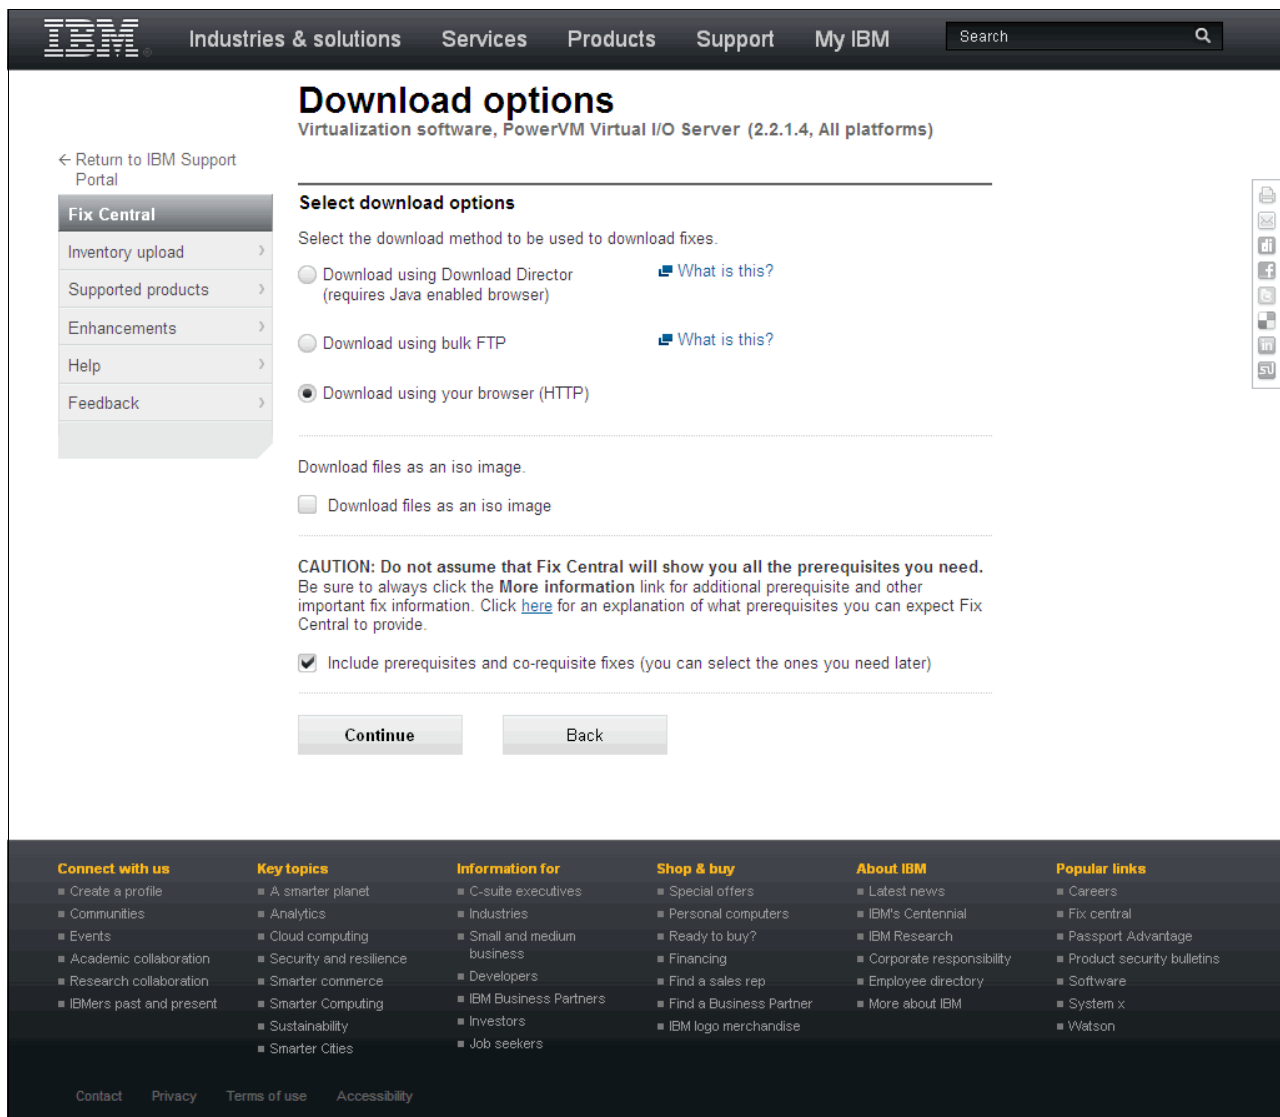


Figure 5-8 Fix Central download options

For the first three download options, all file sets are downloaded into a user-specified directory. When the download completes, the updates can be applied from a directory on your local hard disk. Follow these steps to install the fix pack:

1. Log in to the Virtual I/O Server as the user `padmi n`.
2. Create a directory on the Virtual I/O Server.

```
$ mkdir directory_name
```
3. Using the `ftp` command, transfer the update files to the directory that you created.
4. Apply the update by running the `updateios` command:

```
$ updateios -dev directory_name -install -accept
```

Accept to continue the installation after the preview update is run.
5. Reboot

Verify a successful update by checking the results of the **updateios** command and running the **ioslevel** command. The result of the **ioslevel** command should equal the level of the downloaded package.

Applying updates from a remotely mounted file system

If the remote file system is to be mounted as *read-only*, you must first rename the fix pack file `tableofcontents.txt` to `.toc`; otherwise, you are prevented from installing this fix pack. Use the following steps to apply updates from a remotely mounted file system:

1. Log in to the Virtual I/O Server as user `padmin`.
2. Mount the remote directory onto the Virtual I/O Server:

```
$ mount remote_machine_name:directory /mnt
```

3. Apply the update by running the **updateios** command:

```
$ updateios -dev /mnt -install -accept
```

4. If prompted to remove the `.toc?`, enter `no`.

Verify a successful update by checking the results of the **updateios** command and running the **ioslevel** command. The result of the **ioslevel** command should equal the level of the downloaded package.

Applying updates from the ROM drive

This fix pack can be burned onto a CD by using ISO image files. After the CD is created, do the following steps to apply the update:

1. Log in to the Virtual I/O Server as user `padmin`.
2. Place the update CD into the drive.
3. Apply the update by running the **updateios** command:

```
$ updateios -dev /dev/cdX -install -accept
```

(where X is a device number between 0 and N)

4. Verify a successful update by checking the results of the **updateios** command and running the **ioslevel** command. The result of **ioslevel** command should equal the level of the downloaded package.

Updating from an ioslevel before 1.3.0.0: If you update from an ioslevel before 1.3.0.0, the **updateios** command might indicate several failures (such as missing requisites) while you install the fix pack. This notice is expected. Proceed with the update if you are prompted to Continue with the installation [y/n].

5.1.6 The migration between HMC and IVM

It is important to note that moving between the Hardware Management Console (HMC) and an Integrated Virtualization Manager (IVM) environment requires a certain amount of reconfiguration.

Virtual I/O Server installation: There is no guarantee onto which disk the Virtual I/O Server will install. If the installation takes place on a disk that contains your client volume group, you lose the data and are not able to import it again.

Have a backup of the Virtual I/O Server and virtual I/O clients, as well as profiles for system recovery, before you attempt any migration.

A simple method might be the physical removal of any disks that you want to save, while you do the installation. Then, put the disks back in after installation.

Always make a backup of your environment before you migrate between the HMC and IVM environment.

5.1.7 Recovery after an improper HMC connection

An HMC must not be connected to a system that is running the IVM; otherwise, you cannot do any other operation on the IVM.

If an HMC was connected to a system using the IVM, the following steps explain how to re-enable the IVM capabilities:

1. Power off the system.
2. Remove the system definition from the HMC.
3. Unplug the HMC network cable from the system if directly connected.
4. Connect a TTY console emulator with a serial cross-over cable to one of the system's serial ports.
5. Press any key on the console to open the service processor prompt.
6. Log in as the user `admin` and answer the questions about the number of lines and columns.
7. Reset the service processor.

Type 2 to select **2. System Service Aids**, type 10 to select **10. Reset Service Processor**, and then type 1 to confirm your selection. Wait for the system to reboot.
8. Reset it to the factory configuration (Manufacturing Default Configuration).

Type 2 to select **2. System Service Aids**, type 11 to select **11. Factory Configuration**, then type 1 to select **1. Reset all settings**, and then type 1 to confirm. Wait for the system to reboot.
9. Configure the ASMI IP addresses if needed.

Type 5 to select **5. Network Services**, type 1 to select **1. Network Configuration**, select the wanted IP version (type 1 to select **1. IPv4**, or 2 to select **2. IPV6**), and then configure each Ethernet adapter. For more information, refer to 2.4, "Advanced System Management Interface IP address setup" on page 26.
10. Start the system.

Type 1 to select **1. Power/Restart Control**, type 1 to select **1. Power On/Off System**, type 6 to select **6. Power on**, and press **Enter** to confirm your selection.
11. Go to the system management services (SMS) menu.
12. Update the boot list.

Type 5 to select **5. Select Boot Options**, type 2 to select **2. Configure Boot Device Order**, and select the IVM boot disk.

13. Boot the system.
14. Wait for the IVM to start.
15. Connect to the IVM with the GUI.
16. Restore the partition configuration using the last backup file.

From the Service Management menu in the navigation area, click **Backup/Restore**, and then click **Restore Partition Configuration** in the work area. For more information, refer to 5.1.1, “Backup and restore of the logical partition definitions” on page 152.

This operation updates only the IVM partition configuration and does not restore the LPARs hosted by the IVM.

17. Reboot the IVM. (If changes do not require a reboot, then recovery of the IVM should be done immediately.)
18. Restore the partition configuration using the last backup file.
This time, each LPAR definition is restored.
19. Reboot the IVM.
This reboot is needed to make each virtual device available to the LPARs. (This is also possible by issuing the `cfgdev` command.)
20. Restart each LPAR.

5.1.8 Migration considerations

The following list contains the minimum considerations to migrate between an HMC and the IVM. For a production redeployment, it depends on the configuration of the system:

- ▶ Virtual I/O Server version
- ▶ System firmware level
- ▶ Virtual I/O Server I/O device configuration
- ▶ Backup Virtual I/O Server, virtual I/O clients profile, and virtual I/O devices
- ▶ The mapping information between physical and virtual I/O devices
- ▶ Virtual I/O Server and VIO client backups

Virtual I/O Server version

The `ioslevel` command displays the Virtual I/O Server version. You see output that is similar to the following display:

```
$ ioslevel  
2.2.1.4
```

System firmware level

You can display the system firmware level by using the `lsfware` command. You see output that is similar to the following display:

```
$ lsfware  
system:AL730_087 (t) AL730_087 (p) AL730_087 (t)
```

Virtual I/O Server I/O device configuration

To display I/O devices such as an adapter, disk, or slots, use the `lsdev` command. See Example 5-6 for an example.

Example 5-6 Virtual I/O Server device information

```
$ lsdev -type adapter
name          status      description
ent0          Available  2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent1          Available  2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent2          Available  2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent3          Available  2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent4          Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent5          Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent6          Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent7          Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent8          Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent9          Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent10         Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent11         Available  4-Port 10/100/1000 Base-TX PCI-Express Adapter
(14106803)
ent12         Available  Logical Host Ethernet Port (lp-hea)
ent13         Available  Logical Host Ethernet Port (lp-hea)
ent15         Available  Logical Host Ethernet Port (lp-hea)
ent16         Available  10 Gigabit Ethernet Adapter (ct3)
ent17         Available  Virtual I/O Ethernet Adapter (1-lan)
ent18         Available  Virtual I/O Ethernet Adapter (1-lan)
ent19         Available  Virtual I/O Ethernet Adapter (1-lan)
ent20         Available  Virtual I/O Ethernet Adapter (1-lan)
ent21         Available  Shared Ethernet Adapter
ent22         Available  Virtual I/O Ethernet Adapter (1-lan)
fcs0          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs1          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs2          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs3          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs4          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs5          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs6          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs7          Available  8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
hba0          Available  10 Gigabit Ethernet-SR PCI-Express Host Bus Adapter
(2514300014108c03)
```

```

ibmvmc0      Available  Virtual Management Channel
lhea0       Available  Logical Host Ethernet Adapter (l-hea)
pager0      Available  Pager Kernel Extension
pager1      Available  Pager Kernel Extension
pager2      Available  Pager Kernel Extension
pager3      Available  Pager Kernel Extension
pager4      Available  Pager Kernel Extension
sissas0     Available  PCI Express x8 Ext Dual-x4 3Gb SAS Adapter
sissas1     Available  PCI-X266 Planar 3Gb SAS RAID Adapter
sissas2     Available  PCI Express x1 Auxiliary Cache Adapter
usbhc0      Available  USB Host Controller (33103500)
usbhc1      Available  USB Host Controller (33103500)
usbhc2      Available  USB Enhanced Host Controller (3310e000)
vasi0       Available  Virtual Asynchronous Services Interface (VASI)
vasi1       Available  Virtual Asynchronous Services Interface (VASI)
vasi2       Available  Virtual Asynchronous Services Interface (VASI)
vasi3       Available  Virtual Asynchronous Services Interface (VASI)
vasi4       Available  Virtual Asynchronous Services Interface (VASI)
vbsd0       Available  Virtual Block Storage Device (VBSD)
vbsd1       Available  Virtual Block Storage Device (VBSD)
vbsd2       Available  Virtual Block Storage Device (VBSD)
vbsd3       Available  Virtual Block Storage Device (VBSD)
vbsd4       Available  Virtual Block Storage Device (VBSD)
vfchost0    Available  Virtual FC Server Adapter
vfchost1    Available  Virtual FC Server Adapter
vhost0      Available  Virtual SCSI Server Adapter
vhost1      Available  Virtual SCSI Server Adapter
vhost2      Available  Virtual SCSI Server Adapter
vsa0        Available  LPAR Virtual Serial Adapter
vsa1        Available  LPAR Virtual Serial Adapter
vts0        Available  Virtual TTY Server Adapter
vts1        Available  Virtual TTY Server Adapter
vts2        Available  Virtual TTY Server Adapter

```

```
$ lsvg -lv datapool1
```

```
datapool1:
```

| LV NAME | TYPE | LPs | PPs | PVs | LV STATE | MOUNT POINT |
|-------------|------|------|------|-----|--------------|-------------|
| lpar3rootvg | jfs | 1024 | 1024 | 1 | closed/syncd | N/A |

If you want to display the attribute of each device, use the **lsdev -dev Devicename -attr** command. And you can use the **lsdev -slots** command for the slot information and **lsdev -dev Devicename -child** command for the child devices that are associated with these devices.

Also, you can use the **lsvg -lv volumegroup_name** command to discover system disk configuration and volume group information.

Tip: Note the physical location code of the disk unit that you are using to boot the Virtual I/O Server. To display this code, use the **lsldev -dev Devicename -vpd** command.

To migrate from an HMC to an IVM environment, the Virtual I/O Server must own all of the physical devices. If the server does not, all of the physical devices are assigned to the Virtual I/O Server when it is managed by the IVM.

Backup of Virtual I/O Server and client profiles

Document the information in the virtual I/O clients which have a dependency on the virtual SCSI server and the virtual SCSI client adapter, as shown in Figure 5-9.

Partition Properties - vios1

General Hardware **Virtual Adapters** Settings Other

Actions ▼

Virtual resources allow for the sharing of physical hardware between logical partitions. The current virtual adapter settings are listed below.

Maximum virtual adapters : *200

Number of virtual adapters : 17

| Select | Type | Adapter ID | Server/Client P... | Partner Adapter | Required |
|--------------------------|----------------------|------------|--------------------|--------------------|----------|
| <input type="checkbox"/> | Ethernet | 22 | N/A | N/A | No |
| <input type="checkbox"/> | Ethernet | 3 | N/A | N/A | No |
| <input type="checkbox"/> | Ethernet | 4 | N/A | N/A | No |
| <input type="checkbox"/> | Ethernet | 5 | N/A | N/A | No |
| <input type="checkbox"/> | Ethernet | 6 | N/A | N/A | No |
| <input type="checkbox"/> | Server Fibre Channel | 21 | lpar2_server1(3) | 5 | No |
| <input type="checkbox"/> | Server SCSI | 11 | lpar2_server1(3) | 2 | No |
| <input type="checkbox"/> | Server SCSI | 13 | lpar3_server1(4) | 2 | No |
| <input type="checkbox"/> | Client Serial | 10 | lpar2_server1(3) | 0 | No |
| <input type="checkbox"/> | Client Serial | 12 | lpar3_server1(4) | 0 | No |
| <input type="checkbox"/> | Server Serial | 0 | Any Partition | Any Partition Slot | Yes |
| <input type="checkbox"/> | Server Serial | 1 | Any Partition | Any Partition Slot | Yes |
| <input type="checkbox"/> | VASI | 16 | N/A | N/A | N/A |
| <input type="checkbox"/> | VASI | 17 | N/A | N/A | N/A |
| <input type="checkbox"/> | VASI | 19 | N/A | N/A | N/A |
| <input type="checkbox"/> | VASI | 20 | N/A | N/A | N/A |
| <input type="checkbox"/> | VASI | Z | N/A | N/A | N/A |

Total: 17 Filtered: 17 Selected: 0

OK Cancel Help

Figure 5-9 Virtual SCSI: Client adapter properties

The mapping information between physical and virtual I/O devices

To display the mapping information between the physical I/O devices and virtual I/O devices such as disk, network, and optical media, use the `lsmap -vadaptor vhost#` command. Refer to Example 5-7 for an example.

Example 5-7 Mapping information between physical I/O devices and virtual I/O devices

```
$ lsmap -all
SVSA          Physloc          Client Partition ID
-----
vhost1        U9409.M50.1059C72-V1-C11  0x00000003

VTD          vtscsi0
Status       Available
LUN          0x8200000000000000
Backing device  hdisk1
Physloc      U78A0.001.DNWGPL0-P2-D4
Mirrored     false

SVSA          Physloc          Client Partition ID
-----
vhost2        U9409.M50.1059C72-V1-C13  0x00000004

VTD          vtscsi1
```

| | |
|----------------|---|
| Status | Available |
| LUN | 0x8300000000000000 |
| Backing device | hdisk4 |
| Physloc | U78A0.001.DNWGPL0-P1-C3-T1-W201600A0B829AC12-L2000000000000 |
| Mirrored | false |

Virtual I/O Server and VIO client backups

Before the migration from an HMC to an IVM environment, it is necessary to back up the Virtual I/O Server and virtual I/O client. For more information about backup, refer to 5.1.2, “Backup and restore of the IVM operating system” on page 155.

5.1.9 Migration from an HMC to an IVM environment

For redeployment from HMC to IVM, the managed system must be reset to the Manufacturing Default Configuration by using the ASMI menu function.

Tip: The suggested method is a *complete reinstallation*. The only reason to save the virtual I/O installation is if there is client data on the rootvg (not recommended; use different storage pools for client data). If the client data is on other VGs (volume groups), export the VG and remove the disk to ensure that it does not get installed over. You need to reconfigure all of your device mappings, and so on. The outcome is that this type of installation is more complex and time-consuming than starting with a “fresh” installation.

This migration has the following requirements:

- ▶ Virtual I/O Server of the HMC-managed environment owns all the physical I/O devices
- ▶ Backup of a Virtual I/O Server and a virtual I/O client
- ▶ Virtual I/O Server with version 1.2, or above
- ▶ System firmware level with SF230_120, or above

Figure 5-10 shows the general migration procedure from an HMC to an IVM environment. There is some dependency on system configuration.

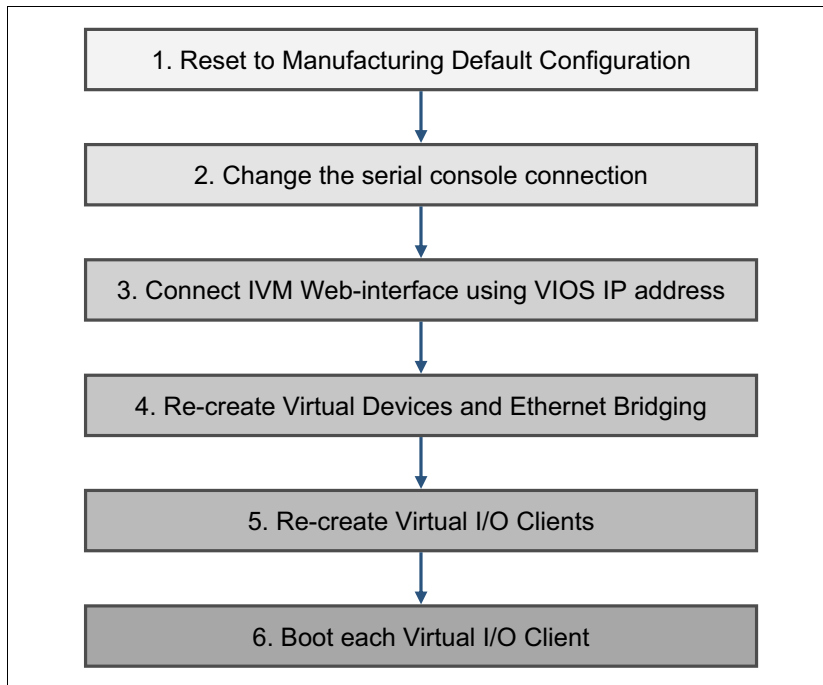


Figure 5-10 General migration procedure from an HMC to an IVM environment

The following steps, as shown in Figure 5-10, describe the general migration procedure from an HMC to an IVM environment:

1. Reset to the *Manufacturing Default Configuration*

If you decide to do this migration, it is necessary to restore the firmware setting, network configuration, and passwords to their factory defaults. When you reset the firmware, it removes all partition configuration and any personalization that was made to the service processor. A default full system partition is created to handle all the hardware resources. Without an HMC, the system console is provided through the internal serial ports. Connections are made using a serial ASCII console and cross-over cable that is connected to the serial port.

If you perform the firmware reset after detaching the HMC, the HMC retains information about the server as a managed system. You can remove this by using the HMC GUI.

When a console session is opened to the reset server, at the first menu, select **1.Power/Restart Control** → **1.Power On/Off system**, as shown in Example 5-8.

Example 5-8 Power On/Off System

```
Power On/Off System
Current system power state: Off
Current firmware boot side: Temporary
Current system server firmware state: Not running
```

1. System diagnostic level for the next boot
 Currently: Normal
2. Firmware boot side for the next boot
 Currently: Temporary
3. System operating mode

- Currently: Normal
- 4. Server firmware start policy
 - Currently: Running (Auto-Start Always)
- 5. System power off policy
 - Currently: Automatic
- 6. Power on
- 7. Default Partition Environment
 - Currently: AIX
- 98. Return to previous menu
- 99. Log out

S1>

Example 5-8 on page 175 shows that the *Power on* menu is 6. This means that the firmware reset is not done and the system is still managed by an HMC. If the firmware reset is done and the system is no longer managed by an HMC, then the *Power on* menu is 8. You can reset the service processor or put the server back to factory configuration through the System Service Aids menu in ASMI.

2. Change the serial connection for IVM

When you change the management system from HMC to IVM, you can no longer use the default console connection through vty0. You change the console connection, as shown in Example 5-9. This change is effective after the Virtual I/O Server reboot, and you will change the physical serial connection from SPC1 to SPC2 for using the vty1 console connection.

Example 5-9 Serial connection change for IVM

```
# lscons
NULL

# lsdev -Cc tty
vty0 Available Asynchronous Terminal

# lsdev -Cl vty0 -F parent
vsa0

# lsdev -Cl vsa0
vsa0 Available LPAR Virtual Serial Adapter

# chcons /dev/vty0
chcons: console assigned to: /dev/vty0, effective on next system boot
```

3. Connect the IVM web-interface by using the Virtual I/O Server IP address

The first web-interface pane that opens after the login process is View/Modify Partitions, as shown in Figure 5-11. You can see only a Virtual I/O Server partition. IVM does not have any information about other virtual I/O clients because the service process is reset to the *Manufacturing Default Configuration*.

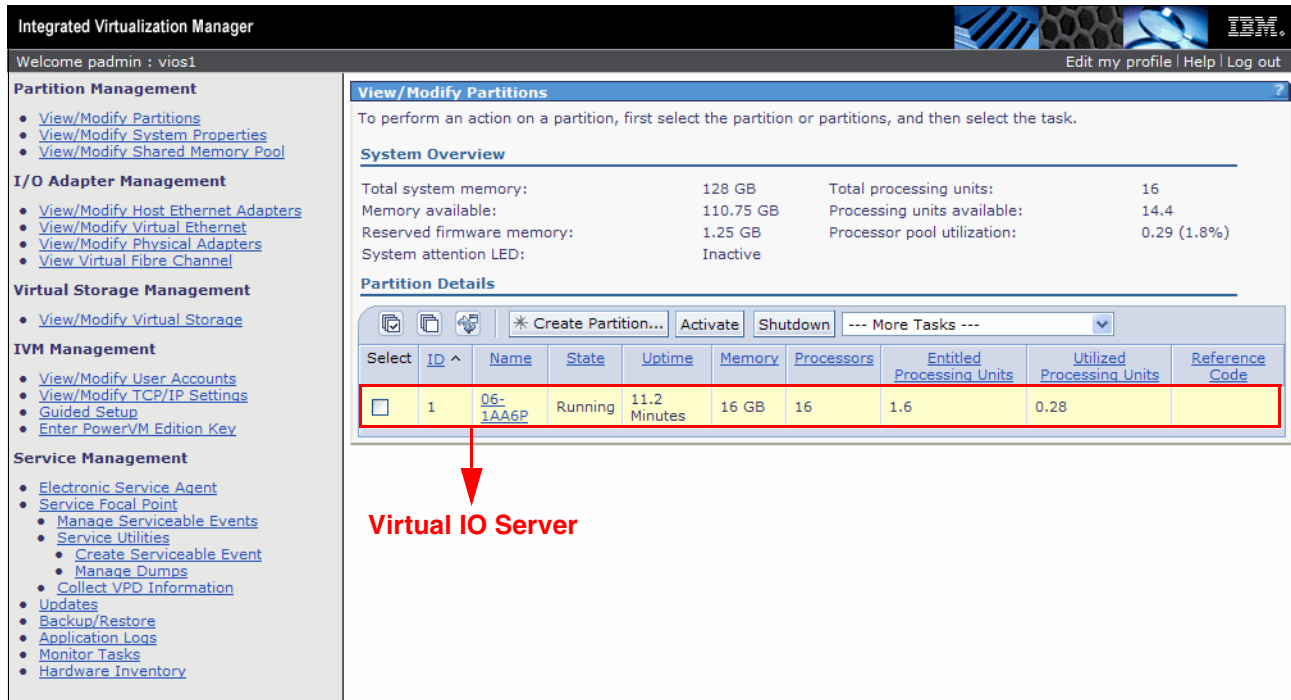


Figure 5-11 View/Modify Partitions

4. Re-create virtual devices and Ethernet bridging

When changed to an IVM environment, the Virtual I/O Server (now *Management Partition*) still has virtual device information left over from the HMC environment. There is the virtual SCSI, virtual Ethernet, shared Ethernet, and virtual target device information, but their status is changed to *defined* after migrating to an IVM environment.

Because these virtual devices no longer exist, remove them before creating the virtual I/O clients in IVM. You can remove the virtual devices, as shown in Example 5-10.

If you define virtual disks for clients from the Management Partition, the virtual SCSI server and client devices are created automatically for you.

Example 5-10 Remove the virtual device

```
$ rmdev -dev vhost1 -recursive
vtscsi0 deleted
vhost1 deleted
$ rmdev -dev vhost2 -recursive
vtscsi1 deleted
vhost2 deleted
$ rmdev -dev vfchost1 -recursive
vfchost1 deleted
$ rmdev -dev vfchost1 -recursive
vfchost1 deleted
$ rmdev -dev ent16
ent16 deleted
```

```
$ rmdev -dev en16
en16 deleted
$ rmdev -dev et16
et16 deleted
$ rmdev -dev ent17
ent17 deleted
$ rmdev -dev en17
en17 deleted
$ rmdev -dev et17
et17 deleted
$ rmdev -dev ent18
ent18 deleted
$ rmdev -dev en18
en18 deleted
$ rmdev -dev et18
et18 deleted
$ rmdev -dev ent19
ent19 deleted
$ rmdev -dev en19
en19 deleted
$ rmdev -dev et19
et19 deleted
$ rmdev -dev ent20
ent20 deleted
$ rmdev -dev en20
en20 deleted
$ rmdev -dev et20
et20 deleted
$ rmdev -dev ent22
ent22 deleted
$ rmdev -dev en22
en22 deleted
$ rmdev -dev et22
et22 deleted
```

After you remove the virtual devices, you can re-create virtual devices by using the `cfgdev` command. Or, you can re-create the devices through the IVM GUI and the virtual Ethernet bridge for virtual I/O clients in the **View/Modify Virtual Ethernet** panel, by following these steps:

- a. In the **View/Modify Virtual Ethernet** panel, click **Initialize Virtual Ethernet**, as shown in Figure 5-12.




Figure 5-12 Initialize Virtual Ethernet

b. Then, create the **Virtual Ethernet Bridge**, as shown in Figure 5-13. Click **Apply**.

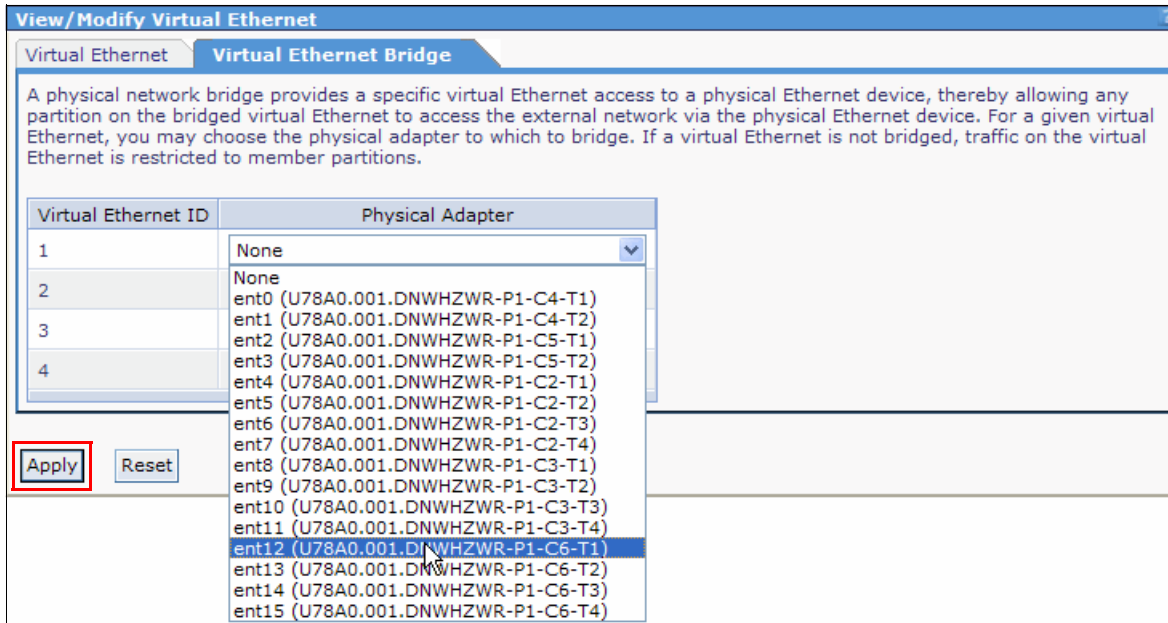


Figure 5-13 Creating a Virtual Ethernet Bridge

5. Re-create the virtual I/O clients

Because the IVM does not have information for the virtual I/O clients, you must re-create them by using the IVM web interface. For more information about creating LPARs, refer to 4.2, “IVM graphical user interface” on page 50.

When you choose **Storage Type**, select **Assign existing virtual disks and physical volumes**, as shown in Figure 5-14.

You can also allow the IVM to create a virtual disk for you by selecting **Create virtual disk**, when needed.

Volume group export: Export any volume group that contains client data by using the `exportvg` command. After migrating, import the volume groups by using the `importvg` command. This is a more efficient method to migrate the client data without loss.

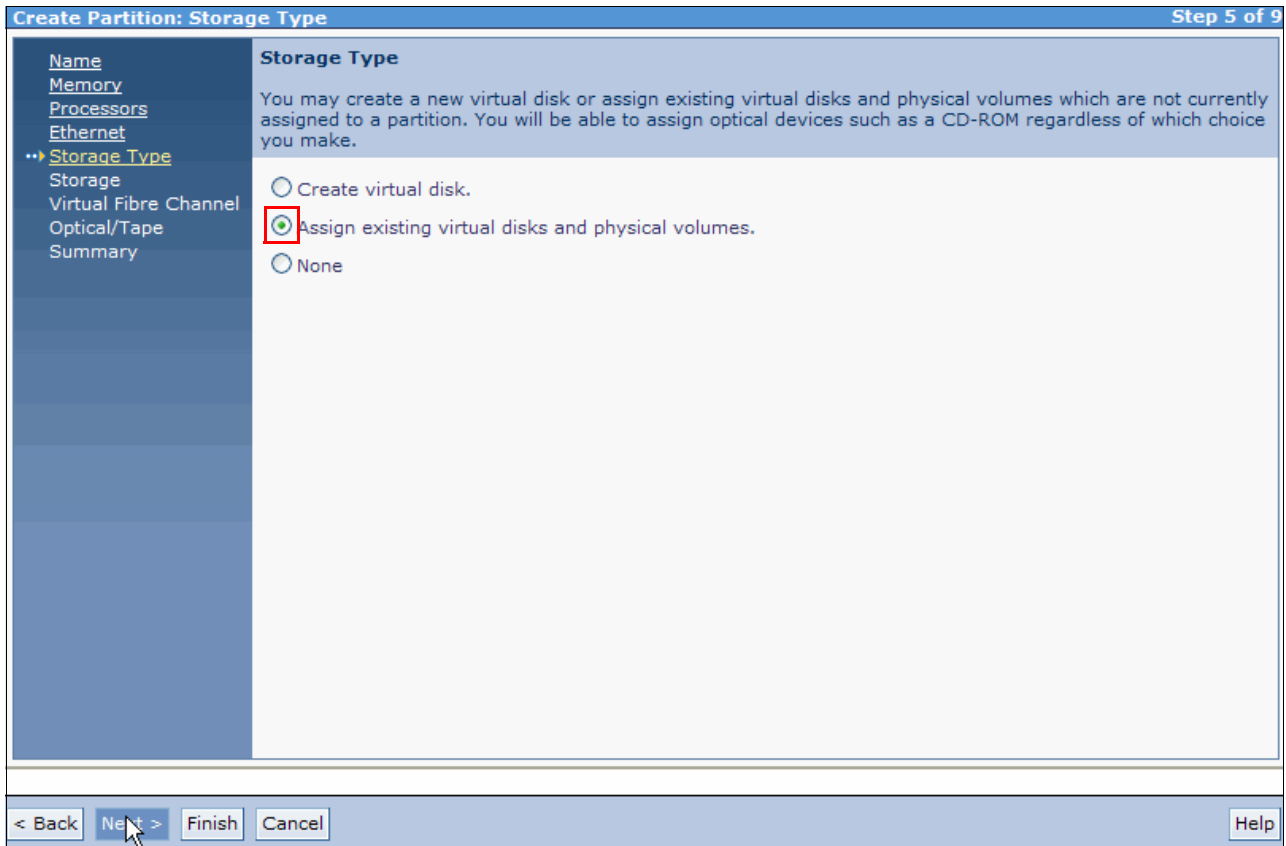


Figure 5-14 Create LPAR: Storage Type

6. Boot each virtual I/O client

5.1.10 Migration from an IVM environment to an HMC

There is no officially announced procedure to migrate from an IVM environment to the HMC.

If an HMC is connected to the system, the IVM interface is disabled immediately, effectively making it just a Virtual I/O Server partition. The managed system goes into recovery mode. After recovery completes, the HMC shows all of the LPARs without a profile. You must create one profile for each LPAR.

Record configuration information: Ensure that you carefully record all configuration information before you do this migration.

This can be migrated more easily compared to the reverse. The IVM environment must own all of the physical devices and there can be only one *Management Partition* per server, so there are no restrictions on server configuration that could affect a possible migration.

Figure 5-15 shows the general migration procedure from an IVM environment to the HMC.

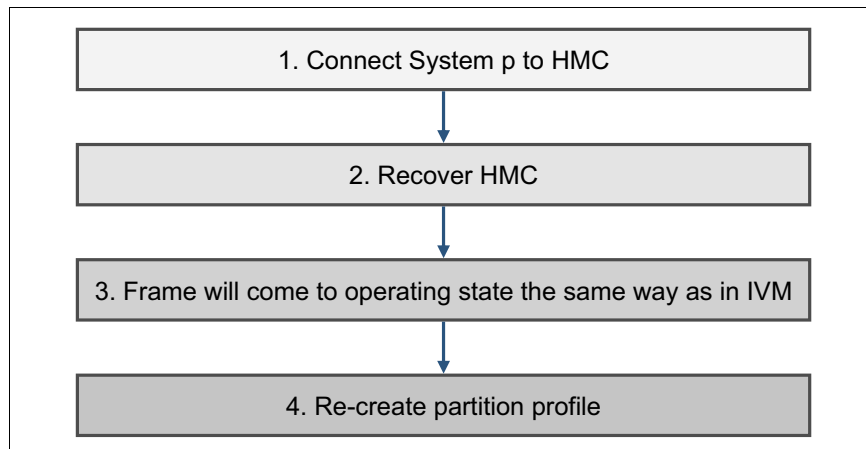


Figure 5-15 General migration procedure from an IVM environment to the HMC

The following steps, as shown in Figure 5-15 on page 181, describe the general migration procedure from an IVM environment to the HMC:

1. Connect the Power Systems server to an HMC

The server is connected and recognized by the HMC, and the IVM interface is disabled immediately, effectively making it just a Virtual I/O Server partition, as shown in Figure 5-16.

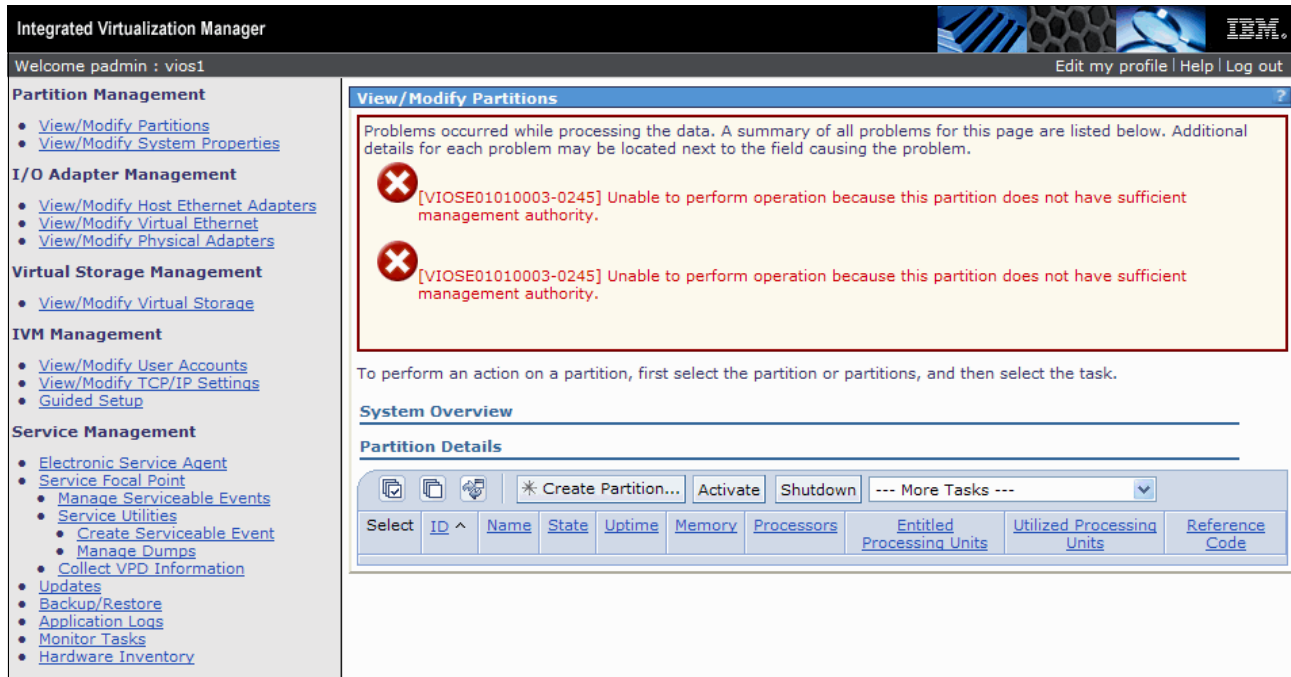


Figure 5-16 IVM management after connecting the HMC

2. Recover server configuration data to the HMC

Add the managed system to the HMC, then the managed system goes into Recovery mode, as shown in Figure 5-17. Right-click the managed system, then select **Recover Partition Data** → **Restore profile data from HMC backup data**. Ensure that at least one of the LPARs is up and running; otherwise, the HMC might delete all of the LPARs.

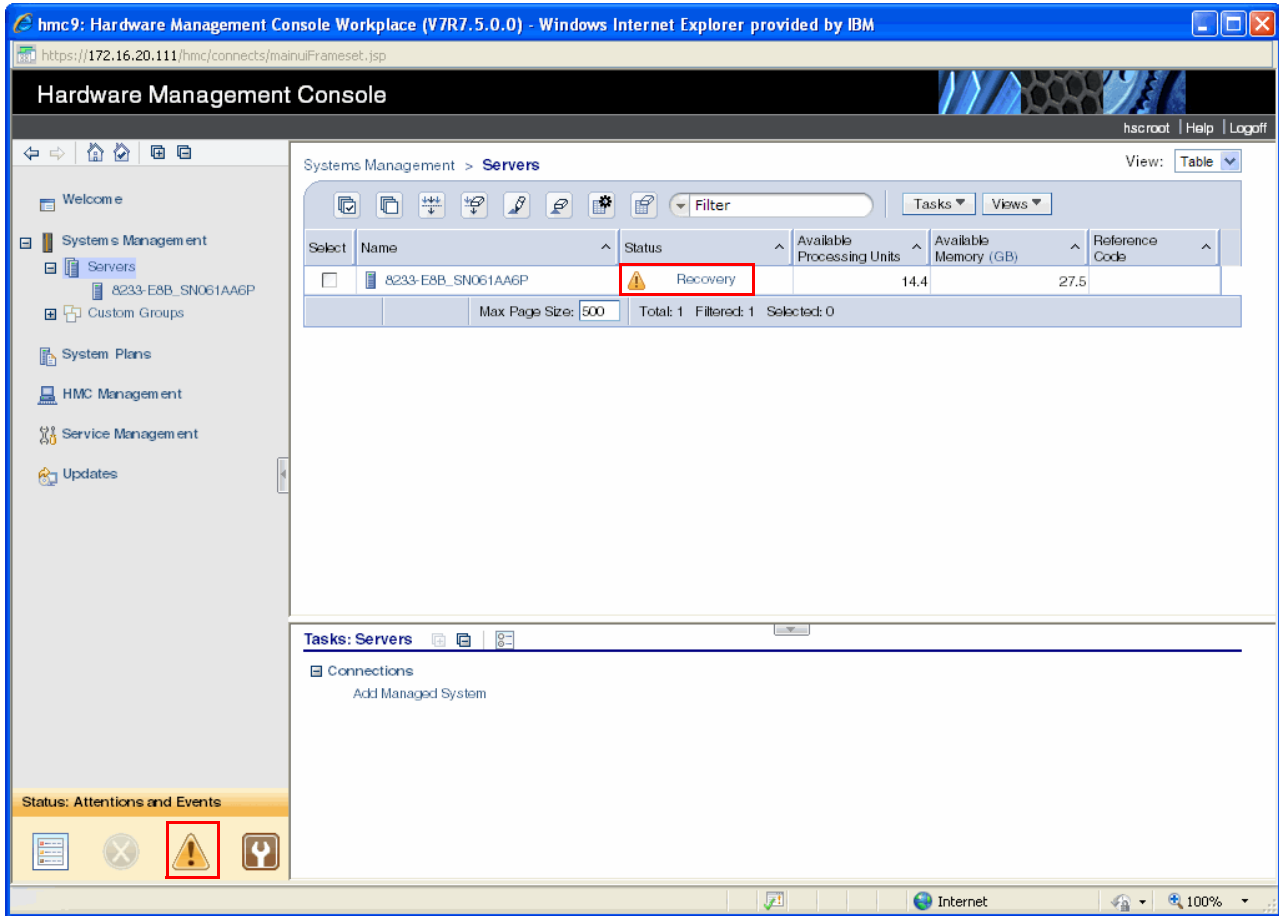


Figure 5-17 HMC recovery mode

3. After the recovery completes, the HMC displays all partitions, in the same state as they were in IVM. The status resolves to the Operating state, as shown in Figure 5-18 and Figure 5-19 on page 184.




Figure 5-18 Managed System going to Operating state

Systems Management > Servers > 8233-E8B_SN061AA6P

| Select | Name | ID | Status | Processing Units | Memory (GB) | Active Profile | Environment | Reference Code |
|--------------------------|---------------|----|---------------|------------------|-------------|----------------|--------------------|----------------|
| <input type="checkbox"/> | vios1 | 1 | Running | 1.3 | 20 | | Virtual I/O Server | |
| <input type="checkbox"/> | lpar2_server1 | 3 | Not Activated | 0.1 | 60 | | AIX or Linux | 00000000 |
| <input type="checkbox"/> | lpar3_server1 | 4 | Not Activated | 0.2 | 10 | | AIX or Linux | 00000000 |

Max Page Size: 500 Total: 3 Filtered: 3 Selected: 0

Figure 5-19 Virtual I/O Server at Operating state

4) Re-creating the partition profile

When the IVM is connected to the HMC, there is no profile so you must create one for both the Virtual I/O Server and the virtual I/O client.

To create the profile, do the following steps:

- a. In the HMC, select the LPAR. Select **Configuration** → **Managed Profiles**, as shown in Figure 5-20.




Figure 5-20 Selecting Manage Profiles in the HMC

- b. In the new pop-up window, select **Actions** → **Save Current Config**, as shown in Figure 5-21.

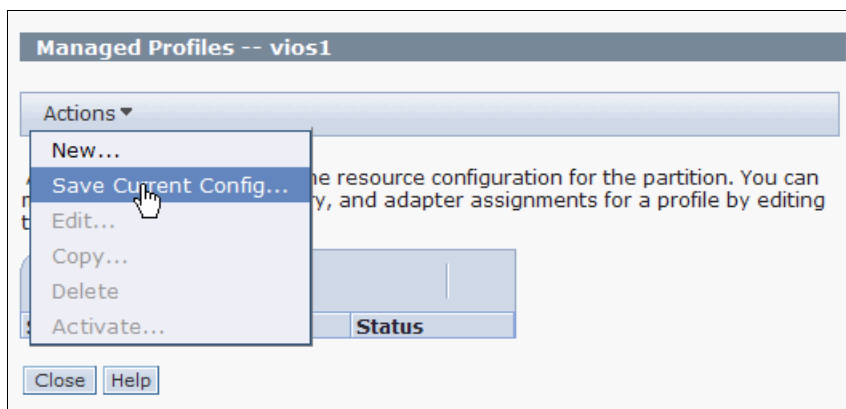


Figure 5-21 Saving the current configuration to a profile in the HMC

- c. Then, select **New profile** and give the profile a name that you want, then select **OK**, as shown in Figure 5-22. This action creates a profile for the LPAR with current configurations.

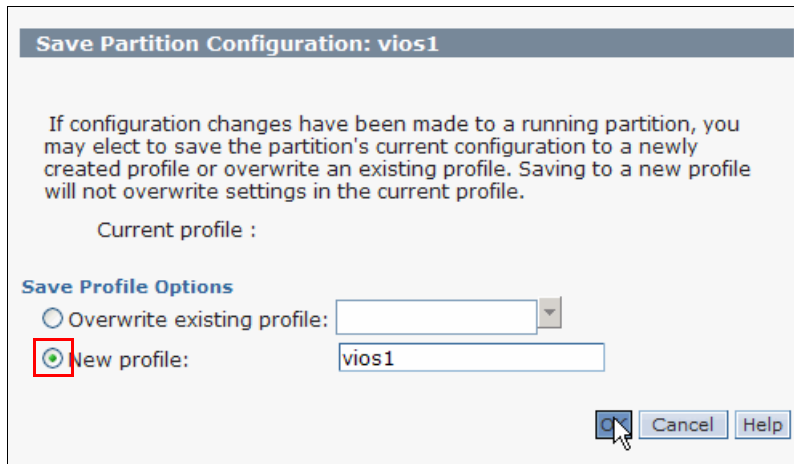


Figure 5-22 Creating a profile with a current configuration

Create a backup before migration: Before migration from an IVM environment to an HMC, it is necessary to back up the Virtual I/O Server and virtual I/O client.

5.2 System maintenance

Operations such as *microcode updates* and *capacity on demand (CoD)* are available for the system that is hosting the IVM.

5.2.1 Microcode update

The IVM provides a convenient interface to generate a microcode survey of the managed system and to download and upgrade microcode.

The following steps describe how to update the device firmware:

1. From the Service Management menu in the navigation area, click **Updates**, then click the **Microcode Updates** tab in the work area.

Updating with IBM BladeCenter: If you are using BladeCenter, use the `diag command` menu to do the update.

2. Click **Generate New Survey**. This action generates a list of devices, as shown in Figure 5-23.

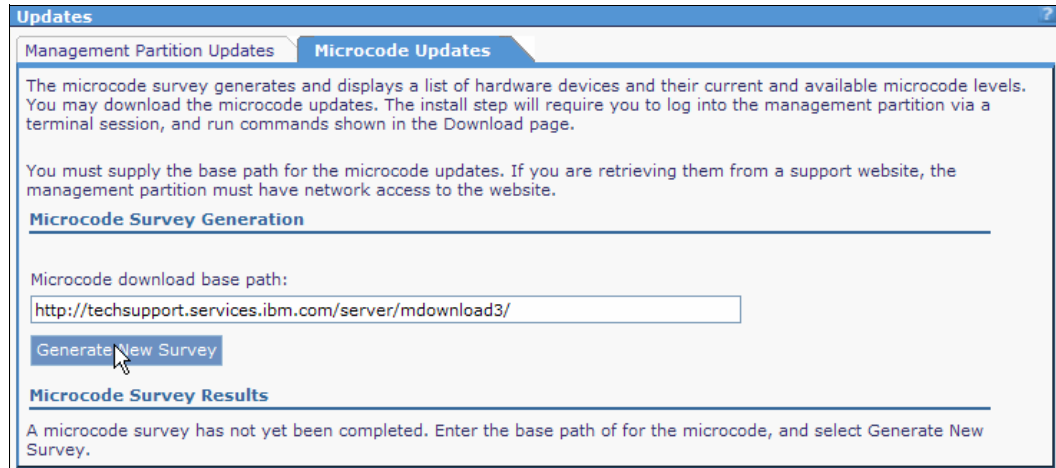


Figure 5-23 Microcode Survey Results

3. From the **Microcode Survey Results** list, select one or more items to upgrade. Click the **Download** link in the task area.
4. Information is displayed about the selected devices, such as the available microcode level and the commands that you need to install the microcode update, as shown in Figure 5-24. Select the **Accept license** check box in the work area, and click **OK** to download the selected microcode and store it on the disk.

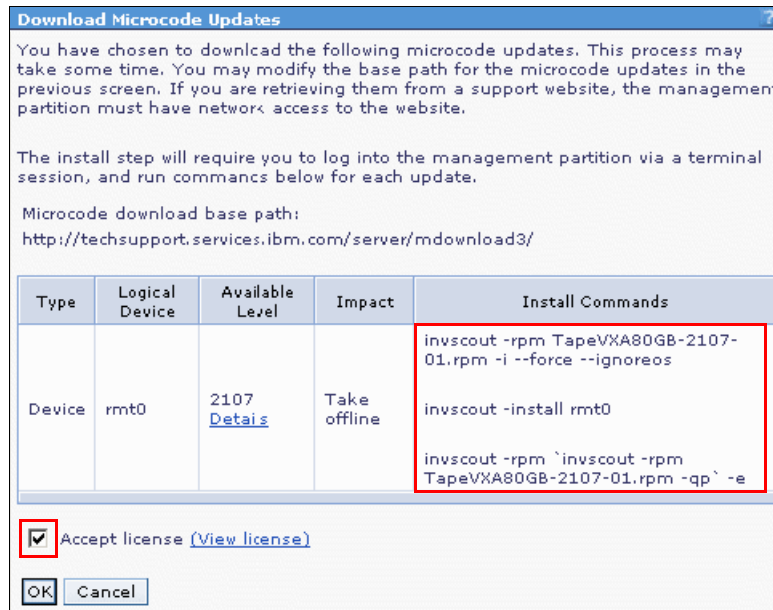


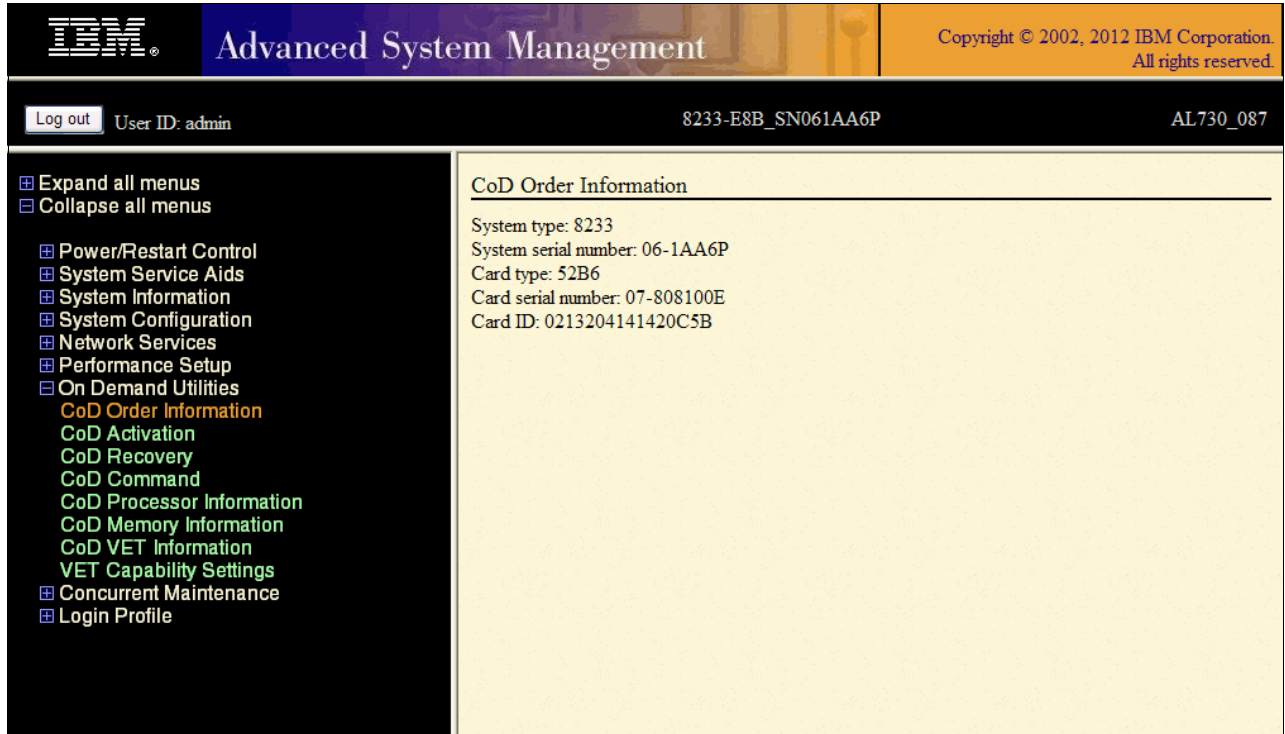
Figure 5-24 Download Microcode Updates

5. Log in to the IVM by using a terminal session.
6. Run the installation commands that are provided by the GUI in step 4.

If you are not able to connect to the GUI of the IVM and a system firmware update is needed, refer to 2.3, “Microcode update” on page 23 for the update procedure with a diagnostic CD.

5.2.2 Capacity on demand operations

Operations for *capacity on demand (CoD)* are available only through the ASMI menu. You must provide the necessary information to get a CoD code. The general information is available in the ASMI menu, as shown in Figure 5-25.



The screenshot displays the ASMI (Advanced System Management Interface) web application. At the top, the IBM logo is on the left, followed by the title "Advanced System Management" in a purple banner. To the right of the banner, it says "Copyright © 2002, 2012 IBM Corporation. All rights reserved." Below the banner, a navigation bar shows "Log out" and "User ID: admin" on the left, and system identifiers "8233-E8B_SN061AA6P" and "AL730_087" on the right. The main content area is split into two columns. The left column is a dark sidebar menu with expandable items. The right column displays the "CoD Order Information" page, which lists system details.

| Expand all menus | CoD Order Information |
|--|--------------------------------|
| <input type="checkbox"/> Collapse all menus | |
| <input type="checkbox"/> Power/Restart Control | System type: 8233 |
| <input type="checkbox"/> System Service Aids | System serial number: 06-1AA6P |
| <input type="checkbox"/> System Information | Card type: 52B6 |
| <input type="checkbox"/> System Configuration | Card serial number: 07-808100E |
| <input type="checkbox"/> Network Services | Card ID: 0213204141420C5B |
| <input type="checkbox"/> Performance Setup | |
| <input type="checkbox"/> On Demand Utilities | |
| <input type="checkbox"/> CoD Order Information | |
| <input type="checkbox"/> CoD Activation | |
| <input type="checkbox"/> CoD Recovery | |
| <input type="checkbox"/> CoD Command | |
| <input type="checkbox"/> CoD Processor Information | |
| <input type="checkbox"/> CoD Memory Information | |
| <input type="checkbox"/> CoD VET Information | |
| <input type="checkbox"/> VET Capability Settings | |
| <input type="checkbox"/> Concurrent Maintenance | |
| <input type="checkbox"/> Login Profile | |

Figure 5-25 CoD menu using ASMI

You can get a CoD code for a processor and memory, as shown in Figure 5-26 and Figure 5-27.

The screenshot displays the IBM Advanced System Management interface. At the top, the IBM logo is on the left, followed by the title "Advanced System Management" in a blue banner. To the right of the banner is the copyright notice: "Copyright © 2002, 2012 IBM Corporation. All rights reserved." Below the banner, a navigation bar shows a "Log out" button, the user ID "admin", the system identifier "8233-E8B_SN061AA6P", and the asset tag "AL730_087".

The main content area is split into two columns. The left column is a dark sidebar menu with the following items:

- Expand all menus
- Collapse all menus
- Power/Restart Control
- System Service Aids
- System Information
- System Configuration
- Network Services
- Performance Setup
- On Demand Utilities
 - CoD Order Information
 - CoD Activation
 - CoD Recovery
 - CoD Command
 - CoD Processor Information**
 - CoD Memory Information
 - CoD VET Information
 - VET Capability Settings
- Concurrent Maintenance
- Login Profile

The right column displays the "CoD Processor Information" page. The data is as follows:

- System type: 8233
- System serial number: 06-1AA6P
- Card type: 52B6
- Card serial number: 07-808100E
- Card ID: 0213204141420C5B
- Resource ID: 7714
- Activated Resources: 0016
- Sequence number: 0041
- Entry check: 21
- Installed processors: 0016
- Permanent processors: 0016
- Inactive processors: 0000
- Configuration index value: 0000
- Processor CCIN: 535B

Figure 5-26 CoD Processor Information

The screenshot displays the IBM Advanced System Management interface, similar to Figure 5-26. The top navigation bar and sidebar menu are identical. In the sidebar menu, "CoD Memory Information" is highlighted in orange.

The right column displays the "CoD Memory Information" page. The data is as follows:

- System type: 8233
- System serial number: 06-1AA6P
- Card type: 52B6
- Card serial number: 07-808100E
- Card ID: 0213204141420C5B
- Resource ID: MODD
- Activated Resources: 0000
- Sequence number: 0040
- Entry check: 32
- Installed memory (GB): 0128
- Permanent memory (GB): 0128
- Inactive memory (GB): 0000

Figure 5-27 CoD Memory Information

You can activate the *virtualization* feature through CoD. For more information about this feature, refer to 2.5, “Virtualization feature activation” on page 29.

When the CoD code is ready, you can apply it either from the ASMI or from IVM to activate.

Figure 5-28 shows the way to activate through ASMI.

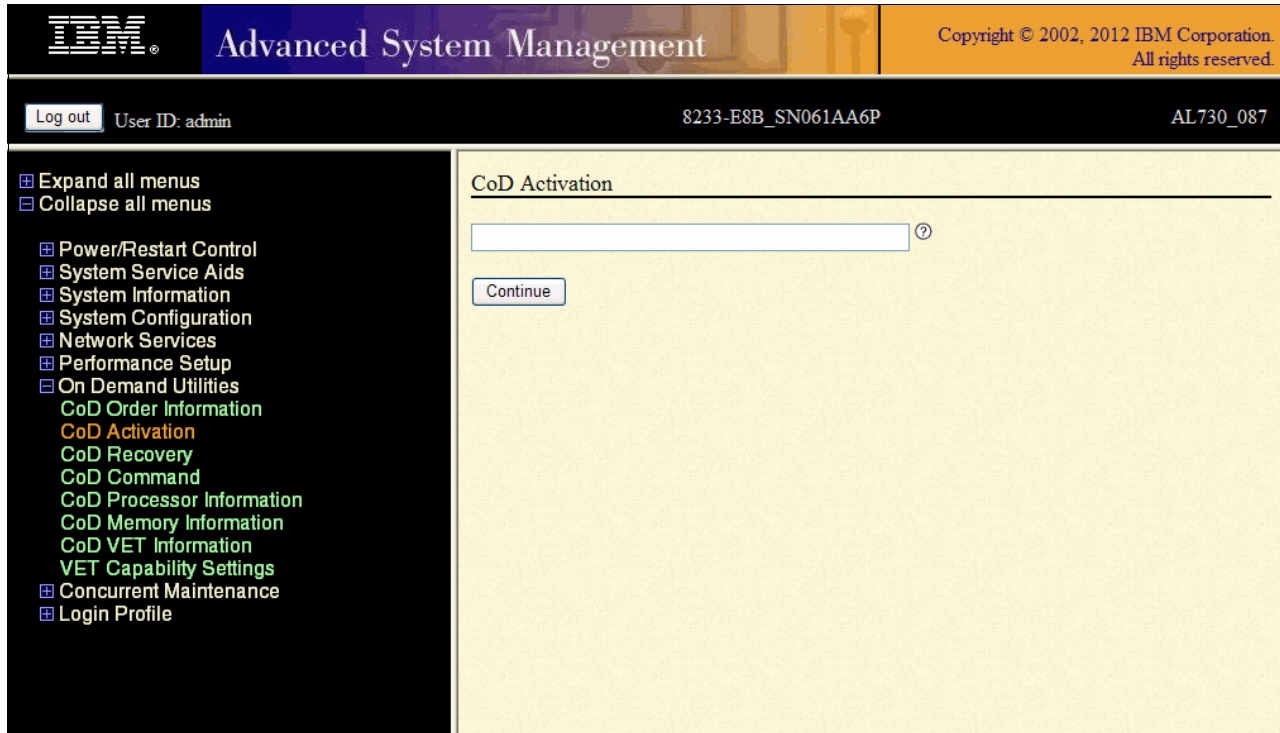


Figure 5-28 CoD Activation through ASMI

Figure 5-29 shows the way to activate through IVM. This process works only to change from the PowerVM Standard to the Express or Enterprise editions, in the **IVM Management** → **Enter PowerVM Edition Key** panel.



Figure 5-29 CoD activation through IVM

5.2.3 Logical partition maintenance

Each LPAR that is hosted by the IVM works like a stand-alone system. For example, you can back up and restore, boot in maintenance mode, and do an operating system update or a migration.

5.2.4 Backup of the operating system

There are many ways to back up LPARs that are hosted by the IVM, depending on which operating system is installed.

The main possibilities for the AIX operating system include the following scenarios:

- ▶ In general, the **mksysb** command creates a bootable image of the **rootvg** volume group either in a file or onto a tape. If your system does not have sufficient space, you can use NFS to mount space from another server system to create a system backup to file. However, the file systems must be writable. Because there is no virtual tape device, tape backup cannot be done locally for the client partitions, but only by a remotely operated tape device. Such a backup can also be done by using more software such as the IBM Tivoli® Storage Manager.
- ▶ The **mkcd** command creates a system backup image (*mksysb*) to a CD-recordable (CD-R) or DVD-recordable (DVD-RAM) media from the system rootvg or from a previously created mksysb image. Multiple volumes are possible for backups over 4 GB. You can create a /mkcd file system that is very large (1.5 GB for CD or 9 GB for DVDs). The /mkcd file system can then be mounted onto the clients when they want to create a backup CD or DVD for their systems.

Large backups: When you create large backups (DVD-sized backups larger than 2 GB) with the **mkcd** command, the file systems must be large-file enabled and this requires that the **ulimit** values are set to *unlimited*.

- ▶ *Network Installation Management (NIM)* creates a system backup image from a logical partition rootvg using the network.

For more information about creating system backups, see this website:

http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.install/doc/insgdrf/create_sys_backup.htm

5.2.5 Restoration of the operating system

The restoration process is the same as on stand-alone systems. Follow these main steps:

1. Log in to the IVM.
2. Open a virtual terminal for the LPAR to be installed with the **mkvt** command, providing the ID of the LPAR to be restored.
3. Start the LPAR in SMS mode.
4. Select the boot device that was used for the backup such as CD, DVD-RAM, or network.
5. Boot the LPAR.
6. Follow the specific restore procedures of the operating system.

5.3 Command logs

All IVM actions are logged in to the system. The log contains all the commands that the IVM web GUI runs and all IVM-specific commands that are issued by an administrator on the command line.

The log contains the following information for each action:

- ▶ User name
- ▶ Date and time
- ▶ The command, which includes all the parameters

The following steps describe how to access the log:

1. Under the Service Management menu in the navigation area, click **Application Logs**.
2. In the work area, use the provided filters to restrict the log search and then click **Apply**. This action generates the selected log entries, as shown in Figure 5-30.

The screenshot shows the IVM interface with the following components:

- Navigation Menu (Left):**
 - Partition Management
 - View/Modify Partitions
 - View/Modify System Properties
 - View/Modify Shared Memory Pool
 - I/O Adapter Management
 - View/Modify Host Ethernet Adapters
 - View/Modify Virtual Ethernet
 - View/Modify Physical Adapters
 - View Virtual Fibre Channel
 - Virtual Storage Management
 - View/Modify Virtual Storage
 - IVM Management
 - View/Modify User Accounts
 - View/Modify TCP/IP Settings
 - Guided Setup
 - Enter PowerVM Edition Key
 - System Plan Management
 - Manage System Plans
 - Service Management
 - Electronic Service Agent
 - Service Focal Point
 - Manage Serviceable Events
 - Service Utilities
 - Create Serviceable Event
 - Manage Dumps
 - Collect VPD Information
 - Updates
 - Backup/Restore
 - Application Logs
 - Monitor Tasks
 - Hardware Inventory

- Main Content Area:**
- Application Logs:** The application logs are events and errors generated by the management partition. To modify the selection criteria, select the filters and then select Apply. To perform an action, first select a log and then the task.
- Select Filters:**
 - Log type: Fatal Error Warning Information Audit
 - Client type: Command line interface Access process Web interface
 - Application:
 - Time period: 10 Minutes
 - Buttons: Apply, Reset
- Selected Log Entries:**

| Select | User ID | Date and Time | Log Type | Application | Client Type | Description |
|--------------------------|---------|--------------------|----------|--------------------|------------------------|---|
| <input type="checkbox"/> | root | 4/19/12 9:45:15 AM | Audit | lpmutil (11141350) | Command line interface | lpmutil -o isapready |
| <input type="checkbox"/> | root | 4/19/12 9:45:15 AM | Audit | lpmutil (11141350) | Command line interface | Command lpmutil returned 12. Additional messages: |
| <input type="checkbox"/> | root | 4/19/12 9:45:14 AM | Audit | lpmutil (8061158) | Command line interface | lpmutil -o isapready |

Figure 5-30 Application command logs

5.4 Integration with IBM Director

The IVM includes an agent that allows full integration and management through the *IBM Director* console. From this centralized console, you can monitor critical resources and events with automated alerts or responses to predefined conditions. You also have control over the hardware to remotely start, stop, and reset the servers and LPARs. You can also examine the software and hardware inventory and deploy new applications or updates across the environment. All processes can be monitored by the console too.

In a few steps, you can look after your IVM managed system and have access to these options:

- ▶ Access the IBM Systems Director for the Power Systems console. Click the **Resource Explorer** link in the left side menu or in the **Plug-ins** tab below the **Discovery Manager** options, as shown in Figure 5-31.

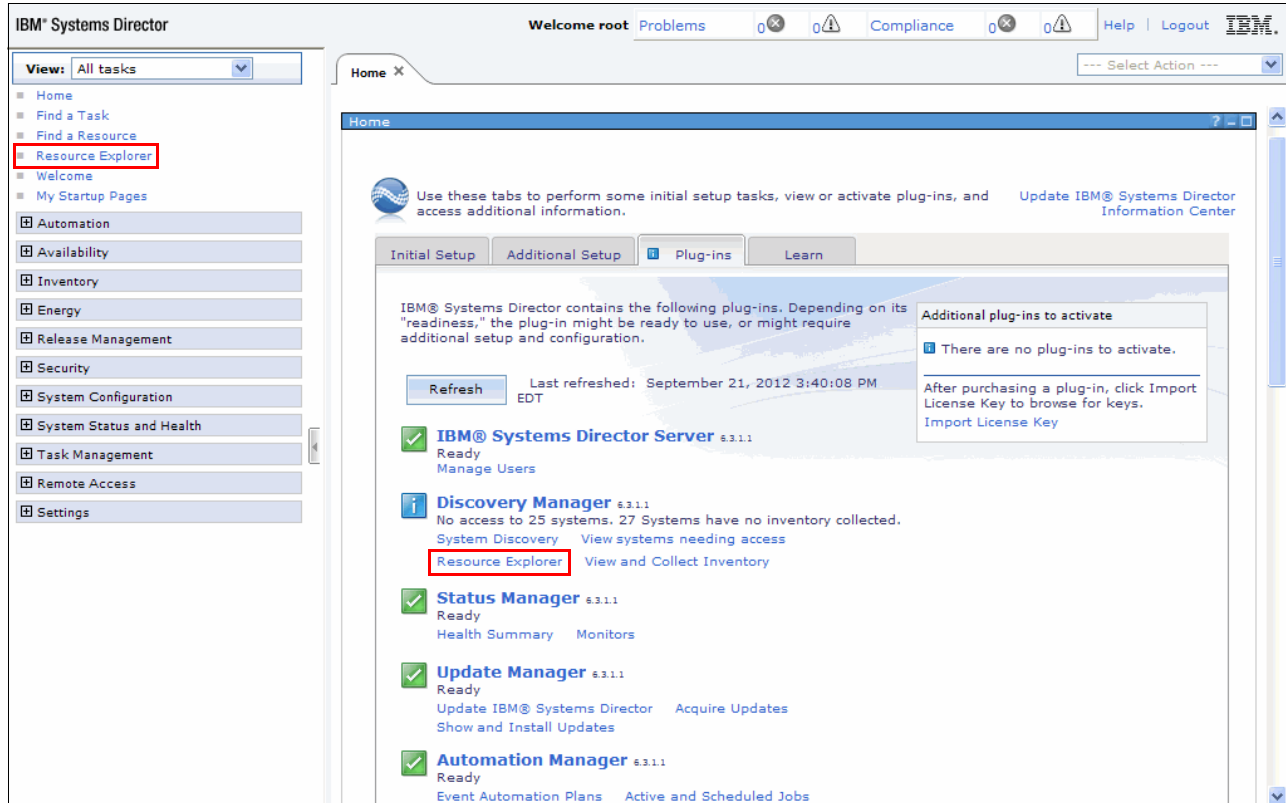


Figure 5-31 Systems Director: Plug-ins tab

- ▶ In the **Resource Explorer** panel, select **Groups by System Type**, as shown in Figure 5-32.

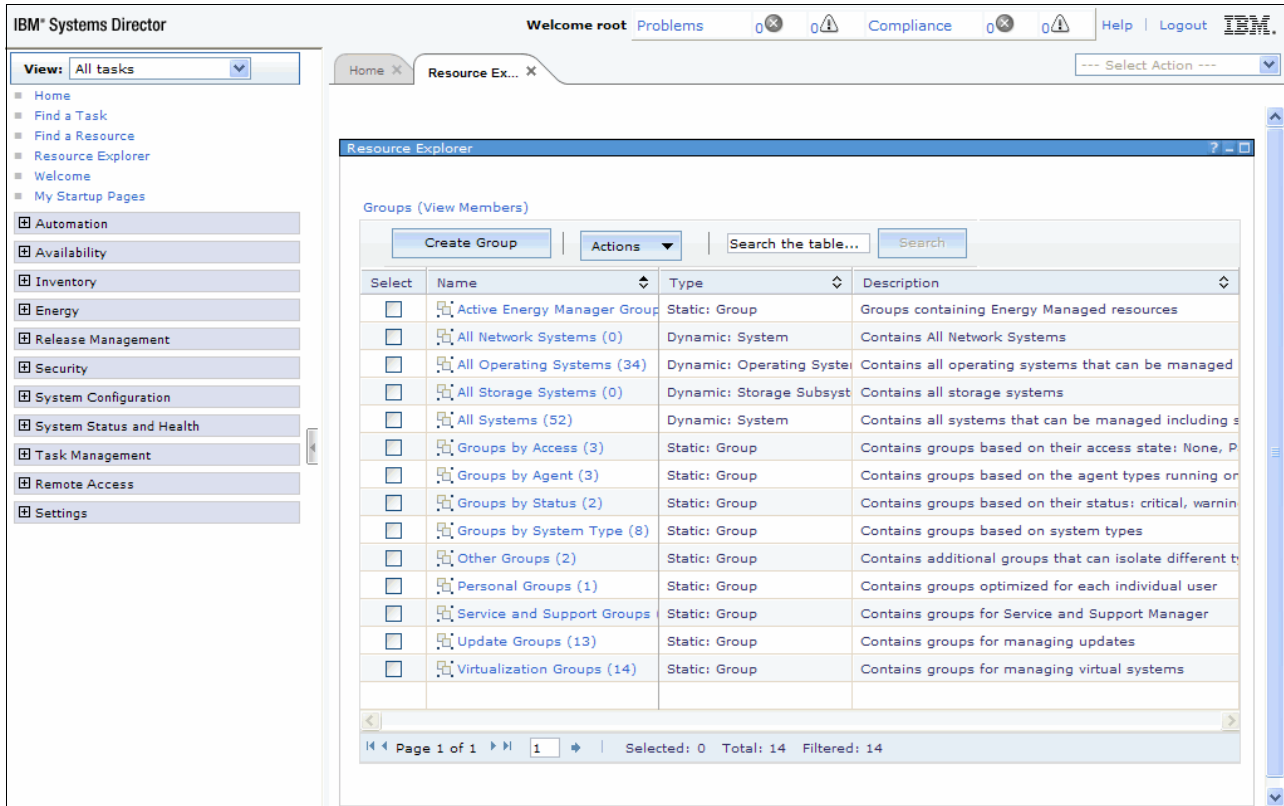


Figure 5-32 Resource Explorer window

- ▶ In the next view, select **Power Systems** servers, as shown in Figure 5-33.

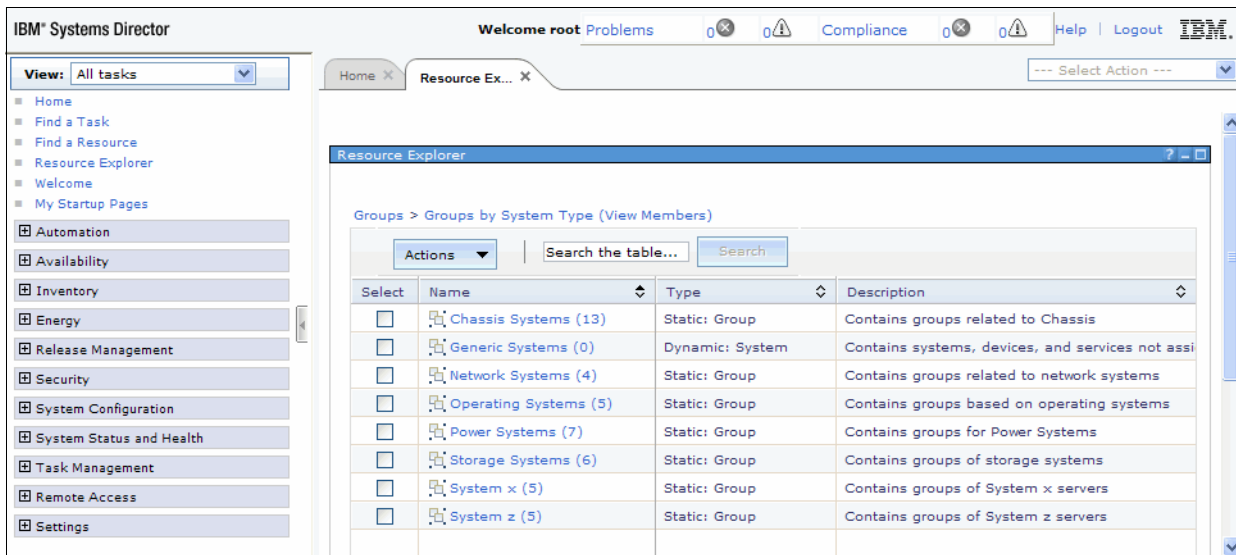


Figure 5-33 Resource Explorer: Groups by System Type

- ▶ Then, select the **IVM and Managed Power Systems Servers**, as shown in Figure 5-34.

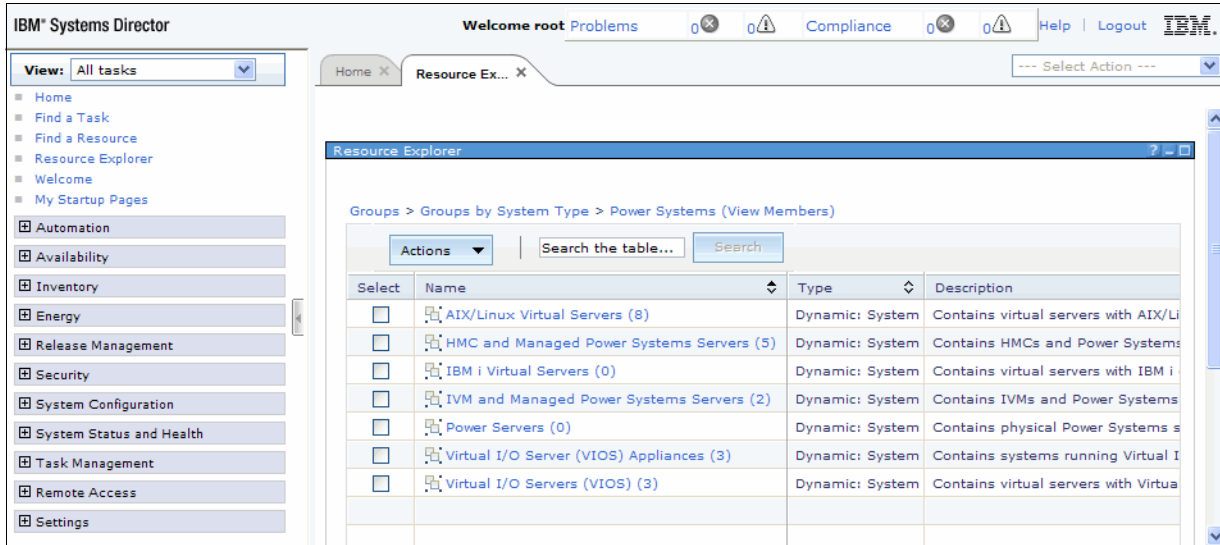


Figure 5-34 Resource Explorer: Power Systems view

- ▶ Now you are able to see all the servers that are managed by the IVM, as shown in Figure 5-35. It is possible to have access to their partitions (and do most of the IVM operations), validating it with a **padmin** user and password.

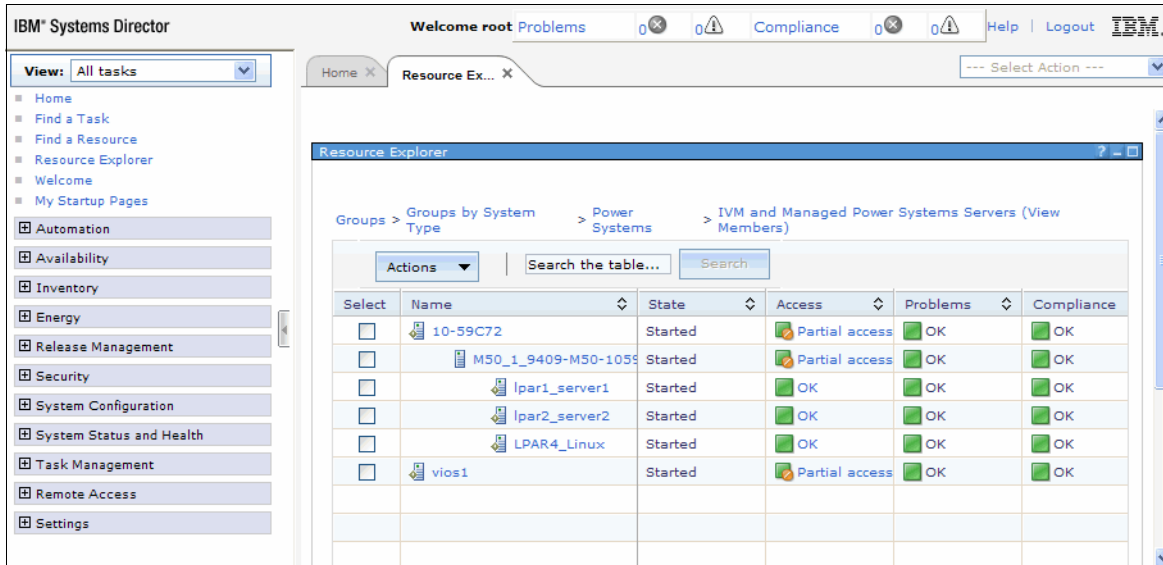


Figure 5-35 Resource Explorer: Power Systems servers view

Systems controlled by HMC and IVM: For more information about managing the systems that are controlled by HMC and IVM, refer to this website:

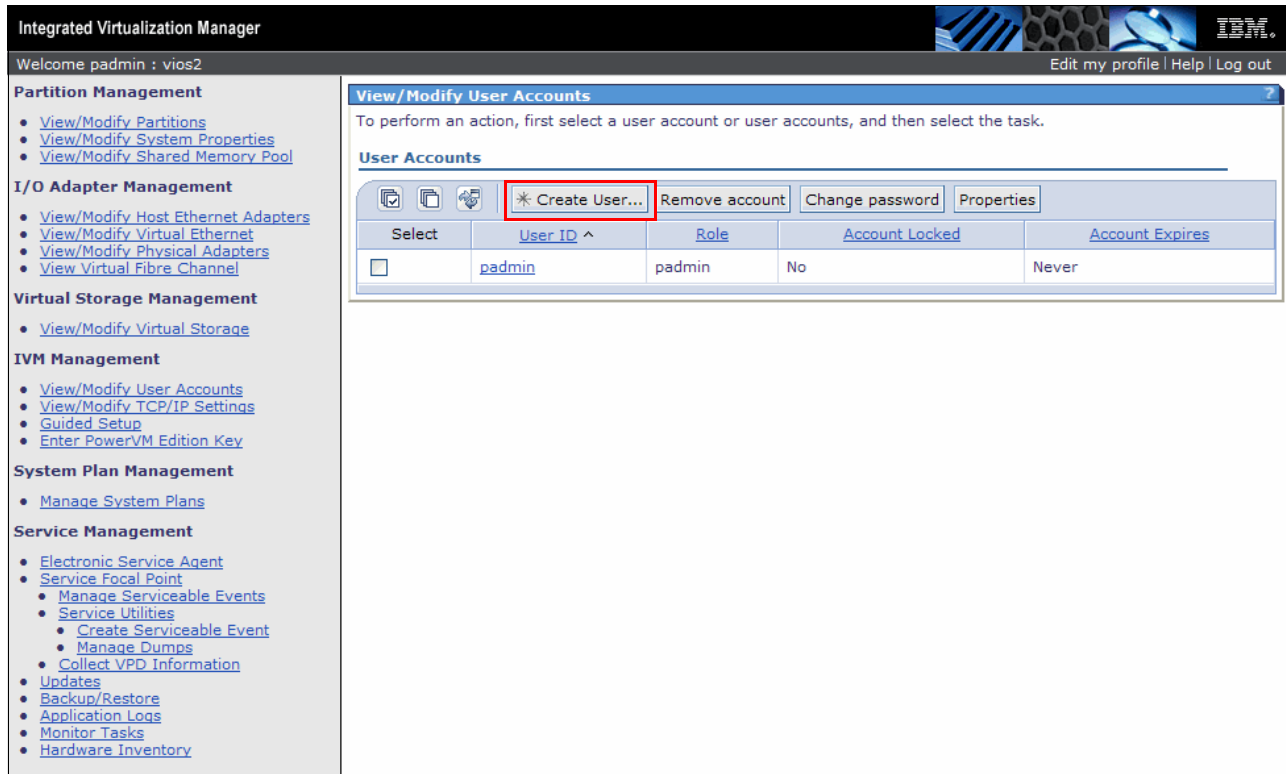
http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.power.helps.doc%2Ffqm0_t_managing_hmc_ivm.html

5.5 User Accounts

After IVM installation, you can access the IVM web browser and also the Virtual I/O Server command line by logging in with a padmin user ID, which is the default user for the IVM/Virtual I/O Server.

It is possible to create other users for IVM by following these steps:

- ▶ Access the **IVM Management** menu and click **View/Modify User Accounts**, as shown in Figure 5-36, and click **Create User**.



The screenshot shows the IVM web interface. The top navigation bar includes the title 'Integrated Virtualization Manager' and the IBM logo. Below the navigation bar, there is a welcome message 'Welcome padmin : vios2' and links for 'Edit my profile | Help | Log out'. The left sidebar contains a navigation menu with categories: Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, and Service Management. The main content area is titled 'View/Modify User Accounts' and contains a table of user accounts. The 'Create User...' button is highlighted with a red box.

| Select | User ID ^ | Role | Account Locked | Account Expires |
|--------------------------|-----------|--------|----------------|-----------------|
| <input type="checkbox"/> | padmin | padmin | No | Never |

Figure 5-36 View/Modify User Accounts window

- ▶ In the opened window, as shown in Figure 5-37, type in the User ID and Password. Then, choose the user's role and click **OK**.

Figure 5-37 Create User Account menu

- ▶ After that, you can see the new user, as shown in Figure 5-38.

| Select | User ID ^ | Role | Account Locked | Account Expires |
|--------------------------|-----------|-------------|----------------|-----------------|
| <input type="checkbox"/> | padmin | padmin | No | Never |
| <input type="checkbox"/> | New_User | View/Modify | No | Never |

Figure 5-38 View/Modify User Accounts panel with the new user

- ▶ After the user creation, it is possible to edit **User Properties**, under the User Settings and Password Settings tabs, as shown in Figure 5-39 and Figure 5-40.

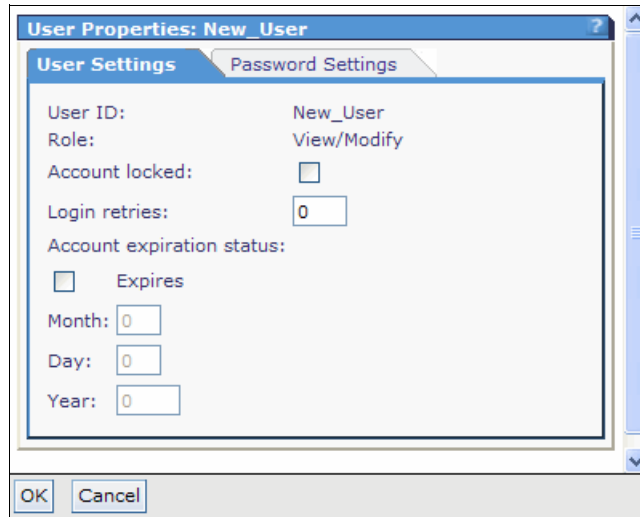


Figure 5-39 User Properties window: User Settings tab

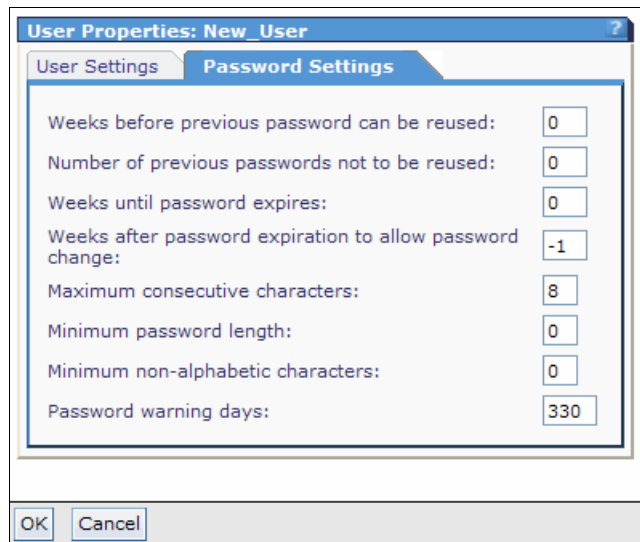
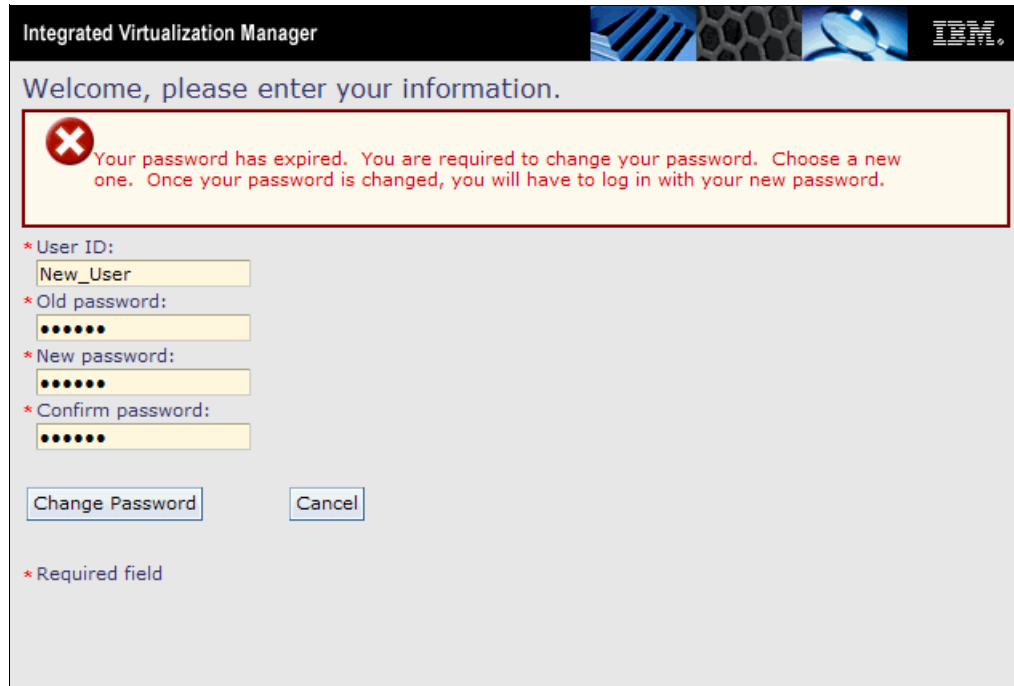


Figure 5-40 User properties window: Password Settings tab

- ▶ By the first time this new user logs in to the IVM, it prompts the user to change the password, as shown in Figure 5-41.



The screenshot shows the Integrated Virtualization Manager (IVM) interface. At the top, the title bar reads "Integrated Virtualization Manager" and the IBM logo is visible on the right. Below the title bar, the text "Welcome, please enter your information." is displayed. A prominent red-bordered warning box contains a red 'X' icon and the message: "Your password has expired. You are required to change your password. Choose a new one. Once your password is changed, you will have to log in with your new password." Below the warning box, there is a form with the following fields and labels:

- * User ID:
- * Old password:
- * New password:
- * Confirm password:

At the bottom of the form, there are two buttons: "Change Password" and "Cancel". A legend at the bottom left indicates "* Required field".

Figure 5-41 New user first login: change password request

5.6 IBM Electronic Service Agent

The *IBM Electronic Service Agent™ (ESA)* is used to monitor hardware problems and send the information automatically to IBM support. It has the following attributes:

- ▶ Submits the problems to IBM through the network
- ▶ Is disabled by default

To use the ESA feature, follow these steps:

1. Configure and start the ESA by using **cfgassist** and then select **Electronic Service Agent**, as shown in Figure 5-42. Press **Enter**.

```
Config Assist for VIOS

Move cursor to desired item and press Enter.

Set Date and TimeZone
Change Passwords
Set System Security
VIOS TCP/IP Configuration
Install and Update Software
Storage Management
Devices
Performance
Role Based Access Control (RBAC)
Shared Storage Pools
Electronic Service Agent

F1=Help          F2=Refresh       F3=Cancel        F8=Image
F9=Shell         F10=Exit         Enter=Do
```

Figure 5-42 Config Assist for VIOS menu to select Electronic Service Agent

2. Then, select **Configure Electronic Service Agent**, as shown in Figure 5-43. Press **Enter**.

```
Electronic Service Agent

Move cursor to desired item and press Enter.

Configure Electronic Service Agent
Configure Service Connectivity
Start Electronic Service Agent
Stop Electronic Service Agent
Verify Electronic Service Agent Connectivity

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit    Enter=Do
```

Figure 5-43 Electronic Service Agent menu to configure the ESA

3. Enter the company name and contact details for configuring the Electronic Service Agent, as shown in Figure 5-44. Then, press **Enter** to confirm the configuration.

```

                                Configuring Electronic Service Agent

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

                                [Entry Fields]
* Company name                                []

Service contact
* Name of the contact person                    []
* Telephone number of the contact person        []
* Email address (myuserid@mycompany.com)       []
* Country or region of contact person          +
  IBM ID                                        []

System location
* Telephone number where the system is located []
* Country or region where the system is located +
* Street address where the system is located   []
* City where the system is located             []
* State or province where the system is located []
* Postal code where the system is located      []
* Building, floor, and office of the system   []

Operational setting
* Port number on which to receive connections [5024] #

F1=Help          F2=Refresh          F3=Cancel          F4=List
F5=Reset         F6=Command          F7=Edit            F8=Image
F9=Shell         F10=Exit             Enter=Do
  
```

Figure 5-44 Details of company to configure the Electronic Service Agent

4. When configured, start the Electronic Service Agent by selecting `cfgassist` → **Electronic Service Agent** → **Start Electronic Service Agent**, as shown in Figure 5-45. This action starts the service.

```
Electronic Service Agent

Move cursor to desired item and press Enter.

Configure Electronic Service Agent
Configure Service Connectivity
Start Electronic Service Agent
Stop Electronic Service Agent
Verify Electronic Service Agent Connectivity

F1=Help          F2=Refresh      F3=Cancel      F8=Image
F9=Shell         F10=Exit       Enter=Do
```

Figure 5-45 Starting Electronic Service Agent

5.7 Service Focal Point

IVM has a facility to manage all the hardware errors that occur in the system. It logs only one error for all the common errors occurred across the LPARs so that you can track errors easily.

You can do the following tasks when you use Service Focal Point. These tasks are described in more detail in the sections that follow:

- ▶ Manage serviceable events
- ▶ Service utilities
 - Create serviceable events
 - Manage system and service processor dumps
- ▶ Collect vital product data (VPD) information

5.7.1 Manage serviceable events

You can *manage all serviceable events* easily in a customized manner. This feature lists the problems with a problem number, reference code, ESA request number, status, reported times, and failing unit.

You can customize the view in the following ways to easily track each problem:

- ▶ Event status
- ▶ Reference code
- ▶ Problem number
- ▶ Number of days
- ▶ Reporting unit
- ▶ Failing unit
- ▶ Field-replaceable unit (FRU) part number, and location code

Figure 5-46 describes the *Manage Serviceable Events* window.

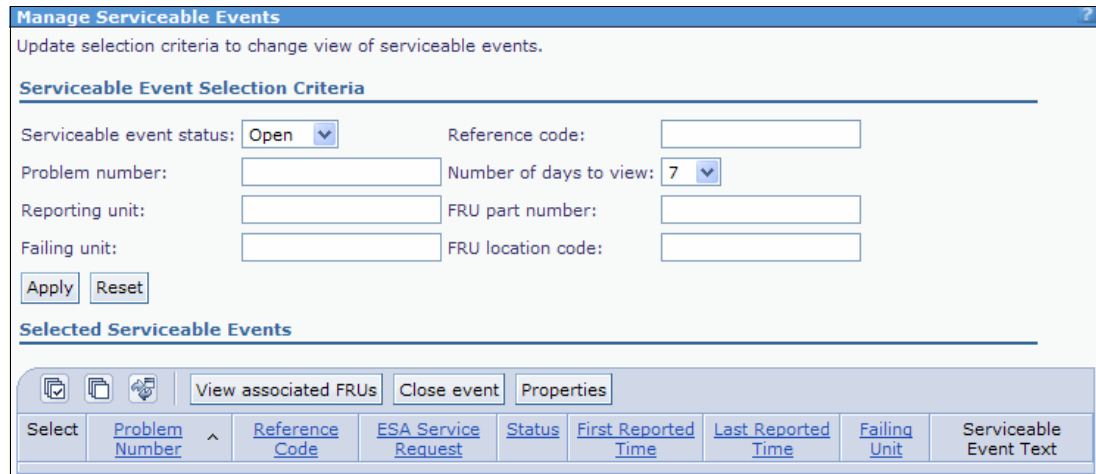


Figure 5-46 Window which is used to Manage Serviceable Events

5.7.2 Service utilities

The following *service utilities* are available in the IVM:

- ▶ Use the *create serviceable event* feature for any problem that you notice.
- ▶ Use the option to connect to the *Advanced System Management Interface (ASMI)*. The following tasks can be handled with the ASMI:
 - Easy Identification of the problem, based on the reference *light-emitting diodes (LEDs)*.
 - Check for correctness of the hardware parts by doing a *lamp test*.
 - Easily manage *system and service processor dumps*.
- ▶ A *diagnostic tool* is available in the Virtual I/O Server restricted shell.

Create serviceable event

You can report any hardware problem to IBM through this window. This feature uses internet access to reach IBM support. View this feature in Figure 5-47.

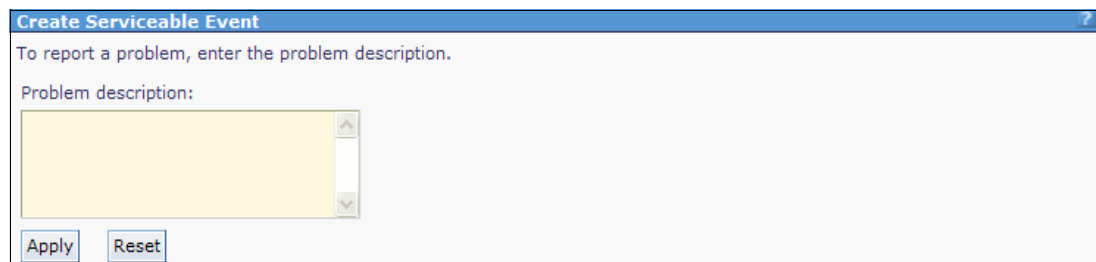


Figure 5-47 Create Serviceable Event window to reach IBM support

Manage system and service processor dumps

You can manage *system and service processor dumps* in this view. This feature allows you to download, check the size and type of the dump, and see when the dump occurred. See Figure 5-48 for a view of the Manage Dumps window.




Figure 5-48 Window to manage dumps

5.8 Application logs

Events and errors of applications that are generated in the management partition can be easily managed in the IVM.

The logs can be customized by the *log type* or *where* you got the error (whether at the command line or web interface). The log can also be customized by *application* and by the *time gap*. See Figure 5-49 for a view of the Application Logs window.

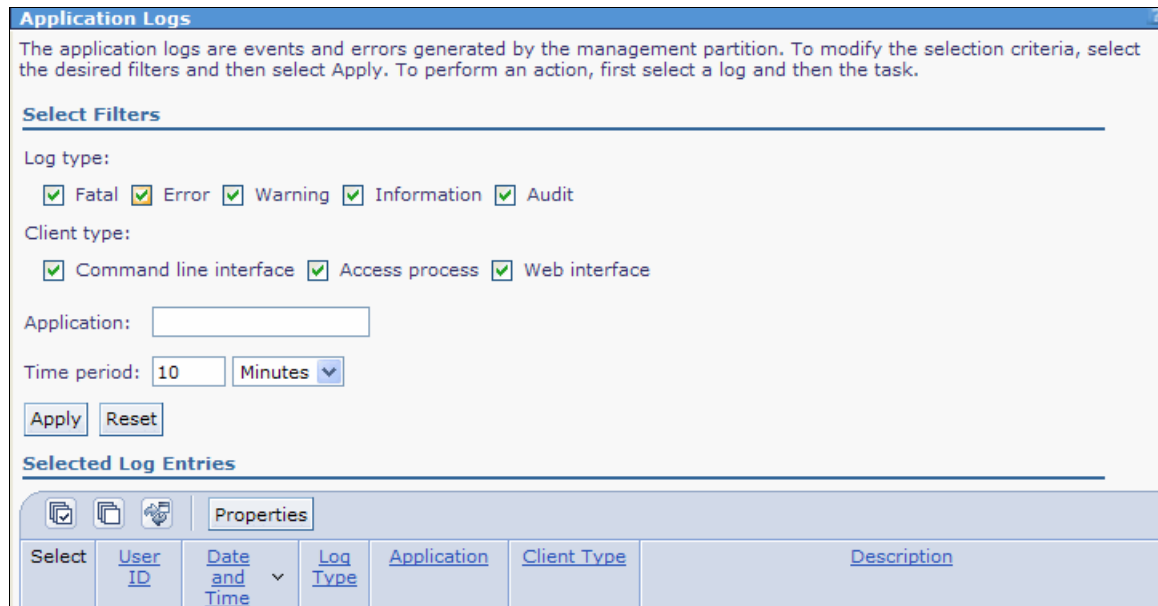


Figure 5-49 Window to customize the application logs

5.9 Monitoring tasks

In IVM, you can view and monitor the most recent 40 tasks. All actions that a user can do in the GUI become *tasks*. All tasks are audited at the task level. Each task can have subtasks; the status of each subtask is managed. When you perform a task, the user gets a *Busy* dialog indicating that the task is running. You can go away from the page, and do other tasks.

To view the properties of the tasks, follow these steps:

1. In the Service Management menu, click **Monitor Tasks**. The Monitor Tasks panel opens.
2. **Select** the task for which you want to view the properties, as shown in Figure 5-50.

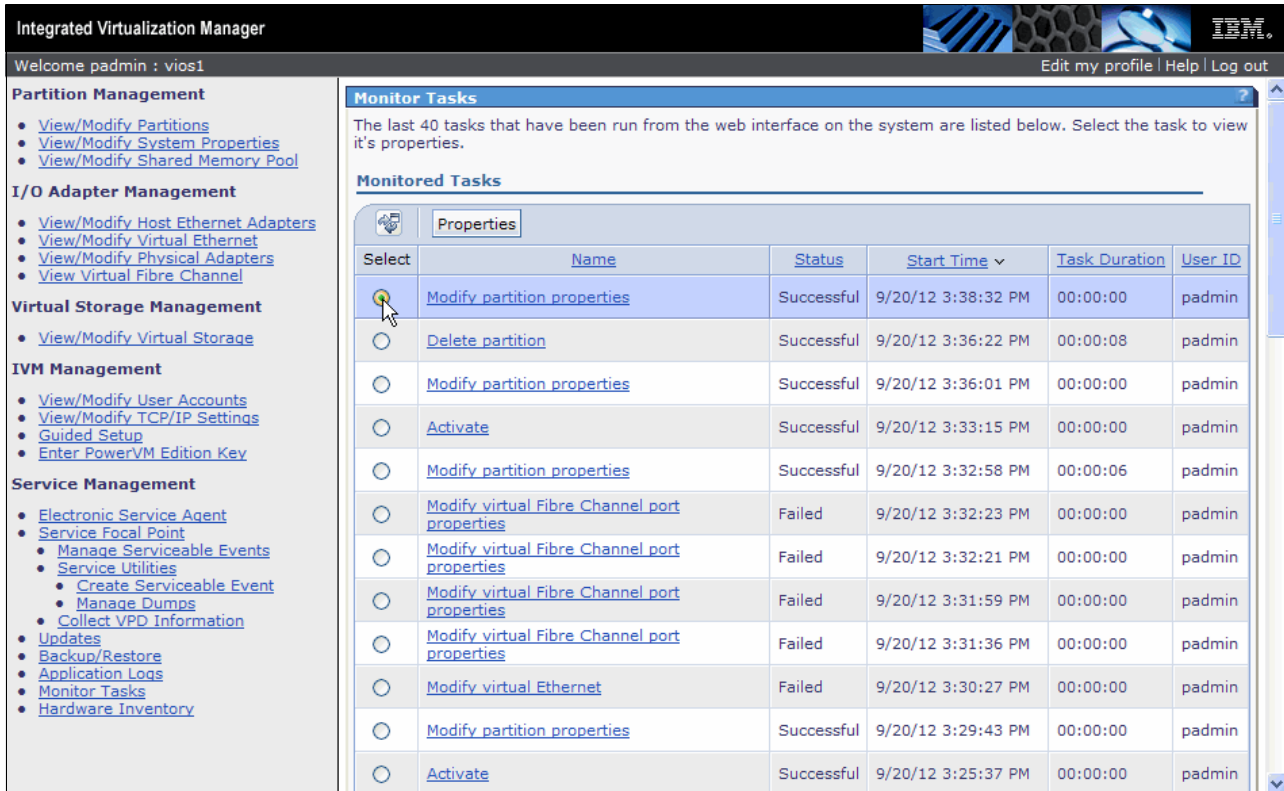


Figure 5-50 Monitor Tasks: the last 40 tasks

3. Click **Properties** (shown in Figure 5-50). The **Task Properties** window opens, as shown in Figure 5-51.

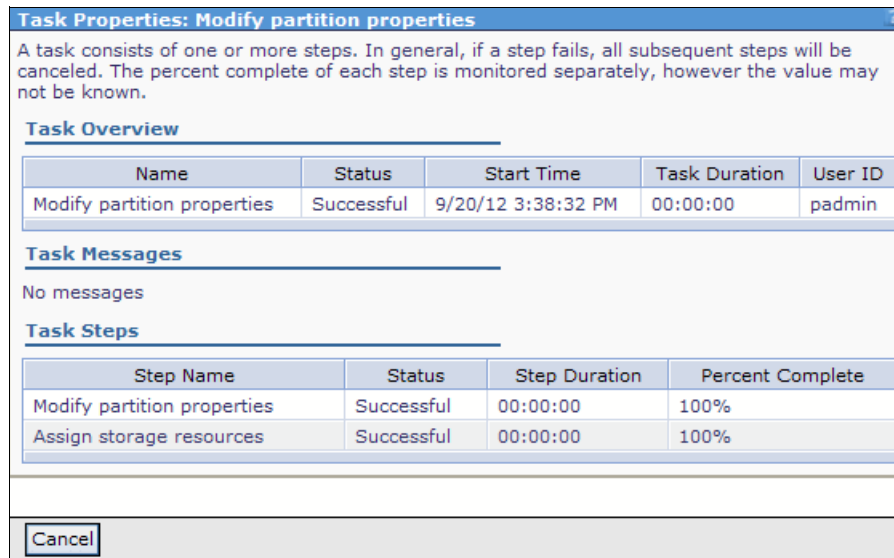


Figure 5-51 Task Properties

4. Click **Cancel** to close the Task Properties window. The Monitor Tasks panel opens.

You can also click the hyperlink of the task for which you want to view the properties. See more about hyperlinks in the following section.

5.10 Hardware inventory


From the *Hardware Inventory* menu, it is possible to see all the physical and virtual devices from the server, as shown in Figure 5-52. If you click **Configure Devices**, you can update the window in case there are any new devices that are not yet shown.

The screenshot shows the 'Hardware Inventory' window in the Integrated Virtualization Manager. The window title is 'Hardware Inventory' and it contains the following text: 'Below is a listing of hardware devices visible to the management partition. Click on **Configure Devices** to add any available devices that are not visible to the management partition.'

Below the text is a 'Configure Devices' button and a table with the following data:

| Device Name | Status | Description | Physical Location Code |
|-------------|-----------|---|----------------------------|
| cd0 | Available | SATA DVD-RAM Drive | U78A0.001.DNWGPL0-P2-D2 |
| Client_data | Defined | Volume group | |
| cluster0 | Available | Cluster Node | |
| en0 | Defined | Standard Ethernet Network Interface | |
| en1 | Defined | Standard Ethernet Network Interface | |
| en2 | Available | Standard Ethernet Network Interface | |
| en3 | Defined | Standard Ethernet Network Interface | |
| en5 | Defined | Standard Ethernet Network Interface | |
| en6 | Defined | Standard Ethernet Network Interface | |
| en7 | Defined | Standard Ethernet Network Interface | |
| en8 | Defined | Standard Ethernet Network Interface | |
| en9 | Defined | Standard Ethernet Network Interface | |
| en10 | Defined | Standard Ethernet Network Interface | |
| en11 | Defined | Standard Ethernet Network Interface | |
| en12 | Defined | Standard Ethernet Network Interface | |
| en13 | Defined | Standard Ethernet Network Interface | |
| ent0 | Available | 2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902) | U78A0.001.DNWGPL0-P1-C4-T1 |
| ent1 | Available | 2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902) | U78A0.001.DNWGPL0-P1-C4-T2 |
| ent2 | Available | Shared Ethernet Adapter | |
| ent3 | Available | EtherChannel / IEEE 802.3ad Link Aggregation | |

Figure 5-52 Hardware Inventory window



IVM and HMC feature summary

Table A-1 provides a comparison between the Integrated Virtualization Manager (IVM) and the Hardware Management Console (HMC).

Table A-1 IVM and HMC comparison at a glance

| | Integrated Virtualization Manager (IVM) | Hardware Management Console (HMC) |
|-------------------------------------|--|--|
| Physical footprint | Integrated into the server | A desktop or rack-mounted appliance |
| Installation | Installed with the Virtual I/O Server (optical or network). Preinstall option available on some systems. | Appliance is preinstalled. Reinstall using optical media or network is supported. |
| Managed operating systems supported | AIX, Linux, and IBM i | AIX, Linux, and IBM i |
| Virtual console support | AIX, and Linux virtual console support | AIX, Linux, and IBM i virtual console support |
| User security | Password authentication with support for either full or ready-only authorities | Password authentication with granular control of task-based authorities and object-based authorities |
| Network security | -Firewall support via command line -Web server SSL support | -Integrated firewall -SSL support for clients and for communications with managed systems |
| Multiple system support | One IVM per server | One HMC can manage multiple servers |

| | Integrated Virtualization Manager (IVM) | Hardware Management Console (HMC) |
|---|---|---|
| Redundancy | One IVM per server | Multiple HMCs can manage the same system for HMC redundancy |
| Maximum number of partitions supported | Firmware maximum | Firmware maximum |
| Uncapped partition support | Yes | Yes |
| Partition Mobility | Yes, between 2 IVM managed systems | Yes, between two systems managed by a management console |
| Host Ethernet Adapter (Integrated Virtual Ethernet) | Yes | Yes |
| Multiple Shared Processor Pool | No, default pool only | Yes |
| Other POWER7 functions (share while active, partition priority, etc.) | Yes | Yes |
| Dynamic Resource Movement (dynamic LPAR) | Yes - full support | Yes - full support |
| I/O Support for AIX and Linux | Virtual and Direct | Virtual and Direct |
| I/O Support for i/OS | Virtual only | Virtual and Direct |
| Maximum # of virtual LANs | 4096 | 4096 |
| Fix/update process for Manager | Virtual I/O Server fixes and updates | Management Console e-fixes and release updates |
| Adapter microcode updates | Inventory scout | Inventory scout |
| Firmware updates | Virtual I/O Server firmware update tools (not concurrent) | Service Focal Point with concurrent firmware updates |
| I/O concurrent maintenance | Virtual I/O Server support for slot and device level concurrent maintenance via the diag hot plug support | Guided support in the Repair and Verify function on the HMC |
| Scripting and automation | Virtual I/O Server command-line interface (CLI) and HMC-compatible CLI | HMC command-line interface |
| Capacity on Demand | Entry of PowerVM codes only | Full support |
| User interface | Web browser (no local graphical display) | Web browser (local or remote) |
| Workload Management (WLM) groups supported | One | 254 |
| LPAR configuration data backup and restore | Yes | Yes |

| | Integrated Virtualization Manager (IVM) | Hardware Management Console (HMC) |
|---|--|--|
| Support for multiple profiles per partition | No | Yes |
| Serviceable event management | Service Focal Point Light: Consolidated management of firmware and management of partition detected errors | Service Focal Point support for consolidated management of operating system and firmware detected errors |
| Hypervisor and service processor dump support | Dump collection with support to do manual dump downloads | Dump collection and call home support |
| Remote support | No remote support connectivity | Full remote support for the HMC and connectivity for firmware remote support |



System requirements

Table B-1 shows the currently supported systems for the IBM Integrated Virtualization Manager (IVM).

Table B-1 Supported server models for IVM

| IBM Power Systems | Supported server models |
|--------------------------|---|
| IBM POWER5 Systems | All server models, except p5-575, p5-590, and p5-595 |
| IBM POWER6 Systems | All server models, except Power 575, Power 590, and Power 595 |
| IBM POWER7 Systems | All server models, except Power 755, Power 770, Power 775, Power 780, and Power 795 |
| IBM BladeCenter | All POWER6 and POWER7 processor-based blades, and IBM BladeCenter JS21 |

The minimum required firmware levels for supporting IVM include the following requirements:

- ▶ SF235 or later for POWER5 processor-based servers
- ▶ All firmware levels for POWER6 and POWER7 processor-based servers

The software minimum supported levels include the following requirements:

- ▶ IBM AIX 5L™ V5.3 or later
- ▶ SUSE Linux Enterprise Server 9 for POWER (SLES 9) or later
- ▶ Red Hat Enterprise Linux AS 3 for POWER, Update 2 (RHEL AS 3) or later
- ▶ IBM i 6.1.1 and later

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks related to PowerVM

For information about ordering these publications, see “How to get IBM Redbooks” on page 214. Some of the documents referenced here might be available in softcopy only.

- ▶ *PowerVM Best Practices*, SG24-8062
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.
- ▶ *IBM PowerVM Getting Started Guide*, REDP-4815.
- ▶ *IBM PowerVM Virtualization Active Memory Sharing*, REDP-4470.
- ▶ *IBM PowerVM Live Partition Mobility*, SG24-7460.

IBM Redbooks related to hardware

These publications are also relevant as further information sources:

- ▶ *IBM BladeCenter PS700, PS701, and PS702 Technical Overview and Introduction*, REDP-4655
- ▶ *IBM BladeCenter PS703 and PS704 Technical Overview and Introduction*, REDP-4744
- ▶ *IBM Power 795 (9119-FHB) Technical Overview and Introduction*, REDP-4640.
- ▶ *IBM Power 770 and 780 (9117-MMD, 9179-MHD) Technical Overview and Introduction*, REDP-4924
- ▶ *IBM Power 770 and 780 (9117-MMC, 9179-MHC) Technical Overview and Introduction*, REDP-4798
- ▶ *IBM Power 770 and 780 (9117-MMB, 9179-MHB) Technical Overview and Introduction*, REDP-4639
- ▶ *IBM Power 750 and 755 (8233-E8B, 8236-E8C) Technical Overview and Introduction*, REDP-4638
- ▶ *IBM Power 720 and 740 (8202-E4B, 8205-E6B) Technical Overview and Introduction*, REDP-4637
- ▶ *IBM Power 720 and 740 (8202-E4C, 8205-E6C) Technical Overview and Introduction*, REDP-4797
- ▶ *IBM Power 710 and 730 (8231-E2B) Technical Overview and Introduction*, REDP-4636.
- ▶ *IBM Power 710 and 730 (8231-E1C, 8231-E2C) Technical Overview and Introduction*, REDP-4796
- ▶ *IBM Power 570 and IBM Power 595 (POWER6) System Builder*, REDP-4439
- ▶ *IBM Power 520 and Power 550 (POWER6) System Builder*, SG24-7765

Online resources

These websites and URLs are also relevant as further information sources:

- ▶ IBM TotalStorage Mid-range/Enterprise disk Systems
<http://www.ibm.com/systems/storage/disk/index.html>
- ▶ IBM Virtualization Engine
<http://www.ibm.com/systems/virtualization/>
- ▶ IBM PowerVM
<http://www.ibm.com/systems/power/>
- ▶ Virtual I/O Server supported environments
<http://www.software.ibm.com/webapp/set2/sas/f/vios/home.html>
- ▶ Hardware Management Console, Microcode, AIX Fix pack support information
<http://www.ibm.com/support/fixcentral/>
- ▶ SUMA on AIX
<http://www.software.ibm.com/support/customercare/sas/f/genunix/suma.html>
- ▶ SUSE Enterprise Linux
<https://www.suse.com/products/server/>
- ▶ Red Hat Enterprise Linux
<http://www.redhat.com/rhel/features/>
- ▶ AIX 6.1 Information Center
<http://pic.dhe.ibm.com/infocenter/aix/v6r1/index.jsp>
- ▶ AIX 7.1 Information Center
<http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.ntl%2FRELNOTES%2FGI11-9815-00.htm>
- ▶ Power Systems Hardware Information Center
<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Integrated Virtualization Manager for IBM Power Systems Servers



**No dedicated
Hardware
Management Console
required**

**Powerful integration
for entry-level servers**

**Key administration
tasks explained**

The *Virtual I/O Server (VIOS)* is part of the IBM PowerVM feature on IBM Power Systems and part of the IBM POWER Hypervisor. The VIOS is also supported on IBM BladeCenter. The server is a single-function appliance that is in the *logical partition (LPAR)* of the Power Systems server. The VIOS facilitates the sharing of physical I/O resources between client partitions (IBM AIX, IBM i, and Linux) within the server.

The Virtual I/O Server provides a virtual SCSI target, N_Port ID Virtualization (NPIV) target, and Shared Ethernet Adapter (SEA) virtual I/O function to client LPARs.

The Virtual I/O Server has the capability of a hardware management function, the *Integrated Virtualization Manager (IVM)*. IVM is a simplified hardware management solution that inherits most of the Hardware Management Console (HMC) features. The console manages a single server, avoiding the need of a dedicated personal computer.

IVM provides the following functions:

- ▶ Shared Storage Pool
- ▶ IBM Active Memory Sharing (AMS)
- ▶ Live Partition Mobility (LPM)
- ▶ Task manager monitor for long-running tasks
- ▶ Security additions such as viosecure and firewall, and other improvements

This IBM Redpaper publication provides an introduction to IVM by describing its architecture and showing how to install and configure a partitioned server using its capabilities. This document is intended for IT personnel who have a complete understanding of partitioning before reading this document.

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