



IBM BladeCenter JS23 and JS43 Implementation Guide

Featuring installation techniques for the IBM
AIX, IBM i, and Linux operating systems

Showing Live Partition Mobility
scenarios

Detailing AMS, IVM, and
power management



Alex Zanetti de Lima
Kerry Anders
Nahman Cohen
Steven Strain
Vasfi Gucer

ibm.com/redbooks

Redbooks



International Technical Support Organization

IBM BladeCenter JS23 and JS43 Implementation Guide

July 2009

Archived

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

First Edition (July 2009)

This edition applies to IBM BladeCenter JS23, IBM BladeCenter JS43, IBM AIX Version 6.1, IBM Vi 6.1, Red Hat Enterprise Linux for POWER Version 5.3, SUSE Linux Enterprise Server 11 for POWER.

© Copyright International Business Machines Corporation 2009. 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	xi
Trademarks	xii
Preface	xv
The team that wrote this book	xv
Acknowledgements	xvi
Become a published author	xvii
Comments welcome	xviii
Part 1. The IBM BladeCenter JS23 and JS43 servers	1
Chapter 1. Introduction to IBM BladeCenter	3
1.1 Highlights of BladeCenter	4
1.2 IBM BladeCenter is right choice, open, easy, green	6
Chapter 2. General description	9
2.1 Overview of JS23 and JS43 Express blade servers	10
2.2 Features and specifications	11
2.3 Blade server features	14
2.4 Physical specifications BladeCenter JS23	16
2.4.1 Minimal and optional features of BladeCenter JS23	16
2.4.2 Processor features	17
2.4.3 Memory features	18
2.4.4 Memory DIMMs	18
2.4.5 Internal disk	18
2.5 Physical specifications of BladeCenter JS43	19
2.5.1 Minimal and optional features of BladeCenter JS43	19
2.5.2 Processor features	20
2.5.3 Memory features	21
2.5.4 Memory DIMMs	21
2.5.5 Internal disk	21
2.6 IBM BladeCenter chassis	22
2.6.1 BladeCenter H	23
2.6.2 BladeCenter S	26
2.6.3 BladeCenter HT	30
2.6.4 Number of JS23 and JS43 blades in a chassis	34
2.6.5 IBM Director	35

Chapter 3. Technical description of the hardware architecture	37
3.1 POWER6 processor	38
3.1.1 Decimal floating point	39
3.2 Simultaneous multi-threading	40
3.3 AltiVec SIMD	41
3.4 IBM EnergyScale technology	42
3.4.1 Thermal Power Management Device	43
3.5 Flexible Service Processor	44
3.6 Management Card	44
3.6.1 Anchor chip on Management Card	45
3.7 Memory subsystem	45
3.7.1 Memory description of JS23 and JS43 Express	45
3.7.2 Memory placement rules for JS23 and JS43	46
3.7.3 Memory RAS	47
3.8 I/O subsystem	48
3.8.1 PCI Express	48
3.8.2 I/O Expansion Cards	48
3.8.3 Universal Serial Bus subsystem	50
3.8.4 Integrated Virtual Ethernet	50
3.8.5 Integrated Video Card	51
3.8.6 Serial Attached SCSI storage subsystem	51
3.9 PowerVM	56
3.10 Operating system support	57
3.10.1 AIX	57
3.10.2 Linux	58
3.10.3 IBM i	59
3.11 Systems management	59
3.11.1 BladeCenter Advanced Management Web interface	59
3.11.2 IBM Director	60
3.11.3 Cluster Systems Management	61
Part 2. System configuration and management	63
Chapter 4. System planning and configuration using VIOS with IVM	65
4.1 Planning considerations	67
4.1.1 General considerations	67
4.1.2 Internal and external storage considerations	70
4.2 VIOS system management using IVM	77
4.2.1 VIOS installation considerations	78
4.2.2 IVM user interface	78
4.2.3 VIOS/IVM command-line interface	79
4.3 First VIOS login	80
4.3.1 Password set	80

4.3.2	License acceptance	80
4.3.3	Initial network setup	81
4.4	First IVM connection	87
4.4.1	Connecting to IVM.	87
4.4.2	Verify and modify VIOS partition memory and processors	89
4.5	VIOS network management and setup	94
4.5.1	Host Ethernet Adapters.	95
4.5.2	Virtual Ethernet Adapters and SEA.	97
4.5.3	Physical adapters	103
4.5.4	VLANs and virtual networks	107
4.6	VIOS Storage Configuration and Management.	116
4.6.1	Physical volumes	117
4.6.2	Storage pools	119
4.6.3	Virtual disks.	122
4.6.4	Optical and tape devices.	127
4.7	Partition configuration for Virtual I/O Client	139
4.7.1	Live Partition Mobility considerations	139
4.7.2	Partition name and environment	140
4.7.3	Partition name	141
4.7.4	Partition memory	142
4.7.5	Partition processors	143
4.7.6	Partition Ethernet	144
4.7.7	Partition storage types and assignments	145
4.7.8	Optical and tape devices.	147
4.7.9	Physical adapters	148
4.7.10	Partition summary	149
4.7.11	Partition properties changes and DLPAR operations	151
4.8	Console access and activating a partition	161
4.8.1	Opening a virtual terminal	161
4.8.2	Activating a partition	163
4.8.3	Shut down a VIO Client partition	166
4.8.4	Shutdown the VIO Server	170
Chapter 5.	Active Memory Sharing configuration using IVM	173
5.1	Planning considerations	174
5.1.1	Prerequisites	175
5.1.2	Workload	175
5.1.3	Paging devices	176
5.1.4	Collaborative Memory Manager	176
5.1.5	Processor resource sizing.	176
5.1.6	Memory weight	177
5.1.7	Consolidation factors.	178
5.2	Configuration of Active Memory Sharing using IVM	178

5.2.1	Defining the shared memory pool and paging storage pool	178
5.2.2	Creating dedicated paging devices for partitions	187
5.2.3	Creating shared memory LPARs	189
5.2.4	Shared memory partitions and dedicated paging devices	199
5.2.5	Active Memory Sharing DLPAR operations	203
5.3	Active Memory Sharing summary	204
Chapter 6.	IBM AIX V6.1 installation	205
6.1	Introduction to installing in PowerVM client partition	206
6.2	Creating a virtual media library	206
6.3	Preparing the PowerVM client partition	212
6.4	Installing AIX 6.1 in a logical partition of VIOS	225
Chapter 7.	IBM i V6.1 installation	235
7.1	Preparing for installation	236
7.1.1	Software installation process	236
7.1.2	Hardware environments	238
7.1.3	BladeCenter hardware preparation	241
7.1.4	VIO Server software environments	244
7.1.5	Network considerations	244
7.1.6	Storage consideration BladeCenter H	245
7.1.7	Disk consideration in BladeCenter S	245
7.1.8	Disk configuration in BladeCenter S	247
7.1.9	Individual BladeCenter S disk configuration	252
7.2	IBM System Access for Windows V6R1	253
7.2.1	Preparation of System i Access for Windows	254
7.2.2	Installing System i Access for Windows	254
7.3	Creating an IBM i V6.1 partition	263
7.3.1	IBM i V6.1 minimum requirements	264
7.3.2	VIO Server configuration	264
7.3.3	Creating an IBM i partition	264
7.3.4	Partition properties	270
7.3.5	IBM i V6.1 installation of media preparation	272
7.3.6	Connect the System i LAN console	273
7.3.7	IBM i V6.1 IPL types	284
7.3.8	Completing the partition	285
7.4	Installing and setting up IBM i V6.1	296
7.4.1	IBM i V6.1 installing PTFs	304
7.4.2	IBM i V6.1 TCP/IP setup	305
7.4.3	IBM i V6.1 Navigator for i	305
7.5	IBM i V6.1 backup and restore	307
7.5.1	Creating a virtual media library for backup	309
7.5.2	Creating virtual media library using IVM	315

7.5.3 Adding image files to media library	317
7.5.4 Attaching a remote PC file or media device	320
7.5.5 IBM Tivoli Storage Manager	324
7.5.6 IBM i V6.1 shutdown and restart	325
Chapter 8. Red Hat Enterprise V5.3 Linux installation	327
8.1 Red Hat operating systems and prerequisites	328
8.1.1 Supported Red Hat operating system versions	328
8.1.2 Considerations and prerequisites	328
8.2 Linux LPAR installation using DVD	329
8.3 Linux network installation	333
8.4 Native Red Hat Enterprise Linux 5.3 installation	346
8.5 Red Hat Enterprise Linux 5.3 automated installation	346
8.5.1 Methods for creating a Kickstart file	347
8.5.2 Creating Kickstart file with Kickstart Configurator tool	348
8.5.3 Performing a Kickstart installation	357
8.6 IBM service and productivity tools	358
Chapter 9. SUSE Linux Enterprise Server V11 installation	359
9.1 Operating systems and prerequisites	360
9.1.1 Supported operating systems	360
9.1.2 Considerations and prerequisites	360
9.2 Linux LPAR installation using DVD	361
9.3 Linux network installation (detailed)	366
9.4 Native SLES 11 installation	387
9.5 SLES 11 automated installation	387
9.6 IBM service and productivity tools	389
Chapter 10. JS23 and JS43 power management using EnergyScale technology	391
10.1 Power management through the AMM	392
10.1.1 Using the AMM Web GUI for blade power management	392
10.1.2 Using the AMM CLI for blade power management	398
10.2 Power management through AEM	401
10.2.1 Active Energy Manager console	402
10.2.2 Active Energy Manager energy properties	404
10.2.3 BladeCenter Energy Properties	407
10.2.4 Creating power policies	415
10.2.5 Conclusion	426
Chapter 11. Performing Live Partition Mobility	427
11.1 Requirements	428
11.1.1 Hardware	428
11.1.2 Firmware	428

11.1.3	VIOS version	428
11.1.4	PowerVM Enterprise	430
11.1.5	LPAR OS versions	431
11.2	Preparation	432
11.2.1	VIOS source and target requirements.	432
11.2.2	Networking	436
11.2.3	Partition requirements	437
11.3	Migrating the LPAR	442
11.3.1	Using the IVM GUI	442
11.3.2	From the command line	450
Chapter 12. System maintenance and diagnostics		453
12.1	Firmware updates	454
12.1.1	Committing new firmware to the TEMP side	454
12.1.2	Starting the firmware image from the TEMP side	456
12.1.3	Verifying current firmware level	462
12.1.4	Interpreting the firmware file name	463
12.1.5	Getting the latest firmware image from IBM	464
12.2	System diagnostics	465
12.2.1	Diagnostic tools	465
12.2.2	Reference codes	472
12.2.3	Diagnostics tasks through the Advanced Management Module	473
12.2.4	IBM i partition diagnostics and errors	476
Part 3. Appendixes		485
Appendix A. Consoles, SMS, and Open Firmware		487
Consoles of the IBM BladeCenter JS23 and JS43		488
Graphical console		488
Serial Over LAN		494
System Management Services menu		495
Powering on JS23 or JS43 using AMM Web interface		496
Powering on JS23 or JS43 using Telnet or SSH into AMM		497
Choosing the active console after first power on		500
IBM BladeCenter JS23 and JS43 power-on self-test		501
Selecting active console using the SMS menu		502
Open Firmware interface		503
Getting access to the firmware prompt		504
QLogic host bus adapter		507
Emulex host bus adapter		510
Appendix B. SUSE Linux Enterprise Server AutoYaST		515
AutoYaST introduction		516
AutoYaST profile creation methods		516

Creating an AutoYaST profile with YaST Control Center	516
Appendix C. Additional Linux installation configuration options	529
Basic preparations for a Linux network installation	530
Services required for installing Linux using the network	530
Configuring a BOOTP or DHCP service	531
Configuring a Trivial File Transfer Protocol service	533
Configuring a Network File System Protocol service	535
Virtual optical device setup and installation	538
Appendix D. Service and productivity tools for Linux	539
Overview	540
Install tools for Red Hat Enterprise Linux 5 or SLES 11	542
Installing tools on BladeCenter servers	542
Installing tools on IVM-managed servers	545
Abbreviations and acronyms	549
Related publications	555
IBM Redbooks	555
Online resources	556
How to get Redbooks	560
Help from IBM	560
Index	561

Notices

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

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

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

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

1350™	IBM Systems Director Active	POWER®
Active Memory™	Energy Manager™	Redbooks®
AIX 5L™	IBM®	Redbooks (logo)  ®
AIX®	iSeries®	RETAIN®
BladeCenter®	Micro-Partitioning™	ServerProven®
DS4000®	OS/400®	Solid®
DS6000™	POWER Hypervisor™	System i®
DS8000®	Power Systems™	System p5®
EnergyScale™	POWER5™	System p®
Enterprise Storage Server®	POWER6™	System Storage™
eServer™	POWER6+™	Tivoli®
i5/OS®	PowerVM™	

The following terms are trademarks of other companies:

Advanced Micro Devices, AMD, ATI, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

BLADE Network Technologies, BNT, and the BLADE logo are registered trademarks of BLADE Network Technologies, Inc.

Emulex, and the Emulex logo are trademarks or registered trademarks of Emulex Corporation.

InfiniBand, and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

QLogic, and the QLogic logo are registered trademarks of QLogic Corporation. SANblade is a registered trademark in the United States.

Red Hat, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

ABAP, SAP NetWeaver, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Java, JRE, Power Management, Solaris, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

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

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

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

Archived

Preface

This IBM® Redbooks® publication provides a detailed technical guide for configuring and using the IBM BladeCenter® JS23 and IBM BladeCenter JS43 servers. These IBM Power Blade servers feature the latest IBM POWER6™ processor technology. This book teaches you how to set up the latest Power Blade servers to run AIX®, i, and Linux® operating systems in the IBM BladeCenter architecture.

This book is useful for data center system architects; network, storage and facilities engineers; system administrators; and application architects that want a detailed system start-up guide and integration guide for IBM Power Blade servers hosting virtualized operating system partitions. We include procedures for installation and configuration of Virtual Input/Output Server (VIOS), Integrated Virtualization Manager (IVM), IBM AIX Version 6.1, IBM i V6.1 (i5/OS® V6R1), Red Hat® Enterprise Linux, and SUSE® Enterprise Linux.

The team that wrote this book

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

Alex Zanetti de Lima is a Software Engineer at the Linux Technology Center in Brazil since 2007 and leads the development of the IBM Installation Toolkit project. He has experience developing, debugging, and deploying SUSE Linux Enterprise Server and Red Hat Enterprise Linux on IBM BladeCenter JS23 and JS43. He is also experienced in developing and deploying server consolidation solutions using BladeCenter servers. Related to this area, Alex holds two patents in the United States Patent and Trademark Office.

Kerry Anders is a Consultant in System p® Lab Services for the IBM Systems and Technology Group, based in Austin, Texas. He supports clients in implementing IBM POWER processor blades using Virtual I/O Server, Integrated Virtualization Manager, and AIX. He was previously the Systems Integration Test Team Lead for the IBM BladeCenter JS21blade with IBM SAN storage using AIX and Linux. His prior work includes test experience with the JS20 blade, also using AIX and Linux in SAN environments. Kerry began his career with IBM in the Federal Systems Division supporting NASA at the Johnson Space Center as a Systems Engineer. He transferred to Austin in 1993.

Nahman Cohen is an IT Manager in Memorex Telex Israel for 10 years. He has 18 years of experience in the network support and hardware fields. His areas of expertise include Windows®, Linux, Sun™ Solaris™, and networks.

Steve Strain is a Software Engineer/Advisory Education Specialist for the Rochester Support Center in Rochester, MN. He is responsible for developing and delivering education that is based on IBM i and the POWER platforms. He has worked on many HMC, POWER5™ and POWER6 courses that are tailored towards hardware maintenance, integration, and system configuration. Prior to his education position, he was a member of the Rochester Support Center Hardware Team specializing in DASD systems diagnostics and repair. He has been working for IBM since 1992.

Vasfi Gucer is an IBM Certified Consultant IT Specialist at the ITSO Austin Center. He has more than 15 years of experience teaching and implementing systems management, networking hardware, and distributed platform software. Vasfi is also an ITIL® Expert and Project Management Professional (PMP).

Acknowledgements

This IBM Redbooks publication would not have been possible without the generous support of many IBM employees around the world. For their contributions, guidance, patience, and technical feedback in support of this project, we gratefully acknowledge the following IBM employees:

Chris Austen, John E Biebelhausen, David Engebretsen, Kaena Freitas, Anita Govindjee, Carol Hernandez, Tenley Jackson, Hoa Nguyen, Vess Natchev, Jonathan Van Niewaal, Gary Rohret, Todd Rosedahl, Steven Royer, Mark Smolen, Mike Stys, Padmavathy Tamirisa, Mark Vanderwiel, Mary Wigham, Kyle Wurgler

We would also like to thank the authors of these previous IBM Redbooks publications: *IBM BladeCenter JS21: The POWER of Blade Innovation*, SG24-7273, *IBM BladeCenter JS21 Technical Overview and Introduction*, REDP-4130, *Techniques to Build a Diskless Boot Linux Cluster of JS21 Blades*, REDP-4238, *IBM BladeCenter JS12 and JS22 Implementation Guide*, SG247655, *IBM PowerVM Live Partition Mobility*, SG247460:

Chris Almond, Kerry Anders, John E Bailey, Trina Bunting, Benjamin Ebrahimi, Hendrik Ernst, Stephen Hochstetler, Tom Junkin, Wayne Kimble, Gregor Linzmeier, Bob Maher, Ameet Paranjape, Thomas Prokop, Guido Somers, Scott Vetter, Bernhard Zeller

Lab systems setup in support of this project required the availability of multiple IBM BladeCenter chassis, POWER® blade servers, plus various adapters and access to disk storage subsystems. We are very thankful for the lab systems support we received from:

Ned Gamble and **Erich J Hauptli**

Finally, the team would also like to acknowledge the support for this project provided by **Scott Vetter**, ITSO System p Team Leader, and our book editor, **Diane Sherman**, also contributed to our production and review efforts.

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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 books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

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

ibm.com/redbooks

- ▶ Send your comments in an e-mail 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



Part 1

The IBM BladeCenter JS23 and JS43 servers

This part provides general and technical descriptions of the BladeCenter products covered in this publication.

Introduction to IBM BladeCenter

This chapter provides an introduction to IBM BladeCenter and blade servers JS23 Express and JS43 Express, and discusses the business benefits of blade servers in general.

This chapter contains the following topics:

- ▶ “Highlights of BladeCenter” on page 4
- ▶ “IBM BladeCenter is right choice, open, easy, green” on page 6

The JS23 Express and JS43 Express blade servers are represented in this book in any of the following ways:

- ▶ JS23 and JS43
- ▶ JS23/JS43 Express
- ▶ JS23 and JS43 Express
- ▶ JS23 Express and JS43 Express

1.1 Highlights of BladeCenter

Blade servers are thin servers that insert into a single rack-mounted chassis, which supplies shared power, cooling, and networking infrastructure. Each server is an independent server with its own processors, memory, storage, network controllers, operating system, and applications. Blade server design is optimized to minimize physical space. Standard rack-mount servers require a number of power cords and network cables, but in the case of blade servers, blade enclosures provide services such as power, cooling, networking for multiple blade servers, thereby reducing the space and cable requirements. See Figure 1-1 for a computing environment with and without blade servers.




Figure 1-1 With and without blade servers

Blade servers came to market around 2000, initially to meet client requirements for greater ease of administration and increased server density in the data center environment. When IBM released the IBM BladeCenter in November 2002, it quickly changed the industry with its modular design. The IBM BladeCenter provides complete redundancy in a chassis, enables network and storage integration. IBM BladeCenter is ideal for infrastructure consolidation, virtualization, and demanding applications that require scalable performance and high-memory capacity. The IBM BladeCenter blade server processors use POWER6 processor technology and have the ability to run AIX, IBM i, and Linux operating systems simultaneously. Overall, IBM BladeCenter provides a secure, resilient, and dynamic infrastructure solution that helps drive down cost, reduces risk, improves energy efficiency, and enhances flexibility.

Built on the promise of the IBM BladeCenter family of products that are easy-to-use and integrated platforms with a high degree of deployment flexibility, energy efficiency, scalability, and manageability, the BladeCenter JS23 and JS43 Express are the premier blades for 64-bit applications. They represent one of the most flexible and cost-efficient solutions for UNIX®, IBM i, and Linux deployments available in the market. Further enhanced by their ability to be installed in the same chassis with other IBM BladeCenter blade servers, the

JS23 and JS43 can deliver the rapid return on investment that clients and businesses demand. Delivering on the promise of a truly dynamic infrastructure the BladeCenter JS23 and JS43 help in delivering superior business and IT services with agility and speed, all in a simple to manage, highly efficient way.

The JS23 and JS43 Express blades have been preconfigured and tested by IBM and are based on proven technology. Utilizing a 4.2 GHz 64-bit POWER6 processor and available in a 4-core or 8-core configuration, they are designed to deliver outstanding performance and capabilities at compelling prices. With faster and more reliable double data rate 2 (DDR2) memory options and support for eight to sixteen memory dual in-line memory module (DIMM) slots along with Serial Attached SCSI (SAS) disk subsystem, the BladeCenter JS23 and JS43 Express blades are designed for increased uptime and enhanced performance. Exploiting the newest in high-performance and energy-efficient solutions like IBM Solid® State Disk Technology, which offers remarkable performance for I/O-intensive applications and those that require fast and frequent data access, the JS23 and JS43 offer the utmost in easy-to-manage, tough-to-break solutions. Along with built-in support for PowerVM™ Editions for advanced virtualization, the JS23 and JS43 offers an ideal blade server solution for driving your most demanding performance and memory-intensive workloads such as virtualization, databases and high-performance computing (HPC) applications. By consolidating and virtualizing on BladeCenter with the JS23 and JS43, you can increase the utilization of your hardware and decrease the number of physical assets you have to watch over. This translates into real dollar savings through better energy conservation and IT resource usage across the data center.

Meeting today's cost challenges while planning for tomorrow's growth can be difficult. When you demand more performance, this blade server is ready to scale to four processors with the addition of a multiprocessor expansion unit (MPE). Simply snap the MPE to the original blade and it becomes a 4-processor 8-core, approximately 60 mm blade server with the additional performance, memory, and double the I/O expansion you require in order to deliver a flexible, resilient and highly scalable IT infrastructure. This expansion capability is highly cost effective and provides the opportunity to standardize on a single blade platform for all your requirements.

With IBM BladeCenter and IBM i, clients can realize innovation with proven 6th generation POWER technology for enhanced performance, efficiency and reliability; cut costs and consolidate workloads with leadership virtualization as a built-in feature; go green with better energy management by using EnergyScale™ technology; and manage growth, complexity and risk with the flexibility of IBM BladeCenter technology. With IBM i support, small and medium clients can consolidate their IBM i and Intel® processor-based servers into a single chassis, leveraging the management, space and power savings provided by IBM BladeCenter solutions. Large or small enterprises can now consolidate

their older i5/OS applications into a centralized BladeCenter environment with a choice of BladeCenter chassis and blade configurations to fit their requirements.

Simplify, cut costs, boost productivity, go green. These are all priorities for IT, and they are all driving organizations to rethink their server strategies and become more receptive to new ways to use IT. Blades are the next-generation solution, promising improvements across the board. The IBM BladeCenter innovative, open design offers a true alternative to today's sprawling racks and overheated server rooms. Migrating to the blade solution will give you a solution that uses less energy and more choices and control with less complexity.

1.2 IBM BladeCenter is right choice, open, easy, green

IBM BladeCenter is the right choice, open, easy, and green.

The *right* choice, tailored to fit your diverse requirements:

- ▶ It is flexible and modular. As requirements evolve, a one-size-fits-all solution does not work. It can help you:
 - Meet your needs with BladeCenter. Everything from a high-performance data center to a small office with limited IT skills, IBM has you covered.
 - Get flexibility with five compatible chassis and five blade types supporting multiple I/O fabrics, all managed from a common point.
- ▶ It is robust and reliable, providing redundancy throughout and the information you require to keep your business up and running. It can:
 - Provide redundancy for no single point of failure with IBM BladeCenter.
 - Preserve application uptime with IBM Predictive Failure Analysis and light path diagnostics.
 - Make decisions based on accurate data for quick problem diagnosis with first failure data capture (FFDC).

Open and innovative, for a flexible business foundation:

- ▶ It is comprehensive, providing broad, fast, and reliable networking and storage I/O with BladeCenter Open Fabric. It can help you:
 - Match your data center requirements and the appropriate interconnect using a common management point, and five I/O fabrics to choose from.
 - Extract the most from your third-party management solutions by utilizing the BladeCenter Open Fabric Manager.
- ▶ It is collaborative, enabling you to harness the power of the industry to deliver innovation that matters:

- Get flexibility from a myriad of solutions created by the blade computing community (<http://blade.org>) members and industry leaders that have downloaded our open specification.

Easy to deploy, integrate and manage:

- ▶ It enables efficient integrated management, which allows you to minimize costs with the tools you require for effective management to:
 - Automate OS installation and BIOS updates remotely with IBM Director tools.
 - Administer your blades at the chassis or rack level with the Advanced Management Module.
 - Plug into your enterprise management software.
- ▶ It enables deployment simplicity without trade-offs by speeding the deployment of new hardware in minutes rather than days, using BladeCenter Open Fabric Manager to:
 - Get significantly faster deployment of servers and I/O than from rack solutions.
 - Reduce costly downtime with integrated failover capability.
 - Manage from a single point of control with the Advanced Management Module.
 - Use with virtually all IBM switches, blades and chassis

Green today for a better tomorrow:

- ▶ It offers control through powerful tools that help you optimize your data center infrastructure so you can be responsive, and:
 - Understand your power requirements with IBM Power Configurator.
 - Monitor, control, and virtualize your power with IBM Systems Director Active Energy Manager™ for x86.
 - Reduce data center hot spots with the IBM Rear Door Heat eXchanger.
 - Optimize and future-proof your data center with IBM Data Center Energy Efficiency services.
- ▶ Our eco-friendly servers and services can help you be environmentally responsible so you can:
 - Become more energy efficient with IBM expertise.

General description

The newest release of the IBM BladeCenter POWER6 processor-based blade family consists of two new models: The JS23 and JS43 Express blade servers.

This chapter contains the following topics:

- ▶ “Overview of JS23 and JS43 Express blade servers” on page 10
- ▶ “Features and specifications” on page 11
- ▶ “Blade server features” on page 14
- ▶ “Physical specifications BladeCenter JS23” on page 16
- ▶ “Physical specifications of BladeCenter JS43” on page 19
- ▶ “IBM BladeCenter chassis” on page 22

2.1 Overview of JS23 and JS43 Express blade servers

The newest release of the IBM BladeCenter POWER6 processor based blade family consists of two new models: The JS23 and JS43 Express blade servers. This family allows processor scalability starting with a 2 processor (4-core single wide) blade and adds the ability to upgrade to a 4 processor (8-core) blade with the addition of a second blade making it a double-wide package.

The blades continue to support AIX, IBM i, and Linux operating systems. Also supported are the IBM EnergyScale technologies, IBM PowerVM Virtualization and AltiVec SIMD acceleration functionality.

The optional IBM Systems Director Active Energy Manager (AEM) for POWER V5.1 and browser-based Integrated Virtualization Manager (IVM) software make it easier than ever to achieve increased utilization and energy efficiency using the PowerVM and EnergyScale capabilities of the system.

The IBM BladeCenter JS23 and JS43 Express blade servers refresh every major subsystem to the latest industry-standard functionality. With the energy-efficient 4.2 GHz high performance dual-core, 64-bit POWER6 processors along with the improved reliability of a SAS disk drive bay supporting multiple disk sizes, the JS23 and JS43 combine leadership performance per watt and enterprise-class reliability features. The BladeCenter JS23 Express supports up to 64 GB of DDR2-based memory in eight DIMM slots. The BladeCenter JS43 Express supports up to 128 GB of DDR2-based memory in sixteen DIMM slots. In addition to standard support for error correction code (ECC) and IBM Chipkill technology designed for improved memory fault protection, the BladeCenter JS23/JS43 Express also offers memory running at up to 667 MHz for outstanding performance.

The BladeCenter JS23 Express (7778-23X), BladeCenter JS43 Express (7778-23X + FC8446), BladeCenter S Chassis (7779-BCS), and BladeCenter H Chassis (7989-BCH) are part of the Cluster 1350™ portfolio.

The JS23 and JS43 blades can be configured in the IBM eConfigurator similar to other IBM System Power Systems servers.

2.2 Features and specifications

This section lists features and specifications of the IBM BladeCenter JS23 and JS43 blades.

Microprocessor

Features include:

- ▶ JS23: Two dual-core (4-core) 64-bit POWER6 microprocessors, 4.2 GHz
- ▶ JS43: Two additional dual-core (total 8-core) 64-bit POWER6 microprocessors, 4.2 GHz
- ▶ Support for Energy Scale thermal management for power management/oversubscription (throttling) and environmental sensing.

Memory

Features include:

- ▶ JS23: Dual-channel (DDR2) with eight very low profile (VLP) DIMM slots (maximum 64 GB)
- ▶ JS43: Dual-channel (DDR2) with 16 VLP DIMM slots (maximum 128 GB)

Supports:

- ▶ 2 GB, 4 GB, and 8 GB DDR2 DIMMs
- ▶ 2-way interleaved, DDR2, PC2-4200 or PC2-5300, ECC SDRAM registered x4, memory scrubbing, Chipkill, and bit steering DIMMs

Storage

Features include:

- ▶ JS23: Supports one internal small-form-factor (SFF) Serial Attached SCSI (SAS) drive or Solid State Drive (SSD) in the base unit
- ▶ JS43: Supports one additional internal SFF SAS drive or SSD in the expansion unit for a total of two drives

Virtualization

PowerVM Standard Edition hardware feature supports Integrated Virtualization Manager (IVM) and Virtual I/O Server. In addition:

- ▶ No HMC support
- ▶ Partition migration
- ▶ No CUoD

Predictive failure analysis

Performed on:

- ▶ Microprocessor
- ▶ Memory

Environment considerations

Environmental specifications include

- ▶ Electrical input is 12 V dc.
- ▶ Air temperature:
 - Blade server on: 10 - 35°C (50 - 95°F)
Altitude: 0 - 914 m (3000 ft)
 - Blade server on: 10 - 32°C (50 - 90°F)
Altitude: 914 - 2133 m (3000 - 7000 ft)
 - Blade server off: -40° to 60°C (-40° to 140°F)
- ▶ Humidity:
 - Blade server on: 8 - 80%
 - Blade server off: 8 - 80%

Physical characteristics

JS23 size characteristics are:

- ▶ Height: 24.5 cm (9.7 inches)
- ▶ Depth: 44.6 cm (17.6 inches)
- ▶ Width: 30 mm (1.14 inches) single-slot blade
- ▶ Maximum weight: 5.0 kg (11 lb)

JS43 size characteristics are:

- ▶ Height: 24.5 cm (9.7 inches)
- ▶ Depth: 44.6 cm (17.6 inches)
- ▶ Width: 60 mm (2.28 inches) double-slot blade
- ▶ Maximum weight: 10.0 kg (22 lb)

Supported I/O options

I/O adapter card options:

- ▶ Up to two PCIe High Speed adapters on JS43. Only one supported on JS23
- ▶ Up to two PCIe CIOv adapters on JS43. Only one on JS23

Integrated functions

Functions include:

- ▶ JS23: Two 1 GB Ethernet controllers connected to the BladeCenter chassis fabric through the 5-port integrated Ethernet switch
- ▶ JS43: Two additional 1 GB Ethernet controllers, connecting directly to BladeCenter Ethernet switch modules
- ▶ Expansion card interface
- ▶ The baseboard management controller (BMC) is a flexible service processor with Intelligent Platform Management Interface (IPMI) firmware and SOL support
- ▶ PCI attached ATI™ RN 50 graphics controller
- ▶ SAS controller
- ▶ Light path diagnostics RS-485 interface for communication with the management module
- ▶ Automatic server restart (ASR)
- ▶ Serial Over LAN (SOL)
- ▶ Four Universal Serial Bus (USB) buses on base planar for communication with keyboard and removable-media drives
- ▶ Support for USB-attached local keyboard, video, and mouse (KVM)
- ▶ Transferable Anchor function (Renesas Technology HD651330 microcontroller) in the management card

Supported operating systems

Supported operating systems include:

- ▶ Linux SLES10 SP2 or later versions
- ▶ Red Hat RHEL 5.2 or later versions
- ▶ Red Hat RHEL 4.6 or later versions
- ▶ AIX 5.3.S, 6.1.F
- ▶ IBM i 6.1

2.3 Blade server features

This section lists the main features of JS23 and JS43 blade servers.

Baseboard management controller

The enhanced baseboard management controller (BMC) for the JS23 blade server and JS43 blade server is a flexible service processor that provides support for the following functions:

- ▶ Alert Standard Format (ASF) and RS-485 interfaces to the management modules
- ▶ Intelligent Platform Management Interface (IPMI)
- ▶ The operating system
- ▶ Power control and advanced power management
- ▶ Reliability, availability, and serviceability (RAS) features
- ▶ Serial Over LAN (SOL)
- ▶ Continuous health monitoring and control
- ▶ Configurable notification and alerts
- ▶ Event logs that are time stamped and saved in nonvolatile memory and that can be attached to e-mail alerts
- ▶ Point-to-Point Protocol (PPP) support
- ▶ Remote power control
- ▶ Remote firmware update and access to critical blade server settings

Disk drive support

The JS23 blade server supports one 2.5-inch hard disk drive. The JS43 blade server can support up to two 2.5-inch hard disk drives. The disk drives can be either the small-form-factor (SFF) Serial Attached SCSI (SAS) or the Solid state drive (SSD).

IBM Director

IBM Director is a workgroup hardware management tool that you can use to centrally manage the JS23 blade server and JS43 blade server, including updating the JS23 and JS43 firmware.

Integrated network support

The blade server has two integrated (onboard) 1 Gbps Host Ethernet Adapter (HEA) controllers that provide advanced acceleration features, such as

checksum offload, TCP large send, and jumbo frames. Capabilities include virtualized adapter sharing among logical partitions that does not require the shared Ethernet adapters of Virtual I/O Server. TCP advanced features include hardware de-multiplexing and per connection queues.

I/O expansion

The blade server has connections on the system board for two optional PCIe expansion cards, such as Fibre Channel and InfiniBand® expansion cards, for adding more network communication capabilities to the blade server.

Large system memory capacity

The memory bus in the JS23 blade server supports up to 64 GB of system memory. The memory bus in the JS43 blade server supports up to 128 GB of system memory.

Light path diagnostics

Light path diagnostics provides light-emitting diodes (LEDs) to help you diagnose problems. An LED on the blade server control panel is lit if an unusual condition or a problem occurs. If this happens, you can look at the LEDs on the system board to locate the source of the problem.

Power throttling

If your BladeCenter unit supports power management, the power consumption of the blade server can be dynamically managed through the management module.

Reliability, availability, and serviceability

The blade server incorporates reliability, availability, and serviceability (RAS) features that provide capability to monitor, diagnose and correct errors on some of the components within the blade server. Several features are:

- ▶ Automatic service processor reset and reload recovery for service processor errors
- ▶ Automatic server recovery and restart that provides automatic reboot after boot hangs or detection of checkstop conditions
- ▶ Automatic server restart (ASR)
- ▶ Built-in monitoring for temperature, voltage, hard disk drives, and flash drives
- ▶ Checkstop analysis
- ▶ Customer-upgradeable basic input/output system (BIOS) code (firmware code)
- ▶ Degraded boot support (memory and microprocessors)

- ▶ Extended Error Handling (EEH) for adapter failures
- ▶ Emergency power off (EPOW) for the blade server and expansion cards
- ▶ Environmental monitors and alerts
- ▶ First-failure data capture (FFDC) for determining failure root cause
- ▶ Service processor communication with the management module to enable remote blade server management
- ▶ Light-emitting diodes (LEDs) for identifying failing customer replaceable units
- ▶ Light path diagnostics
- ▶ POWER6 Hypervisor (PHYP) partition recovery or partition termination when unrecoverable errors occur
- ▶ Power-on self-test (POST)
- ▶ Vital product data (VPD) unique identifiers on blade server and all major electronic components with information stored in nonvolatile memory for remote viewing

2.4 Physical specifications BladeCenter JS23

This section provides more details related to the features of the JS23 blade servers.

2.4.1 Minimal and optional features of BladeCenter JS23

The BladeCenter JS23 blade is based on a modular design of two dual-core modules that contain 64-bit POWER6 processors and integrated L2 cache soldered directly to the system planar board. Additionally, a 32-MB L3 cache is integrated into each dual-core module (DCM). The JS23 is contained in a single wide package.

Table 2-1 lists the JS23 configuration options.

Table 2-1 JS23 standard configuration (#7778-23X)

Processor	L2/L3	Memory	Ethernet	Disk
2-socket, 4-core, 4.2 GHz, POWER6	L2 - 4 MB per core L3 32 MB per DCM	4 GB - 64 GB	Dual gigabit	73,146 SAS, 300 GB SAS, or 69 GB SSD

Figure 2-1 shows the physical layout of the JS23 blade including memory slots, disk, and the expansion option connectors.




Figure 2-1 JS23 Blade physical layout

2.4.2 Processor features

The key processor features are:

- ▶ The BladeCenter JS23 blade provides the support for a 4-socket, 8-core, POWER6 4.2 GHz processor implementation.
- ▶ Each processor is directly mounted to the system planar board, providing multi-processing capability.
- ▶ Each processor core includes a 64 KB Instruction-Cache, 64 KB Data-Cache, and 4 MB of L2 cache.
- ▶ Each dual-core module contains a 32 MB L3 cache.

Table 2-2 shows the supported processor on a BladeCenter JS23 blade.

Table 2-2 BladeCenter JS23 processor support

Feature	Description
#7778-23X	IBM BladeCenter JS23 4-core 64 bit 4.2 GHz

2.4.3 Memory features

The integrated memory controller supports eight pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 4 GB (2x2 GB) and the maximum is 64 GB (8x8 GB). All the memory features support memory scrubbing, error correction, Chipkill, and bit steering. Memory is packaged in Very Low Profile (VLP) RDIMM packages. This newer style memory is much lower in height allowing more memory in the same space.

2.4.4 Memory DIMMs

Table 2-3 provides a list of supported memory on a BladeCenter JS23 blade.

Table 2-3 BladeCenter JS23 memory support

Feature	Description
#8229	4 GB (2x2 GB) DDR2 667 MHz VLP RDIMMs
#8239	8 GB (2x4 GB) DDR2 667 MHz VLP RDIMMs
#8245	16 GB (2x8 GB) DDR2 400 MHz VLP RDIMMs

2.4.5 Internal disk

Table 2-4 lists supported disks on a BladeCenter JS23 blade. Disk drives are not required on the base offering.

Table 2-4 BladeCenter JS23 disk support

Feature	Description
#8237	73 GB SAS 10K SFF hard disk drive
#8236	146 GB SAS 10K SFF hard disk drive
#8274	300 GB SAS 10K SFF hard disk drive
#8273	69 GB Solid State Disk (SSD)

2.5 Physical specifications of BladeCenter JS43

In this section, we discuss the physical specifications of BladeCenter JS43.

2.5.1 Minimal and optional features of BladeCenter JS43

The BladeCenter JS43 blade is based on a modular design of four dual-core modules (DCM) that contain 64-bit POWER6 processors and integrated L2 cache soldered directly to the system planar board. Additionally, a 32 MB L3 cache is integrated into each of the dual-core modules. The JS43 is contained in a double wide package. The JS43 can be obtained by either starting with a JS23 (7778-23x) and adding a Feature 8446 which contains an additional blade to make a double wide package or ordering a Feature 7778-43X which contains the JS23 (7778-23X) and Feature 8446 assembled at the factory.

Table 2-5 shows the JS43 configuration options.

Table 2-5 JS43 configuration (7778-23X; Plus #8446 add-on expansion module)

Processor	L2/L3	Memory	Ethernet	Disk
4-socket, 8-core, 4.2 GHz, POWER6	L2 4 MB per core L3 32 MB per DCM	8 - 128 GB	Dual gigabit	73,146 SAS, 300 GB SAS, or 69 GB SSD

Figure 2-2 on page 20 shows the physical layout of the JS43 blade multiprocessor expansion unit (MPE) including memory slots, disk, and the expansion option connectors. The MPE stacks on top of the single wide JS23 making a double-wide blade. Each section has its own processors, memory, disk, and adapter cards.

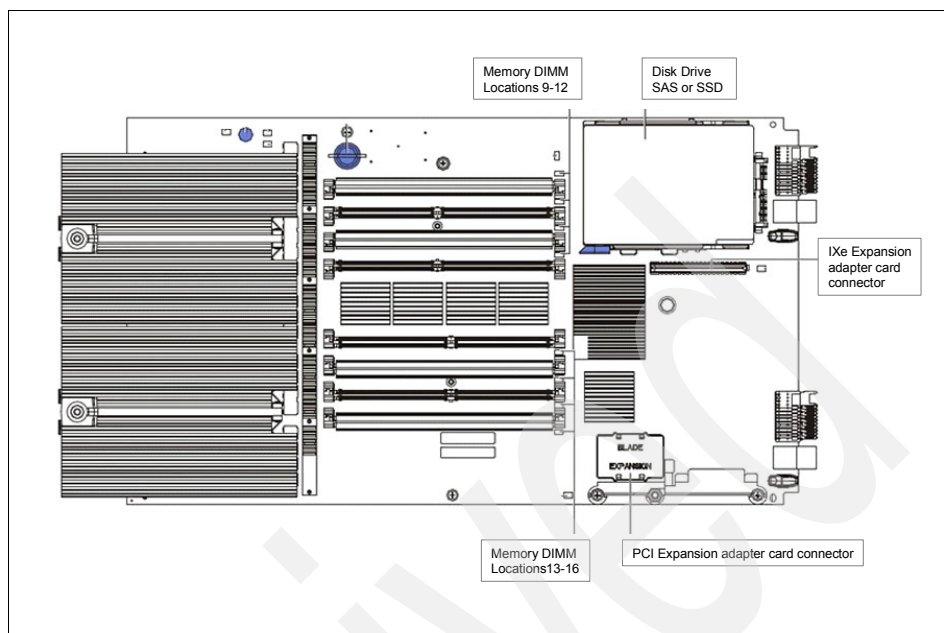


Figure 2-2 JS43 multiprocessor expansion unit (MPE)

2.5.2 Processor features

The key processor features are:

- ▶ The BladeCenter JS43 blade provides the support for a 2-socket, 4-core, POWER6 4.2 GHz processor implementation. Each of the two blade boards contain processors.
- ▶ Each processor is directly mounted to the system planar board, providing multi-processing capability.
- ▶ Each processor core includes a 64 KB Instruction-Cache, 64 KB Data-Cache, and 4 MB of L2 cache.

Each dual-core module contains a 32 MB L3 cache. Table 2-6 shows the supported processor on a BladeCenter JS43 blade.

Table 2-6 BladeCenter JS43 processor support

Feature	Description
#7778-23X Plus 8446	IBM BladeCenter JS43 8-core 64 bit 4.2 GHz

2.5.3 Memory features

The integrated memory controller supports sixteen pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 4 GB (2x2 GB) and the maximum is 128 GB (16x8 GB). All memory features support memory scrubbing, error correction, Chipkill, and bit steering.

2.5.4 Memory DIMMs

Table 2-7 lists supported memory on a BladeCenter JS43 blade.

Table 2-7 BladeCenter JS43 memory support

Feature	Description
#8229	4 GB (2x2 GB) DDR2 667 MHz VLP RDIMMs
#8239	8 GB (2x4 GB) DDR2 667 MHz VLP RDIMMs
#8245	16 GB (2x8 GB) DDR2 400 MHz VLP RDIMMs

2.5.5 Internal disk

Table 2-8 lists supported disks on a BladeCenter JS43 blade. Disk drives are not required on the base offering.

Table 2-8 BladeCenter JS43 disk support

Feature	Description
#8237	73 GB SAS 10K SFF hard disk drive
#8236	146 GB SAS 10K SFF hard disk drive
#8274	300 GB SAS 10K SFF hard disk drive
#8273	69 GB Solid State Disk (SSD)

When you have a JS43 and both halves contain disk drives, RAID is supported. You can utilize either RAID 0 (striping) or RAID 1 (mirroring). Either the SAS drives or the SSD disk units can be configured for RAID, however the drives must be of the same type. Although having drives of the same capacity is preferred, RAID can be performed using dissimilar capacities. If differing capacities are used you will only have the effective capacity of the smaller drive.

2.6 IBM BladeCenter chassis

The BladeCenter JS23 and BladeCenter JS43 Express blade are supported in the BladeCenter chassis as shown in Table 2-9. Note that operating system selection (in particular IBM i) should be taken into account for overall system support.

Table 2-9 BladeCenter JS23 and JS43 chassis support

Blade	BCH	BCS	BCHT	BCT	BCE
JS23	YES	YES	YES	NO	NO
JS43	YES	YES	YES	NO	NO

The table indicates three chassis are available in the BladeCenter chassis family:

- ▶ IBM BladeCenter H (BCH) delivers high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments.
- ▶ IBM BladeCenter S (BCS) combines the power of blade servers with integrated storage, all in an easy-to-use package designed specifically for the office and distributed enterprise environment.
- ▶ IBM BladeCenter HT (BCHT) models are designed for high-performance flexible telecommunication environments by supporting high-speed inter-networking technologies such as 10 Gb Ethernet.

Table 2-10 provides a list of supported BladeCenter chassis and the total number of JS23 or JS43 blades installable into a chassis. This table describes the physical limitations of blades in chassis.

Note: The number of blade servers that can be installed into the chassis is dependent on the power supply configuration, power supply input (110/220 V) and power reduction/redundancy options. See 2.6.4, “Number of JS23 and JS43 blades in a chassis” on page 34 for more information.

Table 2-10 BladeCenter support

Chassis	Number of JS23 blades	Number of JS43 Blades
BladeCenter S chassis	6	3
BladeCenter H chassis	14	7
BladeCenter HT chassis	12	6

2.6.1 BladeCenter H

IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility to even the most demanding IT environments. In a 9U rack space, the BladeCenter H chassis can contain up to 14 blade servers, 10 I/O modules, and four power supplies to provide the necessary I/O network switching, power, cooling, and control panel information to support the individual servers.

The chassis supports up to four traditional fabrics using networking switches, storage switches, or pass-through devices. The chassis also supports up to four high-speed fabrics for support of protocols such as 4X InfiniBand or 10 Gigabit Ethernet. The built-in media tray includes light path diagnostics, two front USB inputs, and a DVD drive.

Figure 2-3 shows the front view of an IBM BladeCenter H (and Figure 2-4 on page 24 shows the rear view).




Figure 2-3 Front view of BladeCenter H

The key features on the front of the BladeCenter H are:

- ▶ A media tray on the right, with a DVD drive, two USB v2.0 ports, and a system status LED panel
- ▶ One pair of 2900 W power modules. An additional power module option (containing two 2900 W power modules) is available.
- ▶ Two hot-swap fan modules (two extra hot-swap fan modules are included with the additional power module option)
- ▶ Fourteen hot swap blade server bays supporting various blade server types

Figure 2-4 shows the rear view of an IBM BladeCenter H.




Figure 2-4 Rear view of BladeCenter H

The key features on the rear of the BladeCenter H are:

- ▶ Two hot-swap blower modules as standard
- ▶ Two hot-swap management module bays, with one management module as standard
- ▶ Four traditional fabric switch modules
- ▶ Four high-speed fabric switch modules

The BladeCenter H chassis allows for either 14 single-slot blade servers or seven double-slot blade servers. However, you can mix the blade server models in one chassis to meet your requirements.

The BladeCenter H chassis ships standard with one Advanced Management Module. This module provides the ability to manage the chassis as well as providing the local KVM function. The optional redundant Advanced Management Module provides the IBM BladeCenter H with higher levels of resiliency. While in the chassis, the second module is in passive or standby mode. If the active or primary module fails, the second module is automatically enabled with all of the configuration settings of the primary module. This function provides clients with easy remote management and connectivity to the BladeCenter H chassis for their critical applications.

BladeCenter H does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity requirements. An Ethernet Switch Module (ESM) or Passthrough Module will be required in I/O module bays 1 and 2, to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3 and 4 depend on the I/O Expansion Card installed in the blade servers. The I/O modules required in the high-speed I/O module bays 7, 8, 9, and 10 depend on the HSFF or CFFv I/O Expansion Adapter cards installed in the blade servers.

The BladeCenter H chassis ships with no power cord; power cords have to be ordered separately.

The BladeCenter H chassis ships standard with:

- ▶ One advanced management module
- ▶ Two blower modules
- ▶ Two power supply modules (one pair of 2900 W power modules)
- ▶ Two hot-swap power supply fan modules
- ▶ Two USB v2.0 ports
- ▶ One DVD-ROM drive

The chassis does not have a diskette drive. An optional USB-attached 1.44 MB diskette drive is available. The feature code for ordering the external diskette drive is FC 2591 of FRU P/N 03N4962.

Table 2-11 on page 26 lists the BladeCenter H (BCH) specifications.

Table 2-11 BladeCenter H specifications

Feature	Specification
Machine type	8852-4XU
Rack dimension	9U x 28 inches (711 mm)
DVD/CD standard drive	1 x DVD ROM (in media tray)
Diskette drive	None
Number of blade slots	14 (30 mm blade servers)
Number of switch module slots	10 hot-swap (4 x high speed, 4 x standard, 2 x bridge)
Switch module standard	None (in standard chassis offerings)
Power supply	2900 W ac
Number of power supplies (standard/maximum)	2 / 4 ^a
Number of blowers (standard/maximum)	2 / 2
Dimensions	Height: 400 mm (15.75 inches) Width: 422 mm (17.40 inches) Depth: 711 mm (28.00 inches)

a. Four power supplies are required to use high-speed bays 7 - 10, and any blade server in slots 8 - 14.

2.6.2 BladeCenter S

The BladeCenter S chassis is a robust and flexible physical platform. Its modular tool-free design allows easy access and maintenance. All external components (except running blade servers) are hot swappable and release levers/handles are clearly marked.

Figure 2-5 on page 27 shows the front view of an IBM BladeCenter S (Figure 2-6 on page 28 shows the rear view).




Figure 2-5 BladeCenter S front view

The key features on the front of the BladeCenter S are:

- ▶ A media tray at the front, with a DVD drive, two USB 2.0 ports, a system status LED panel, and two bay locations that are reserved for future use.
- ▶ Six hot-swap blade server bays supporting different blade server types.
- ▶ Two bays for disk storage modules; each storage module can house up to six 3.5-inch disk drives of internal storage. No storage modules are standard with the BladeCenter S chassis.

Figure 2-6 on page 28 shows the rear view of the IBM BladeCenter S.




Figure 2-6 BladeCenter S rear view

The key features on the rear of the BladeCenter S are:

- ▶ Four hot-swap blower modules, as standard
- ▶ One hot-swap management-module bay with one management module, as standard
- ▶ Four I/O bays for standard switch modules (bays 1, 3, and 4 can be used for installing I/O modules, bay 2 is reserved for future use)
- ▶ One pair of 950/1450-watt power modules. An additional power module option (configured in pairs of two 950/1450 W feature 4548 power modules) is available.
- ▶ One bay for a serial pass-through module to give a direct serial connection to installed blades (for those blades with the functionality)

The BladeCenter S chassis allows for either six single-slot blade servers or three double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements.

The BladeCenter S chassis ships standard with an Advanced Management Module. This module provides the ability to manage the chassis as well as providing the local KVM function. Only one management module is supported with a BladeCenter S chassis.

The BladeCenter S chassis does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity requirements. An Ethernet Switch Module (ESM) is required in I/O module bay 1, to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3

and 4 depend on the I/O Expansion Card installed in the blade servers. Bay 2 is reserved for future use. The chassis does not ship with any storage modules.

The BladeCenter S chassis uses either 100 - 127 V or 200 - 240 V ac power and can be attached to standard office power outlets.

The BladeCenter S chassis ships with:

- ▶ One advanced management module
- ▶ Four blower modules
- ▶ Two power supply modules (one pair of 950/1450-watt power modules)
- ▶ Two 2.5 m rack jumper power cords (IEC 320 C19 - C20)
- ▶ Four country-specific line cords
- ▶ One CD-RW/DVD-ROM drive

The chassis does not have a diskette drive. An optional USB-attached 1.44 MB diskette drive is available.

Table 2-12 lists the BladeCenter S (BCS) specifications.

Table 2-12 BladeCenter S specifications

Feature	Specification
Machine type	8886-1MY
Rack dimension	7U x 28.9 inches (733.4 mm)
DVD/CD standard drive	1 x DVD ROM (in media tray)
Diskette drive	None
Number of blade slots	6 (30 mm blade servers)
Number of switch module slots	4 hot-swap (3 x standard, 1 x standard for future use)
Switch modules (standard/maximum)	0 / 4
Storage modules (standard/maximum)	0 / 2
Power supply	950 W connected to 100 V ac 1450 W connected to 220 V ac
Number of power supplies (standard/maximum)	2 / 4

Feature	Specification
Number of blowers (standard/maximum)	4 / 4
Dimensions	Height: 306.3 mm (12.00 inches) Width: 440 mm (17.50 inches) Depth: 733.4 mm (28.90 inches)

2.6.3 BladeCenter HT

The IBM BladeCenter HT is a 12-server blade chassis designed for high-density server installations, typically for telecommunications use. It offers high performance with the support of 10 Gb Ethernet installations. This 12U high chassis with DC or AC power supplies provides a cost-effective, high-performance, high-availability solution for telecommunication network and other rugged non-telecommunication environments. The IBM BladeCenter HT chassis is positioned for expansion, capacity, redundancy, and carrier-grade NEBS level/ETSI compliance in DC models.

IBM BladeCenter HT provides a solid foundation for Next Generation Networks (NGN), enabling service providers to become on demand providers. Coupled with technological expertise within the enterprise data center, IBM leverages the industry know-how of key Business Partners to jointly deliver added value within service provider networks.

The BladeCenter HT brings significant new capabilities to the broad IBM ecosystem of hundreds of NGN applications already being deployed on BladeCenter. A key example is the introduction of the Nortel 10 Gb Ethernet Switch Module for BladeCenter, which delivers 10 Gb to each blade server deployed in the BladeCenter H or BladeCenter HT chassis, and six 10 Gb Ethernet uplinks. This capability helps to greatly reduce the cost of implementing IPTV and other high bandwidth NGN applications.

The key features of the BladeCenter HT include:

- ▶ Support for up to 12 blade servers, compatible with the other chassis in the BladeCenter family
- ▶ Four standard I/O and four high-speed I/O module bays, compatible with the other chassis in the BladeCenter family
- ▶ A media tray at the front with light path diagnostics, two USB 2.0 ports, and optional compact flash memory module support
- ▶ Two hot-swap management module bays (one management module standard)

- Four hot-swap power-module bays (two power modules standard)
- New serial port for direct serial connection to installed blades
- Compliance with the NEBS 3 and ETSI core network specifications

Figure 2-7 shows the front view of the IBM BladeCenter HT (Figure 2-8 on page 32 shows the rear view).




Figure 2-7 IBM BladeCenter HT front view




Figure 2-8 IBM BladeCenter HT rear view

Table 2-13 lists the features of the IBM BladeCenter HT.

Table 2-13 BladeCenter HT specifications

Feature	Specification
Machine type	8740-1RY (DC) 8750-1RY (AC)
Rack dimension	12U x 27.8 inches (706 mm)
DVD/CD standard drive	None
Diskette drive	None
Number of blade slots	12 (30 mm blade servers)
Number of switch module slots	4

Feature	Specification
Number of high-speed switch module slots	4
Switch modules (standard/maximum)	None
Number of power supplies (standard/maximum)	2 / 4 ^a
Number of blowers (standard/maximum)	4 / 4
Dimensions	Height: 528 mm (21.00 inches) Width: 440 mm (17.50 inches) Depth: 706 mm (27.8 inches)

- a. Four power supplies are required to use the high-speed bays 7 to 10, and any blade servers in slots 7 to 12.

The BladeCenter HT chassis allows for either 12 single-slot blade servers or six double-slot blade servers. However, you can mix different blade server models in one chassis. The BladeCenter HT chassis ships standard with one Advanced Management Module. This module provides the ability to manage the chassis, as well as providing the local KVM function. The optional redundant Advanced Management Module provides the IBM BladeCenter HT with higher levels of resiliency. This module provides clients with easy remote management and connectivity to the chassis.

The BladeCenter HT does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity requirements. An Ethernet Switch Module (ESM) is required in I/O module bays 1 and 2 to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3 and 4 depend on the I/O Expansion Card installed in the blade servers.

High-speed switch modules can be installed into I/O bays 7 through 10, and are used together with high-speed expansion cards installed into blade servers. I/O bays 1 through 4 can also be used for bridge modules.

For more information related to IBM BladeCenter technology and BladeCenter products, see *IBM BladeCenter Products and Technology*, SG24-7523.

2.6.4 Number of JS23 and JS43 blades in a chassis

This section describes the number of IBM BladeCenter JS23 and JS43 Express blades in the supported BladeCenter chassis. IBM BladeCenter JS23 Express and JS43 Express each have their own power consumption characteristics. The amount of power requirements for each type of blade dictates the number of blades supported in each BladeCenter chassis.

Table 2-14 shows the power consumption values for the IBM BladeCenter JS23 and JS43 Express, running with maximum configuration (8 x 8 GB, HSDC, 1Xe and HDD on the JS23, and this same configuration on each JS43's planar).

Table 2-14 Power Consumption for IBM BladeCenter JS23 and JS43 Express

Blades	Power consumption (W)	
	Full power	-30% P6 and L3, Static Power Save
JS23	400	334
JS43	777	654

Figure 2-9 and Figure 2-10 on page 35 show the maximum number of blades, running in the maximum configuration, for each supported BladeCenter chassis.

	BCS				BCH				BCH-T			
	Total of 6 slots				Total of 14 slots, 7 in each Power Domain (PD)				Total of 12 slots, 6 in each Power Domain (PD)			
	110VAC PS		220VAC PS		Only PD1 (No PD2)		PD1 and PD2		AC Power Supply		DC Power Supply	
	2PS	4PS	2PS	4PS	PD1	None	PD1	PD2	PD1	PD2	PD1	PD2
Fully Redundant without Performance Reduction	1	5	2	5	6	-	6	6	5	6	4	4
Redundant with Performance Reduction	1	6	2	6	7	-	7	7	6	6	5 (*)	5 (*)
Basic Power Mode (Max Power Capacity)	3	6	5	6	7	-	7	7	6	6	6	6
(*) It is possible to select a special JS23 configuration such that all 12 blades can be supported in the Redundant with Performance Reduction mode on a BCH-T chassis. For instance, JS23 with one HSDC card, 4 x 8GB of memory and 1 HDD.												

Figure 2-9 Maximum number of IBM BladeCenter JS23 Express blades per IBM BladeCenter chassis

	BCS				BCH				BCH-T			
	Total of 6 slots				Total of 14 slots, 7 in each Power				Total of 12 slots, 6 in each Power			
	110VAC PS		220VAC PS		Only PD1 (No PD2)		PD1 and PD2		AC Power Supply		DC Power Supply	
	2PS	4PS	2PS	4PS	PD1	None	PD1	PD2	PD1	PD2	PD1	PD2
Fully Redundant without Performance Reduction	1 JS23	2 JS43 + 1 JS23	2 JS43	2 JS43 + 1 JS23	3 JS43	-	3 JS43	3 JS43	2 JS43 + 1 JS23	3 JS43	2 JS43	2 JS43
Redundant with Performance Reduction	1 JS23	3 JS43	2 JS43	3 JS43	3 JS43 + 1 JS23	-	7 JS43		3 JS43	3 JS43	2 JS43 + 1 SJ23 (*)	2 JS43 + 1 SJ23 (*)
Basic Power Mode (Max Power Capacity)	1 JS43 + 1 JS23	3 JS43	2 JS43 + 1 JS23	3 JS43	3 JS43 + 1 JS23	-	7 JS43		3 JS43	3 JS43	3 JS43	3 JS43
	(*) It is possible to select a special JS43 configuration such that all 6 blades can be supported in the "Redundant with Performance Reduction" mode on a BCH-T chassis. For instance, JS43 with one HSDC card, 4 x 8GB of memory on each planar and 1 HDD.											

Figure 2-10 Maximum number of IBM BladeCenter JS23 and JS43 Express per IBM BladeCenter chassis.

Note: A best practice is to use the BladeCenter chassis with all power supplies installed to get full advantage of all performance enhancements made in the IBM BladeCenter JS23 and JS43 Express. This practice can also maximize the number of blades supported within the chassis.

2.6.5 IBM Director

Use IBM Director to perform network and system management tasks.

With IBM Director, a network administrator can perform the following tasks:

- ▶ View the hardware configuration of remote systems, in detail.
- ▶ Monitor the usage and performance of critical components, such as microprocessors, disks, and memory.
- ▶ Centrally manage individual or large groups of IBM and other x86 processor-based servers, desktop computers, workstations, and notebook computers on a variety of platforms.

IBM Director provides a comprehensive entry-level workgroup hardware manager. It includes the following key features:

- ▶ Advanced self-management capabilities for maximum system availability
- ▶ Multiple operating-system platform support, including Microsoft® Windows Server® 2003, Windows 2000 Server, Windows XP Professional, Windows Vista®, Red Hat Linux, SUSE Linux, VMware®, Novell® NetWare, AIX, and IBM i operating system (formerly known as i5/OS)
- ▶ Support for IBM and other servers, desktop computers, workstations, and notebook computers
- ▶ Support for systems-management industry standards

- ▶ Integration into leading workgroup and enterprise systems-management environments
- ▶ Ease-of-use, training, and setup

IBM Director also provides an extensible platform that supports advanced server tools, which are designed to reduce the total cost of managing and supporting networked systems. By deploying IBM Director, you can achieve reductions in ownership costs through the following benefits:

- ▶ Reduced downtime
- ▶ Increased productivity of IT personnel and users
- ▶ Reduced service and support costs
- ▶ Reduced power consumption

Technical description of the hardware architecture

IBM BladeCenter JS23 Express is a single-wide blade; the IBM BladeCenter JS43 Express is a double-wide blade, consisting of the JS23's Base planar and a *multiprocessor expansion unit* planar (MPE). The MPE planar design is similar to the base planar, but with reduced functions. In this chapter, we present the technical details of JS23's Base planar, highlighting differences from the MPE planar.

This chapter discusses hardware architecture, including the following topics:

- ▶ “POWER6 processor” on page 38
- ▶ “Simultaneous multi-threading” on page 40
- ▶ “Altivec SIMD” on page 41
- ▶ “IBM EnergyScale technology” on page 42
- ▶ “Flexible Service Processor” on page 44
- ▶ “Management Card” on page 44
- ▶ “Memory subsystem” on page 45
- ▶ “I/O subsystem” on page 48
- ▶ “PowerVM” on page 56
- ▶ “Operating system support” on page 57
- ▶ “Systems management” on page 59

3.1 POWER6 processor

The POWER6 processor capitalizes on the enhancements brought by the POWER5 processor.

Two enhancements of the POWER6 processor are the ability to do processor instruction retry and to alternate processor recovery. These significantly reduce exposure to both hard (logic) and soft (transient) errors in the processor core.

Enhancements include:

- ▶ Processor instruction retry

Soft failures in the processor core are transient errors. When an error is encountered in the core, the POWER6 processor first automatically retries the instruction. If the source of the error was truly transient, the instruction succeeds and the system continues as before. On predecessor IBM systems, this error would have caused a checkstop.

- ▶ Alternate processor retry

Hard failures are more difficult, being true logical errors that are replicated each time the instruction is repeated. Retrying the instruction does not help in this situation because the instruction will continue to fail. Systems with POWER6 processors introduce the ability to extract the failing instruction from the faulty core and retry it elsewhere in the system, after which the failing core is dynamically deconfigured and called out for replacement. The entire process is transparent to the partition owning the failing instruction. Systems with POWER6 processors are designed to avoid what would have been a full system outage.

- ▶ POWER6 single processor checkstopping

Another major advancement in POWER6 processors is single processor checkstopping. A processor checkstop would result in a system checkstop. A feature in the BladeCenter JS23 Express server is the ability to contain most processor checkstops to the partition that was using the processor at the time. This significantly reduces the probability of any one processor affecting total system availability.

- ▶ POWER6 cache availability

In the event that an uncorrectable error occurs in L2 cache, the system is able to dynamically remove the offending line of cache without requiring a reboot. In addition, POWER6 utilizes an L1/L2 cache design and a write-through cache policy on all levels, helping to ensure that data is written to main memory as soon as possible. It can also make use of an additional 32 MB chip-shared L3 cache.

- POWER6 processor modules on JS23 Express and JS43 Express

IBM BladeCenter JS23 Express comes with 2 POWER6 processor modules (4-core), and IBM BladeCenter JS43 Express comes with two additional POWER6 modules (total 8-core).

Each POWER6 modules is 2-core Dual Chip Module (DCM), containing one 64bit 2-core POWER6 processors (4.2GHz) and one 32MB L3 cache.

Figure 3-1 shows a high-level view of the POWER6 module present in the JS23 and JS43 Express servers.

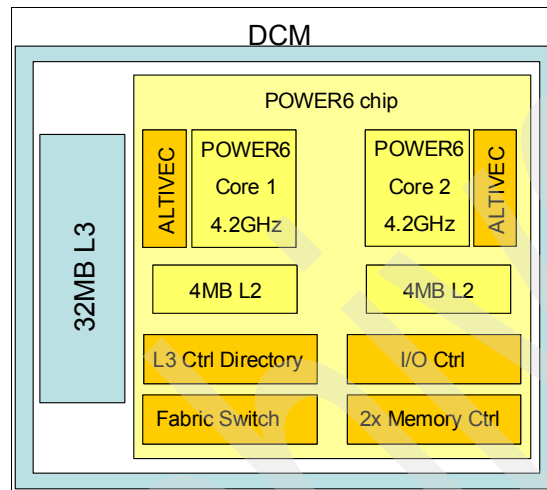


Figure 3-1 POWER6 processor module (2-core DCM)

3.1.1 Decimal floating point

This section describes the behavior of the decimal floating-point processor, the supported data types, formats, and classes, and the usage of registers. The decimal floating-point (DFP) processor shares the 32 floating-point registers (FPRs) and the floating-point status and control register (FPSCR) with the binary floating-point (BFP) processor. However, the interpretation of data formats in the FPRs, and the meaning of some control and status bits in the FPSCR are different between the BFP and DFP processors.

The DFP processor supports three DFP data formats: DFP32 (single precision), DFP64 (double precision), and DFP128 (quad precision). Most operations are performed on the DFP64 or DFP128 format directly. Support for DFP32 is limited to conversion to and from DFP64. For some operations, the DFP processor also supports operands in other data types, including signed or unsigned binary fixed-point data, and signed or unsigned decimal data.

DFP *instructions*, described in the following list, are provided to perform arithmetic, compare, test, quantum-adjustment, conversion, and format operations on operands held in FPRs or FPR pairs:

Arithmetic	Perform addition, subtraction, multiplication, and division operations.
Compare	Perform a comparison operation on the numerical value of two DFP operands.
Test	Test the data class, the data group, the exponent, or the number of significant digits of a DFP operand.
Quantum-adjustment	Convert a DFP number to a result in the form that has the designated exponent, which may be explicitly or implicitly specified.
Conversion	Perform conversion between different data formats or data types.
Format	Facilitate composing or decomposing a DFP operand.

For example, the SAP® NetWeaver 7.10 ABAP™ kernel introduces a new SAP ABAP data type named DECFLOAT to enable more accurate and consistent results from decimal floating point computations. The decimal floating point (DFP) support by SAP NetWeaver® leverages the built-in DFP feature of POWER6 processors. This approach allows for highly simplified ABAP-coding while increasing numeric accuracy, and with a potential for significant performance improvements.

3.2 Simultaneous multi-threading

The POWER6 processor core has been designed to support both enhanced Simultaneous multi-threading (SMT) and single-threaded (ST) operation modes. Both IBM BladeCenter JS23 and JS43 Express support the SMT technology.

Enhanced SMT features

To improve SMT performance for various workloads and provide robust quality of service, POWER6 provides two features:

- **Dynamic resource balancing**

The objective of dynamic resource balancing is to ensure that the two threads executing on the same processor flow smoothly through the system.

Depending on the situation, the POWER6 processor resource balancing logic has a different thread throttling mechanism (a thread reached threshold of L2

cache misses and will be throttled to allow other threads to pass the stalled thread).

- Adjustable thread priority

Adjustable thread priority lets software determine when one thread should have a greater (or lesser) share of execution resources. POWER6 supports eight software-controlled priority levels for each thread.

Single-threaded operation

Not all applications benefit from SMT. Having threads executing on the same processor does not increase the performance of applications with execution unit limited performance or applications that consume all the chip's memory bandwidth. For this reason, POWER6 supports the ST execution mode. In this mode, the POWER6 processor gives all the physical resources to the active thread. Highly optimized scientific codes are one example where single-threaded operation is ideal.

3.3 AltiVec SIMD

IBM Semiconductor's advanced Single Instruction, Multiple Data (SIMD) technology based on the AltiVec instruction set is designed to enable exceptional general-purpose processing power for high-performance POWER processors. This leading-edge technology is engineered to support high-bandwidth data processing and algorithmic-intensive computations, all in a single-chip solution.

With its computing power, AltiVec technology also enables high-performance POWER processors to address markets and applications in which performance must be balanced with power consumption, system cost, and peripheral integration.

The AltiVec technology is a well-known environment for software developers who want to add efficiency and speed to their applications. A 128-bit vector execution unit was added to the architecture. This engine operates concurrently with the existing integer and floating-point units and enables highly parallel operations, up to 16 operations in a single clock cycle. By leveraging AltiVec technology, developers can optimize applications to deliver acceleration in performance-driven, high-bandwidth computing.

3.4 IBM EnergyScale technology

IBM EnergyScale technology is featured in IBM POWER6 processor-based systems. It provides functions to help you understand and control IBM server power and cooling usage.

In this section, we describe IBM EnergyScale features and hardware and software requirements:

- Power Trending

EnergyScale provides continuous power usage data collection. This feature provides administrators with the information to predict power consumption across their infrastructure and to react to business and processing requirements. For example, an administrator could adjust server consumption to reduce electrical costs. To collect power data for the IBM BladeCenter JS23 and JS43 Express, you do not require any additional hardware because it collects the information internally.

- Power Saver Mode

Power Saver Mode can be static or dynamic. *Static Power Saver Mode* lowers the processor frequency and voltage on a system a fixed amount, reducing the power consumption of the system while still delivering predictable performance. This percentage is predetermined to be within a safe operating limit and is not user configurable. Under current implementation, this is a 14% frequency drop. When processor utilization is low, Power Saver Mode has no impact on performance. Power Saver Mode can reduce the processor usage up to 30%. Power Saver Mode is not supported during boot or reboot although it is a persistent condition that will be sustained after the boot when the system starts executing instructions.

Dynamic Power Saver Mode, however, varies processor frequency and voltage based on the utilization of the system's POWER6 processors. When dynamic power savings mode is enabled, the firmware of the system continuously monitors the utilization of the system, and adjusts the processor clock speed and voltage to provide enough power to run the current workload. The less the system is utilized, the more power savings are achieved. In addition, you can specify whether you want to favor performance or favor power when enabling dynamic power savings mode. With favor performance, the peak frequency of the processors can be greater than 100%. With favor power, the processors are limited to 95% of nominal frequency under full utilization.

- Power Capping

Capping enforces a user-specified limit on power usage. Power Capping is not a power saving mechanism. It enforces power caps by actually throttling

the processors in the system, degrading performance significantly. The idea of a power cap is to set something that should never be reached but frees up margined power in the data center. The margined power is the amount of extra power that is allocated to a server during its installation in a data center. It is based on the server environmental specifications that usually are never reached because server specifications are always based on maximum configurations and worst case scenarios.

- **Processor Core Nap**

The IBM POWER6 processor uses a low-power mode called Nap that stops processor execution when there is no work to do on that processor core, that is, both threads are idle. Nap mode allows the hardware to clock off most of the circuits inside the processor core. Reducing active power consumption by turning off the clocks allows the temperature to fall, which further reduces leakage (static) power of the circuits causing a cumulative effect. Unlicensed cores are kept in core Nap until they are licensed and return to core Nap when they are unlicensed again.

For more information about EnergyScale Technology on POWER6 processor systems, see:

<http://www.ibm.com/systems/power/hardware/whitepapers/energyscale.html>

3.4.1 Thermal Power Management Device

The implementation of performance-aware power and thermal management for POWER6 processor-based systems is called the EnergyScale architecture, which meets a number of basic requirements for system-level power. IBM BladeCenter JS23 and JS43 Express implementation uses an integrated circuit called Thermal Power Management™ Device (TPMD), placed on the management card.

On IBM BladeCenter JS43 Express there is only one TPMD processor, located in the Base planar. This TPMD processor manages power consumption and thermal conditions of both Base and MPE planars.

The TPMD card can dynamically optimize the processor performance depending on processor power and system workload. The EnergyScale design supports a number of power and thermal management policies.

- **Benchmark**

Benchmark maximizes the single-threaded performance of the system by putting one core of each processor into a low-power state, as in POWER6 the nap mode.

- ▶ Maximum performance

The EnergyScale implementation regulates the system in such a way as to use the maximum performance possible without violating the power or thermal limits of the system.

- ▶ Power cap

The system is set to use a previously defined power usage limit. Maximum power savings attempts to save as much power as possible for a given workload.

- ▶ Optimal power/performance

In this mode the EnergyScale implementation changes the system to the most optimal power/performance settings on the basis of workload characteristics and the power and thermal environment.

3.5 Flexible Service Processor

Also known as the *Service Processor* or *Flexible Support Processor*, the *Flexible Service Processor* (FSP) is integrated on the system board. It provides control busses for all major components on the system. The FSP-1 also has a 128MB flash memory used to store a compressed image of the system's firmware image.

Note: The IBM BladeCenter JS43 Express has two Service Processors, one in the Base planar, and one in the MPE planar. The Service Processor located in the MPE planar has only I/O functions, and does not provide redundancy nor backup support to the FSP in the Base planar.

3.6 Management Card

The Management Card provides a means for making the Anchor system information chip pluggable. Management Card's plug is located on Base planar, just below the DIMMs (see Figure 3-4 on page 49). The Management Card is *always present* in its specially designated connector on the Base planar.

Both IBM BladeCenter JS23 and JS43 Express have only one Management Card, located in the Base planar.

Management Card contains the TPMD processor, RN50 video controller and associated SDRAM, RISCWatch, Async and debug Ethernet ports, Anchor chip and P6 characterization connector, and it is a *customer-replaceable unit* (CRU).

3.6.1 Anchor chip on Management Card

The Anchor Smartchip resides on Management Card. This Anchor chip stores system-specific information. The pluggable Management Card enables the transfer of system-specific information from a faulty system CRU to the replacement CRU.

3.7 Memory subsystem

For IBM BladeCenter JS23 and JS43 Express, the memory controller is internal to the POWER6 processor. It is interfaced to memory buffer chips located on the system board. The memory buffer chip in the POWER6 processor-based server is the next generation Synchronous Memory Interface ASIC that connects memory controllers to the *Industry Standard Very Low Profile (VLP) Registered DIMMs (RDIMM)*.

3.7.1 Memory description of JS23 and JS43 Express

IBM BladeCenter JS23 Express has two memory channels per POWER6 processor module (four channel total), and each memory channel connects to a memory buffer chip. This same configuration is present on the MPE planar of a IBM BladeCenter JS43 Express, for a total of eight channels.

Each memory buffer chip connects to two Registered DIMMs, giving a total of eight DIMMs in the IBM BladeCenter JS23 Express, and 16 DIMMs in the BladeCenter JS43 Express.

Figure 3-2 on page 46 show the memory wiring layout for the base planar on an IBM BladeCenter JS23 Express. This same wiring layout is used in the MPE planar of an IBM BladeCenter JS43 Express.




Figure 3-2 IBM BladeCenter JS23 Express memory wiring layout (Base planar - P1)

3.7.2 Memory placement rules for JS23 and JS43

The following memory placement rules are for JS23 and JS43:

- ▶ Install DIMM fillers in unused DIMM slots for proper cooling.
- ▶ DIMMs are to be installed in pairs. First filling BusA then BusB of each planar, as shown in Figure 3-2:
 - Base planar (P1): (C1, C3), (C6, C8), (C2, C4), (C5, C7)
 - MPE planar (P2): (C1, C3), (C6, C8), (C2, C4), (C5, C7)

Important: Both IBM BladeCenter JS23 and JS43 require a minimum of 4 GB (2x 2 GB DIMM), and we suggest to plug them in slots P1-C1 and P1-C3 (BusA), as shown in Figure 3-3 on page 47.

- ▶ Both DIMMs in a pair must be of same size, speed, and technology. Mixing compatible DIMMs from different manufacturers is possible.
- ▶ All DIMMs controlled by the same POWER6 processor module must be the same size and speed.

Important: Install only supported Registered DIMMs. For a complete list of supported Registered DIMMs, see 2.1, “Overview of JS23 and JS43 Express blade servers” on page 10. You can also find the most recent list of supported DIMMs in the IBM ServerProven® Web site:

<http://www.ibm.com/servers/eserver/serverproven/compat/us/>

Figure 3-3 shows the memory placement layout for both IBM BladeCenter JS23 and JS43 Express.

RDDIM Count	JS23 Base blade planar (P1) RDDIM slots															
	P1-C1	P1-C2	P1-C3	P1-C4	P1-C5	P1-C6	P1-C7	P1-C8								
2	X		X													
4	X		X			X		X								
6	X	X	X	X		X		X								
8	X	X	X	X	X	X	X	X								
RDDIM Count	JS43 Base blade planar (P1) RDDIM slots								JS43 Expansion unit planar (P2) RDDIM slots							
	P1-C1	P1-C2	P1-C3	P1-C4	P1-C5	P1-C6	P1-C7	P1-C8	P2-C1	P2-C2	P2-C3	P2-C4	P2-C5	P2-C6	P2-C7	P2-C8
2	X		X													
4	X		X						X		X					
6	X		X			X		X	X		X					
8	X		X			X		X	X		X			X		X
10	X	X	X	X		X		X	X		X			X		X
12	X	X	X	X		X		X	X	X	X	X		X		X
14	X	X	X	X	X	X	X	X	X	X	X	X		X		X
16	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Figure 3-3 Memory placement for IBM BladeCenter JS23 and JS43 Express

3.7.3 Memory RAS

IBM BladeCenter JS23 and JS43 Express supports Memory Scrubbing, ECC, Chipkill Correction, and Bit Steering. You can find more details about these and other POWER Systems RAS technologies in the following white papers:

- *IBM POWER Systems: Designed for Reliability*

http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03019USEN&attachment=POW03019USEN.PDF&appname=STGE_PO_PO_USEN_WH

- *IBM POWER Systems: Designed for Availability*

http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03020USEN&attachment=POW03020USEN.PDF&appname=STGE_PO_PO_USEN_WH

- *IBM POWER Systems: Designing and Implementing Serviceability*

http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03021USEN&attachment=POW03021USEN.PDF&appname=STGE_PO_PO_USEN_WH

3.8 I/O subsystem

In this section, we discuss the I/O subsystem of JS23 and JS43 blades.

3.8.1 PCI Express

Both IBM BladeCenter JS23 and JS43 Express has support for *PCI Express* (PCIe) cards only. Each planar has connectors for one *PCIe High Speed Daughter Card* (HSDC), and one *PCIe Combination I/O Form Factor – Vertical expansion adapter* (CIOv) Card.

PCIe is the successor to Peripheral Component Interconnect (PCI) and PCI-x bus systems, realized by point-to-point implementation with the following general I/O interface standards:

- ▶ Point-to-point serial interconnect with packetized, layered protocol
- ▶ Speed 2.6 Gbits per pin pair in each direction
- ▶ Dual simplex connection
- ▶ Scalable bus widths
- ▶ Embedded clocking technique using 8-bit/10-bit encoding
- ▶ Isochronous data transfer support
- ▶ Compatible with PCI at the software layers

3.8.2 I/O Expansion Cards

IBM BladeCenter JS43 Express has two HSDC 450-pin connectors, one in each planar, and two CIOv 160-pin connectors, one in each planar also.

Figure 3-4 on page 49 shows how the HSDC and CIOv cards fit together inside the Base planar of an IBM BladeCenter JS23 Express.

Note: IBM BladeCenter JS23 and JS43 Express supports only *Combined Form Factor* High Speed Daughter Cards. Not supported are *Small Form Factor* (SFF) and *Large Form Factory* (LFF) High Speed Daughter cards.

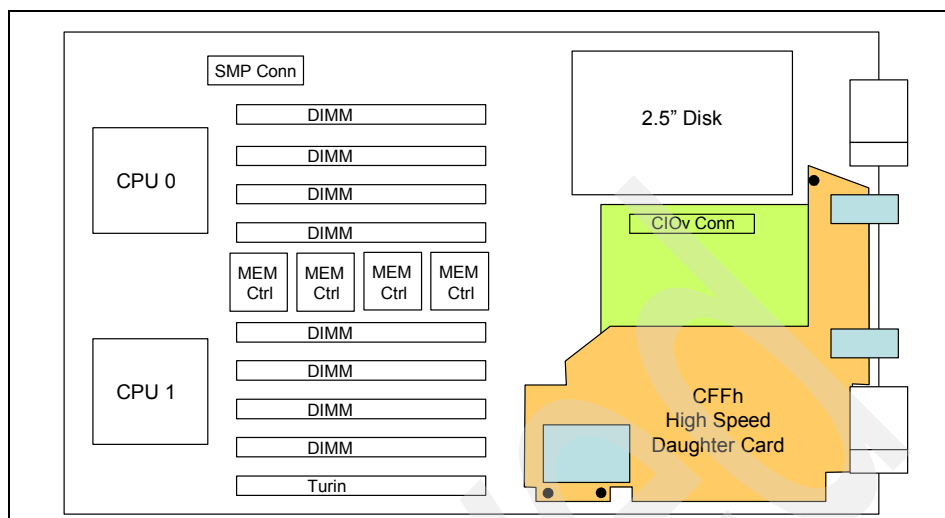


Figure 3-4 JS23 Base planar with a CFFh High Speed Daughter Card and CIOv Card and 2.5"HDD

Table 3-1 and Table 3-2 provide lists of several HSDC and CIOv daughter cards supported by the IBM BladeCenter JS23 and JS43 Express.

Table 3-1 Supported High Speed Daughter Cards (HSDC)

FRU Name	Feature	Supported OS
QLogic® 4 Gb FChannel and Broadcom 1 Gb Enet Combo	8252	AIX, Linux, IBM i
Mellanox 4X InfiniBand Dual Port DDR Expansion Card	8258	AIX, Linux
Qlogic 8 Gb FChannel	8271	Linux

Table 3-2 Supported CIOv PCIe Expansion Cards

FRU Name	Feature	Supported OS
Emulex® 8 Gb Fibre Channel Expansion card	8240	AIX, Linux, IBM i
QLogic 4 Gb FC Expansion Card	8241	AIX, Linux, IBM i
Qlogic 8 Gb Fibre Channel Expansion card	8242	AIX, Linux, IBM i
3 Gb SAS Passthrough Expansion Card	8246	AIX, Linux, IBM i

3.8.3 Universal Serial Bus subsystem

The Universal Serial Bus (USB) connects USB devices to a USB host. The USB subsystem used in the BladeCenter chassis conform to the Universal Serial Bus 2.0 specification.

The IBM BladeCenter JS23 and JS43 Express use the USB subsystem for shared resources such as keyboard, mouse, CD-ROM and floppy drive USB devices between the processor blades. These devices are connected to the processor blades through the Mid-plane.

The USB subsystem routed to the Management Module is used for the keyboard on the IBM BladeCenter JS23 and JS43 Express. This approach allows the Management Module to connect the keyboard to the IBM BladeCenter JS23 and JS43 Express, and control the routing of keyboard inputs to one of fourteen blades installed in a BladeCenter chassis.

The other USB subsystem is routed to the media tray, which connects to the DVD-ROM drive and diskette drive. Both USB subsystems are controlled by the Management Module independent of each other. This provides the capability to assign the CD and floppy disk drive to one blade while the keyboard is assigned to another blade. However, it does not restrict the assignment of shared resources of both USB subsystems to a single IBM BladeCenter JS23 or JS43 Express.

3.8.4 Integrated Virtual Ethernet

Integrated Virtual Ethernet (IVE) is the name given to the collection of hardware components (including the Host Ethernet Adapter (HEA), the software, and the hypervisor functions that provide the integrated Ethernet adapter ports with hardware assisted virtualization capabilities.

The IVE was developed to meet general market requirements for better performance and better virtualization for Ethernet. It offers:

- ▶ Two 1-Gbps ports
- ▶ External network connectivity for LPARs using dedicated ports without requiring a Virtual I/O server
- ▶ The speed and performance of the GX+ bus, faster than PCI Express x16

The IVE consists of a physical Ethernet adapter that is connected directly to the GX+ bus of the POWER6 processor instead of being connected to a PCIe bus. This method provides IVE with the high throughput and low latency of a bus

embedded in the I/O controller. IVE also includes special hardware features that provide logical Ethernet adapters.

The IVE feature implemented on the IBM BladeCenter JS23 Express provides two 1-Gbps physical ports. IBM BladeCenter JS43 Express has two additional 1-Gbps physical ports, located in the MPE planar.

Prior to IVE, virtual Ethernet provided a connection between LPARs. The use of an SEA and the Virtual I/O server allowed connection to an external network. The IVE replaces the requirement for both the virtual Ethernet and the SEA. It provides most of the functionality of each.

Therefore, this eliminates having to move packets (using virtual Ethernet) between partitions and then through a shared Ethernet adapter (SEA) to an Ethernet port. LPARs can share IVE ports with improved performance.

The terms IVE and HEA are used interchangeably in this document.

3.8.5 Integrated Video Card

Both IBM BladeCenter JS23 and JS43 Express have an integrated ATI RN50 video controller in the Base planar. This video controller is used to support 2D graphics on an SVGA video display. This is the same controller as in the IBM BladeCenter JS21 Express, but with DDR2 SDRAM.

3.8.6 Serial Attached SCSI storage subsystem

IBM BladeCenter JS23 and JS43 Express uses an embedded Serial Attached SCSI (SAS) controller that operates at 32-bit PCI-X at 133MHz.

Note: The SAS Drive in the JS23 Base planar is not hotpluggable.

On IBM BladeCenter JS23 Express are four SAS ports. Two of them are wired to the SAS hard drive, and the other two go to the CIOv PCIe connector, connecting to the BladeCenter SAS Switch bay 3 and bay 4, when a SAS paddle card is used in the CIOv connector.

In BladeCenter S, these two SAS ports are connected to the two *Disk Storage Modules* (DSM) through the SAS Switch in bay 3 and bay 4. In BladeCenter H, external storage are connected through the four external ports of the SAS Switches in bay 3 and bay 4.

IBM BladeCenter JS43 Express has four SAS ports that goes from the Base planar to the MPE planar, and they are used in similar functions as those on

Base planar. IBM BladeCenter JS43 Express has only one SAS controller, located in the Base planar.

The boot HDD can either be on the Base or MPE planar of an IBM BladeCenter JS43 Express. The boot disk can also be from a BladeCenter HDD through SAS CIOv paddle card.

Tip: Install the boot HDD on the Base planar of the IBM BladeCenter JS43 Express.

SAS Drive

IBM BladeCenter JS23 and JS43 Express supports SAS 2.5" *Hard Disk Drive* (HDD) or SAS 2.5-inch *solid disk drive* (SDD).

For a complete list of the supported SAS disks, refer to "Storage" on page 11.

Important: SSD drivers are formatted in 528-byte sector, and the SAS controller does not write to them if it finds configuration sectors that are not clean. This situation happens when the SSD was previously used in RAID application. To be able to write to the SSD again, you have to reformat it.

RAID support

IBM BladeCenter JS23 Express has no RAID available.

IBM BladeCenter JS43 Express has support for RAID functions when more than one SAS disk is installed in the system. If there is only one drive then there is no RAID function.

For two drives in the IBM BladeCenter JS43 Express, the supported RAID functions are:

RAID 0 Striping

RAID 1 Mirroring

The drives on the Base planar and MPE planar can be either rotating hard drives (HDD) or solid state drives (SSD). However, the SAS controller will not place an SSD and an HDD in the same RAID set.

Important: For RAID operations, all drives must be of the same type. Either HDD on both Base and MPE planars, or SDD on both Base and MPE planars.

Additionally, RAID can be done with different drives capacities, but you get only the effective capacity of the smaller device.

Configuring a SAS RAID array

Two SAS disk drives in the IBM BladeCenter JS43 Express can be used to implement and manage RAID level-0 and RAID level-1 arrays in operating systems that are on the ServerProven list.

For the blade server, you must configure the SAS RAID array through the command `smit sasdam`, which is the *SAS Disk Array Manager for AIX*.

The *SAS Disk Array Manager* is packaged with the Diagnostics utilities on the Diagnostics CD. Use the command `smit sasdam` to configure the disk drives for use with the SAS controller.

Tip: Refer to Chapter 3 of *IBM Power Systems SAS RAID controller for AIX* for more details on how to use the IBM SAS Disk Array Manager available at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/arebj/arebj.pdf>

Important: Depending on your RAID configuration, you must create the array before you install the operating system in the blade server.

Before you can create a RAID array, you must reformat the hard disk drives so that the sector size of the drives changes from 512 bytes to 528 bytes. If you later decide to remove the hard disk drives, delete the RAID array before you remove the drives.

If you decide to delete the RAID array and reuse the hard disk drives, you must reformat the drives so that the sector size of the drives changes from 528 bytes to 512 bytes.

When the IBM BladeCenter JS43 Express is installed in a BladeCenter Storage (BCS) and all or some of the BCS DASD are configured to that blade slots, those devices, if formatted to 528-Byte sector, can also be candidates for RAID array members with the JS43 Express drives themselves.

SAS CIOv paddle card

When the SAS CIOv paddle card is installed on an IBM BladeCenter JS23 or JS43 Express, accessing the hard drives installed on a BladeCenter S or external hard drive array is possible.

Supported SAS topology through CIOv paddle card

When the SAS CIOv paddle card is installed on Base or MPE planar, the SAS controller can access the hard drives installed on a BCS or external hard drive

array. The CIOv Paddle connects the SAS controller to the SAS switch bay 3 and bay 4. Each SAS Switch bay may be populated with a *Non-RAID SAS Switch Module (NSSM)*.

Note: At the time of this writing, *RAID SAS Switch Modules (RSSM)* is not supported.

In BCS, the switch expands the SAS controller ports to the two internal DSMs and to storage enclosures that are attached to 4 external connections of the switch modules. In BladeCenter H (BCH) or BladeCenter HT (BCHT), only external storage enclosures are attached to the 4 external ports of the SAS switch. No internal storage exists in BCH nor BCHT.

Configurations are allowed with either one or both switch bays populated. Each NSSM contains four mini-SAS connectors. The two SAS devices that can be attached to these connectors are SAS Tape devices and DS3200.

DS3200 is a RAID box that appears as an SAS End device/Target with one or more LUNs. SAS Tape devices are also end devices.

Furthermore, when an IBM BladeCenter JS23 Express is installed in a BCS, *internal storage* Disk Storage Modules in BCS attach to the NSSM through the BladeCenter Mid planar. Each DSM has two SAS expanders with each expander connecting to the six DASDs, one DSM connects to the primary ports of the DASD and the other expander connects to the secondary port of the DASD.

The A side expander of each DSM is wired to NSSM in switch bay 3 and the B side expander is wired to the NSSM in switch bay 4.

Figure 3-5 on page 55 shows the supported SAS topology for the IBM BladeCenter JS23 and JS43 Express on the BCS.

Figure 3-6 on page 56 shows the supported SAS topology for the IBM BladeCenter JS23 and JS43 Express on the BCH and BCHT.

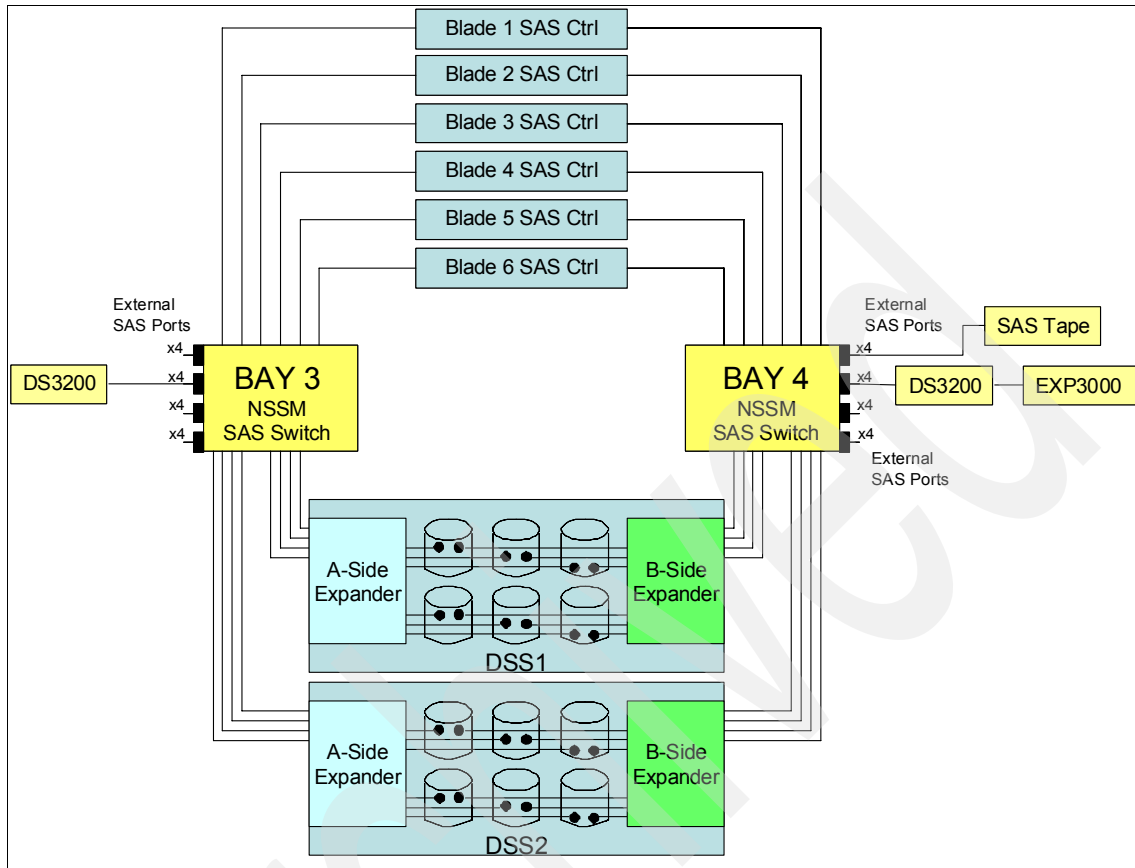


Figure 3-5 IBM BladeCenter JS23 and JS43 Express BCS SAS Topology

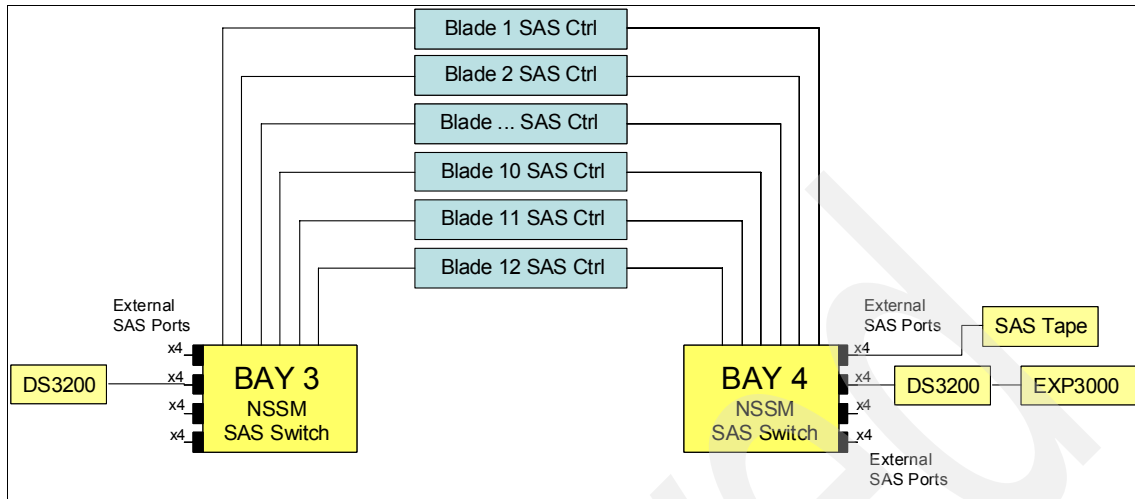


Figure 3-6 IBM BladeCenter JS23 and JS43 Express BCH and BCHT SAS Topology

3.9 PowerVM

The PowerVM platform is the family of technologies, capabilities, and offerings that provide virtualization capabilities on the IBM BladeCenter JS23 and JS43 Express systems. PowerVM is the new umbrella branding term for IBM Power Systems™ Virtualization (Logical Partitioning, IBM Micro-Partitioning™, Hypervisor, Virtual I/O Server, Advanced Power Virtualization, Live Partition Mobility, Workload Partitions, and so on).

As with Advanced Power Virtualization in the past, PowerVM is a combination of hardware enablement and value-added software. Table 3-3 on page 57 shows the licensed features of each of the two different editions of PowerVM for IBM BladeCenter JS23 and JS43 Express.

PowerVM standard edition offers a Web-based interface for managing virtualization within a single blade. The Integrated Virtualization Manager (IVM) component of Virtual I/O server (VIOS) provides the setup of management capabilities of logical partitions (LPARs). It manages Virtual I/O and Virtual Ethernet so that storage and communications adapters can be shared among all the LPARs running on the IBM BladeCenter JS23 and JS43 Express.

Upgrade from PowerVM Standard Edition to Enterprise Edition is possible and completely undistruptive. The upgrade does not even require the installation of additional software. Just enter a valid activation key code in the hypervisor to unlock the next level of functions.

Table 3-3 PowerVM editions for IBM BladeCenter JS23 and JS43 Express

Description	Standard Edition	Enterprise Edition
Maximum LPARs	10 / core	10 / core
Virtual I/O server	YES	YES
Integrated Virtualization Manager	YES	YES
Shared Dedicated Capacity	YES	YES
Live Partition Mobility	NO	YES
IBM Active Memory™ Sharing	NO	YES

3.10 Operating system support

The IBM BladeCenter JS23 and JS43 Express supports the IBM AIX, Novell SUSE Linux Enterprise Server, Red Hat Enterprise Linux, and IBM i operating systems.

When using the PowerVM Standard Edition or Enterprise Edition on an IBM BladeCenter JS23 and JS43 Express running a VIOS, a combination of IBM i, IBM AIX, and Linux partitions can be deployed.

This section describes in detail the supported operating system versions and levels.

3.10.1 AIX

The following versions of IBM AIX 5L™ are supported on the IBM BladeCenter JS23 and JS43 Express:

- ▶ AIX V5.3 with the 5300-07 Technology Level with Service Pack 9, or later
- ▶ AIX V5.3 with the 5300-08 Technology Level with Service Pack 7, or later
- ▶ AIX V5.3 with the 5300-09 Technology Level with Service Pack 4, or later
- ▶ AIX V5.3 with the 5300-10 Technology Level, or later
- ▶ AIX V6.1 with the 6100-03 Technology Level, or later
- ▶ AIX V6.1 with the 6100-02 Technology Level with Service Pack 4, or later
- ▶ AIX V6.1 with the 6100-01 Technology Level with Service Pack 5, or later
- ▶ AIX V6.1 with the 6100-00 Technology Level with Service Pack 9, or later

IBM periodically releases maintenance packages for the AIX 5L operating system. These packages are available on CD-ROM, or you can download them from:

<http://www.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>

The Web page provides information about how to obtain the CD-ROM. You can also get individual operating system fixes and information about obtaining AIX 5L service at this site.

In AIX 5L V5.3, the **suma** command is also available, and helps the administrator to automate the task of checking and downloading operating system downloads. For more information about the **suma** command functionality, visit:

<http://www.ibm.com/systems/p/os/aix/whitepapers/suma.html>

3.10.2 Linux

The following versions of Linux for IBM POWER Series are supported on BladeCenter JS23 and JS43 Express:

- ▶ SUSE Linux Enterprise Server 10 Service Pack 2 for POWER Systems, or later
- ▶ SUSE Linux Enterprise Server 11 for POWER Systems, or later
- ▶ Red Hat Enterprise Linux 4.6 for POWER
- ▶ Red Hat Enterprise Linux 4.7 for POWER, or later
- ▶ Red Hat Enterprise Linux 5.1 for POWER
- ▶ Red Hat Enterprise Linux 5.2 for POWER
- ▶ Red Hat Enterprise Linux 5.3 for POWER, or later

For information about the features and external devices supported by Linux for IBM POWER on the IBM BladeCenter JS23 and JS43 Express server, visit:

<http://www.ibm.com/systems/bladecenter/hardware/servers/index.html>

For information about SUSE Linux Enterprise Server 10, visit:

<http://developer.novell.com/yesssearch/Search.jsp>

For information about Red Hat Enterprise Linux, visit:

<https://hardware.redhat.com/?pagename=hcl&view=certified&vendor=4&class=8>

Many of the features described in this document are operating system dependent and might not be available on Linux. For more information, visit:

http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html

3.10.3 IBM i

IBM i V6.1 is supported on both IBM BladeCenter JS23 and JS43 Express. It uses IBM PowerVM Standard Edition, which includes the IBM POWER Hypervisor™, Micro Partitioning, and Virtual I/O server with Integrated Virtualization Manager (IVM). PowerVM Standard Edition is available for no additional charge on the BladeCenter JS23 and JS43 Express.

For more information about IBM i operating systems running on IBM BladeCenter JS23 and IBM BladeCenter JS43 Express, see:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

3.11 Systems management

Several tools are available to manage IBM BladeCenter JS23 and JS43 Express. The BladeCenter Web interface can effectively manage the blades. In addition, IBM Director and Cluster Systems Management (CSM) are management tools that can be used to manage your environment.

3.11.1 BladeCenter Advanced Management Web interface

The BladeCenter Web interface allows system administrators to easily and effectively manage up to 14 blades from an integrated interface. From trivial tasks such as powering blades on or off, to more complex tasks such as firmware management, the Web interface allows powerful control over all blades and I/O modules that are attached to the BladeCenter chassis.

Important: IBM BladeCenter JS23 Express and JS43 Express do not support firmware updates through the Advanced Management Module. Their server-enhanced service processor has a larger firmware image, therefore downloading and installing over the RS-485 bus of the management module is impractical.

For more details about updating firmware, see 12.1, “Firmware updates” on page 454.

The BladeCenter Web interface enables the System Administrator to:

- ▶ Easily and effectively manage up to 14 blade servers from an integrated interface.
- ▶ Power on or off the IBM BladeCenter JS23 and JS43 Express.

- ▶ Have control over all blade servers and input/output (I/O) modules that are attached to the BladeCenter chassis even with a mixed environment.
- ▶ Manage other BladeCenter resources such as I/O modules and retrieval of system health information.
- ▶ Configure SOL for the IBM BladeCenter JS23 and JS43 Express.

3.11.2 IBM Director

IBM Director is a GUI-based management tool that provides management functions for the BladeCenter. IBM Director enables you to remotely manage many IBM and other servers, including the IBM BladeCenter JS23 and JS43 Express. The IBM Director console allows System Administrators to manage multiple BladeCenter chassis in a heterogeneous environment or environments where a Director infrastructure exists.

Visit the following Web site to download the IBM Director 6.1, and get the latest information about IBM Director 6.1 to IBM BladeCenter JS23 and JS43 Express:

<http://www.ibm.com/systems/management/director/downloads/>

IBM Systems Director Active Energy Manager

IBM Systems Director Active Energy Manager is the strategic power management interface for all IBM server products. From Active Energy Manager a user can display power trending information, display thermal information, enable/disable power savings modes or set a power cap. An administrator may connect a single instance of Active Energy Manager to multiple platforms, including IBM BladeCenter chassis with one or more IBM BladeCenter JS23 and JS43 Express blades.

As the strategic power management interface, Active Energy Manager supports the most comprehensive feature set of all the EnergyScale user interfaces, and is delivered as an extension to IBM Director.

Active Energy Manager has now been integrated into the Web-based interface of Director 6.1. The result is tighter integration of Active Energy Manager and IBM Director, eliminating the separately installable Active Energy Manager console interface of previous Active Energy Manager releases.

In addition to the Web interface and ability to install IBM Director server on AIX, Active Energy Manager leverages Director 6.1 to provide the following features:

- ▶ Dynamic Power Save for supported POWER6 servers
- ▶ Threshold-driven events for power, input temperature, and effective processor speed

- ▶ Enhanced power policies for individual servers, and groups of servers
- ▶ Expanded power capping ranges through *soft* power caps on supported servers
- ▶ Complete command-line interface (CLI) support
- ▶ Support for additional hardware, including:
 - SynapSense Wireless Sensor Nodes
 - Power Distribution Units (PDUs) from Eaton and Raritan
 - Uninterruptible power supplies from Eaton
 - Monitoring of equipment managed by Emerson-Liebert's SiteScan, including Computer Room Air Conditioning (CRAC) units, chillers, uninterruptible power supplies, and PDUs. This includes getting real-time data on power consumption, alerts when problems occur

For more information about IBM Active Energy Manager, see the following resources:

<http://www.ibm.com/systems/management/director/extensions/actengmrg.html>
<http://www.ibm.com/systems/power/hardware/whitepapers/energyscale.html>

3.11.3 Cluster Systems Management

IBM Cluster Systems Management (CSM) provides several useful functions to manage a cluster from a single point-of-control. These include resource monitoring, automated monitoring and operation, remote hardware control, remote command execution, security, configuration file management, parallel network installation, and diagnostics.

CSM V1.7 supports the following functions on the IBM BladeCenter JS23 and JS43 Express:

- ▶ Hardware control
- ▶ Install and update software on nodes
- ▶ Distributed command execution
- ▶ File synchronization across cluster
- ▶ Monitoring synchronization across cluster
- ▶ Monitoring and automated response
- ▶ Automatic security configuration
- ▶ Management of node groups (static and dynamic)
- ▶ Diagnostics tools

For more information about CSM, visit:

<http://www.ibm.com/systems/clusters/software/csm/>
<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>



Part 2

System configuration and management

Having covered the basic system information and architecture in Part 1, we expand on that to include how to get BladeCenter JS23 and JS43 running the supported operating systems, and several other management-oriented topics.

System planning and configuration using VIOS with IVM

This chapter describes how to perform basic system planning prior to and configuration after you install Virtual Input/Output Server (VIOS). The configuration can be done by using the command-line interface (CLI) and graphical user interface (GUI). The Web browser-based GUI is an integral part of the Integrated Virtualization Manager (IVM) and is included in the VIOS.

This chapter contains the following topics:

- ▶ “Planning considerations” on page 67
- ▶ “VIOS system management using IVM” on page 77
- ▶ “First VIOS login” on page 80
- ▶ “First IVM connection” on page 87
- ▶ “VIOS network management and setup” on page 94
- ▶ “VIOS Storage Configuration and Management” on page 116
- ▶ “Partition configuration for Virtual I/O Client” on page 139
- ▶ “Console access and activating a partition” on page 161

Complete VIOS configuration and maintenance information can be found in *System i and System p Using the Virtual I/O Server*, which is available at:

http://publib.boulder.ibm.com/infocenter/iseriess/v1r3s/en_US/info/iphb1/iphb1.pdf

The PowerVM Editions Web site also contains useful information:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm>

Additional IVM information can be found in *Integrated Virtualization Manager on IBM System p5*, REDP-4061.

4.1 Planning considerations

When planning your system environment for an IBM BladeCenter JS23 or JS43, a complete overview of the BladeCenter, blades, network, and storage should be reviewed. Crafting an overall solution will help to eliminate expensive rework.

4.1.1 General considerations

We start with the general considerations.

BladeCenter

An understanding of how network and other I/O traffic is routed through the mid-plane of a BladeCenter chassis is necessary so you can plan external network and storage connections. Figure 4-1 on page 68 and Figure 4-2 on page 69 show the connections between a blade's integrated ports and expansion card ports to the individual switch module bays.

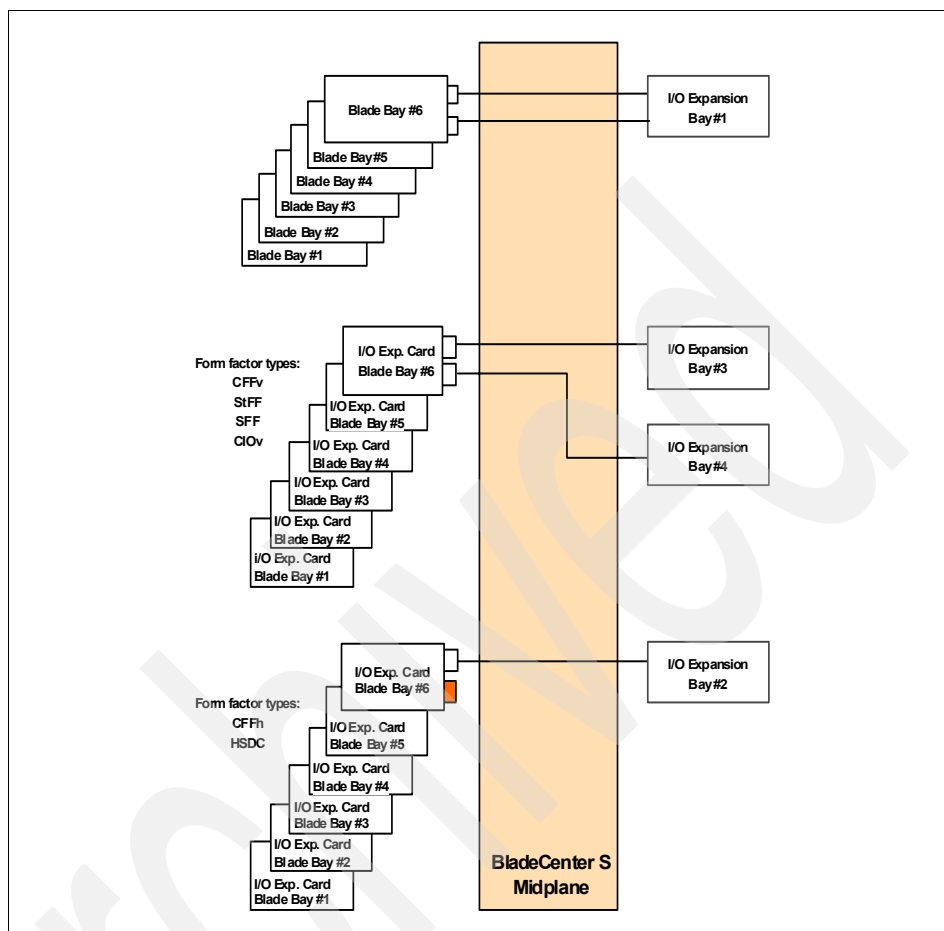


Figure 4-1 Expansion card port path through BladeCenter S mid plane

Note: An RPQ is required when a Qlogic Ethernet and 4 Gb Fibre Channel CFFh expansion card (*combo card*) is used in a BladeCenter S; only the two Ethernet ports are directed to the switch module in I/O bay 2.

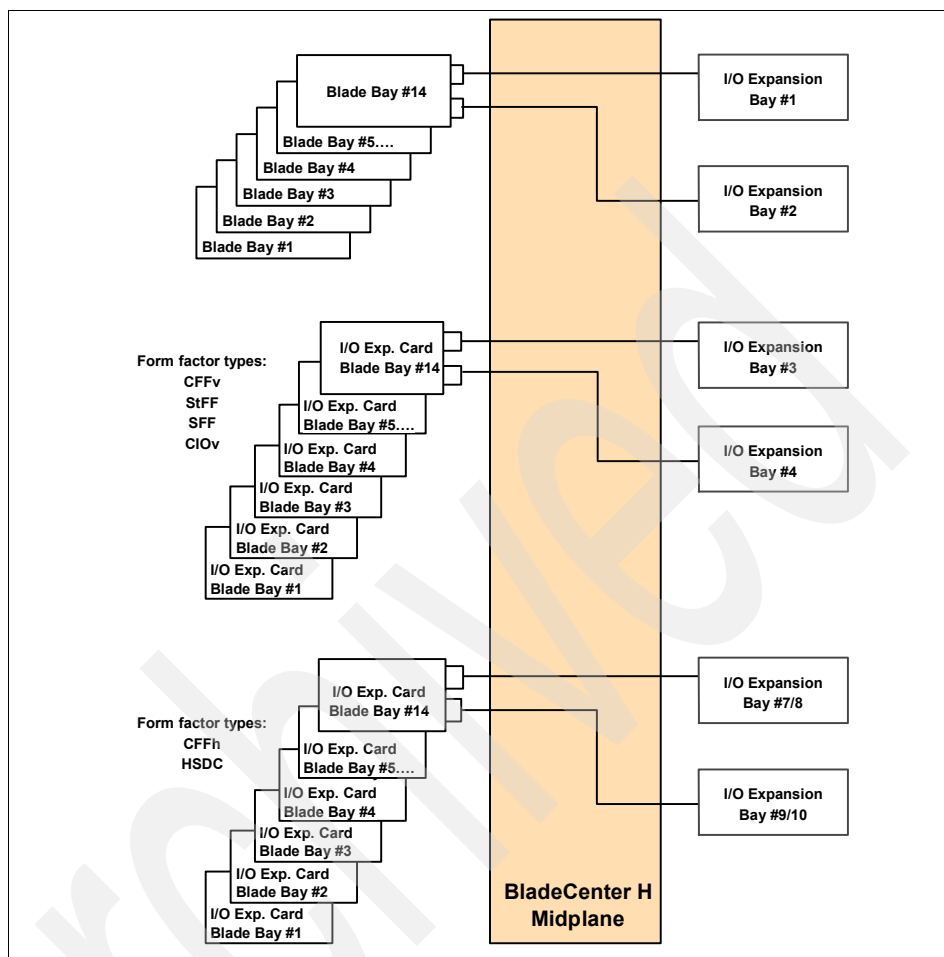


Figure 4-2 Expansion card port path through BladeCenter H mid plane

Internal storage

The IBM BladeCenter JS23 can have a single internal SAS drive. The IBM BladeCenter JS43 can have two internal drives. A best practice is to use disk mirroring for the operating system. The mirror can be between the internal drive and a SAN LUN, between two SAN LUNs, or between two internal drives when available.

Virtualization

After the physical connection planning aspects, the logical partitions or LPARs (assuming PowerVM virtualization) must be designed so that their requirements fit the processor and memory available on the blade. Additional considerations

are required. If you want to use advanced operations available under PowerVM Enterprise Edition, such as Live Partition Mobility (LPM) and Active Memory Sharing (AMS).

The decision regarding whether to use a shared processor pool or dedicated processors should be made prior to configuring an LPAR. Changing from one mode to the other with the IVM GUI requires the deletion of the LPAR and the creation of a new one, the VIOS CLI can use the **chsyscfg** command. The **chsyscfg** command can be used to switch when the partition is in the Not Activated state.

Network

EtherChannel configuration (if used), and BladeCenter switch module type and bay location, require pre-planning. HEA ports cannot be mixed with the PCI-X adapter ports available on the CFFh combo cards to create an EtherChannel device. The use of the HEA ports to create an EtherChannel requires the use of pass-thru modules, the Intelligent Copper Pass-thru Module for IBM BladeCenter in switch module bay one to retain Serial Over LAN (SOL) capability, or Ethernet modules that provide for logical grouping of switches. network interface backup (NIB), can mix HEA and other adapter ports types.

Note: The HEA adapter implementation on IBM BladeCenter JS23 or JS43 blades always show a link status of Up, as shown using the **entstat** command:

```
$ entstat -all ent1 |grep Link
Logical Port Link State: Up
Physical Port Link State: Up
```

When an HEA adapter is used as primary in a NIB configuration, the link state cannot be used as the failover determination. NIB should be set up using the *ping* option.

4.1.2 Internal and external storage considerations

Different types of storage can be used with the Virtual I/O Server. Each type of storage offers advantages and disadvantages. This section explains possible configuration scenarios with the IBM BladeCenter JS23 or JS43. For information about supported I/O expansion cards, along with their feature codes and option part numbers, refer to 3.8.2, “I/O Expansion Cards” on page 48.

Currently, two form factor types of I/O expansion cards are supported in the IBM BladeCenter JS23 or JS43:

- ▶ Combinational I/O vertical (CIOv) I/O expansion cards
- ▶ Combined Form Factor horizontal (CFFh) I/O expansion cards

CIOv adapter cards ports are always connecting to bays 3 and four of a BladeCenter chassis when installed in an IBM BladeCenter JS23 or JS43.

Figure 4-3 show an Active SAS Pass-through *paddle* expansion card in CIOv form factor. A QLogic 4 Gb Fibre Channel HBA, and Qlogic and Emulex 8 Gb Fibre Channel HBAs are also available in the same form factor.




Figure 4-3 Active SAS Pass-through Expansion Card CIOv

The CFFh combo card has a dual-port 4 Gb Fibre Channel Qlogic host bus adapter and a dual-port Broadcom gigabit Ethernet adapter over a PCI-x bridge connected to the PCIe bus of the blade. Figure 4-4 on page 72 shows the CFFh form factor. A Qlogic 8 Gb Fibre Channel 2-port HBA in the CFFh format is also available.




Figure 4-4 Qlogic Ethernet and 4 Gb Fibre Channel “combo card” CFFhTh

Together with an installed Qlogic Ethernet and 4 Gb Fibre Channel combo card, installing the CIOv I/O expansion card is also possible. Using a BladeCenter H with a JS23 combination gives, in addition to the two onboard network ports, six more I/O ports. These six additional ports are four Fibre Channel ports and two 1 Gb Ethernet ports. The JS43 doubles this port count to four integrated network ports and 12 additional I/O ports or eight Fibre Channel ports and four additional 1 Gb Ethernet ports.

Note: When a CIOv card and a CFFh card are installed together, the CIOv must be installed first.

You must install the matching type of I/O module in the bays for the type of I/O expansion card. Specifically, this means:

- ▶ You must install a module that provides Ethernet connectivity in module bay 1 and bay 2 of a BladeCenter H or bay 1 of a BladeCenter S¹.
- ▶ You must install a module that supports the I/O connectivity of the installed CIOv expansion cards in Bay 3 and Bay 4.
- ▶ You must install a module or modules in a BladeCenter H Bays 7-10 or switch module in Bay 2 of a BladeCenter S that support the CFFh expansion cards.

Note: When the Optical Pass-thru Module is used with a 4 Gbs Fibre Channel host bus adapter, only a transfer rate of 2 Gbs can be used. Therefore, set the Fibre Channel switch ports that connect to the OPM to 2 Gbs.

¹ Both integrated network ports of a blade connect to the module in bay 1. A VLAN configuration is required to separate the different networks and broadcast domains.

When JS23 and JS43 blades with CFFh cards are installed in a BladeCenter H or HT, the cards connect to the high speed bays 7, 8, 9, and 10 depending on the ports on the card. These module bays have a horizontal orientation. (The standard module bays have a vertical orientation.)

When JS23 and JS43 blades with a supported CFFh card are installed in a BladeCenter S, the cards are connected to bay 2

Although certain CFFh cards utilize the high-speed bays, they use standard modules for connectivity. This connectivity is achieved by using the Multi-Switch Interconnect Module (MSIM). One MSIM² gets installed in bays 7 and 8, and one in bays 9 and 10. Each MSIM can be populated with a standard Fibre Channel switch module and an Ethernet switch module for the BladeCenter.

Figure 4-5 shows a diagram of the MSIM and the type of switches that can be installed in the module bays. Refer to *IBM BladeCenter Interoperability Guide* for details about which switches are supported in the MSIM. The guide can be found at:

<http://www.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-5073016&brandind=5000020>




Figure 4-5 Multi-Switch Interconnect Module

² A interposer card is required for each MSIM to install in a BladeCenter HT chassis. Contact your IBM sales representative for information about ordering interposer cards.

JS23/JS43 storage

Currently, four types of storage are available:

- ▶ Internal 73 GB or 146 GB SAS Hard Disk Drive (HDD) disk storage
- ▶ Internal 73 GB SAS Solid State Drive (SSD) storage
- ▶ External SAS/SATA disk storage
- ▶ External Fibre Channel storage

No hardware initiator card is available for the IBM BladeCenter JS23 or JS43 for iSCSI storage system attachment. Software initiators are available for AIX and Linux (no VIOS support).

Internal SAS storage

For internal storage, the JS23 blade can have one SAS HDD or SSD. The JS43 can have up to two SAS HDD or SSD storage devices installed or available. The JS43 supports RAID 0/1 types but both SAS devices must be of the same type HDD or SSD.

Note: Internal SAS drives are not currently supported for IBM i under the VIOS; in this case, they are only for VIOS use.

External SAS/SATA storage

External SAS or SATA storage is currently available in the BladeCenter S chassis. A BladeCenter S can be equipped with up to two Disk Storage Modules (DSM), which are each capable of holding up to six SAS or SATA hard drives.

This storage can be utilized by installing an Active SAS Pass through Expansion Card for IBM BladeCenter in the JS23 or JS43 and, one or two SAS Connectivity modules for configurations that are not RAID-enabled. Two SAS RAID Controller Modules are required for RAID-enabled configurations.

With this setup, you are limited to the capacity that can be provided the two DSMs holding up to the maximum of 12 disks. More external storage could be made available using iSCSI and software initiators, but this is not a good practice because of limited network port availability. You have only the two onboard network interfaces of the blade available.

External Fibre Channel storage

Various IBM and non-IBM Fibre Channel storage subsystems can be used to connect to a JS23 or JS43 installed in a BladeCenter, including models from the following series:

- ▶ IBM Total Storage DS8000® series
- ▶ IBM Total Storage DS6000™ series
- ▶ IBM Total Storage DS5000 series

- ▶ IBM Total Storage DS4000® series
- ▶ IBM Total Storage DS3000 series
- ▶ IBM Total Storage N series

The Virtual I/O Server data sheet provides an overview of supported storage subsystems and the failover driver that is supported with the subsystem. The data sheet can be found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/data-sheet.html>

Verify that your intended operating system supports these storage subsystems.

Also decide which type of I/O expansion card to use, and plan your storage area network. IBM offers Fibre Channel and Ethernet Switch modules from various vendors, and it is always easier to connect switches from the same vendor. In a Fibre Channel SAN environment, zoning limitations exist if different switch vendors are used.

General storage considerations and support matrixes

You have to check several support matrixes to plan your Live Partition Mobility installation. This section provides steps, and points to support matrixes you can use, to build a fully supported solution, as follows:

1. Start with the blade itself. All supported hardware and operating systems are listed on IBM ServerProven Web page:
<http://www.ibm.com/servers/eserver/serverproven/compat/us/>
Click the BladeCenter picture for blade and BladeCenter information. You can find consolidated information from different sources in *IBM BladeCenter Interoperability Guide* at:
<https://www.ibm.com/systems/support/supportsite.wss/docdisplay?ln doc id=MIGR-5073016&brandind=5000020>
2. Verify the supported operating systems on the blade by using the link to NOS Support on the Compatibility for BladeCenter products page of ServerProven:
<http://www.ibm.com/servers/eserver/serverproven/compat/us/eserver.html>
3. Select the blade from the Compatibility for BladeCenter products page that you opened in step 1. A list of supported and tested hardware in combination with the blade itself is shown. Clicking a listed component retrieves details about the tested hardware and limitations that might exist for the component.
4. Using the information you gained about supported HBAs and storage products, you can start checking the Virtual I/O Server Data sheet to verify which components supported by the blade are supported by the Virtual IO Server as well.

The data sheet can be found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>

5. Check the support matrix of the storage subsystem of your choice. In the case of Fibre Channel attached storage, verify the SAN switch support matrix. The following list points to the matrixes of IBM storage products.

DS8000 interoperability matrix:

<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

DS6000 interoperability matrix:

<http://www.ibm.com/servers/storage/disk/ds6000/interop.html>

DS5000 interoperability matrix:

<http://www.ibm.com/systems/storage/disk/ds5000/interop-matrix.html>

DS4000 interoperability matrix:

<http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html>

DS3000 interoperability matrix:

<http://www.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf>

Enterprise Storage Server® (ESS) interoperability matrix:

<http://www.ibm.com/servers/storage/disk/ess/interop-matrix.html>

N series interoperability matrix:

<http://www.ibm.com/systems/storage/nas/interophome.html>

SAN Volume Controller support matrix:

<http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html>

SAN switch interoperability matrix:

<http://www.ibm.com/systems/storage/san/index.html>

6. If you plan to implement IBM i on a JS23 or JS43, then in addition to the previous steps, verify the support matrix of IBM i on blades by going to the following links. Certain considerations apply when IBM i is used.
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>
http://www.ibm.com/systems/power/hardware/blades/supported_environments.pdf
7. After verifying the supported hardware components, check the storage area network cabling that is required for the storage product you plan to implement. The cabling is described in the product documentation of the storage subsystem. Verify which failover drivers are supported by the storage subsystem. In the product documentation, check the recommended zoning configuration.

8. Use the Virtual I/O Server data sheet again to check which failover drivers are included in the Virtual I/O Server and which failover drivers can be installed.

Note: The IBM System Storage™ Interoperation Center (SSIC) helps to identify supported storage environments. You find this Web-based tool at:

<http://www.ibm.com/systems/support/storage/config/ssic>

9. The storage subsystem usually requires a specific FCode, firmware, and driver level that is used with the selected host bus adapter in the blade. This information is not for every subsystem in the same location.

The host bus adapters (HBA) Web-based tool helps to identify the required firmware level. You can check also product documentation or host system attachment guides when available. DS3000 and DS4000 support Web pages usually provide update packages for the supported adapters that contain the settings required for the HBA. They can be found at:

<http://www.ibm.com/systems/support/storage/config/hba/index.wss>

10. Define the boot device that you plan to use. You may boot from internal disk or external disk. Depending on the type of blade and I/O expansion card you plan to install, you might not have internal disks available. Redundancy might be not available at boot time for a boot disk. Manual interaction might be required and an administrator might have to be trained for situations where interaction with the system is required.

4.2 VIOS system management using IVM

Using VIOS is required when two or more partitions are required, and when using IBM i.

IVM provides a unique environment to administer logical partition-capable servers. It provides two ways to configure and manage logical partitions (LPARs):

- ▶ A graphical user interface (GUI) designed to be as simple and intuitive as possible, incorporating partition management, storage management, serviceability, and monitoring capabilities
- ▶ A command-line interface (CLI) enables scripting capabilities

You can use either interface to create, delete, and update the logical partitions and perform dynamic operations on LPARs (DLPAR) including the VIOS itself.

4.2.1 VIOS installation considerations

The Virtual I/O Server (VIOS) installation is performed similar to a native installation of AIX. The basic requirements are:

- ▶ Console access, provided by Serial Over LAN (SOL) or KVM³
- ▶ AIX media or a NIM server
- ▶ A BladeCenter media tray assigned to the target blade, or network access to a NIM server
- ▶ Internal or external storage availability

4.2.2 IVM user interface

The IVM user interface is an HTML-based interface. It enables you to create LPARs on a single managed system, manage the virtual storage and virtual Ethernet on the managed system, perform maintenance, and view service information related to the managed system.

The IVM GUI consists of several elements. The following two elements are used most frequently:

- ▶ Navigation area

This area on the left side of the window displays the tasks and links that you can access in the work area.

- ▶ Work area

This area on the right side of the window contains information related to management tasks that you perform using IVM and related to the objects on which you can perform management tasks.

Figure 4-6 on page 79 shows the navigation and works of the IVM GUI. Help can be obtained from any window by clicking the question mark (?) icon near the upper right corner of any window or view.

Logging in to the GUI is described in 4.4.1, “Connecting to IVM” on page 87

³ Keyboard, video, mouse (KVM)

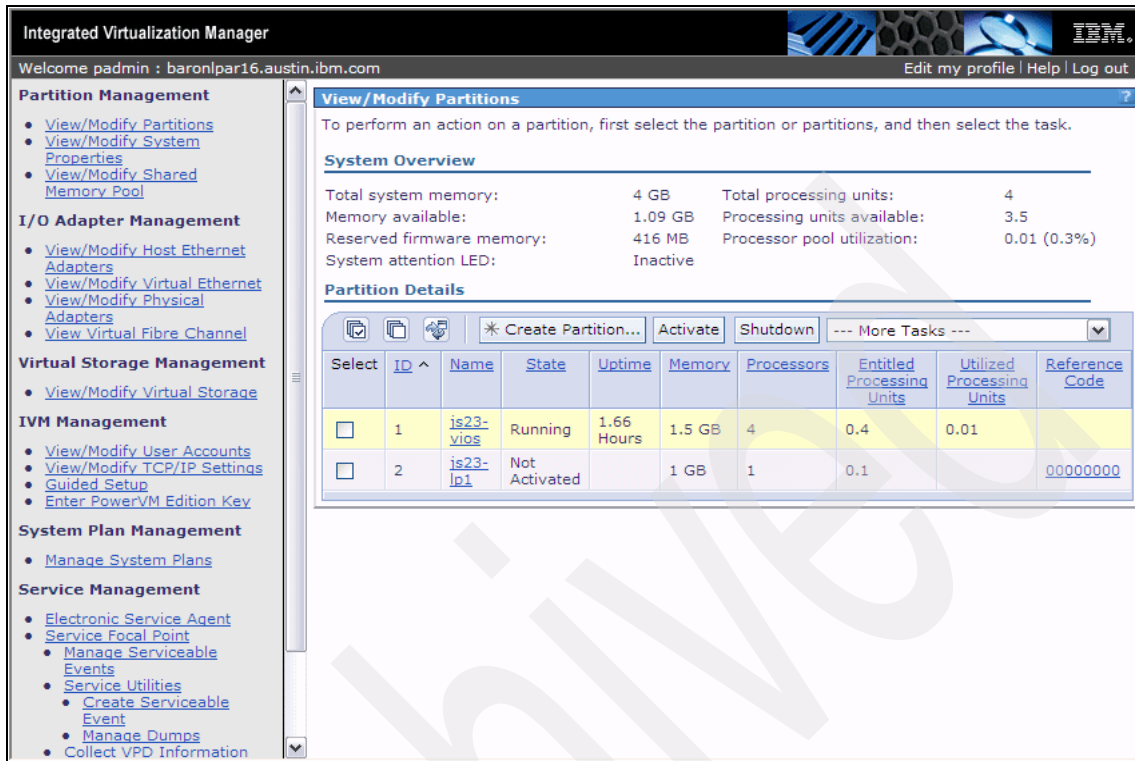


Figure 4-6 IVM navigation and work areas

4.2.3 VIOS/IVM command-line interface

The command-line interface (CLI) requires more experience to master than the GUI, but it offers more possibilities for tuning the partition's definitions. It can also be automated through the use of scripts.

The text-based console with the CLI is accessible through a Serial Over LAN (SOL) terminal connection to the blade through the Advanced Management Module (AMM) or through network connectivity using Telnet or Secure Shell (SSH).

The login to the CLI is to a restricted shell. The restricted shell can be identified by the dollar sign (\$) prompt. Help is available:

- ▶ To display a list of available commands, use the **help** command.
- ▶ To view individual VIOS command help, use the **-help** flag.
- ▶ To view IVM command help, use the **--help** flag.
- ▶ To view detailed command help, use the **man** command.

Note: Not all IVM commands are displayed by using the **help** command. For a complete listing of these commands, refer to *Virtual I/O Server and Integrated Virtualization Manager Command Reference*, available from:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphcg/iphcg.pdf>

4.3 First VIOS login

During the first login to the VIOS running on an IBM BladeCenter JS23 or JS43, you are prompted to complete tasks that are required to make the server operational and IVM accessible through a Web browser. This first connection can be made from an SOL or KVM session.

4.3.1 Password set

From the login prompt, use the default user ID of `padmin`. You enter a new password and then re-enter it for confirmation, as shown in Example 4-1.

4.3.2 License acceptance

The licensing agreements must be accepted prior to starting any configuration tasks. After setting the password, you will be presented three license options as shown in Example 4-1.

Enter `a` to accept, `d` to decline, or `v` to view the license terms.

Example 4-1 Initial login screen showing password set and licensing options

IBM Virtual I/O Server

login: `padmin`

[compat]: 3004-610 You are required to change your password.
Please choose a new one.

`padmin`'s New password:
Enter the new password again:

Indicate by selecting the appropriate response below whether you accept or decline the software maintenance terms and conditions.
[Accept (a)] | Decline (d) | View Terms (v)

After you accept the agreement (by selecting a), enter the **license -accept** command as shown in Example 4-2.

Example 4-2 The license command

```
$ license -accept
```

The status of the license can be verified by using the **license** command with no flags, as shown in Example 4-3.

Example 4-3 The license status

```
$ license
The license has been accepted
en_US Apr 2 2009, 12:33:16 10(padmin)
```

4.3.3 Initial network setup

IVM requires a valid network configuration to be accessed by a Web browser. The VIOS IP address can be set by two methods:

- ▶ By using the **mktcpip** command
- ▶ By starting a SMIT-like tool by using the **cfgassist** command

Note: If you are reinstalling VIOS on the same blade and want to remove the previous logical partition information and return the blade to an original unconfigured state prior to making any configuration changes, execute the following command:

```
lpcfgop -o clear
```

The LPAR information will now be cleared on the next system restart.

The help function in the IVM CLI does not display the **lpcfgop** command. However, you can obtain the command's description by using:

```
man lpcfgop
```

Use the **lstcpip -adapters** command to determine the available network adapters, as shown in Example 4-4 on page 82. When configuring an interface, be sure to pick an adapter that has physical access to the network such as a Logical Host Ethernet Port and not a Virtual I/O Ethernet Adapter.

Example 4-4 The `lstcpip -adapters` command used to display available network adapters

```
$ lstcpip -adapters
```

Ethernet adapters:

```
ent0 Available Logical Host Ethernet Port (lp-hea)
ent1 Available Logical Host Ethernet Port (lp-hea)
ent2 Available Virtual I/O Ethernet Adapter (1-lan)
ent3 Available Virtual I/O Ethernet Adapter (1-lan)
ent4 Available Virtual I/O Ethernet Adapter (1-lan)
ent5 Available Virtual I/O Ethernet Adapter (1-lan)
ent6 Available 05-20 Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent7 Available 05-21 Gigabit Ethernet-SX PCI-X Adapter (14106703)
ibmvmc0 Available Virtual Management Channel
$
```

Choose the corresponding interface on an adapter (ent0 and en0, in this example) that will be the target of the TCP/IP configuration.

Using the `mktcpip` command method

The `mktcpip` command has the following syntax:

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

The `-cabletype` and `-plen` flags are the only two optional parameters.

Example 4-5 shows how to configure a VIOS network interface.

Example 4-5 The `mktcpip` command

```
$ mktcpip -hostname saturn -inetaddr 172.16.1.200 -interface en0 -start
-netmask 255.255.255.0 -gateway 172.16.1.1 -nsrvaddr 172.16.1.199
-nsrvidomain customer.com
```

Using the `cfgassist` menu method

The second method using the `cfgassist` tool requires the same network information, but is entered in a more user-friendly method.

Example 4-6 on page 83 shows the initial `cfgassist` menu. You start the tool by entering the `cfgassist` command.

Example 4-6 The cfgassist menu

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
Electronic Service Agent
```

Esc+1=Help
F9=Shell

Esc+2=Refresh
F10=Exit

Esc+3=Cancel
Enter=Do

F8=Image

By selecting **VIOS TCP/IP Configuration**, you will be presented a list of available network interfaces as shown in Example 4-7.

Example 4-7 cfgassist VIOS Available Network Interfaces

Config Assist for VIOS

```
Mo+-----
      Available Network Interfaces

      Move cursor to desired item and press Enter.

      [TOP]
      en0      Standard Ethernet Network Interface
      en1      Standard Ethernet Network Interface
      en2      Standard Ethernet Network Interface
      en3      Standard Ethernet Network Interface
      en4      Standard Ethernet Network Interface
      en5      Standard Ethernet Network Interface
      en6 05-20 Standard Ethernet Network Interface
      en7 05-21 Standard Ethernet Network Interface
      et0      IEEE 802.3 Ethernet Network Interface
      et1      IEEE 802.3 Ethernet Network Interface
      [MORE...6]

      Esc+1=Help      Esc+2=Refresh      Esc+3=Cancel
      F8=Image        F10=Exit          Enter=Do
Es| /Find            n=Find Next
F9+-----
```

Select the desired interface. On the next screen, shown in Example 4-8, you enter the TCP/IP configuration by pressing the Enter key. This completes the initial TCP/IP configuration of the VIOS.

Example 4-8 cfgassist TCP/IP interface configuration entry page

VIOS TCP/IP Configuration

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

	[Entry Fields]
* Hostname	[saturn]
* Internet ADDRESS (dotted decimal)	[172.16.1.200]
Network MASK (dotted decimal)	[255.255.255.0]
* Network INTERFACE	en0
Default Gateway (dotted decimal)	[172.16.1.1]
NAMESERVER	
Internet ADDRESS (dotted decimal)	[172.16.1.199]
DOMAIN Name	[customer.com]
CableType	tp

Esc+1=Help	Esc+2=Refresh	Esc+3=Cancel	Esc+4=List
Esc+5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Reviewing the network configuration

After you configure a network adapter, you can review the settings by using the **lstcpip** command. The **lstcpip** command has the following syntax:

```
Usage:  lstcpip [-num] [ [-routtable] | [-routinfo] | [-state] | [-arp]
]
           [-sockets [-family inet | inet6 | unix] ]
           [-stored] [-adapters] [-hostname] [-namesrv]
lstcpip [-state [-field FieldName ...] ] |
           [-routtable [-field FieldName ...] ] [-fmt delimiter]
[-num]
lstcpip [-namesrv] | [-interfaces] [-fmt delimiter]
```

Example 4-9 on page 85 shows sample output for the **-adapters**, **-interfaces**, **-routtable**, and **-stored** flags as they are used with the **lstcpip** command to show basic TCP/IP configuration. The example also shows that a shared Ethernet adapter is created.

Example 4-9 lstcpip command sample output

```
$ lstcpip -interfaces
```

Name	Address	Netmask	State	MAC
en1	-	-	detach	00:1a:64:76:00:09
en3	-	-	detach	2a:08:4b:c1:cb:04
en4	-	-	detach	2a:08:4b:c1:cb:05
en5	-	-	detach	2a:08:4b:c1:cb:06
et1	-	-	detach	00:1a:64:76:00:09
et3	-	-	detach	2a:08:4b:c1:cb:04
et4	-	-	detach	2a:08:4b:c1:cb:05
et5	-	-	detach	2a:08:4b:c1:cb:06
en6	172.16.1.200	255.255.255.0	up	00:1a:64:76:00:08
et6	-	-	detach	00:1a:64:76:00:08

```
$ lstcpip -routtable
```

Routing tables

Destination	Gateway	Flags	Refs	Use	If	Exp	Groups
-------------	---------	-------	------	-----	----	-----	--------

Route Tree for Protocol Family 2 (Internet):

default	172.16.1.1	UG	2	288871	en6	-	-
127/8	localhost	U	9	209	lo0	-	-
172.16.1.0	saturn.customer.com	UHSb	0	0	en6	-	-
- =>							
172.16.1/24	saturn.customer.com	U	2	5628596	en6	-	-
-							
saturn.ibm.com	localhost	UGHS	0	39074	lo0	-	-
172.16.1.255	saturn.customer.com	UHSb	0	4	en6	-	-
-							

Route Tree for Protocol Family 24 (Internet v6):

::1	::1	UH	0	232	lo0	-	-
-----	-----	----	---	-----	-----	---	---

```
$ lstcpip -stored
```

saturn

Network Interface Attributes

attributes: en0

IPv4 address = 172.16.1.200

Network Mask = 255.255.255.0

State = detach

attributes: en1

State = down

attributes: en2

State = down

```
attributes: en3
    State = down

attributes: en4
    State = down

attributes: en5
    State = down

attributes: et0
    State = detach

attributes: et1
    State = down

attributes: et2
    State = down

attributes: et3
    State = down

attributes: et4
    State = down

attributes: et5
    State = down

attributes: en6
    IPv4 address = 172.16.1.200
    Network Mask = 255.255.255.0
    State = up

attributes: et6
    State = detach

attributes: en7
    State = down

attributes: en8
    State = down

attributes: et7
    State = down

attributes: et8
    State = down

Static Routes:
```

Route 1:
 hopcount = 0
 default gateway = 172.16.1.1

DNS information:
 nameserver 172.16.1.199
 domain customer.com

To remove all or part of the TCP/IP configuration, use the **rmtcpip** command. The **rmtcpip** command has the following syntax:

Usage: **rmtcpip** [-f] [-nextboot] {-all | [-hostname] [-routing]
 [-interface ifnameList]}
 rmtcpip [-f] {-all | [-namesrv] [-hostname] [-routing]
 [-interface ifnameList]}

To remove all TCP/IP configuration, use the **rmtcpip -f -all** command.

Note: The instructions given in this section are for IPv4. For instructions of how to configure IPv6, refer to the **mktcpip** man page or refer to:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphcg/mktcpip.htm>

4.4 First IVM connection

When the initial networking tasks are completed, the IVM GUI should be accessible through a Web browser. The VIOS Web server supports HTTP and HTTPS connections.

4.4.1 Connecting to IVM

Using a Web browser window, connect using HTTP or HTTPS to the IP address that you assigned to the VIOS during the installation process (see 4.3.3, “Initial network setup” on page 81).

A Welcome window that contains the login and password prompts opens, as shown in Figure 4-7 on page 88. The default user ID is `padmin`, and the password is the one you defined during the VIOS installation.

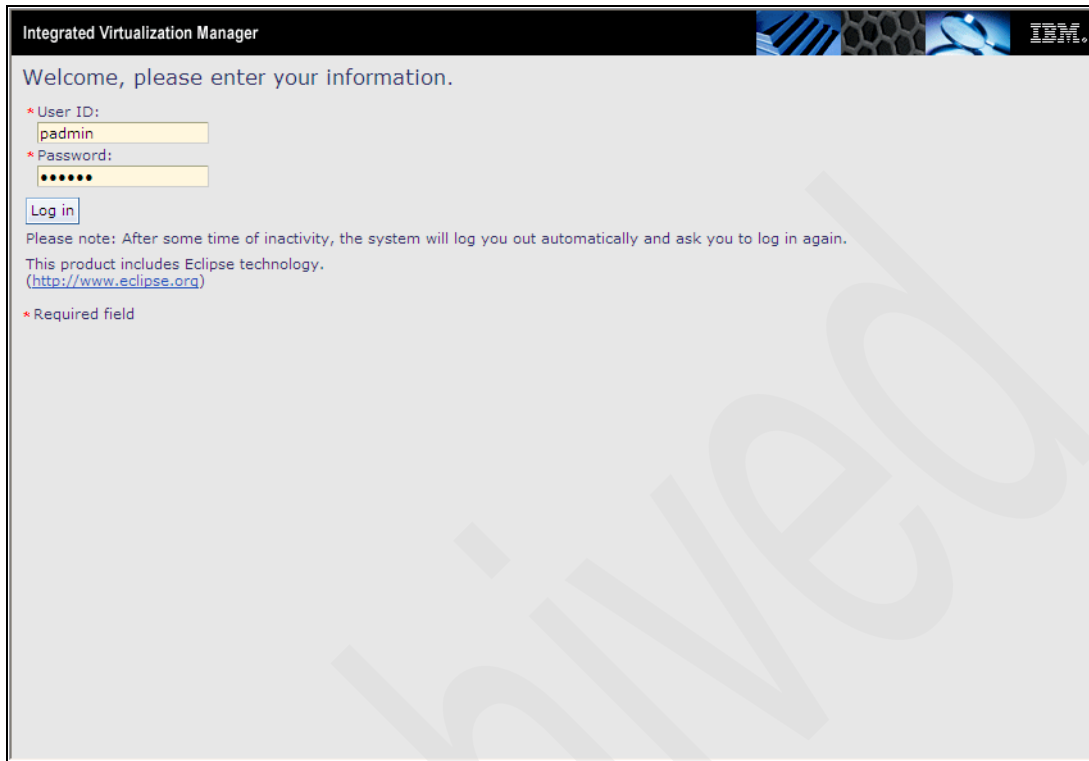


Figure 4-7 The Welcome window

The first connection to the IVM GUI will display the guided setup window as shown in Figure 4-8 on page 89. Expanding the sections on the window provides additional information about configuration and management tasks, with links directly to some of the functions. You can return to this window at any time by clicking the **Guided Setup** link in the navigation area.

Before configuring any additional LPARs or resources, review the default installation values of the VIOS.



Figure 4-8 Guided Setup window

4.4.2 Verify and modify VIOS partition memory and processors

After the initial installation of the VIOS, there is only one LPAR, the VIOS, on the system with the following characteristics:

- ▶ The ID is 1.
- ▶ The name is equal to the system's serial number.
- ▶ The state is Running.
- ▶ The allocated memory is between 1 GB and one-eighth of the installed system memory.
- ▶ The number of (virtual) processors is equal to the number of cores, and the Entitled Processing Units is equal to 0.1 times the number of allocated processors.

Details of all configured partitions, including the VIOS, are summarized on the View/Modify Partitions panel as shown in Figure 4-9 on page 90. This panel can be accessed by the **View/Modify Partitions** link in the navigation area.

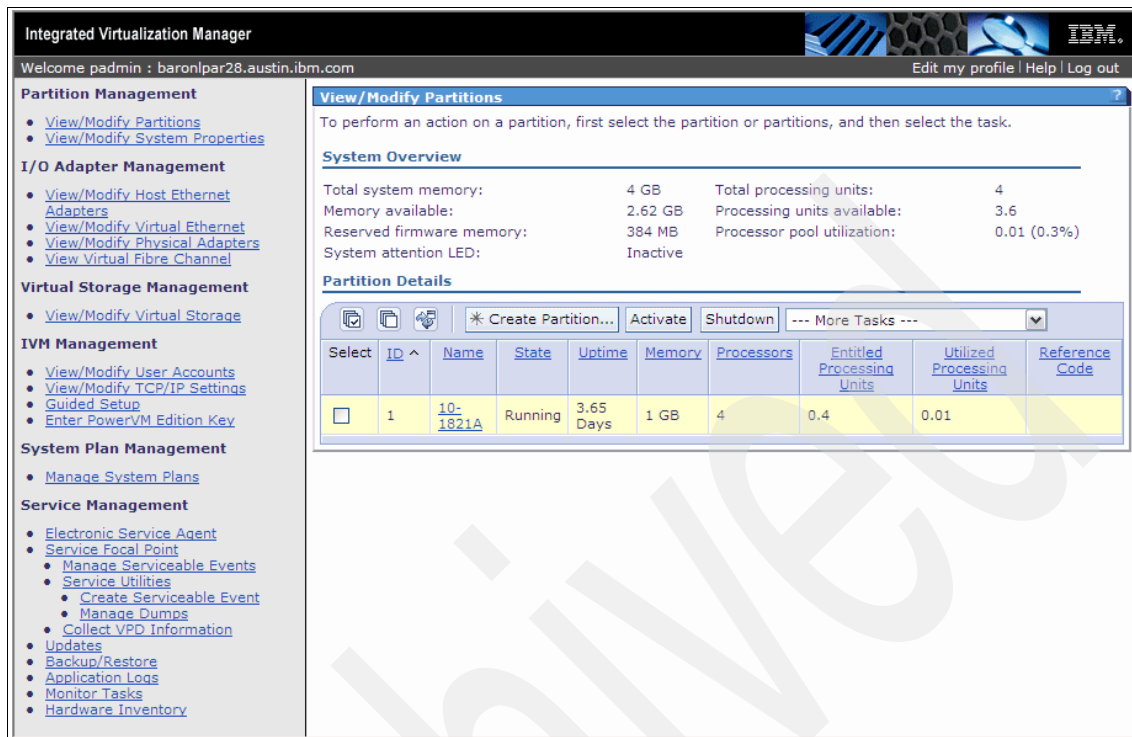


Figure 4-9 View/Modify Partitions window

Administrators can change properties of the VIOS LPAR, including memory or processing units allocation by using the IVM GUI. From the View/Modify Partitions window, click the link in the Name column that corresponds to ID 1 (The VIOS will always be ID or LPAR 1).

The Partition Properties window will be displayed in a new window, as shown in Figure 4-10 on page 91. The name of the VIOS can be changed from the General tab, if desired. The Memory and Processing tabs are used to view or change the allocations. Figure 4-11 on page 92 shows the General tab.

Partition Properties: js23-vios (1)

General | Memory | Processing | Ethernet | Physical Adapters

General

Partition name:

Partition ID:

Environment: Virtual I/O Server

State: Running

Attention LED: ▼

Settings

Boot mode: ▼

Keylock position: ▼

Partition workload group participant: ☐

Automatically start when system starts: ☒

Dynamic Logical Partitioning (DLPAR)

Partition hostname or IP address: 9.3.29.70

Partition communication state: Active

Memory DLPAR capable: Yes

Processing DLPAR capable: Yes

I/O adapter DLPAR capable: Yes

Figure 4-10 Partition Properties, General tab

Figure 4-11 on page 92 shows the Memory tab.

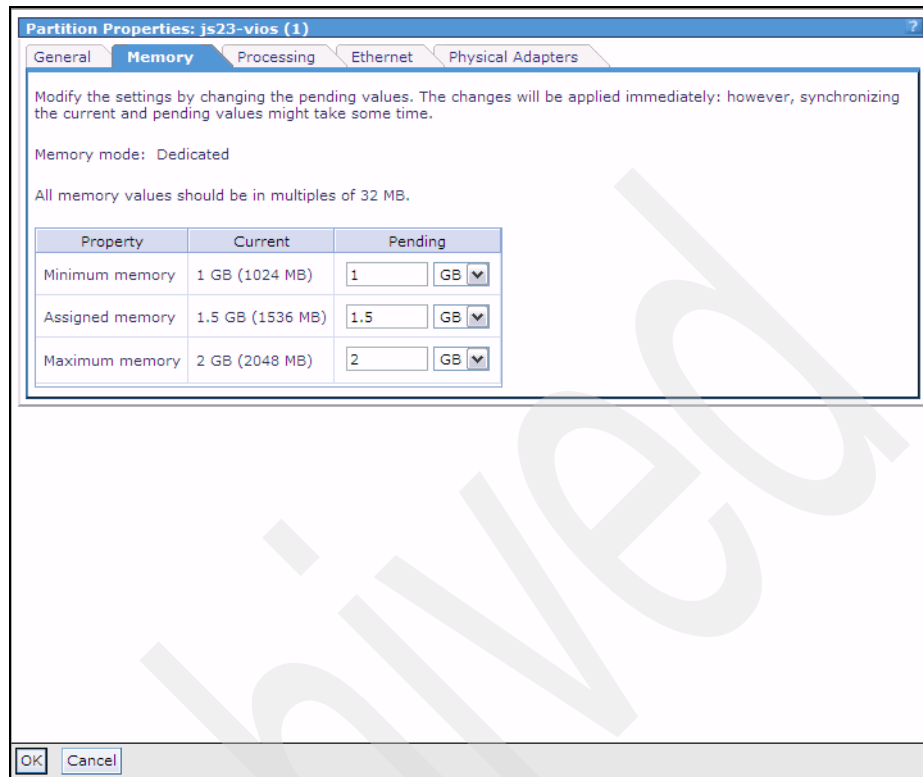


Figure 4-11 Partition Properties, Memory tab

The default memory configuration for the VIOS LPAR, one-eighth of system memory with a minimum value of 1 GB. You may have to increase memory values if you are using additional expansion cards or combinations of expansion cards and EtherChannel configurations, or you plan to have an LPAR supporting IBM i partition.

Tip: You should have approximately 128 MB for each logical port. If you have JS43, that means 4 logical ports by default. VIOS uses 128 MB of memory per port if the port is configured (for example, in the up state, or have an EtherChannel or Shared Ethernet Adapter on top of the device).

Figure 4-12 on page 93 shows the Processing tab.

Partition Properties: js23-vios (1)

General

Memory

Processing

Ethernet

Physical Adapters

Modify the settings by changing the pending values. The changes will be applied immediately; however, synchronizing the current and pending values might take some time.

Processing Units

Property	Current	Pending
Minimum	0.1	<input type="text" value="0.1"/>
Assigned	0.4	<input type="text" value="0.4"/>
Maximum	4	<input type="text" value="4.0"/>

Virtual Processors

Property	Current	Pending
Minimum	1	<input type="text" value="1"/>
Assigned	4	<input type="text" value="4"/>
Maximum	4	<input type="text" value="4"/>

General

Property	Current	Pending
Uncapped weight	Medium - 128	<input type="text" value="Medium - 128"/>

Processor compatibility mode:

Current value: POWER6+

Preferred value:

OK

Cancel

Figure 4-12 Partition Properties, Processing tab

Processing unit allocations for the VIOS should remain at the installation defaults. But you should monitor utilization and adjust the Assigned amount, as required. The Virtual Processor default settings should not be changed.

The `lshwres` and `chsyscfg` commands are used to display and change memory and processor values, as shown in Example 4-10 on page 94.

Example 4-10 CLI commands for displaying and altering memory and processor values

```
$ lshwres -r mem --level lpar --filter "\"lpar_ids=1\"" -F curr_mem
1024

$ chsyscfg -r prof -i "lpar_id=1,desired_mem=1152"

$ lshwres -r mem --level lpar --filter "\"lpar_ids=1\"" -F curr_mem
1152
$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F
curr_proc_units
0.40

$ chsyscfg -r prof -i "lpar_id=1,desired_proc_units=0.5"

$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F
curr_proc_units
0.50

$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F curr_procs
4
$ chsyscfg -r prof -i "lpar_id=1,desired_procs=3"
$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F curr_procs
3
```

The Ethernet tabs are discussed in 4.5.2, “Virtual Ethernet Adapters and SEA” on page 97. Physical Adapters tabs are discussed in 4.5.3, “Physical adapters” on page 103.

4.5 VIOS network management and setup

When you are preparing the IBM BladeCenter JS23 or JS43 and VIOS for additional LPARs, and planning for the ability for those LPARs to reach a physical network, you have to understand the networking options that are available. The three choices are:

- ▶ Logical Host Ethernet Adapters (lp-HEA)
- ▶ Virtual Ethernet Adapters bridged to a physical adapters by a Shared Ethernet Adapter (SEA)
- ▶ Physical Adapters

4.5.1 Host Ethernet Adapters

Using the Host Ethernet Adapters (HEA) can reduce overall system processor usage by the VIOS compare to a SEA, but will not allow partition mobility. In a JS23, 16 individual logical ports in one port group are available across the two physical ports. The JS43 has 32 logical ports over four physical ports and two ports groups. In either case, the VIOS claims one logical port on each physical port, leaving 14 available logical ports on the JS23 for LPAR assignment and 28 logical ports on the JS43. The assignment of a logical port from the HEA can be done during the LPAR creation process. Refer to 3.8.4, “Integrated Virtual Ethernet” on page 50 for additional technical details about the HEA.

You configure the HEA port mode by selecting **View/Modify Host Ethernet Adapters** from the navigation area to open the panel shown in Figure 4-13.

The screenshot displays the 'View/Modify Host Ethernet Adapters' window in the Integrated Virtualization Manager (IVM). The window title is 'View/Modify Host Ethernet Adapters'. Below the title, a description states: '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.' The main content area features a table with the following columns: 'Select', 'Type', 'Link State', 'Physical Location Code ^', 'Connected Partitions', and 'Available Connections'. The table lists four HEA ports, each with a 1 Gbps link state and 14 available connections. The left sidebar contains navigation links for Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, and Service Management.

Select	Type	Link State	Physical Location Code ^	Connected Partitions	Available Connections
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P1-T6	1	0
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P1-T7	2	13
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P2-T6	1	14
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P2-T7	1	0

Figure 4-13 View/Modify Host Ethernet Adapters window

All four HEA ports on a JS43 are shown. The default configuration is port-sharing with 14 logical connections available per port pair. In this example, two ports

have been set to allow bridging (Available Connections 0) one port has an additional partition connected (Connected Partitions 2).

To view and modify the port properties, select one of the ports and click **Properties**. The HEA Physical Port Properties window opens in a new window; Figure 4-14 shows the General tab already selected.

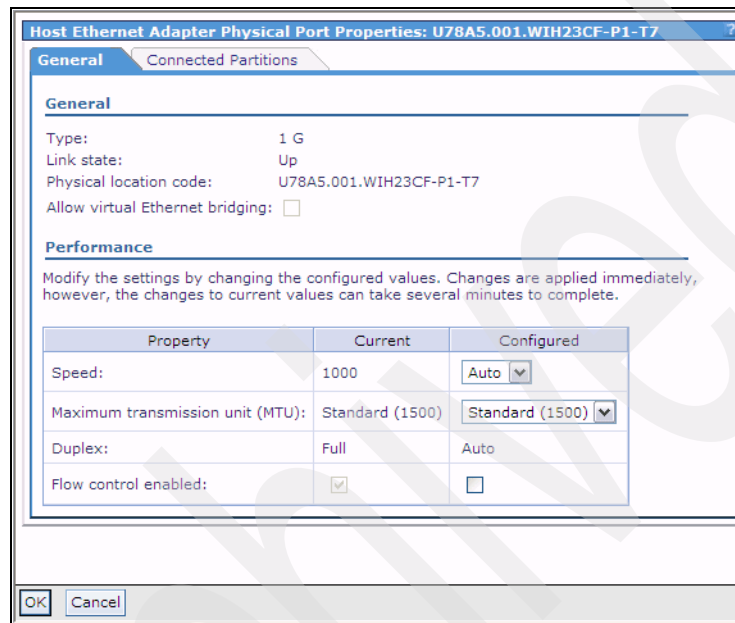


Figure 4-14 HEA Port Properties

You can display a list of connected partitions (if any) and MAC addresses by selecting the **Connected Partitions** tab, as shown in Figure 4-15 on page 97.

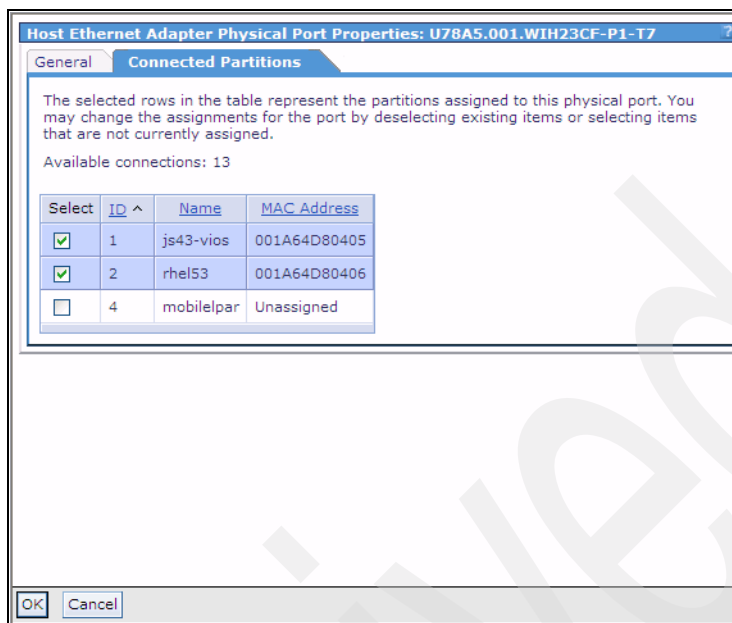


Figure 4-15 HEA Port Properties, Connected Partitions

4.5.2 Virtual Ethernet Adapters and SEA

Virtual adapters exist in the hypervisor that allows LPARs to communicate with each other without requiring a physical network. They can be created for each partition provided by the hypervisor.

Four virtual Ethernet adapters are created by default on the VIOS, and two each for every logical partition. Additional virtual adapters can be created on both the VIOS and logical partitions.

A Shared Ethernet Adapter (SEA) is a new virtual adapter that is created by bridging between a physical adapter (HEA port or expansion card port) and a virtual adapter on the VIOS. A SEA adapter can also be created by bridging between an EtherChannel adapter and a virtual adapter. The SEA interface can then be mapped to logical partitions, thus providing network connectivity outside of the VIOS and logical partition.

For an HEA adapter port to participate in an SEA Ethernet bridge, the configuration of the port must be changed. The configuration is changed by ensuring that the box **Allow virtual Ethernet bridging** on the HEA Physical Port Properties window is checked, as shown in Figure 4-16 on page 98. This setting

allows the port to operate in promiscuous mode. When this mode is enabled, there is only one logical port available and it is assigned to the VIOS LPAR.

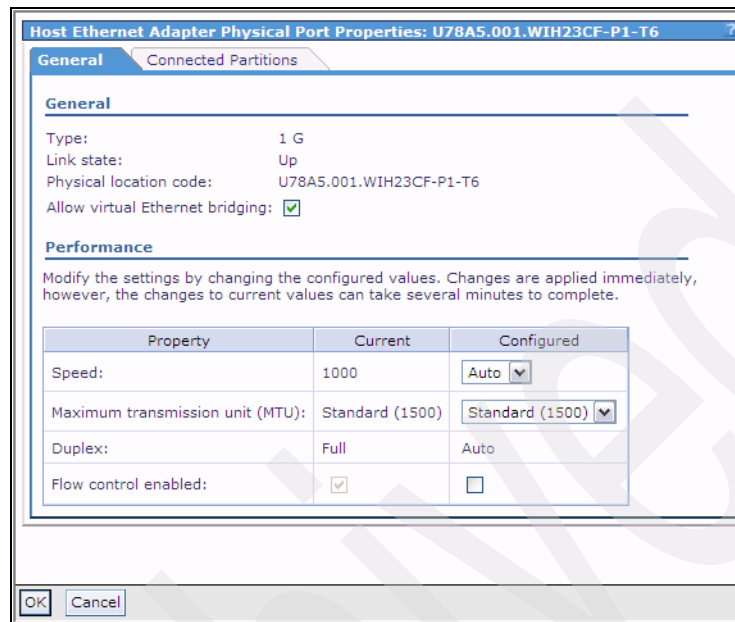


Figure 4-16 HEA port setting for Ethernet bridging

Physical Ethernet ports on an expansion cards do not require configuration prior to being used in a SEA environment.

The SEA adapter is configured by selecting the **View/Modify Virtual Ethernet** link in the navigation area. If the four default virtual Ethernet adapter have not already been created, you see the window as shown in Figure 4-17 on page 99. Click the **Initialize Virtual Ethernet** button to create the default adapters.

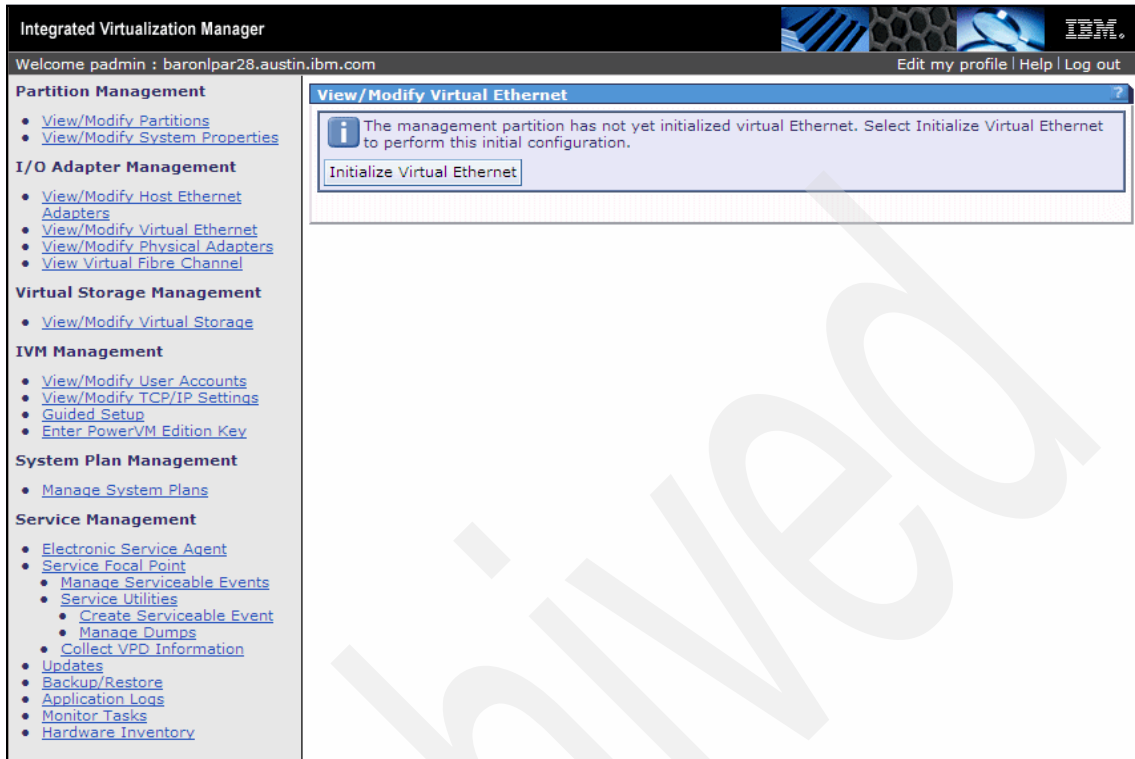


Figure 4-17 View/Modify Virtual Ethernet showing Initialize Virtual Ethernet option

Figure 4-18 on page 100 shows the four virtual Ethernet adapters that are created by default on the VIOS.

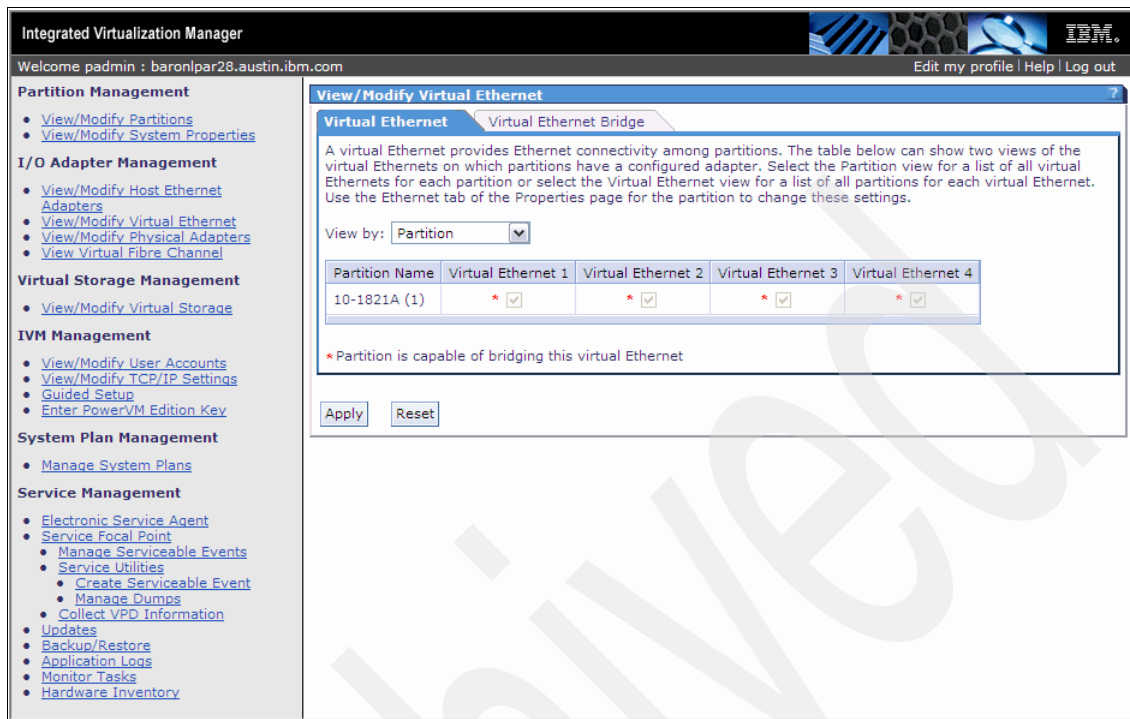


Figure 4-18 View/Modify Virtual Ethernet window

Select the **Virtual Ethernet Bridge** tab to display the virtual to physical options for creating an SEA, as shown in Figure 4-19 on page 101. The drop-down box in the Physical Adapter column lists the adapters that are available for creating the SEA.

Notes:

- ▶ A physical adapter can only be used to create one SEA in combination with a virtual adapter. Although the drop-down menu allows the selection of the same adapter for another virtual Ethernet ID, an error message is generated when you click **Apply**.
- ▶ EtherChannel adapters, if created, will also be listed in the drop-down box.

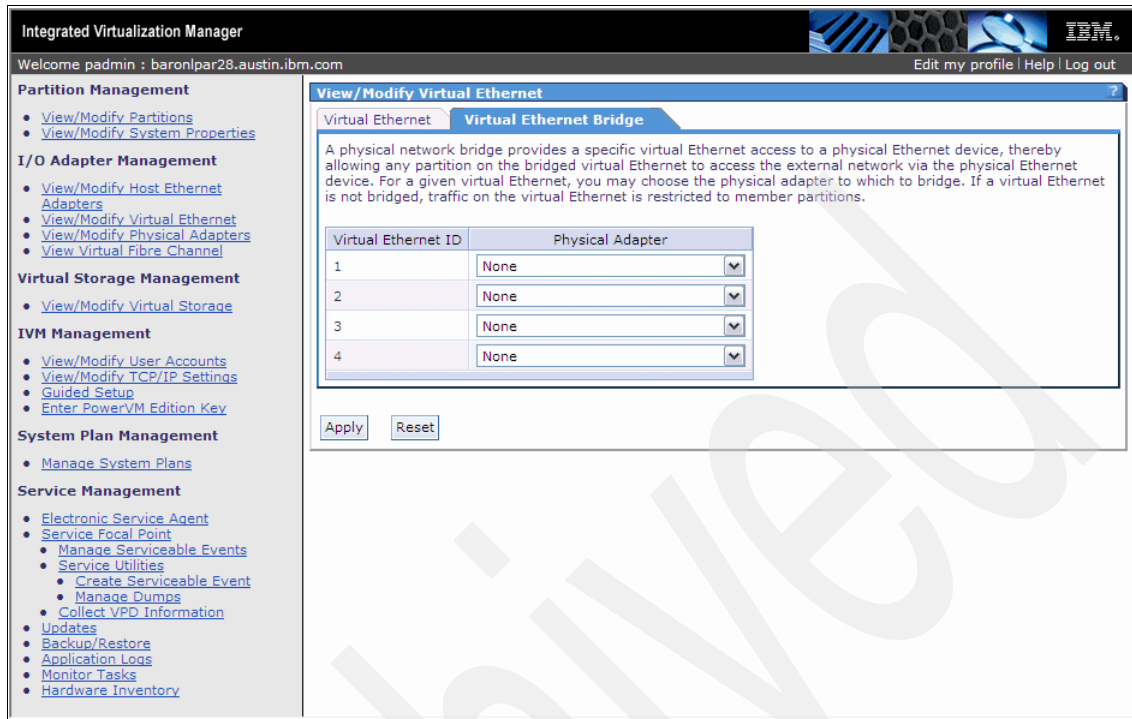


Figure 4-19 View/Modify Virtual Ethernet Bridge tab

Figure 4-20 on page 102 shows a physical adapter selection.

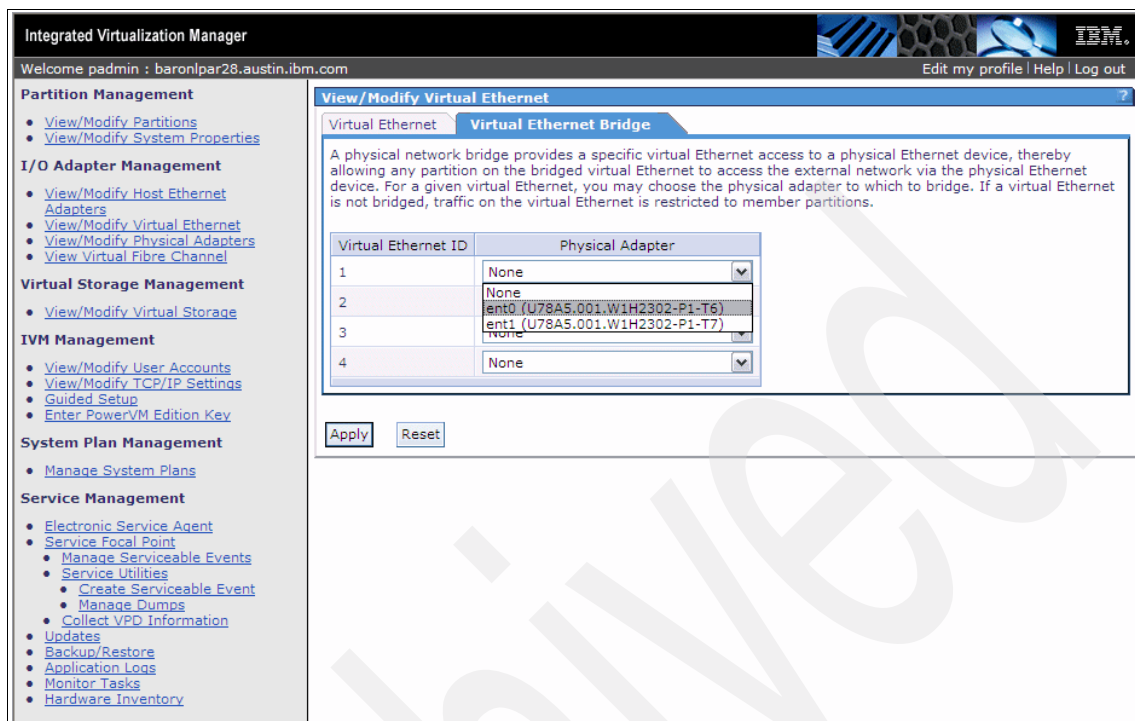


Figure 4-20 Physical adapter selection for SEA creation

Figure 4-21 on page 103 indicates the successful creation of the SEA.

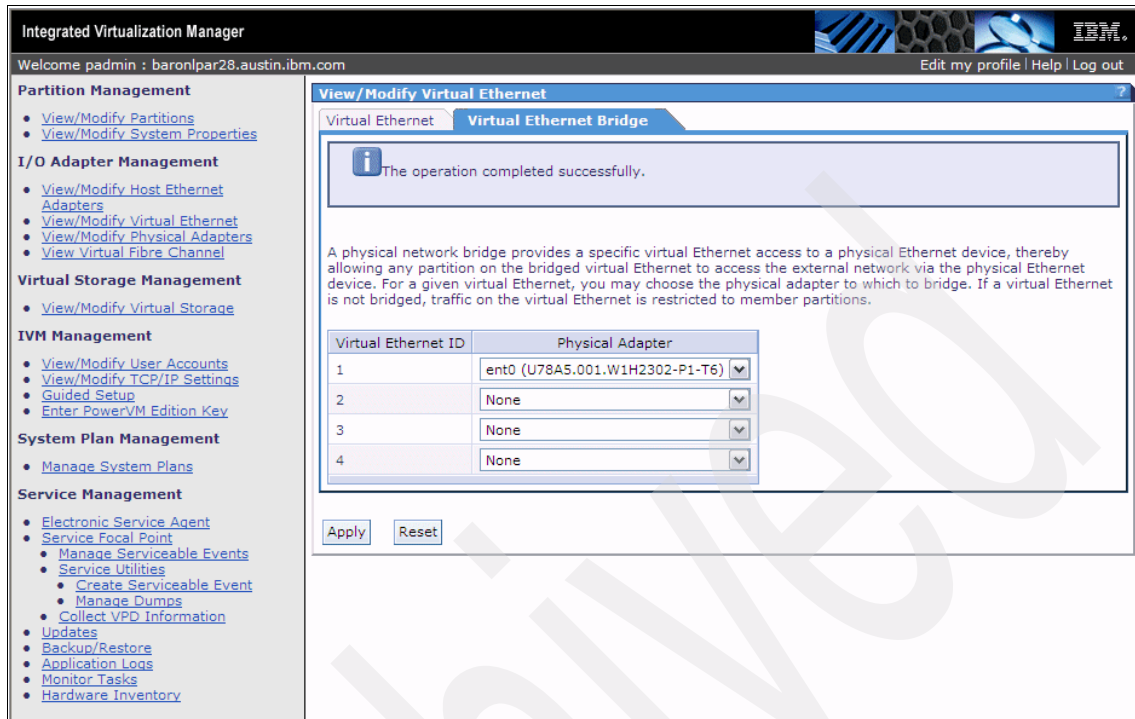


Figure 4-21 Successful SEA creation result

4.5.3 Physical adapters

With the IBM BladeCenter JS23 or JS43, you have the option to assign physical hardware adapters to an LPAR. From a network perspective, only Ethernet expansion cards can be reassigned to an LPAR. The HEA adapter ports cannot be assigned to a logical partition.

Note: When using IBM i and shared memory partitions, the resources must be purely virtual.

To assign a physical adapter, select **View/Modify Physical Adapters** in from the navigation area to display the panel shown in Figure 4-22 on page 104.

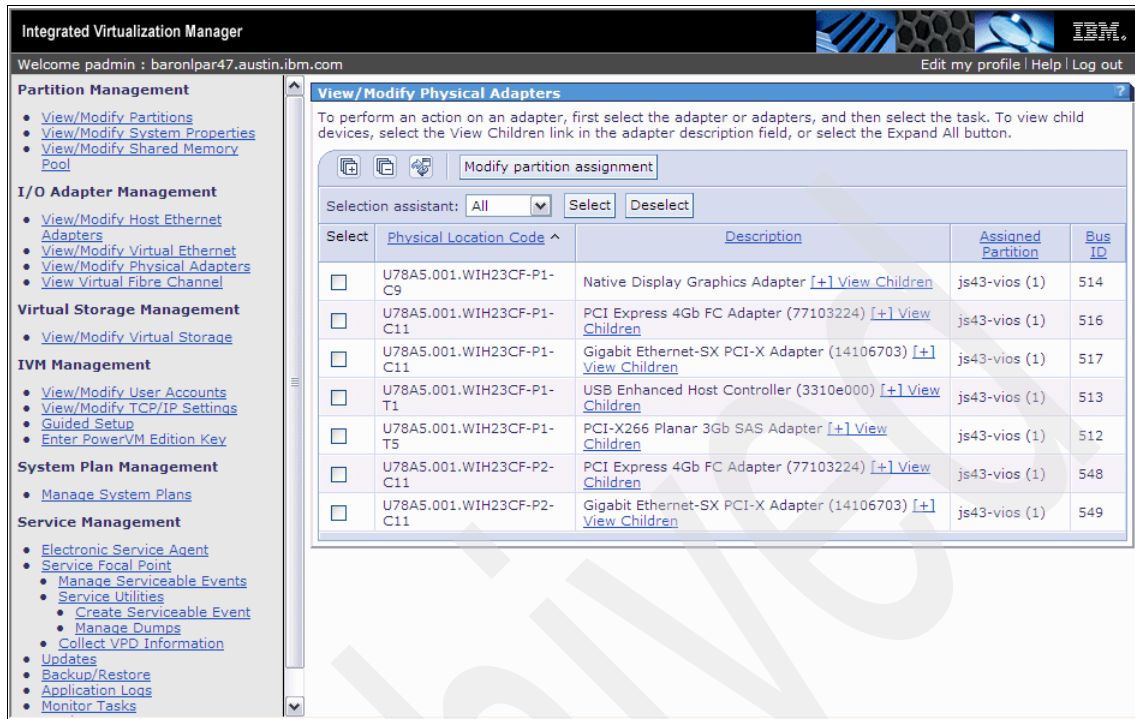


Figure 4-22 View/Modify Physical Adapters window

By default, all physical adapters are owned by the VIOS LPAR. By using the **Modify Partition Assignment** button, you can change the assigned partition.

In the example shown in Figure 4-23 on page 105, the Gigabit Ethernet expansion card ports are being reassigned to partition 2.

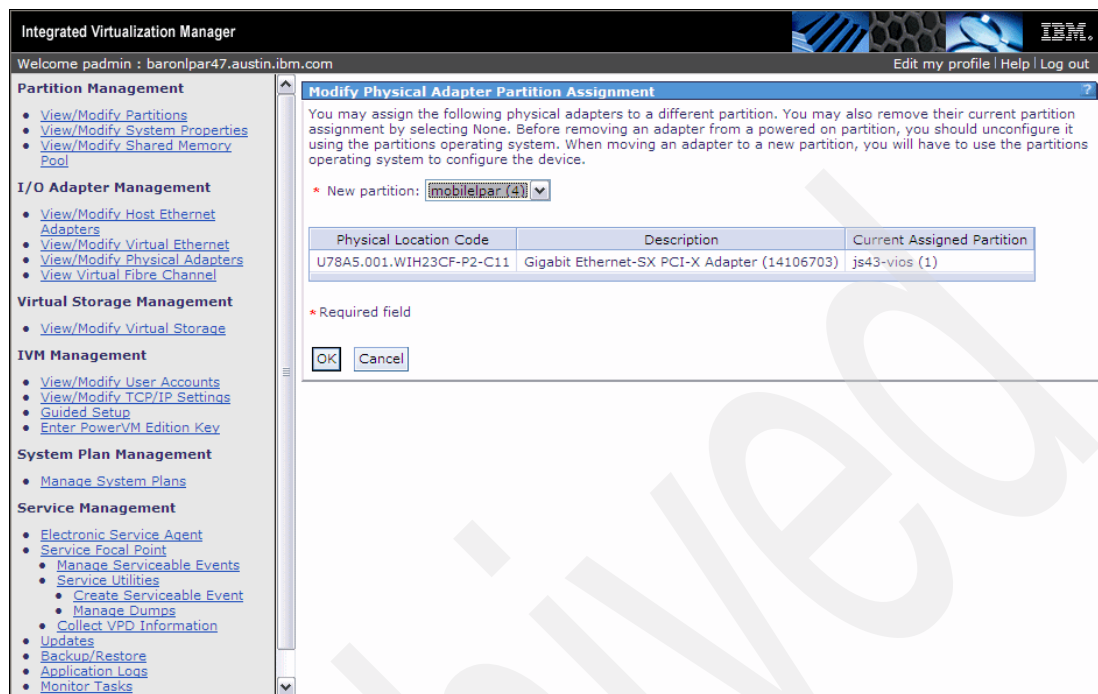


Figure 4-23 Physical Adapter assignment to new partition

Figure 4-24 on page 106 shows the change in partition ownership.

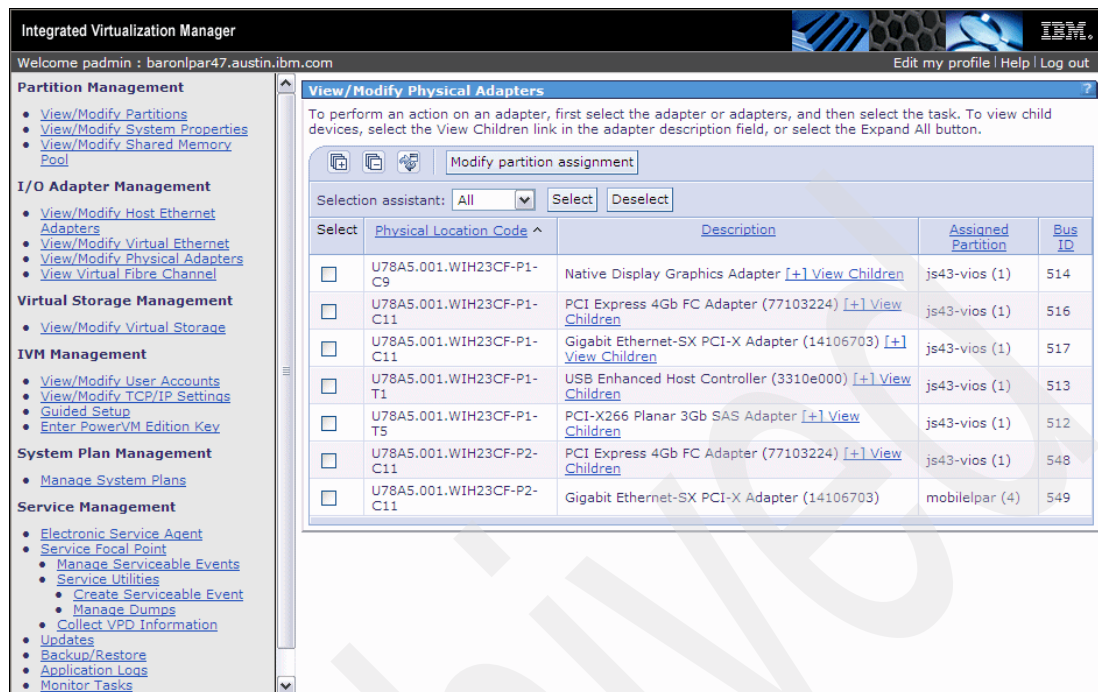


Figure 4-24 View/Modify Physical Adapter panel shows change of ownership of Gigabit Ethernet Adapter

Example 4-11 on page 107 shows the changes in adapter availability in an AIX logical partition, starting with the original virtual Ethernet adapter through the addition of the two physical ports from an IBM BladeCenter JS23 or JS43 expansion card.

Example 4-11 Physical port addition to logical partition

```
# lsdev -Cc adapter
ent0 Available Virtual I/O Ethernet Adapter (1-lan)
vsa0 Available LPAR Virtual Serial Adapter
vscsi0 Available Virtual SCSI Client Adapter
# lsdev -Cc adapter
ent0 Available Virtual I/O Ethernet Adapter (1-lan)
ent1 Defined 01-20 Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent2 Defined 01-21 Gigabit Ethernet-SX PCI-X Adapter (14106703)
vsa0 Available LPAR Virtual Serial Adapter
vscsi0 Available Virtual SCSI Client Adapter
# cfgmgr
# lsdev -Cc adapter
ent0 Available Virtual I/O Ethernet Adapter (1-lan)
ent1 Available 01-20 Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent2 Available 01-21 Gigabit Ethernet-SX PCI-X Adapter (14106703)
vsa0 Available LPAR Virtual Serial Adapter
vscsi0 Available Virtual SCSI Client Adapter
#
```

Note: When removing a physical adapter from an LPAR you might have to remove a PCI bus device by using an `rmdev` command on the LPAR's command line. The IVM interface displays an error message with text indicating the device that must be removed before the change in LPAR assignment can be performed.

4.5.4 VLANs and virtual networks

When the Virtual I/O Server (also referred to as VIO Server, or VIOS) is installed, four virtual Ethernet adapters are created automatically. These adapters have Port VLAN Identifiers (PVID) of 1 - 4. These PVIDs are considered untagged and do not leave the VIO Server. To extend a VLAN environment into the virtual network, the concept of VLAN tagging using IEEE 802.1q protocol is used to make the extension.

IVM does not provide a direct way to create 802.1q aware interfaces but can manage the assignment of these interfaces after they are created using the CLI.

The steps to perform, and the VIO Server user interfaces used are:

- ▶ Create IEEE 802.1q aware virtual Ethernet adapter on the VIOS that will act as a trunking adapter to carry multiple VLANs (CLI)
- ▶ Create a SEA between the physical adapter and the VLAN aware virtual Ethernet adapter on the VIO Sever (IVM)

- ▶ Assign the LPAR virtual Ethernet adapter to the VIO Server virtual Ethernet
- ▶ Run **cfgmgr** on LPAR (if running) and configure new interfaces (client LPAR CLI)

Section 3.6 in *IBM System p Advanced POWER Virtualization (PowerVM) Best Practices*, REDP-4194 contains additional discussions on extending VLANs into virtual networks. The implementation covered in this Redpaper is for a system managed by the Hardware Management Console (HMC) not IVM, but the technical discussions of VLANs in a virtual network are relevant.

Another example of implementation from an IBM BladeCenter and IVM perspective, *Complex networking using Linux on Power blades*, can be found at:

<http://www.ibm.com/developerworks/power/library/1-bladenetconf/index.html?ca=drs->

VLAN configuration of BladeCenter Ethernet switch modules or other Ethernet switches external to the BladeCenter are not covered in this document.

Creating new VIOS virtual Ethernet adapters

The four default virtual adapters that are created by the VIO Server during installation cannot be modified for VLAN tagging use. Therefore new virtual adapters must be created using the CLI with the desired VLAN information.

The **lsdev** command can be used to review the current real and virtual Ethernet adapters on our lab VIO Server, as shown Example 4-12.

Example 4-12 lsdev command used to list current Ethernet adapters

```
$ lsdev -type adapter |grep ent
ent0          Available   Logical Host Ethernet Port (lp-hea)
ent1          Available   Logical Host Ethernet Port (lp-hea)
ent2          Available   Virtual I/O Ethernet Adapter (1-lan)
ent3          Available   Virtual I/O Ethernet Adapter (1-lan)
ent4          Available   Virtual I/O Ethernet Adapter (1-lan)
ent5          Available   Virtual I/O Ethernet Adapter (1-lan)
ent6          Available   Gigabit Ethernet-SX PCI-X Adapter
ent7          Available   Gigabit Ethernet-SX PCI-X Adapter
ent8          Available   Shared Ethernet Adapter
```

Using the **lshwres** command, we can view the existing virtual Ethernet resources. In Example 4-13 on page 109, the four default virtual Ethernet adapters are shown. The output gives details of the adapter characteristics such as LPAR ownership, PVIDs (1 - 4), additional VLANs (none), and IEEE 802.1q capability (0=NO).

Example 4-13 lshwres command showing VIO Server virtual Ethernet adapters

```
$ lshwres -r virtualio --rsubtype eth --level lpar
lpar_name=js23-vios,lpar_id=1,slot_num=3,state=1,ieee_virtual_eth=0,por
t_vlan_id=1,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C03
lpar_name=js23-vios,lpar_id=1,slot_num=4,state=1,ieee_virtual_eth=0,por
t_vlan_id=2,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C04
lpar_name=js23-vios,lpar_id=1,slot_num=5,state=1,ieee_virtual_eth=0,por
t_vlan_id=3,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C05
lpar_name=js23-vios,lpar_id=1,slot_num=6,state=1,ieee_virtual_eth=0,por
t_vlan_id=4,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C06
```

From the **View/Modify Virtual Ethernet** view in IVM, as shown in Figure 4-25, the four default VIO Server Ethernet adapters are displayed.

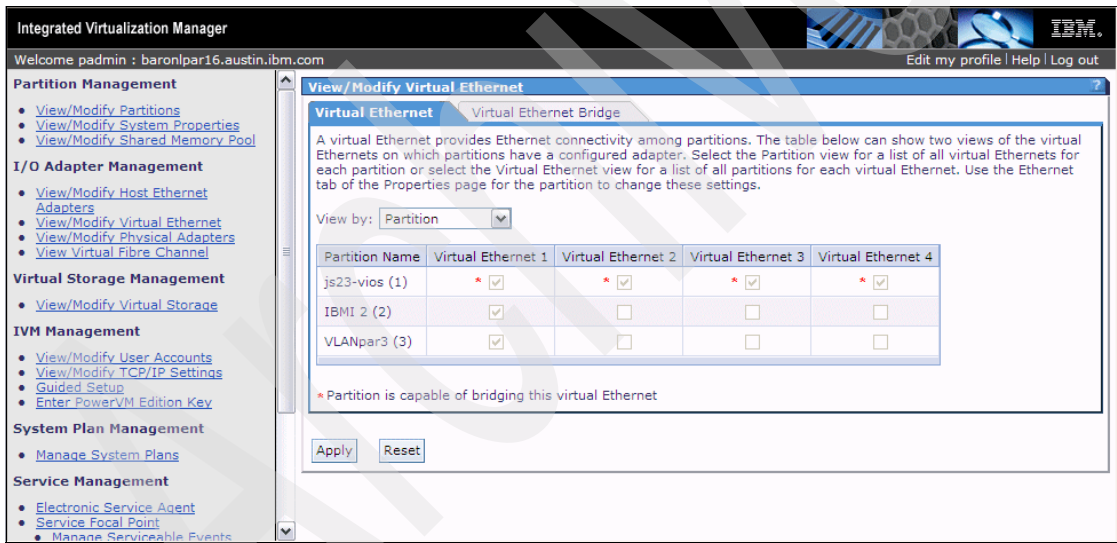


Figure 4-25 Default VIO Server virtual Ethernet Adapters shown by IVM

Note: Figure 4-25 shows additional partitions. Partition creation is not covered until 4.7, “Partition configuration for Virtual I/O Client” on page 139.

When creating a new virtual Ethernet adapter for the VIO Server, we will be assigning PVIDs and additional VLAN IDs. The PVID used should be unique and

not used by any clients in the network or physical Ethernet switch ports. This requirement is to prevent the unintentional removal or stripping of VLAN tags from network packets when entering the VIO Server.

The **chhwres** command is used to create a new VIO Server virtual Ethernet adapter. Example 4-14 creates a new VIOS adapter with a PVID of 555, with IEEE 802.1q enabled and additional VLANs of 20, 30, and 40.

Example 4-14 Using chhwres command to create new VIOS virtual Ethernet adapter

```
$ chhwres -r virtualio --rsubtype eth -o a --id 1 -s 15 -a
port_vlan_id=555,ieee_virtual_eth=1,\"addl_vlan_ids=20,30,40\",is_trunk
=1,trunk_priority=1
```

The flags and their attributes are:

-r virtualio	--rsubtype eth	Type of hardware resource to change
-o a		Perform add operation
--id 1		LPAR ID number
-s 15		Slot number to use
-a		Attributes to add

Attributes for the -a flag are:

port_vlan_id=555	PVID
ieee_virtual_eth=1	Turns on IEEE 802.1q support
add_vlan_ids=20,30,40	Defines additional VLAN ids
is_trunk=1	Must be turned on to pass multiple VLANs
trunk_priority=1	Priority of adapter can be set in range of 1 - 15

The **lsdev** command is repeated and shows the new Virtual I/O Ethernet Adapter at ent9 as shown in Example 4-15.

Example 4-15 The sdev command for listing new ent9 Ethernet adapters

```
$ lsdev -type adapter |grep ent
ent0          Available    Logical Host Ethernet Port (lp-hea)
ent1          Available    Logical Host Ethernet Port (lp-hea)
ent2          Available    Virtual I/O Ethernet Adapter (1-lan)
ent3          Available    Virtual I/O Ethernet Adapter (1-lan)
ent4          Available    Virtual I/O Ethernet Adapter (1-lan)
ent5          Available    Virtual I/O Ethernet Adapter (1-lan)
ent6          Available    Gigabit Ethernet-SX PCI-X Adapter
ent7          Available    Gigabit Ethernet-SX PCI-X Adapter
ent8          Available    Shared Ethernet Adapter
ent9          Available    Virtual I/O Ethernet Adapter (1-lan)
```

The **lshwres** command is run again as shown in Example 4-16 to display the new adapter indicated in bold.

Example 4-16 lshwres command showing new VIOS virtual Ethernet adapter

```
$ lshwres -r virtualio --subtype eth --level lpar
lpar_name=js23-vios,lpar_id=1,slot_num=3,state=1,ieee_virtual_eth=0,por
t_vlan_id=1,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C03
lpar_name=js23-vios,lpar_id=1,slot_num=4,state=1,ieee_virtual_eth=0,por
t_vlan_id=2,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C04
lpar_name=js23-vios,lpar_id=1,slot_num=5,state=1,ieee_virtual_eth=0,por
t_vlan_id=3,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C05
lpar_name=js23-vios,lpar_id=1,slot_num=6,state=1,ieee_virtual_eth=0,por
t_vlan_id=4,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C06
lpar_name=js23-vios,lpar_id=1,slot_num=15,state=1,ieee_virtual_eth=1,po
rt_vlan_id=555,"addl_vlan_ids=20,30,40",is_trunk=1,trunk_priority=1,is
required=0,mac_addr=067E5E2D8C0F
```

With a new VIO Server virtual Ethernet adapter created, we are ready to use IVM for the next steps in configuration.

Create SEA using IEEE802.1q virtual Ethernet adapter

IVM now displays the new virtual Ethernet adapter in the **View/Modify Virtual Ethernet** view as shown in Figure 4-26 on page 112. Notice that the PVID and additional VLAN numbers are shown to be associated with this new virtual adapter.

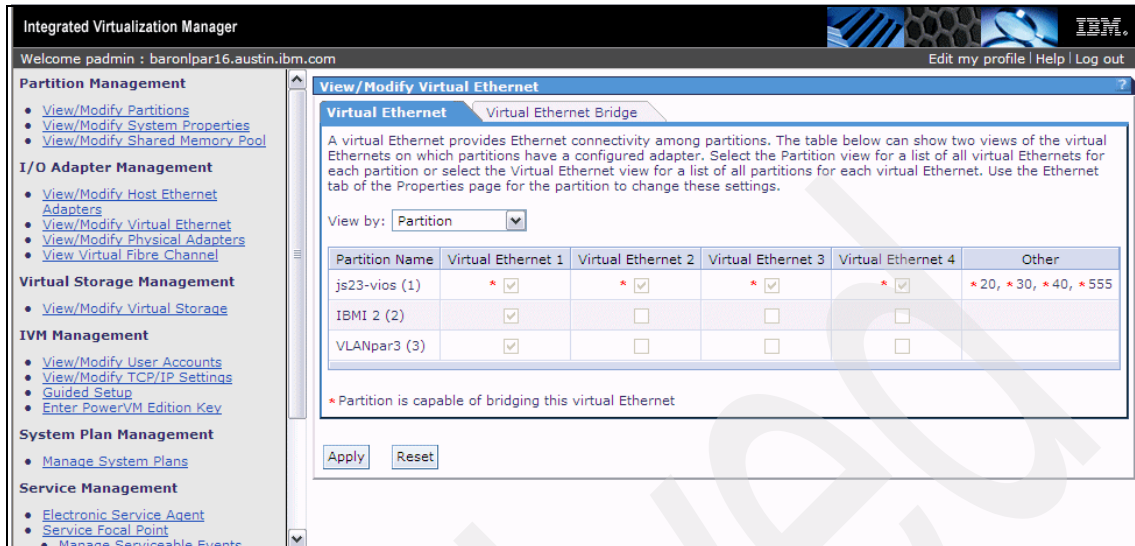


Figure 4-26 View Modify Virtual Ethernet view showing new virtual adapter with multiple VLANs

We now create a SEA or bridge between this new virtual adapter and a physical Ethernet port, in this case a HEA adapter, by first clicking the **Virtual Ethernet Bridge** tab. From the virtual Ethernet list we choose 555(20,30,40) and map it to ent1 as shown in Figure 4-27 on page 113. Click **OK** to complete the assignment and the creation of the SEA.

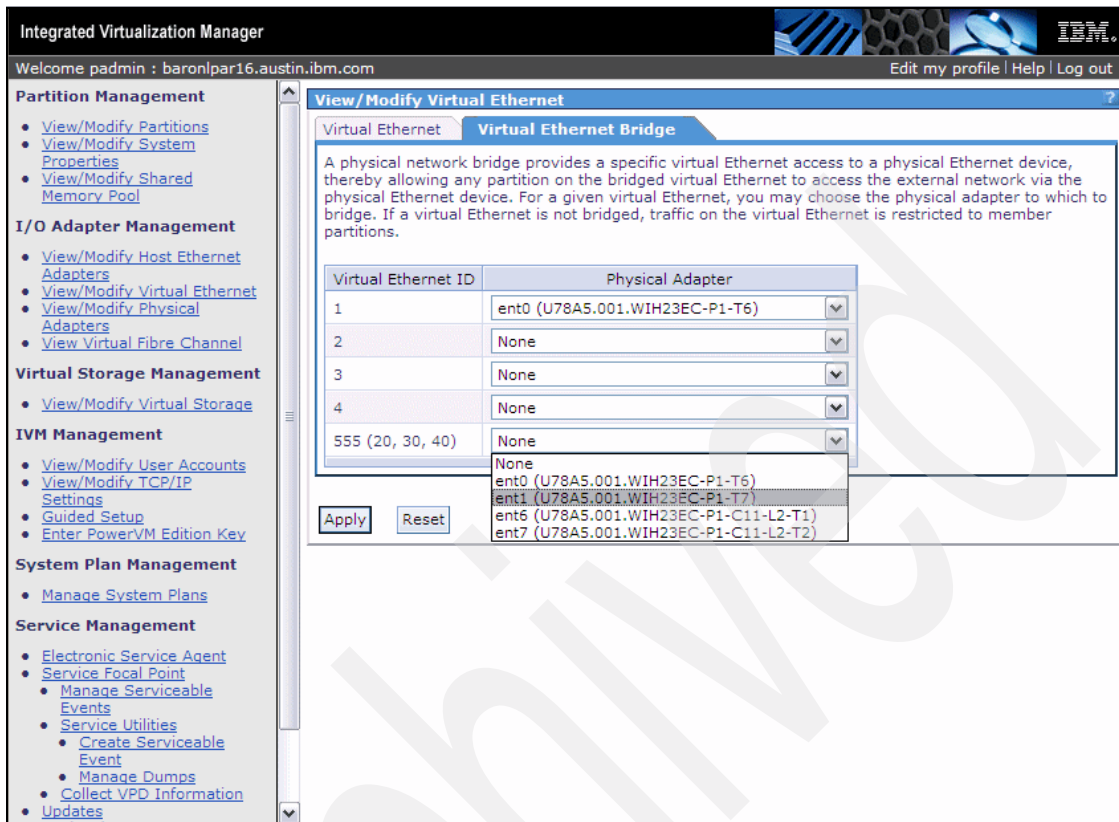


Figure 4-27 Creating a SEA using an IEEE 802.1q enabled virtual adapter

By using the `lsdev` command, we see the new VIO Server virtual Ethernet adapter `ent9`, and the new SEA `ent10` as shown in Example 4-17

Example 4-17 `lsdev` command showing new IEEE 802.1q virtual adapter and SEA

```
$ lsdev -type adapter |grep ent
ent0      Available   Logical Host Ethernet Port (lp-hea)
ent1      Available   Logical Host Ethernet Port (lp-hea)
ent2      Available   Virtual I/O Ethernet Adapter (1-lan)
ent3      Available   Virtual I/O Ethernet Adapter (1-lan)
ent4      Available   Virtual I/O Ethernet Adapter (1-lan)
ent5      Available   Virtual I/O Ethernet Adapter (1-lan)
ent6      Available   Gigabit Ethernet-SX PCI-X Adapter
ent7      Available   Gigabit Ethernet-SX PCI-X Adapter
ent8      Available   Shared Ethernet Adapter
ent9      Available   Virtual I/O Ethernet Adapter (1-lan)
ent10     Available   Shared Ethernet Adapter
```

With the successful creation of the SEA, we can use the **entstat** command on the VIO Server to get additional details of the components of the SEA as shown in Example 4-18.

Example 4-18 entstat command used to provide VLAN details

```
$ entstat -all ent10 |grep VLAN
VLAN Ids :
    VLAN Extract: False
    VLAN tagged filtering mode: Filter according to VLAN permit array
Max number of VLAN IDs per HEA port: 20
Invalid VLAN ID Packets: 0
Port VLAN ID:    555
VLAN Tag IDs:    40    30    20
```

VIO Client LPAR virtual Ethernet mapping

The next step is to map a VIO Client LPAR virtual Ethernet adapter to the VIO Server virtual Ethernet. This procedure is done using IVM. Click the LPAR name from the View/Modify Partitions view to open the Partition Properties window then click the **Ethernet** tab. As shown in Figure 4-28 on page 115 in the Virtual Ethernet Adapter section, select a VIOC (VIO Client) adapter, and from the corresponding pull-down menu, select a VLAN ID. In this example, we used VIOC adapter 2 and VIOS virtual Ethernet corresponding to VLAN 20. When the selection is complete, click **OK**.

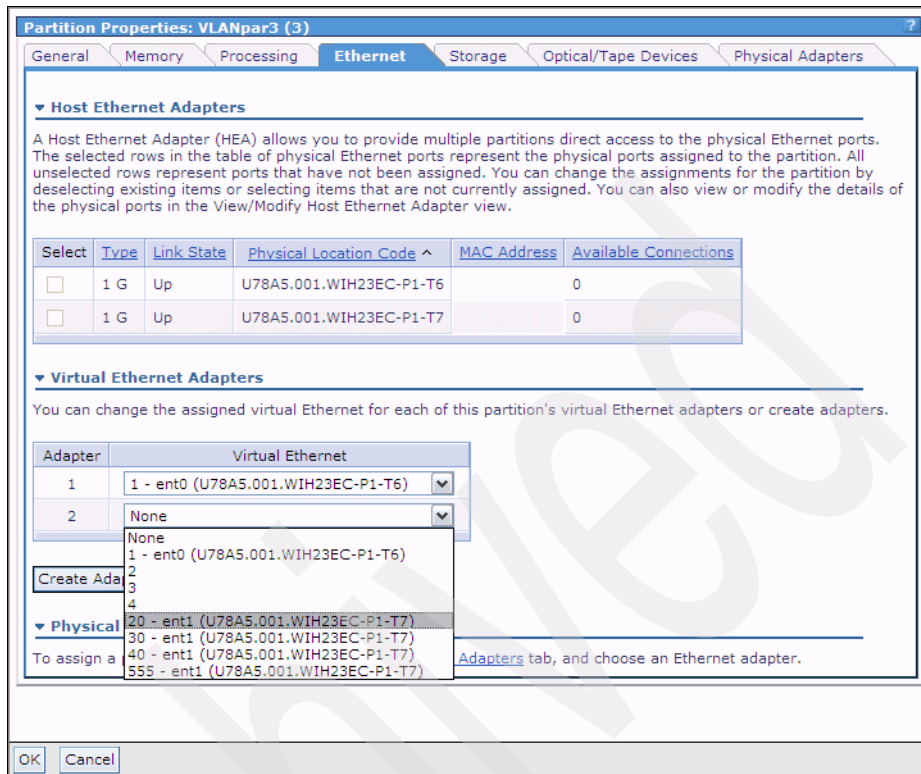


Figure 4-28 VIOC adapter to VIOS virtual Ethernet mapping

VIO Client verification and configuration

If the partition is not active, the new adapter will be discovered upon activation of the LPAR. If the partition is already active, you might have to take additional steps such as run the **cfgmgr** command in AIX. IBM i LPARs with Autoconfig enabled will automatically configure the new adapter.

With the discovery of the new virtual adapter for the LPAR complete, a new virtual Ethernet adapter, ent1, is available in our lab example VIO Client, as shown in Example 4-19.

Example 4-19 lsdev command from VIO Client showing new virtual Ethernet

```
# lsdev |grep ent
ent0      Available      Virtual I/O Ethernet Adapter (1-lan)
ent1      Available      Virtual I/O Ethernet Adapter (1-lan)
```

VLAN details of ent1 can be displayed using the **entstat** command on the VIO Client (assuming an AIX client) as shown in Example 4-20.

Example 4-20 entstat command from VIO Client showing details of new virtual Ethernet

```
# entstat -d ent1 |grep VLAN
Invalid VLAN ID Packets: 0
Port VLAN ID:      20
VLAN Tag IDs:      None
```

In this AIX LPAR example, the interface ent1 on VLAN 20 can now be configured with the desired TCP/IP properties

4.6 VIOS Storage Configuration and Management

Virtual Storage Management allows you to manage the disk storage requirements for logical partitions. The following terms are used to describe storage management concepts:

- ▶ Physical volumes
- ▶ Storage pools
- ▶ Virtual disks
- ▶ Optical and tape devices

These terms are explained in the next sections.

To work with VIOS storage, click **View/Modify Virtual Storage** in the navigation area of the IVM as shown in Figure 4-29.

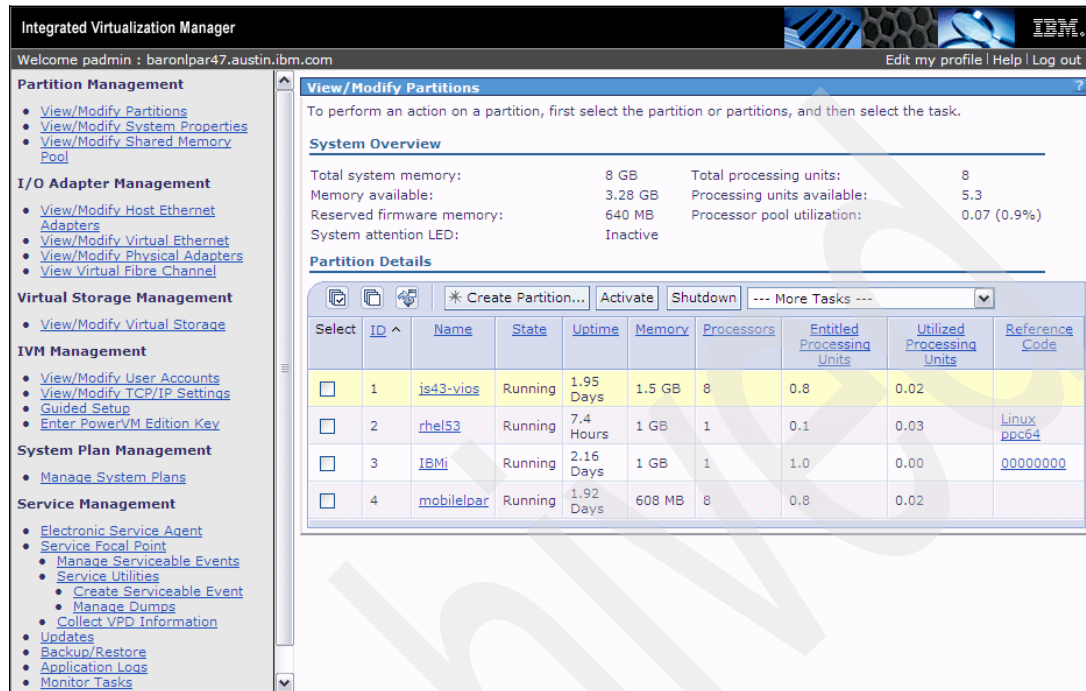


Figure 4-29 View and modify virtual storage

4.6.1 Physical volumes

Physical volumes are the hard drives that are available to the VIOS. They can be installed locally in the IBM BladeCenter JS23 or JS43 blades, SAS drives available from IBM BladeCenter S chassis, or LUNs available from a Fibre Channel storage area network subsystem.

A physical volume is shown as hdisk0, hdisk1 and so on. The Virtual I/O Server LPARs can be assigned complete physical volumes. However, they appear as a virtual SCSI disk drive on the LPAR. This direct assignment is a requirement if you are planning shared memory partitions or using IBM i. Physical volumes can also be used to build storage pools or AMS dedicated paging devices.

To verify the available physical volumes in your VIOS, in the navigation area, click **View/Modify Virtual Storage**. Then click the **Physical Volumes** tab, as shown in Figure 4-30 on page 118. The list of the physical volumes available to the VIOS is displayed.

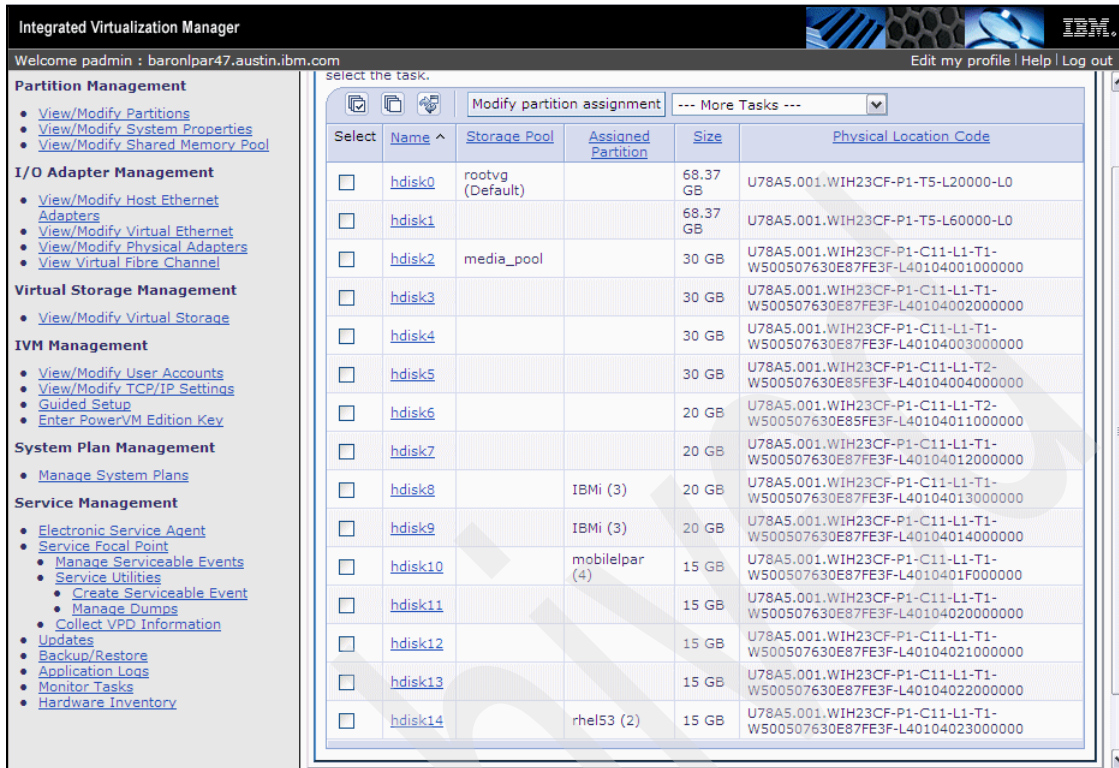


Figure 4-30 Physical volumes shown in IVM

Similar information can be retrieved on the Virtual I/O Server CLI by using the **lsdev** and **lspv** commands. Example 4-21 shows the output of the **lsdev -type disk** command.

Example 4-21 Physical volumes found with *lsdev*

```
$ lsdev -type disk
name          status      description
hdisk0        Available   SAS Disk Drive
hdisk1        Available   SAS Disk Drive
hdisk2        Available   IBM MPIIO FC 1750
hdisk3        Available   IBM MPIIO FC 1750
hdisk4        Available   IBM MPIIO FC 1750
hdisk5        Available   IBM MPIIO FC 1750
hdisk6        Available   IBM MPIIO FC 1750
hdisk7        Available   IBM MPIIO FC 1750
hdisk8        Available   IBM MPIIO FC 1750
hdisk9        Available   IBM MPIIO FC 1750
```

hdisk10	Available	IBM MPIO FC 1750
hdisk11	Available	IBM MPIO FC 1750
hdisk12	Available	IBM MPIO FC 1750
hdisk13	Available	IBM MPIO FC 1750
hdisk14	Available	IBM MPIO FC 1750

Example 4-22 shows the output of the **lspv -size** command.

Example 4-22 Physical volumes found with lspv -size

```
$ lspv -size
```

NAME	PVID	SIZE(megabytes)
hdisk0	000181ca0005e5c6	70006
hdisk1	000181ca6309a681	70006
hdisk2	000181ca7d20d77c	30720
hdisk3	none	30720
hdisk4	none	30720
hdisk5	none	30720
hdisk6	none	20480
hdisk7	none	20480
hdisk8	none	20480
hdisk9	none	20480
hdisk10	000180ea884b6253	15360
hdisk11	000180ea884b8500	15360
hdisk12	000180ea884b98b4	15360
hdisk13	000180ea884baacf	15360
hdisk14	000180ea884bbe1e	15360

4.6.2 Storage pools

A *storage pool* is a single entity that consists of one or more physical volumes or files. Logical volumes or virtual disks are created within a storage pool. Physical volumes or file spaces can only be assigned to a single storage pool.

After installation of the VIOS, a default storage pool is created from space in the volume group rootvg. A new default storage pool should be created if you plan to use storage pools. As a general best practice, the storage pool in volume group rootvg should not be used.

Creating a new storage pool

To create a new storage pool, select the **Storage Pools** tab in the View/Modify Virtual Storage window. Figure 4-31 on page 120 shows a list of all available storage pools.

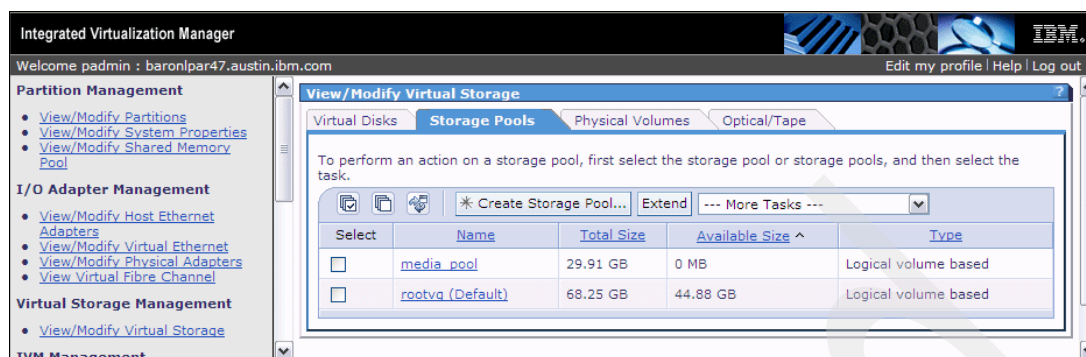


Figure 4-31 Storage pools shown in IVM

Click **Create Storage Pool** to create a new storage pool. A dialog opens that guides you through the setup of the storage pool.

Specify a name (for example, SP-Media-Lib) to use for the storage pool. The name used for the storage pool must be a valid name for volume groups, for example no spaces are allowed and the name cannot exceed 15 characters.

Specify the storage pool type as Logical Volume-based. The File-based option currently requires a local file system. Select one or more available physical volumes to be used for the new storage pool, then click **OK**.

Figure 4-32 shows that, in this case, **hdisk3** was selected.

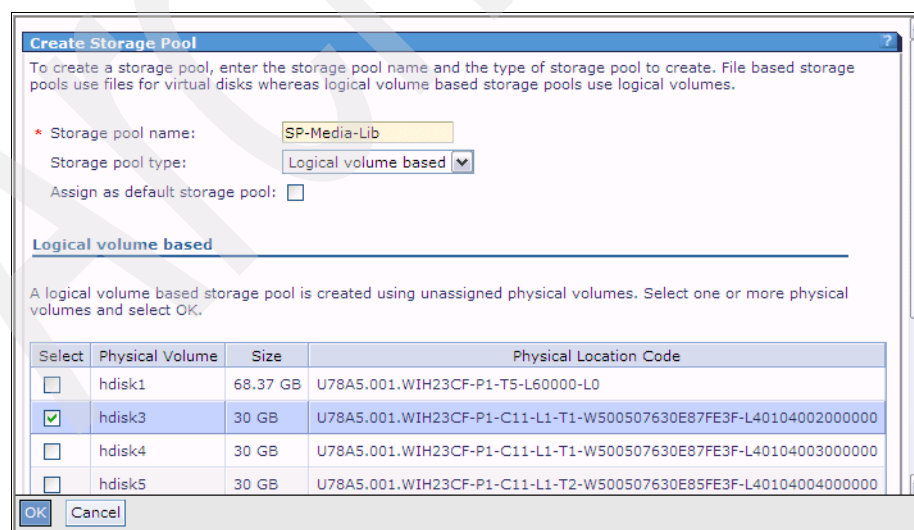


Figure 4-32 Create new storage pool

Figure 4-33 shows the new storage pool.

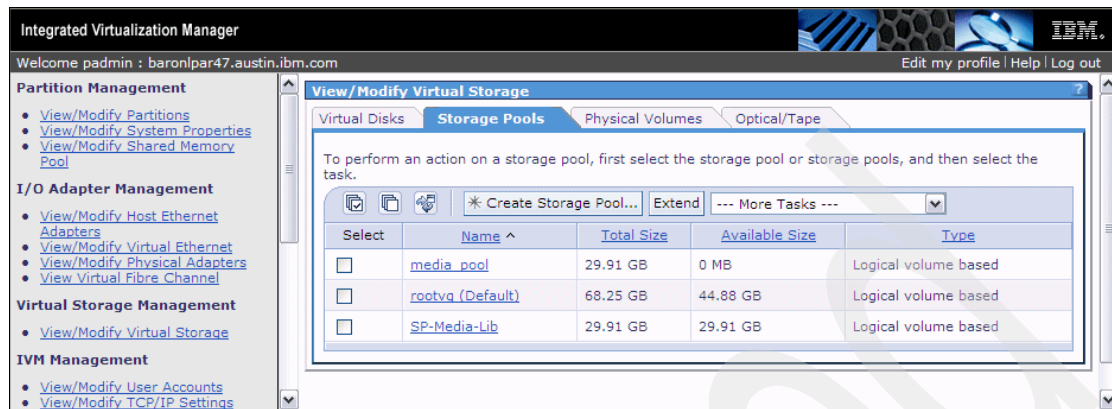


Figure 4-33 Newly created storage pool shown in IVM

Deleting or reducing a storage pool

To delete or reduce a storage pool, start from the Storage Pool tab in the Modify Virtual Storage window.

Select the storage pool you want to delete or reduce. Click **Reduce** from the More Tasks drop-down box as shown in Figure 4-34 on page 121. A dialog opens that guides you through the modification of the storage pool.




Figure 4-34 Reduce or delete a storage pool

Select the physical volumes that you want to remove from the storage pool. The storage pool will be deleted when all physical volumes that are assigned to the storage pool are removed. Click **OK**, as shown in Figure 4-35.




Figure 4-35 Delete storage pool

4.6.3 Virtual disks

Virtual disks are created in storage pools. After they are assigned to a logical partition, they are seen as virtual SCSI disk drives by the LPAR. These assignments are represented in the LPAR as hdisks. Multiple virtual disks can be created in a single storage pool. However, this method does not support Live Partition Mobility and is not recommended for IBM i.

You can create virtual disks from the View/Modify Virtual Storage window by selecting the **Virtual Disks** tab, as described in the following section. The Create Partition Wizard, as described in 4.7.2, “Partition name and environment” on page 140, can also be used to create virtual disks. Both methods require free space in a storage pool.

Creating virtual disks

To create a logical volume, a storage pool must be available. Refer to 4.6.2, “Storage pools” on page 119, for information about how to create a storage pool.

To create a new virtual disk, start with the View/Modify Virtual Storage window and select the **Virtual Disks** tab. From this window, click **Create Virtual Disk**, as shown in Figure 4-36 on page 123.

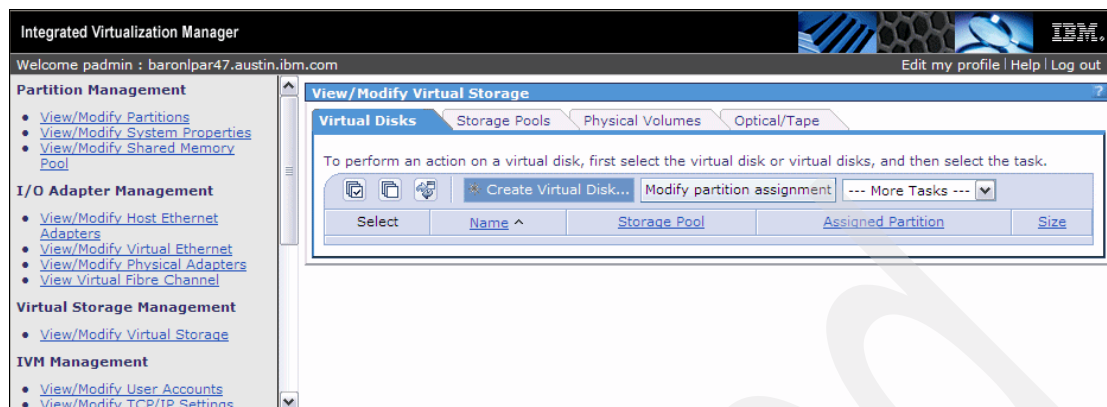


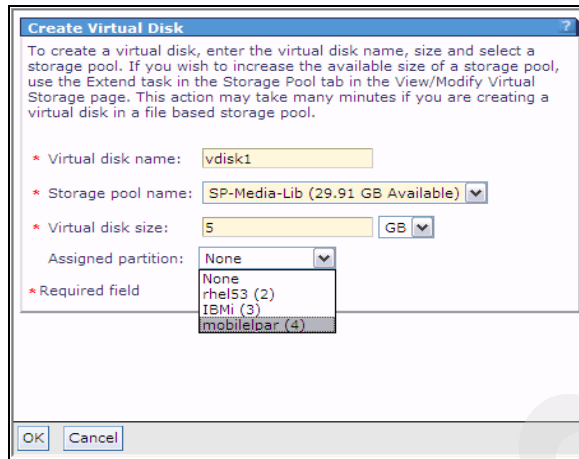
Figure 4-36 Create Virtual Disk

From the Create Virtual Disk window, specify the name of the virtual disk, select a storage pool that will be used from the drop-down box, and specify the virtual disk size.

Optionally, you can make a partition assignment during virtual disk creation as shown in Figure 4-37 on page 124. A new virtual disk will be created when you click **OK**.

Note: When a virtual disk is created during the Create Partition wizard, the default naming schema for virtual disks uses the partition ID and the number of the assigned virtual disk to the LPAR. The result looks like:

`lp{number}vd{number}`



Create Virtual Disk

To create a virtual disk, enter the virtual disk name, size and select a storage pool. If you wish to increase the available size of a storage pool, use the Extend task in the Storage Pool tab in the View/Modify Virtual Storage page. This action may take many minutes if you are creating a virtual disk in a file based storage pool.

* Virtual disk name:

* Storage pool name:

* Virtual disk size: GB

Assigned partition:

* Required field

- None
- rhel53 (2)
- IBMi (3)
- mobilelpar (4)

OK Cancel

Figure 4-37 Virtual disk settings

The newly created virtual disk appears in the list, as shown in Figure 4-38.

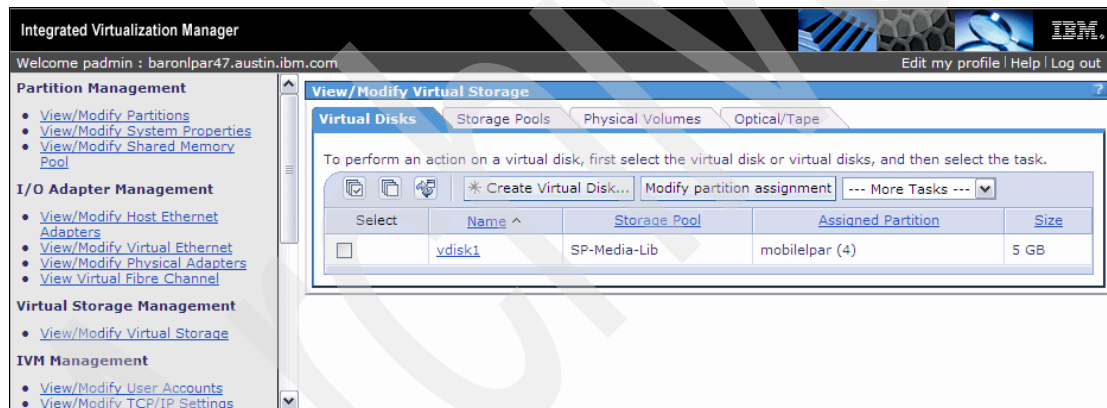


Figure 4-38 The newly created virtual disk

The size of the virtual disk can be extended, as described in the following section.

Extending a virtual disk

You can extend a virtual disk if enough free space is available in the storage pool. To extend a virtual disk, select the virtual disk you plan to extend in the check box. From the More Tasks drop-down menu, select **Extend**, as shown in Figure 4-39 on page 125.

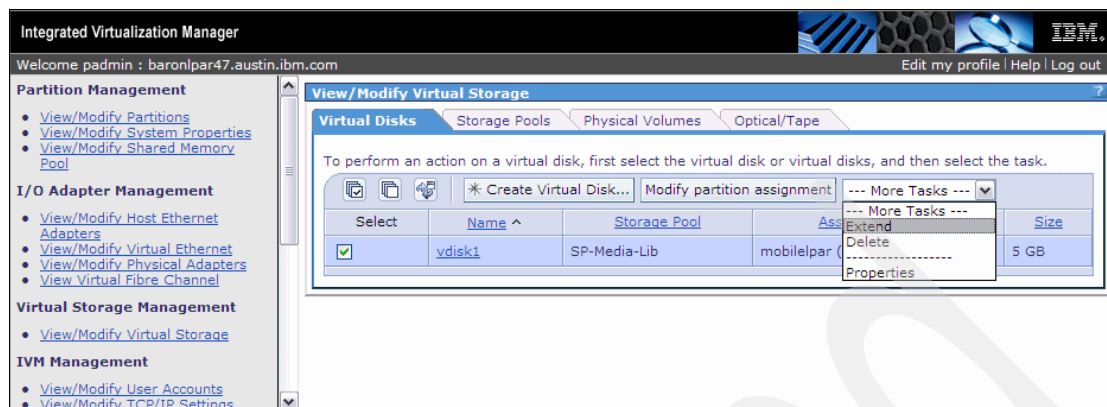


Figure 4-39 Extend virtual disk

Specify the amount of space that the virtual disk will be extended, then click **OK** as shown in Figure 4-40. If the storage pool does not have enough free space, it can be extended from the Storage Pools tab.

Note: When you attempt to extend virtual disk on a running partition, a warning message is generated, alerting the administrator. To continue, select the **Force extend on running partition** check box and click **OK** again.

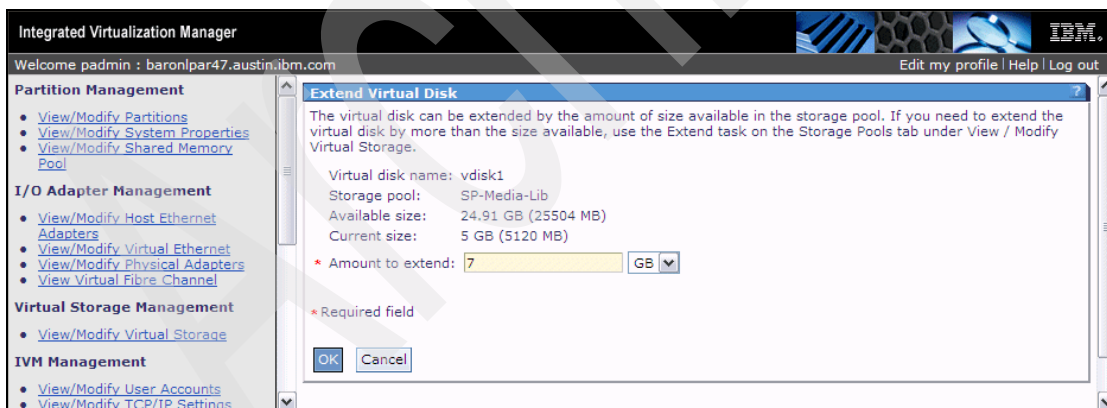


Figure 4-40 Virtual disk extension settings

The new size is shown in the list of available virtual disks when the extension is complete, as shown in Figure 4-41 on page 126.

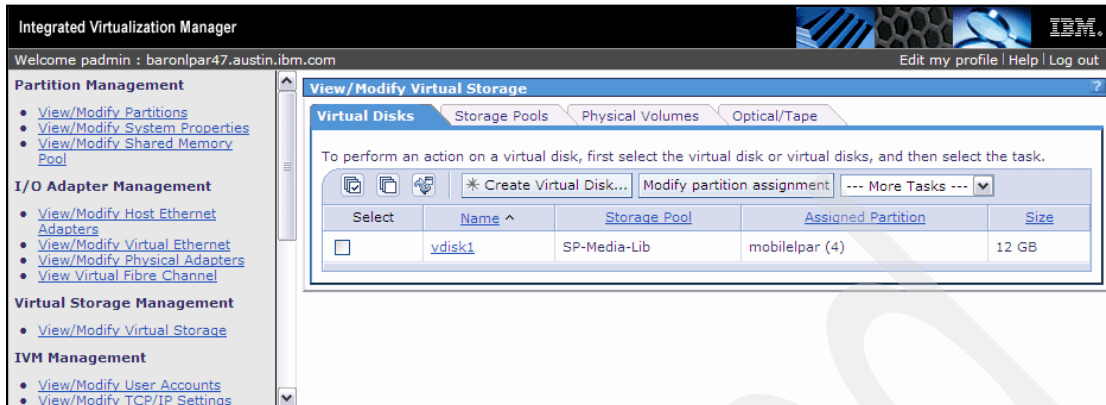


Figure 4-41 Extended virtual disk

Deleting virtual disks

A virtual disk that is assigned to a partition must have that assignment removed before the virtual disk can be deleted.

Note: When you try to delete a virtual disk on a running partition, a warning message is generated, alerting the administrator. To continue, select the **Force device removal from a running partition** check box and click **OK** again.

To delete a virtual disk, click the **Virtual Disks** tab in the View/Modify Virtual Storage window. Select the virtual disk that you want to delete, and then select **Delete** from the More Tasks drop-down menu, as shown in Figure 4-42.

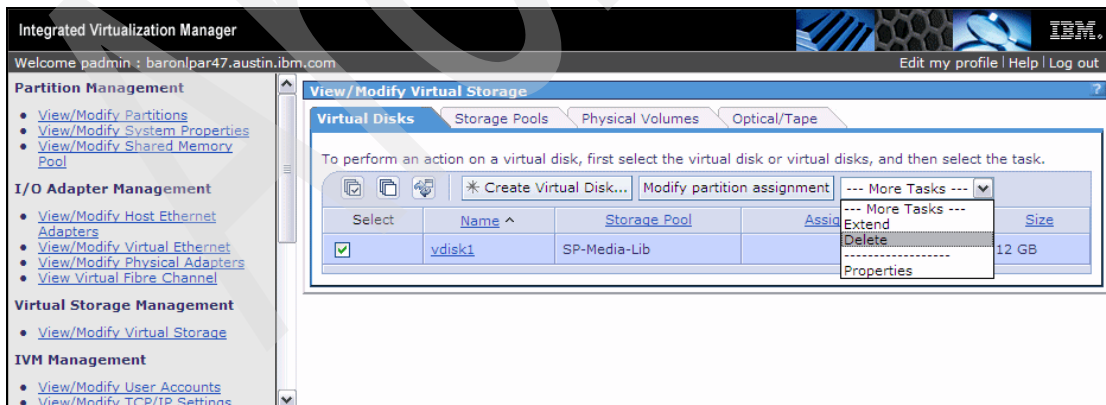


Figure 4-42 Delete virtual disk

Confirm the deletion of the virtual disk by clicking **OK**, as shown in Figure 4-43.

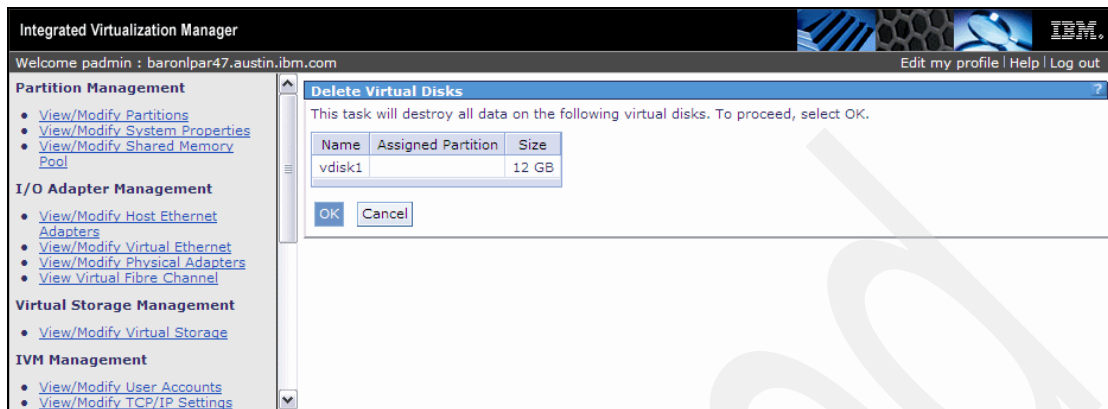


Figure 4-43 Confirm deletion of the virtual disk

The virtual disk will be deleted and the occupied space in the storage pool will become available.

4.6.4 Optical and tape devices

Optical devices are CD or DVD drives. The two types of optical devices are:

- ▶ Physical optical devices
 - Local in BladeCenter media tray
 - Remote media on local desktop or mobile computer
- ▶ Virtual optical devices

Physical tape devices must be Serial Attached SCSI (SAS).

Physical optical devices

Physical optical devices are the CD or DVD drives installed in the media tray of an IBM BladeCenter. Each type of BladeCenter chassis is delivered with a CD drive or a DVD- drive.

The other physical optical device that can be used is remote media. An ISO image or a CD or DVD in your mobile computer or desktop can be assigned to the blade. The Web interface of the Advanced Management Module provides this capability.

As Table 4-1 on page 128 shows, the two different optical drive types can be identified by their location paths.

Table 4-1 Optical drive location paths

Location path	Description
U78A5.001.WIH01AA-P1-T1-L1-L2-L3	CD or DVD drive in the media tray
U78A5.001.WIH01AA-P1-T1-L1-L1	Remote media

The name of the optical drives can vary, depending on the kind of drive or remote media you are using.

Before the BladeCenter physical optical device can be used, the media tray must be assigned to the blade slot you are working with. The physical optical device cannot be shared between LPARs or the VIOS and the LPARs.

“Changing the assignment of physical optical drives” on page 128 describes how to use the Storage Management to change the assignment of physical optical devices.

Changing the assignment of physical optical drives

The assignment of the physical optical device can be changed at any time. The assignment can be made or changed from the Optical Devices tab in the View/Modify Virtual Storage window. This section describes how to use the Storage Management to change the assignment of physical optical devices.

To change the assignment of physical optical drives, click the **Optical/Tape** tab. Figure 4-44 on page 129 shows a list of available physical devices. The table in that figure indicates the physical optical device and shows the LPAR assignment. The example shows that cd1 is assigned to the partition named JS231par2.

The list of physical devices can vary, depending on the media tray assignment and the usage of remote media. Figure 4-44 on page 129 shows two physical CD-ROM drives. The second drive is a remote CD-ROM drive that is mounted over the remote control interface of the Advanced Management Module (AMM) in the BladeCenter chassis.

Note: The remote control function for the IBM BladeCenter JS23 or JS43 is only available to the blade slot that has the media tray assignment.

To change the assignment of a physical optical device, select the check box of the device to be changed and click **Modify partition assignment**. A dialog opens that guides you through the assignment change.

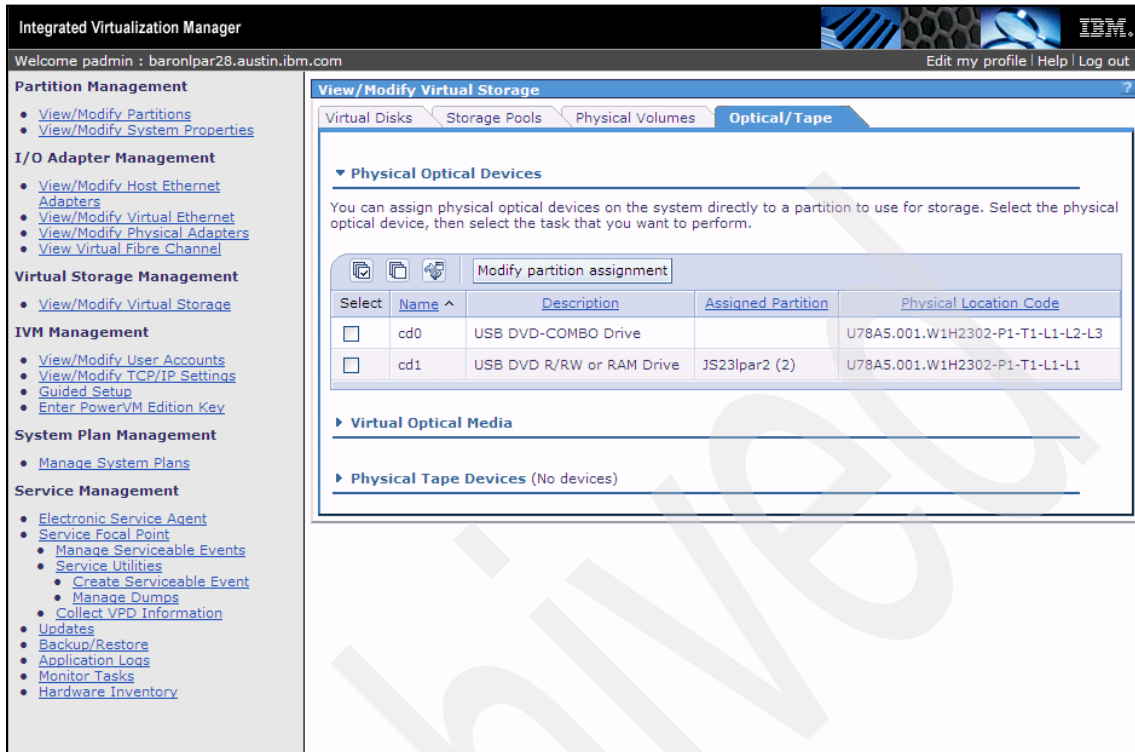


Figure 4-44 Physical optical and tape devices in IVM

To make the physical optical device available to the VIOS itself, select **None**. Otherwise, select the logical partition that the physical optical device will be assigned and then click **OK**, as shown in Figure 4-45 on page 130.

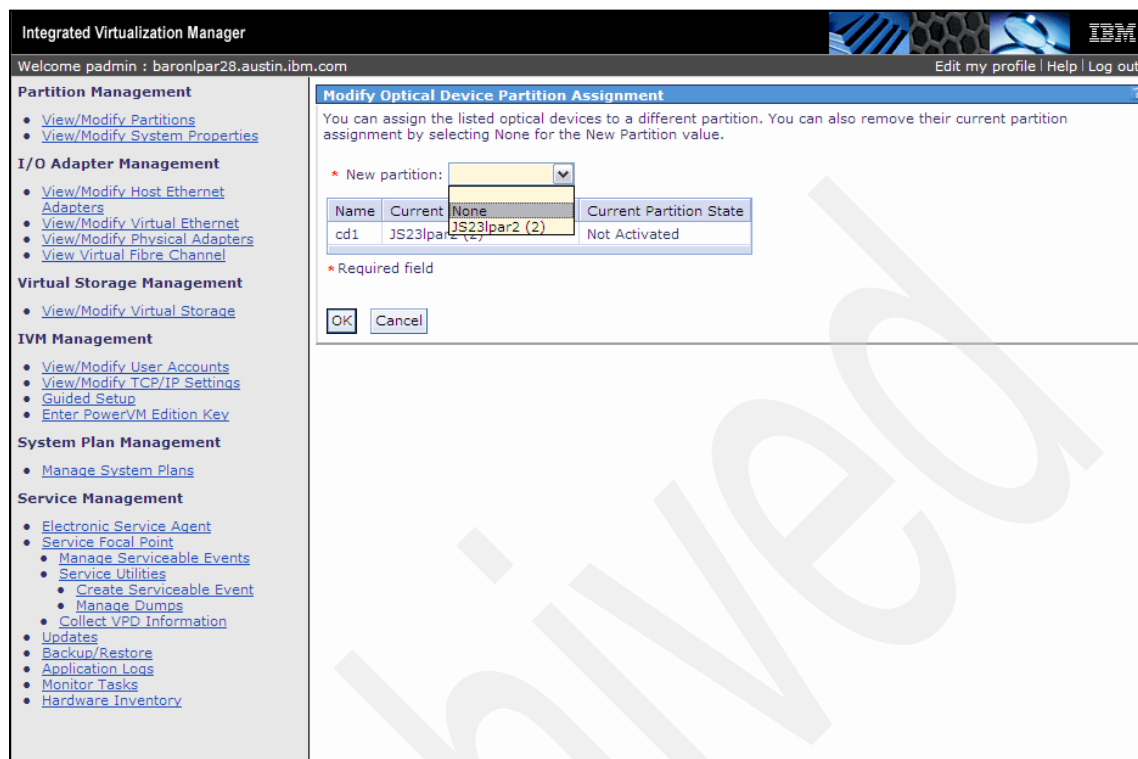


Figure 4-45 Change physical optical device assignment

Virtual optical devices

Virtual optical devices were introduced with Virtual I/O Server V1.5. Together with the Media Library of a Virtual I/O Server, this device is able to virtualize CD or DVD images that are stored in the VIOS media library to one or more logical partitions. Before virtual optical device can be used, you must configure a media library.

Creating a media library

To set up a media library:

1. Ensure a storage pool is available. Refer to 4.6.2, “Storage pools” on page 119, for an explanation of how to set up a storage pool.
2. Click the **Optical/Tape** tab in the View/Modify Virtual Storage window to create a media library. Click **Virtual Optical Media** section to expand it. Then, click **Create Library**, as shown in Figure 4-46 on page 131.

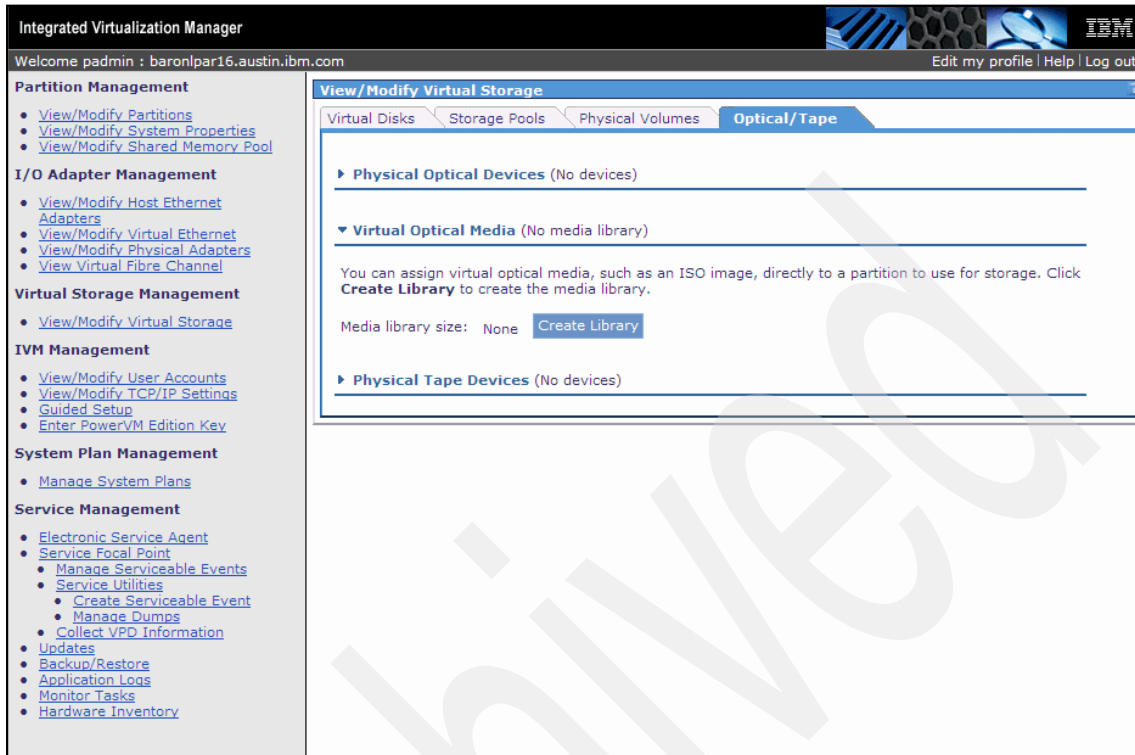


Figure 4-46 Create media library

3. In the Create Media Library panel, shown in Figure 4-47 on page 132, select an available storage pool and the amount of storage space that will be used from this storage pool to create the media library, and then click **OK**.

Important: In production environments, picking rootvg as the storage pool for the media library is not a good practice because reinstalling VIOS can remove all data on rootvg and the library will be lost.

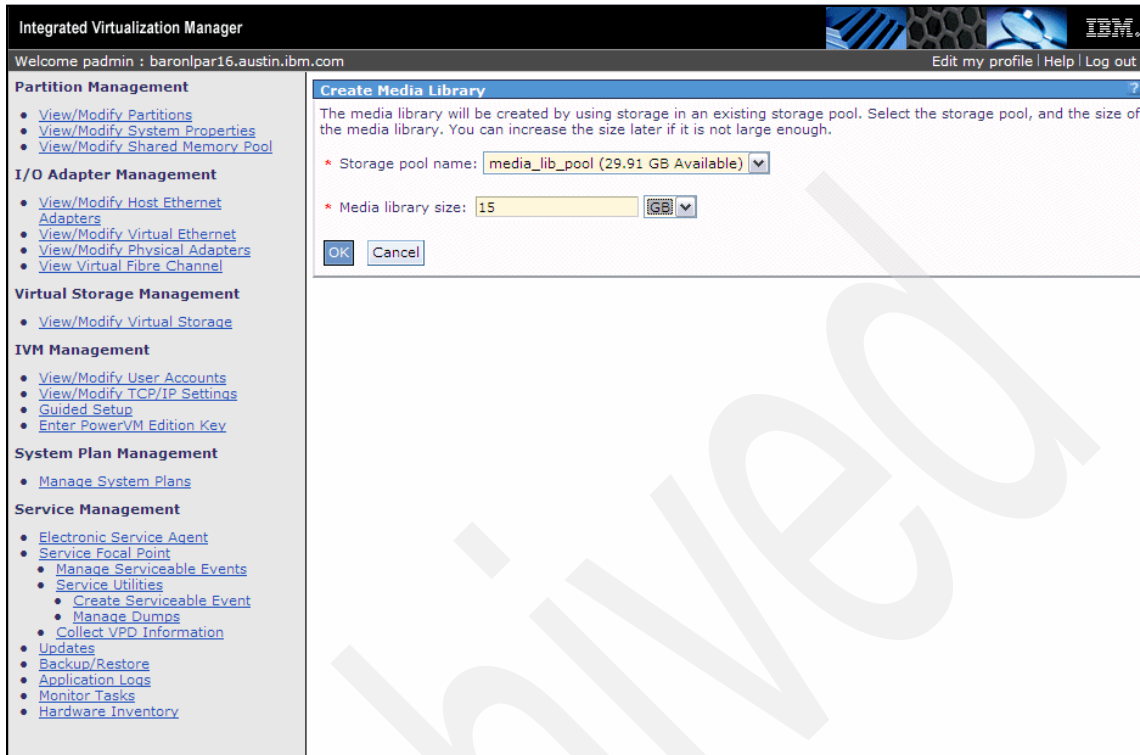


Figure 4-47 Media library size and storage pool

Depending on the size of the media library, the creation time can vary. After the media library is successfully created, the current view in the View/Modify Virtual Storage window will change, showing Media Library options.

The size of media library can be increased at any time by clicking the **Extend Library** button. Selecting the **Delete Library** button allows you to delete the complete media library and all added media and return the space to the storage pool.

Next, you have to add the new media into the media library.

Adding new media into the media library

New media can be added to the media library and later assigned to the virtual optical devices. The new media can consist of image files such as ISO images or copies from physical CDs or DVDs. In addition, you can create blank media that can be written to an LPAR using the virtual optical device.

To add new media in the media library, click **Add Media** as shown in Figure 4-48.

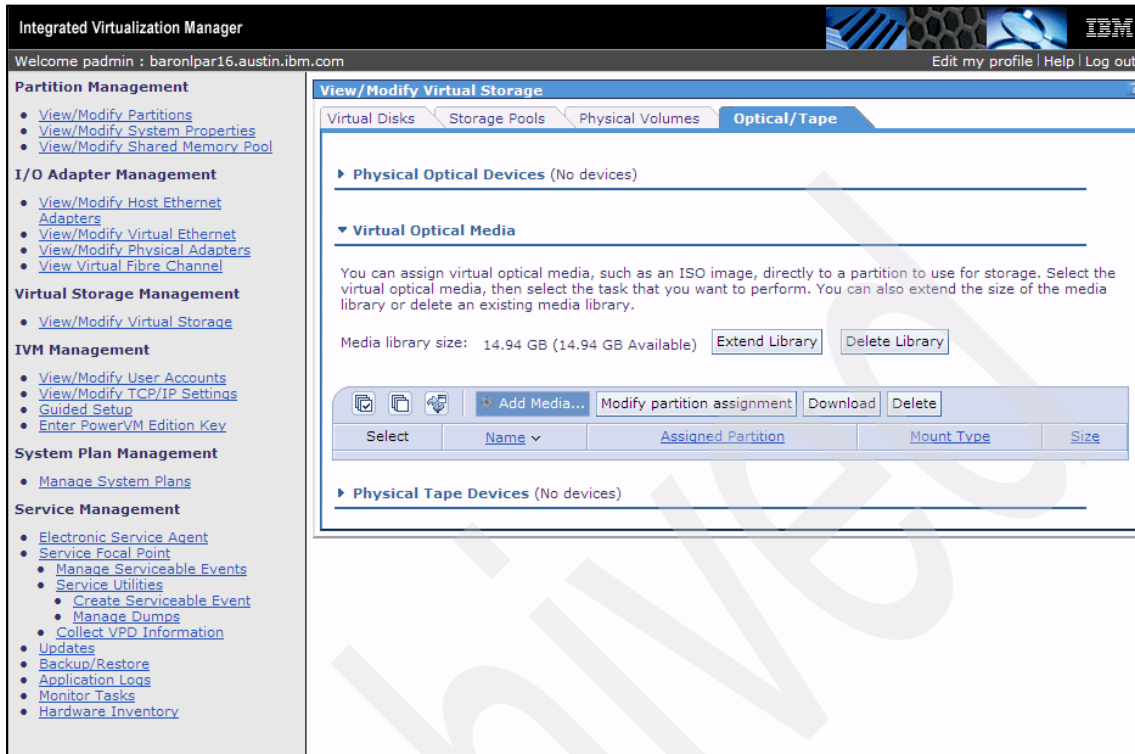


Figure 4-48 Add media to media library

Four options are available for creating new media:

- ▶ Upload media
- ▶ Add existing file
- ▶ Import from physical optical device
- ▶ Create blank media

The Upload media option enables you to transfer files or ISO images from a workstation directly to the media library. File size is limited to 2 GB for this option.

The Add existing file option adds an existing file that is available in a VIOS file system as new media.

The Import from physical optical device option allows you to use a the physical CD or DVD. The data will be copied from the CD or DVD into the media library.

Note: Our testing revealed that the local CD or DVD drive in the media tray of the BladeCenter chassis is a faster option compared to the remote media option with a physical CD or DVD drive.

The Create blank media option allows you to create blank media that may be written to from an LPAR.

Figure 4-49 shows an example that uses Import from physical optical device to create the new media. Click **OK** to start the copy task.

Note: Do not use spaces in the name of the new media. If you use spaces in the name, IVM returns the following error message:

Specify a valid media name. The media name cannot contain blank spaces, be greater than 30 characters in length or begin with a period.

Add Media

You may upload an optical media file from your local workstation, specify an existing file in your home directory, import from a physical optical device, or create a blank media file. Depending on the size of the file, these operations may take several minutes.

☐ Upload media
☐ Add existing file
☒ Import from physical optical device
☐ Create blank media

Media type: Read only ▼

Media name: AIX6.1_install_disk_1

Select an available optical device that contains the media you wish to import, and select OK. This may take some time while the media is copied into the library.

Select	Name	Description	Physical Location Code
<input type="checkbox"/>	cd0	USB DVD-COMBO Drive	U78A5.001.WIH23EC-P1-T1-L1-L2-L3
<input checked="" type="checkbox"/>	cd1	USB CD-ROM Drive	U78A5.001.WIH23EC-P1-T1-L1-L1

OK Cancel

Figure 4-49 Add new media - settings

The copy task takes some time to complete. While the copy task is running, you may proceed with other configuration tasks after you see the message Performing Task - Please Wait displayed, as shown in Figure 4-50 on page 135.




Figure 4-50 Performing task

Click the **Monitor Task** link from the Navigation area to verify the completion of the task. Monitor Tasks contains a list of events and the status, either running, successful, or failed.

Note: An alternative way to monitor the process of creating new media is to review the list under the Optical Devices tab, as shown in Figure 4-51 on page 136.

If your new media is not listed here, click the **Refresh** button. During the copy operation, the new media is shown as Read/Write and the size will increase on a refresh. After the copy operation is finished, the mount type changes to Read only.

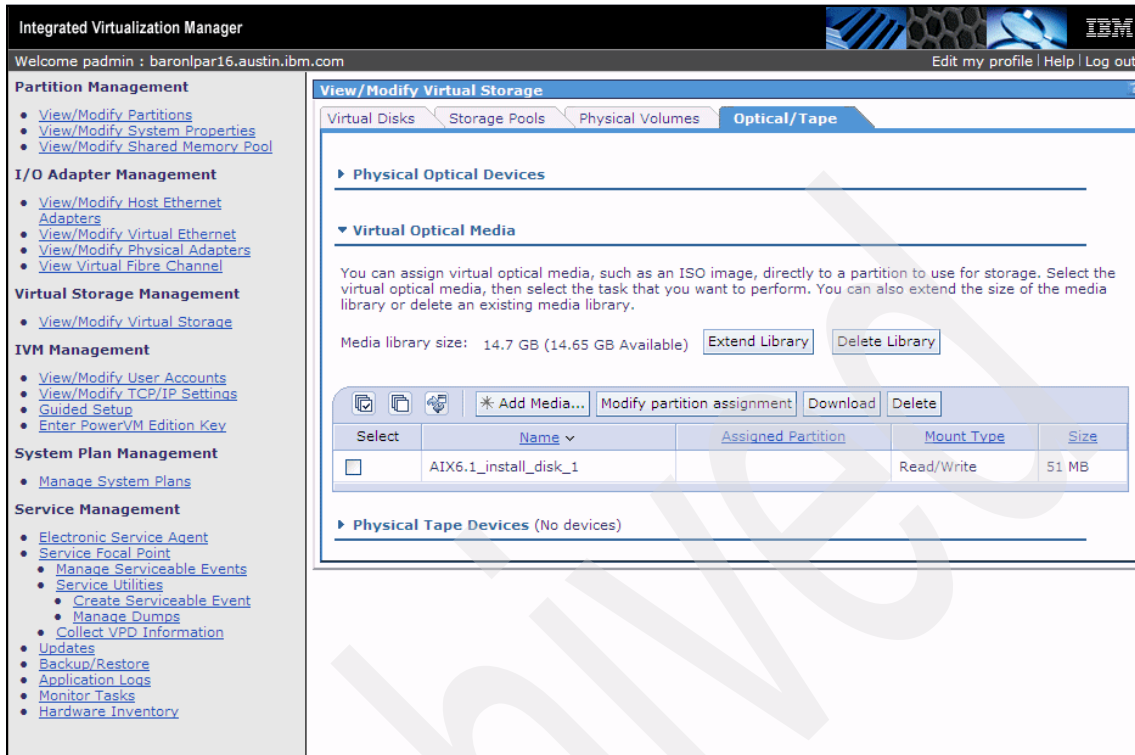


Figure 4-51 Newly created media with the copy operation in progress

Modifying media assignment to virtual optical devices in logical partitions

Media can be assigned from the Optical/Tape tab in the View/Modify Virtual Storage window, when using the Create Partition wizard or from the Partition Properties window. The next step is to modify the partition assignment of the media in the media library.

Note: The logical partition *must* have a virtual optical device assigned prior to assigning or modifying the partition assignment of the selected media.

To modify the virtual media assignment, select the media name and then click **Modify partition assignment** under the Optical/Tape tab from the View/Modify Virtual Storage window.

As shown in Figure 4-52 on page 137, the media AIX6.1_install_disk_1 is not assigned to any LPAR. This media, because it is read-only, is selected and

assigned to two LPARs. Select the check box for the media and click the **Modify partition assignment** button.

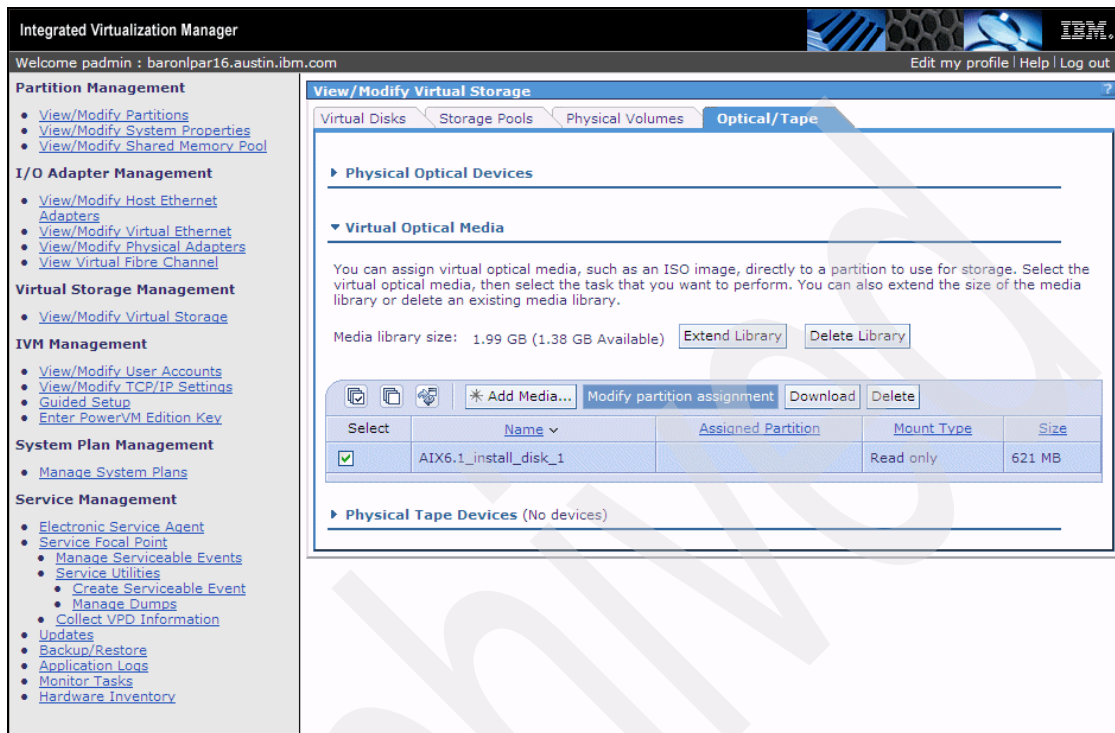


Figure 4-52 Modify partition assignment

As shown in Figure 4-53 on page 138, no LPARs are assigned to the media AIX6.1_install_disk_1.

Next, LPARs JS23DM1par4 and JS23 DP1par5 will be assigned the same media by selecting the check box next to the logical partitions. Choose the Media type **Read only** or **Read/Write** and click **OK**. Only read-only media can be assigned to more than one LPAR.

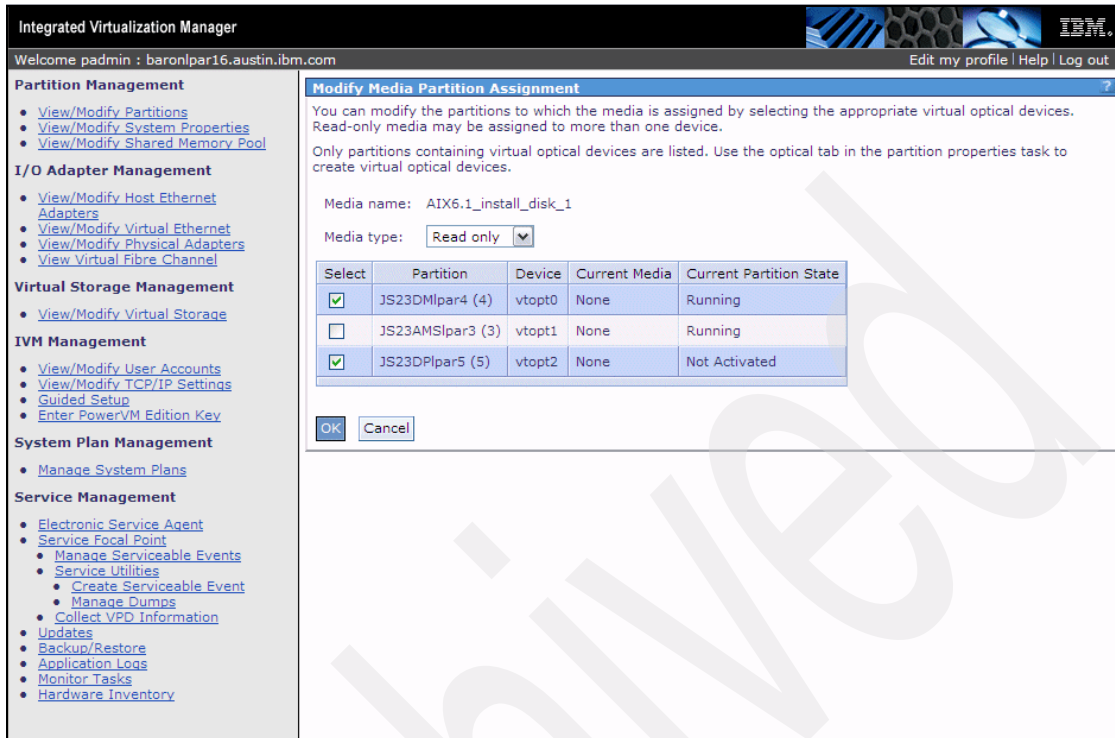


Figure 4-53 Modify media partition assignment

Click **OK** to return to the view of the optical devices.

Notice that the updated table, shown in Figure 4-54 on page 139, now contains the LPARs JS23DMLpar4 and JS23 DPLpar5 in the Assigned Partition column as assigned partitions for the media AIX6.1_install_disk_1.

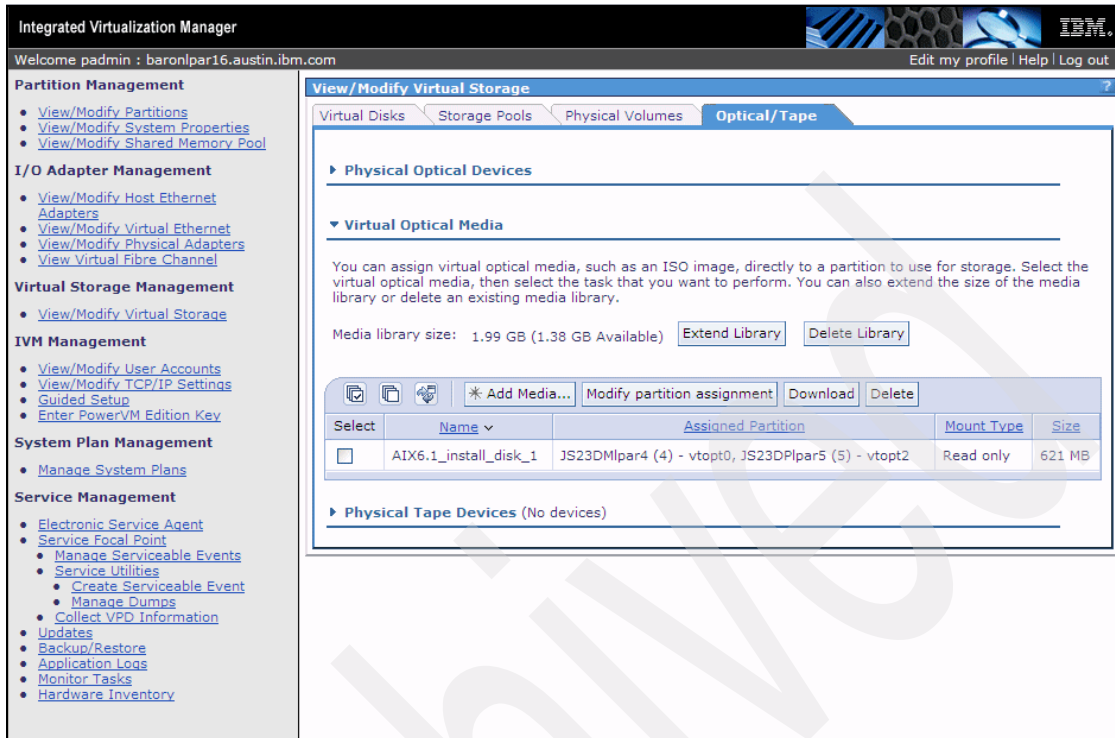


Figure 4-54 New assigned media to partitions

A media can be removed from a partition following the same procedure by deselecting the media that is assigned to the partition.

4.7 Partition configuration for Virtual I/O Client

With networking and storage defined, you can now create additional Virtual I/O Client (VIOC) LPARs for the installation of additional supported operating systems.

4.7.1 Live Partition Mobility considerations

If Live Partition Mobility (LPM) is being considered for a VIOC in logical partitions, you should review the planning considerations (currently unavailable on IBM i) listed in this section.

The planning considerations include:

- ▶ VIOS running on source and target IBM BladeCenter JS23 or JS43:
 - VIOS should be at the latest fix pack.
 - IBM BladeCenter JS23 or JS43 should be at the latest system firmware.
- ▶ All I/O must be virtual to the LPAR:
 - SEA adapters are required. No HEA logical ports can be assigned.
 - No virtual optical drives can be assigned.
 - No physical adapters can be assigned.
- ▶ SAN storage must be properly configured for sharing between the two Virtual I/O Servers.
- ▶ Processor compatibility modes must be between source and target systems.
- ▶ Memory region sizes must match between source and target systems.
- ▶ If Active Memory Sharing (AMS) is being used on the source VIOS/VIOC, it must be available on the target VIOS.
- ▶ Only IVM- to-IVM managed systems are allowed (no HMC-to-IVM or IVM-to-HMC is allowed)

For more information, refer to:

- ▶ Chapter 11, “Performing Live Partition Mobility” on page 427 provides setup and configuration details for Partition Mobility with a JS23 or JS43.
- ▶ *PowerVM Live Partition Mobility*, SG24-7460 provides an overview of Partition Mobility and explains how it applies to other System p configurations.

4.7.2 Partition name and environment

To start the create partition wizard, select **View/Modify Partitions** from the navigation area of IVM interface, shown in Figure 4-55 on page 141. Then, in the View/Modify Partitions panel, click the **Create Partition** button.

Integrated Virtualization Manager

Welcome padmin : baronlpar47.austin.ibm.com

[Edit my profile](#) | [Help](#) | [Log out](#)

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](#)

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:

8 GB

Total processing units:

8

Memory available:

2.22 GB

Processing units available:

5.3

Reserved firmware memory:

704 MB

Processor pool utilization:

0.16 (2.0%)

Available shared memory pool size:

768 MB

System attention LED:

Inactive

Partition Details

Create Partition...

Activate

Shutdown

--- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is43-vios	Running	4.78 Days	1.5 GB	8	0.8	0.12	
<input type="checkbox"/>	2	rhel53	Running	3.14 Days	1 GB	1	0.1	0.03	Linux ppc64
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1	1.0		00000000
<input type="checkbox"/>	4	mobilelpar	Running	4.75 Days	608 MB	8	0.8	0.02	

Figure 4-55 View/Modify Partition

4.7.3 Partition name

When the wizard starts, a new window opens as shown in Figure 4-56 on page 142. This window gives you the opportunity to change the Partition ID number, provide a Partition name, and select an operating system environment. Click **Next** for the memory step.

Create Partition: Name Step 1 of 8

Name

To create a partition complete the following information.

System name: Server-7778-63X-SN10181CA

Partition ID:

* Partition name:

Environment:

* Required field

< Back Next > Finish Cancel Help

Figure 4-56 Create Partition: Name

4.7.4 Partition memory

Figure 4-57 on page 143 shows how to assign memory to the partition. The two memory options are dedicated and shared. In this section, we discuss only dedicated memory. Shared memory is discussed in Chapter 5, “Active Memory Sharing configuration using IVM” on page 173.

Total system memory and the current memory available for a new partition is summarized under Memory Mode selection section. The amount entered in the box determines the initial *Assigned* and *Maximum* values in the LPAR partition properties. After you enter the desired amount of memory, click **Next**.

Note: IVM does not allow you to over-commit *dedicated* memory resources.

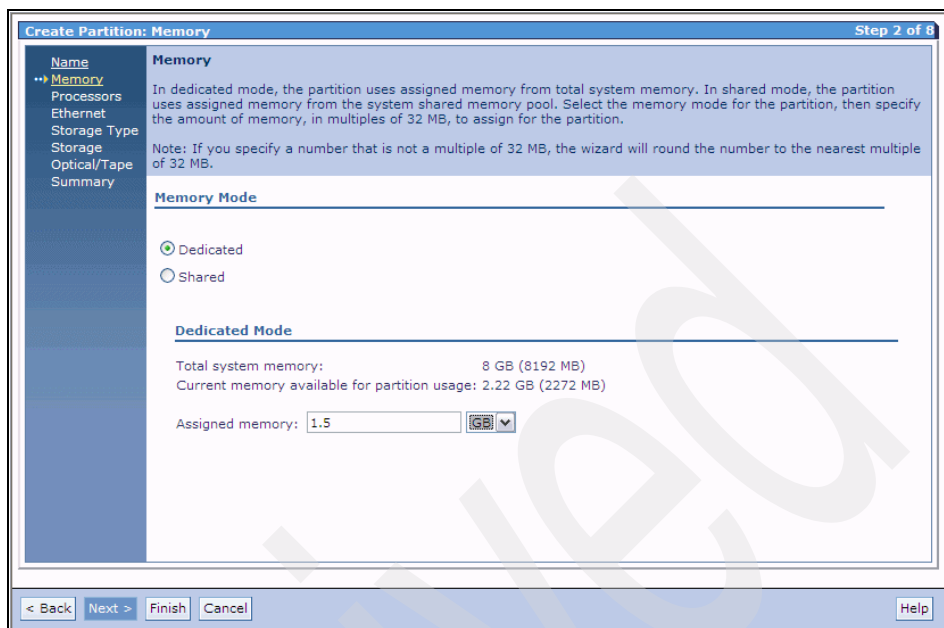


Figure 4-57 Create Partitions: Memory

4.7.5 Partition processors

On the Create Partition: Processors window you have the option of assigning dedicated or shared processors. In shared mode, for each virtual processor, 0.1 processing units are assigned. In dedicated mode, each assigned processor uses one physical processor.

Available processor resources are displayed on the window and, as with dedicated memory resources, they cannot be over-committed. Figure 4-58 on page 144 shows a selection of shared mode and eight assigned processors for this example. After you make your selections, click **Next**.

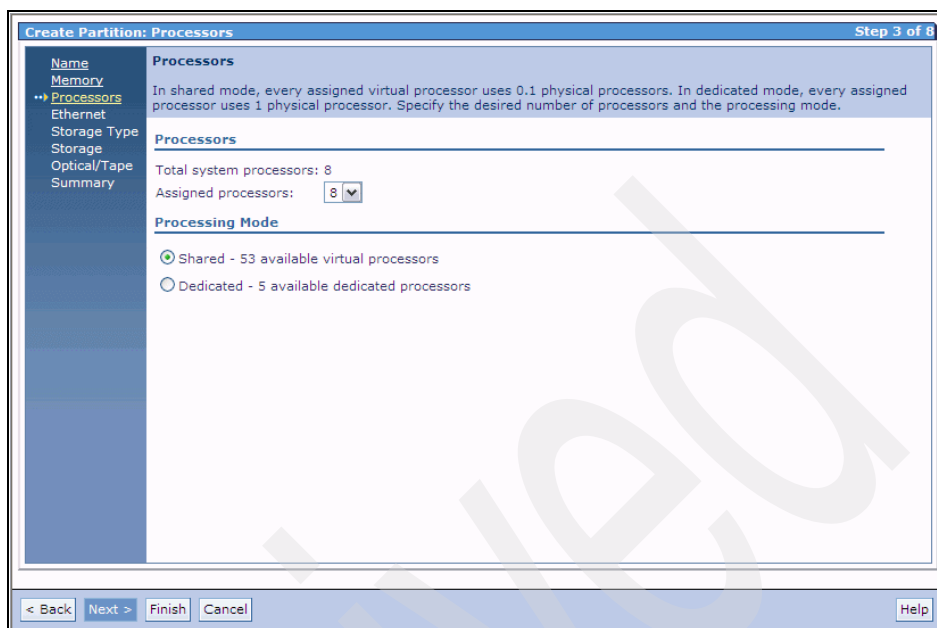


Figure 4-58 Create Partition: Processors

Note: After an LPAR is created, the processor mode cannot be changed from shared to dedicated or dedicated to shared from IVM, only from the VIOS CLI by using the `chsyscfg` command.

4.7.6 Partition Ethernet

The Create Partition: Ethernet window displays the choices for assigning network connectivity. The choices, as previously noted, are:

- ▶ HEA logical port
- ▶ Virtual adapter
- ▶ SEA adapter
- ▶ Physical adapter

HEA logical ports, Virtual adapters, or SEA adapters can be selected during the this step. Physical adapters, if available, are assigned during a later step in the Create partition wizard.

Note: If creating an LPAR with shared memory resources, you only have the option of assigning virtual Ethernet adapters.

Figure 4-59 shows the first three options. The selection in this example is virtual Ethernet adapter 1 on the logical partition assigned to a SEA adapter. Note that you also have an opportunity at this time to create additional virtual Ethernet adapters for the logical partition.

Create Partition: Ethernet

Step 4 of 8

Name

Memory

Processors

Ethernet

Storage Type

Storage

Optical/Tape

Summary

Ethernet

Specify the desired Host Ethernet Adapter ports and specify virtual Ethernet for each virtual Ethernet adapter. Virtual Ethernet requires a bridge to access the external network. Host Ethernet does not require a bridge, but it does not support mobility.

Host Ethernet Adapter Ports

Select	Type	Link State	Physical Location Code ^	MAC Address	Available Connections
<input type="checkbox"/>	1 G	Up	U78A5.001.WIH23CF-P1-T6		0
<input type="checkbox"/>	1 G	Up	U78A5.001.WIH23CF-P1-T7	001A64D80407	13
<input type="checkbox"/>	1 G	Up	U78A5.001.WIH23CF-P2-T6	00215E0B028A	14
<input type="checkbox"/>	1 G	Up	U78A5.001.WIH23CF-P2-T7		0

Virtual Ethernet Configuration

Create Adapter

Adapter	Virtual Ethernet
1	1 - ent0 (U78A5.001.WIH23CF-P1-T6)
2	None

Figure 4-59 Create Partition: Ethernet

Note: HEA logical ports and physical adapter assignments cannot be used on logical partitions that will be considered for Partition Mobility.

4.7.7 Partition storage types and assignments

Logical partition disk storage can be virtual disks from a storage pool or physical volumes. Figure 4-60 on page 146 shows these choices, and also shows the option to create additional virtual disks from a storage pool that has already been defined. Refer to 4.6.2, “Storage pools” on page 119 for an explanation of how to create storage pools.

Note: For logical partitions that will be used in Partition Mobility LPARs, the storage type must be physical volumes.

Chapter 4. System planning and configuration using VIOS with IVM 145

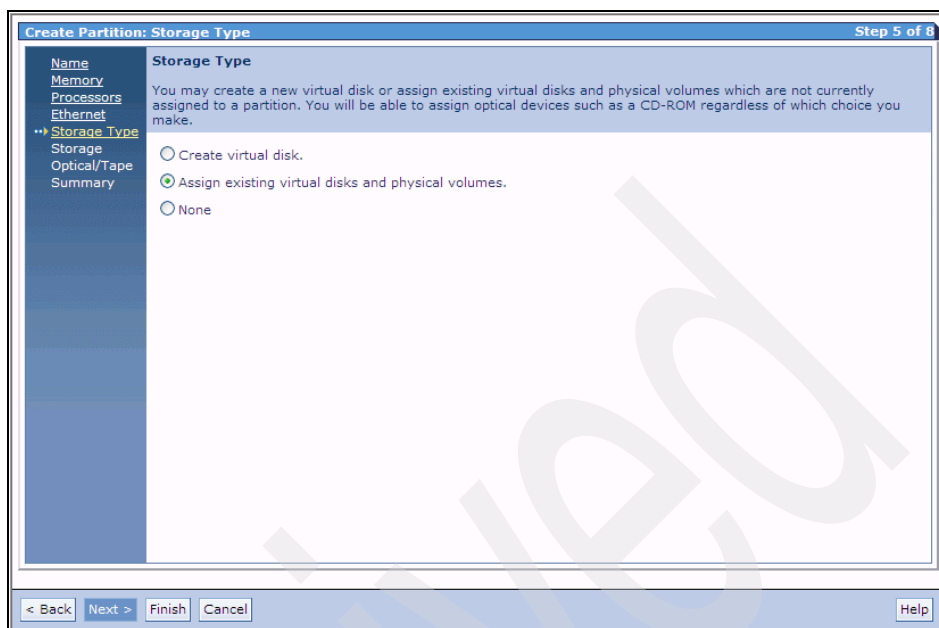


Figure 4-60 Create Partition: Storage Type

In this example we are using physical volumes. Click the option **Assign existing virtual disks and physical volumes**, and then click **Next**.

Figure 4-61 on page 147 shows the available physical volumes. Note that no virtual disks have been defined for this example, so the table under Available Virtual Disks is empty. Select one or more available hdisks, then click **Next**.

Note: For initial LPAR setup and operating system installation, we suggest that you select only the installation target disk at this time.

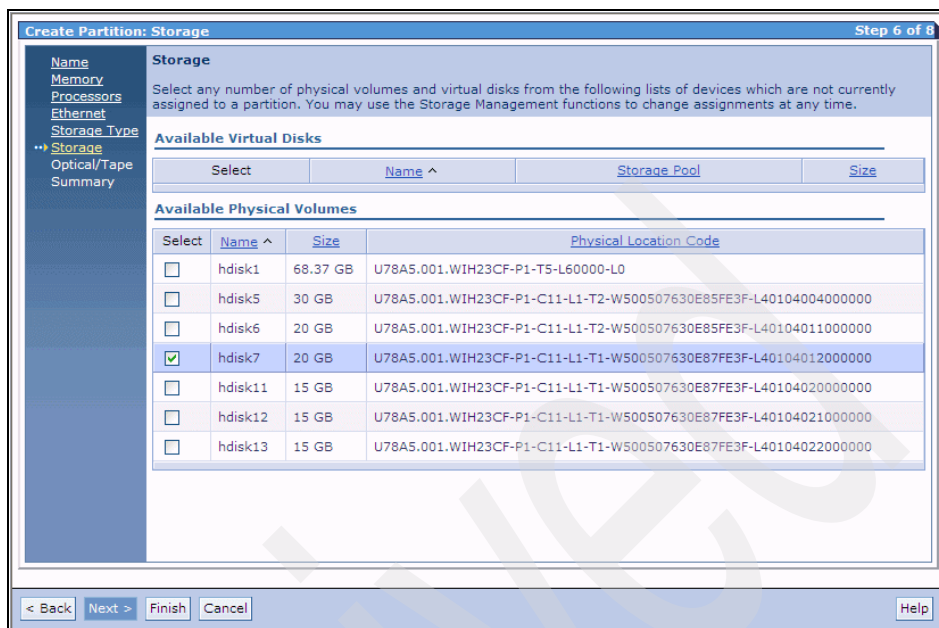


Figure 4-61 Logical Partition: Storage

4.7.8 Optical and tape devices

Optical devices, both physical and virtual, and physical tape devices can be assigned to an LPAR. With an IBM BladeCenter JS23 or JS43, the physical optical device must be available to the BladeCenter slot that you are working with through the media tray assignment before assignment to an LPAR can be made.

Virtual optical devices are not dependent on the media tray assignment. Refer to “Virtual optical devices” on page 130 for an explanation of how to create the media library and virtual optical devices.

Note: Physical and virtual optical devices cannot be used on logical partitions that will be used in Partition Mobility.

A virtual tape or virtual optical device is required for backup for IBM i.

As of this writing, only the IBM System Storage TS2240 SAS attached tape system is supported in a BladeCenter JS23 and JS43 VIOS environment.

Figure 4-62 on page 148 shows the optical device selection window. In this example, no physical optical devices are available. By default, the LPAR wizard

presents and selects a virtual optical device. If you do not want a virtual optical device selected, deselect its check box.



Figure 4-62 Create Partition: Optical

If unassigned physical adapters are available on the system, the next window to open provides the opportunity to assign them to the LPAR being created. If no physical adapter resources are available, you are directed to the summary window. Click **Next** to proceed to the Physical Adapters window (if available) or the Summary window.

4.7.9 Physical adapters

The physical adapter step is shown only if I/O adapters are available for assignment to an LPAR. Figure 4-63 on page 149 shows availability of an Ethernet adapter that could be assigned to the example LPAR.

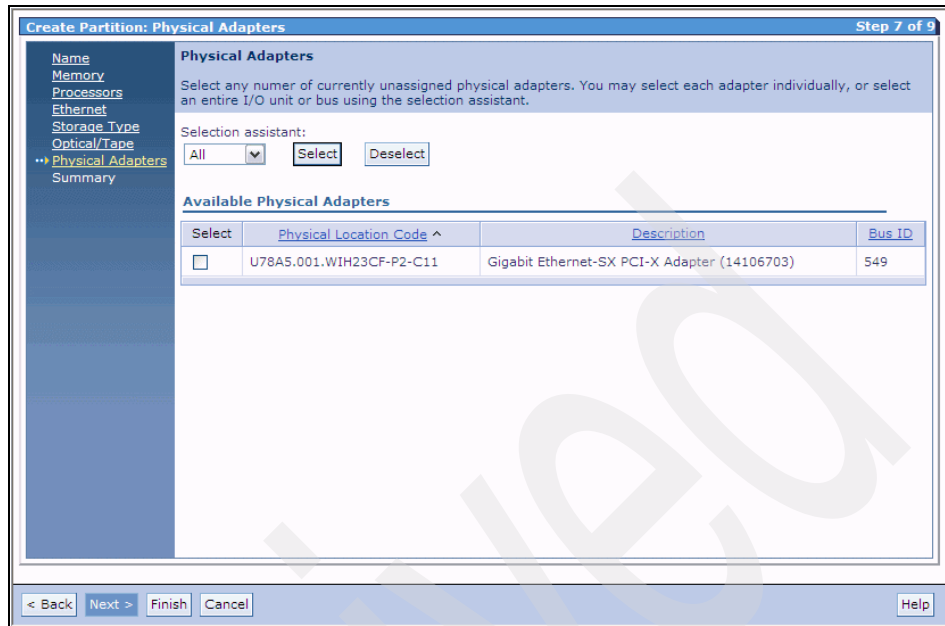


Figure 4-63 Create Partition: Physical Adapters

4.7.10 Partition summary

The final window of the Create Partition wizard is the Create Partition: Summary, as shown in Figure 4-64 on page 150. All previous selections can be reviewed on this window and modified, if necessary, by using the **Back** button.

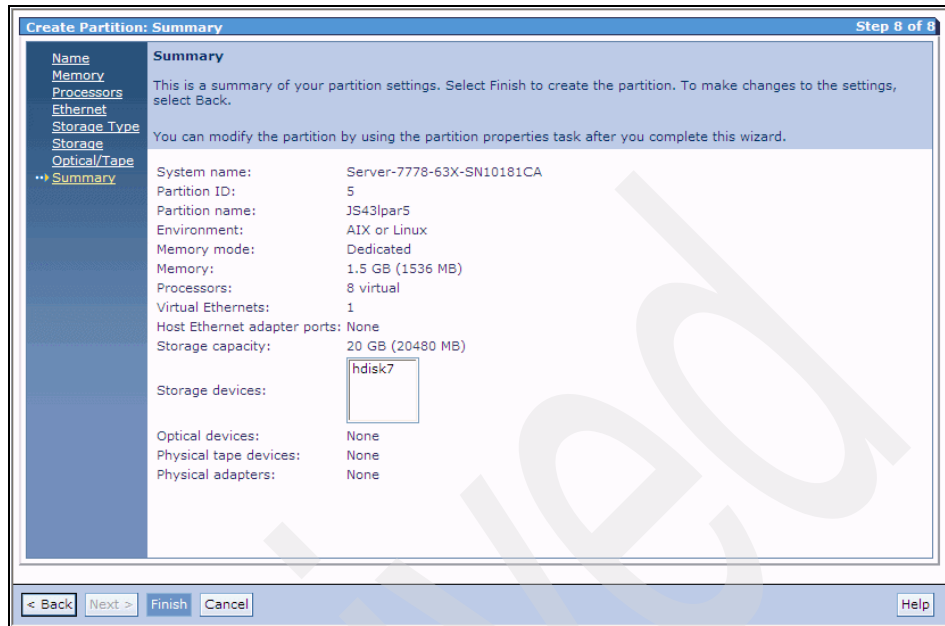


Figure 4-64 Create Partition: Summary

After your review is done and any necessary adjustments have been made, click **Finish** to complete the logical partition creation.

Figure 4-65 on page 151 of the View/Modify Partitions window shows the new logical partition that was created.

Integrated Virtualization Manager

Welcome padmin : baronlpar47.austin.ibm.com

Edit my profile | Help | Log out

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

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:	8 GB	Total processing units:	8
Memory available:	672 MB	Processing units available:	4.5
Reserved firmware memory:	768 MB	Processor pool utilization:	0.11 (1.4%)
Available shared memory pool size:	768 MB		
System attention LED:	Inactive		

Partition Details

Create Partition... Activate Shutdown --- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js43-vios	Running	4.78 Days	1.5 GB	8	0.8	0.06	
<input type="checkbox"/>	2	rhel53	Running	3.14 Days	1 GB	1	0.1	0.03	Linux ppc64
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1	1.0		00000000
<input type="checkbox"/>	4	mobilepar	Running	4.75 Days	608 MB	8	0.8	0.02	
<input type="checkbox"/>	5	JS43lpar5	Not Activated		1.5 GB	8	0.8		00000000

Figure 4-65 View/Modify Partition showing new partition

4.7.11 Partition properties changes and DLPAR operations

The IVM interface provides quick access to change an LPAR's properties and perform dynamic LPAR (DLPAR) operations on an active LPAR. The IBM BladeCenter JS23 or JS43 have the capability to perform DLPAR operations on memory, processors, and real or virtual I/O adapters.

Partition Properties changes and DLPAR operations are accessed by clicking the name of the partition to be modified when seen from the View/Modify Partitions view. When a partition name is selected, a new window will open to the General tab as shown in Figure 4-66 on page 152.

You can change the partition name, attention LED state, boot mode, and keylock position. You can also set participation in a partition workload group, as well as how the LPAR will start in the event of a complete managed system restart. The default setting is for the LPAR to automatically start when the system starts.

Partition Properties: JS23DMLpar4 (4)

General | Memory | Processing | Ethernet | Storage | Optical/Tape Devices | Physical Adapters

General

Partition name: JS23DMLpar4
 Partition ID: 4
 Environment: AIX or Linux
 State: Running
 Attention LED: Inactive

Settings

Boot mode: Normal
 Keylock position: Normal
 Partition workload group participant: ☐
 Automatically start when system starts: ☒

Dynamic Logical Partitioning (DLPAR)

Partition hostname or IP address: 9.3.29.119
 Partition communication state: Active
 Memory DLPAR capable: Unknown **Retrieve Capabilities**
 Processing DLPAR capable: Unknown
 I/O adapter DLPAR capable: Unknown

OK Cancel

Figure 4-66 Partition Properties General tab

DLPAR capabilities can be retrieved by clicking **Retrieve Capabilities**.

Figure 4-67 on page 153 shows the DLPAR capabilities of the IBM BladeCenter JS23 or JS43.

IBM i LPARs have a different Partition Properties General tab view. Refer to 7.3, “Creating an IBM i V6.1 partition” on page 263 for more information.

Note: The AIX or Linux LPARs must have a network interface configured and have an active Resource Monitoring and Control (RMC) connection with the management partition so you can retrieve capabilities or perform DLPAR operations. The VIOS partition and IBM i partitions do not require RMC for DLPAR.

Partition Properties: JS23DMLpar4 (4)

General Memory Processing Ethernet Storage Optical/Tape Devices Physical Adapters

General

Partition name: JS23DMLpar4
 Partition ID: 4
 Environment: AIX or Linux
 State: Running
 Attention LED: Inactive

Settings

Boot mode: Normal
 Keylock position: Normal
 Partition workload group participant: ☐
 Automatically start when system starts: ☒

Dynamic Logical Partitioning (DLPAR)

Partition hostname or IP address: 9.3.29.119
 Partition communication state: Active
 Memory DLPAR capable: Yes
 Processing DLPAR capable: Yes
 I/O adapter DLPAR capable: Yes

OK Cancel

Figure 4-67 DLPAR retrieved capabilities

Selecting the **Memory** tab displays current and pending memory values for the LPAR, as shown in Figure 4-68 on page 154. In addition, if a shared memory pool has been configured, you will have the option to change between dedicated and shared memory. The change between dedicated and shared can be done only on an inactive LPAR.

An active LPAR can have its *Assigned* memory value changed between the range of the minimum and maximum values as a DLPAR operation. The minimum memory and maximum memory pending values can be changed only when the LPAR is *not* activated.

Note: The VIOS or management partition can change the minimum and maximum memory values while active, but does not become effective until a partition shutdown and restart is performed.

Partition Properties: JS23DMLpar4 (4)

General **Memory** Processing Ethernet Storage Optical/Tape Devices Physical Adapters

Modify the settings by changing the pending values. The changes will be applied immediately; however, synchronizing the current and pending values might take some time.

Memory mode: Dedicated **i** You cannot change the memory mode of this partition because the partition is active.

All memory values should be in multiples of 32 MB.

Property	Current	Pending
Minimum memory	128 MB	<input type="text" value="128"/> MB
Assigned memory	1 GB (1024 MB)	<input type="text" value="1"/> GB
Maximum memory	4 GB (4096 MB)	<input type="text" value="4"/> GB

OK Cancel

Figure 4-68 Partition Properties Memory tab

Click the **Processing** tab, to change the processing units, virtual processors, partition priority weighting, and processor compatibility mode for LPARs using a shared processor pool, as shown in Figure 4-69 on page 155. When changing the processor compatibility mode, a partition shutdown and restart is required for an active LPAR to apply the change. If the LPAR is already inactive, an activation is required before the current value is updated.

Note: For Partition Mobility, **Processor compatibility mode** must be selected when moving an LPAR from a JS23 or JS43 to a JS12 or JS22. The earlier blades do not support Power6+ or Power6+ Enhanced modes.

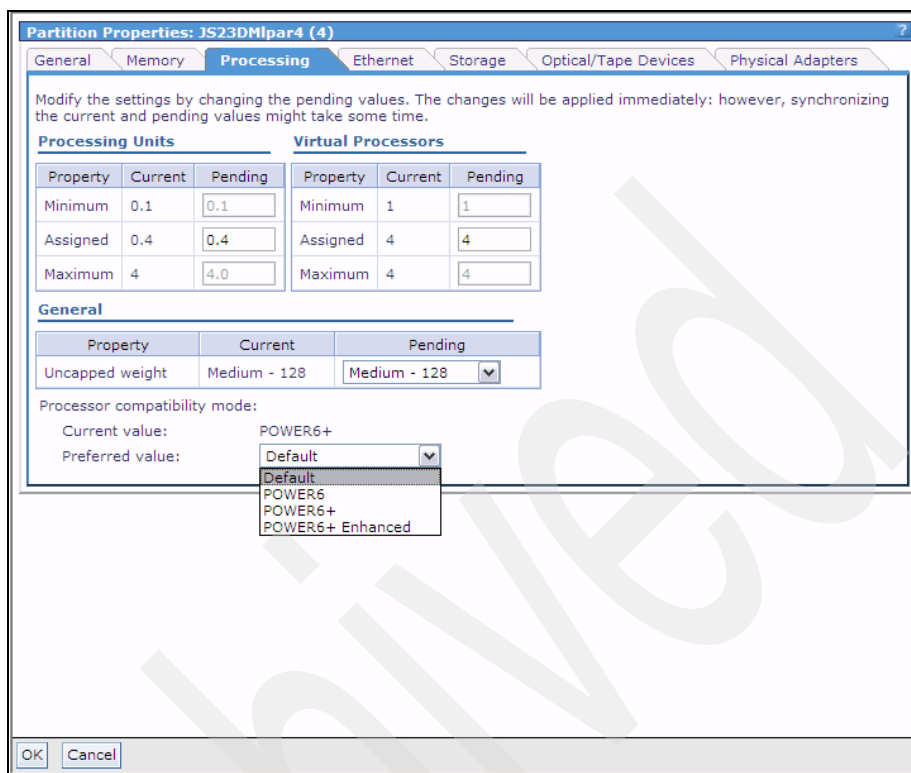


Figure 4-69 Partition Properties, Processing tab for shared pool

If partitions use dedicated processors, the window shown in Figure 4-70 on page 156 opens. This example shows the LPAR state as not activated, and the minimum, assigned, and maximum values can be changed. In an active LPAR, only the assigned value can be altered as a DLPAR operation.

This window also enables changing the mode of sharing idle processors. The four modes are:

- ▶ When inactive (default), shares excess capacity when the LPAR is not active
- ▶ When active, shares excess capacity when the LPAR is active, but not when inactive
- ▶ Always, shares excess capacity when the LPAR is active or inactive
- ▶ Never, no access capacity

The excess capacity that is shared is utilized by the shared pool.

The processor compatibility mode can also be changed when using dedicated processors.

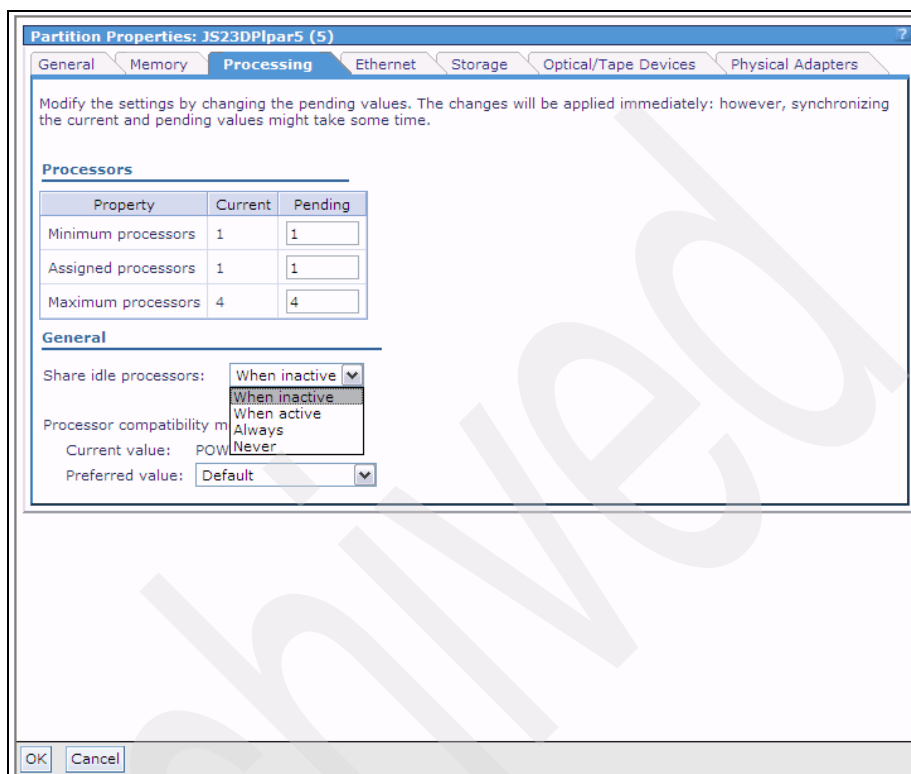


Figure 4-70 Partition Properties, Processing tab for dedicated processors

Click the **Ethernet** tab in Partition Properties, which allows the addition or removal of Ethernet adapters, as shown in Figure 4-71 on page 157.

Note: Before you can dynamically remove LPAR Ethernet adapters from an active AIX LPAR by using the **DLPAR remove** command, first use the **rmdev** command to remove the devices from the LPAR.

HEA virtual ports required the removal of Logical Host Ethernet Adapter (l-hea) and the Logical Host Ethernet Port (lp-hea). Virtual Ethernet adapters can be removed by deleting the Virtual I/O Ethernet Adapter (l-lan). Physical Ethernet adapters require the deletion of the adapter (ent) and the parent. The parent can be determined with the **lsdev** command. For example:

```
# lsdev -Cl ent1 -F parent
pci1
```

If you try to use the alternate method, which is with the **DLPAR** command, IVM returns an error message containing details about the parent and child devices that must be removed.

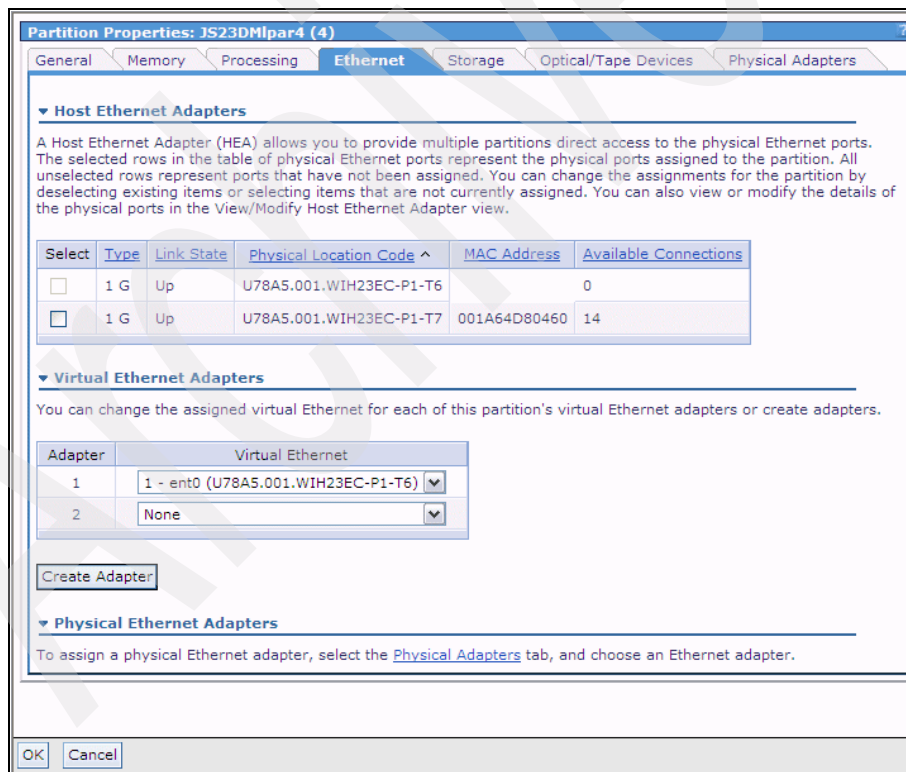


Figure 4-71 Partition Properties, Ethernet tab

Note: Partitions that are configured for shared memory or IBM i partitions cannot own HEAs. Therefore, the Host Ethernet Adapter section of this window will not be shown when viewing the properties of these types of LPARs.

Click the **Storage** tab, which is where you can add or remove storage devices, either physical volumes or virtual disks, as shown in Figure 4-72.

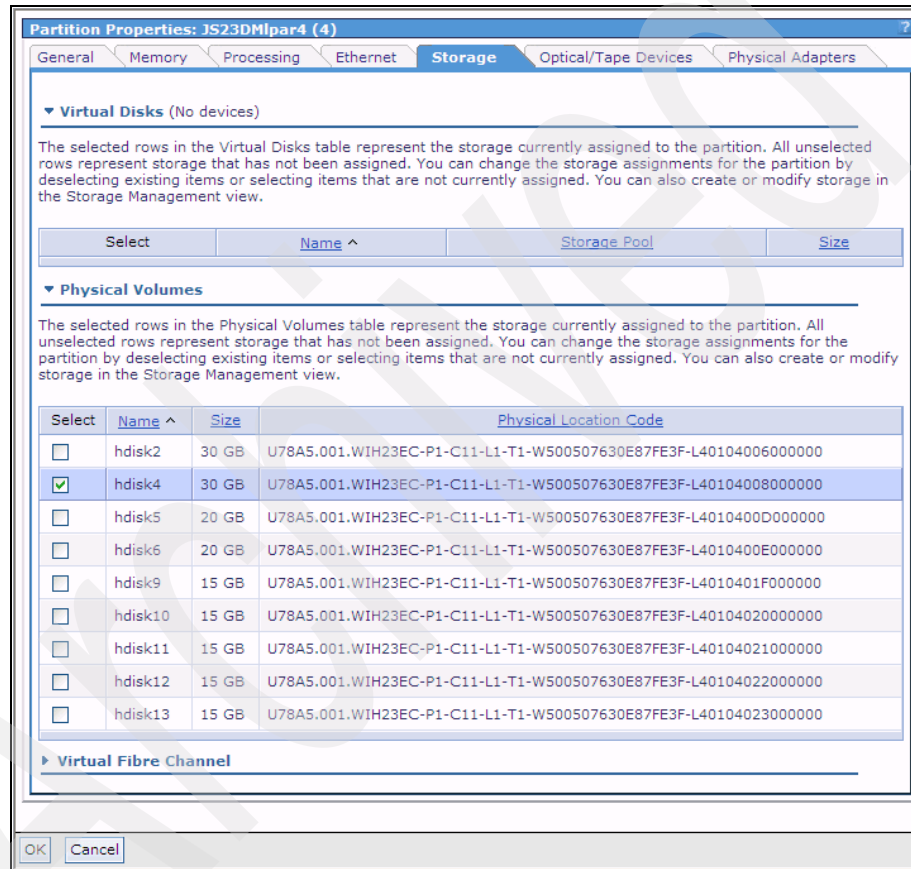


Figure 4-72 Partition Properties, Storage tab

Note: When removing storage from an active partition, IVM requires that you verify that a forced device removal from a running partition is requested.

Optical device assignments, both physical and virtual, and physical tape assignments can be managed from the **Optical /Tape Devices** tab shown in Figure 4-73.

Partition Properties: JS23DMLpar4 (4)

General Memory Processing Ethernet Storage **Optical/Tape Devices** Physical Adapters

▼ **Physical Optical Devices**

You can assign physical optical devices on the system directly to a partition to use for storage. In the Physical Tape Device table, selected rows represent those devices with assignments to the current partition and unselected rows represent devices that do not have an assignment to any partition. To remove a device assignment for the current partition, clear the selection for that device in the table. To assign a device to the current partition, select that device in the table.

Select	Name ^	Description	Physical Location Code
<input type="checkbox"/>	cd0	USB DVD-COMBO Drive	U78A5.001.WIH23EC-P1-T1-L1-L2-L3

▼ **Virtual Optical Devices (No 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 current partition. Selected rows in the Virtual Optical Devices table represent those devices with assignments to the current partition and unselected rows represent devices that do not have an assignment to any partition. To remove a device assignment for the current partition, clear the selection for that device in the table. To assign a device to the current partition, select that device in the table. Click Modify to change the mounted media for a specific optical device.

Select	Name ^	Current Media	Current Media Size	Mount Type
--------	--------	---------------	--------------------	------------

Create Device

▼ **Physical Tape Devices (No devices)**

You can assign physical tape devices directly to a partition to use for storage. In the Physical Tape Device table, selected rows represent those devices with assignments to the current partition and unselected rows represent devices that do not have an assignment to any partition. To remove a device assignment for the current partition, clear the selection for that device in the table. To assign a device to the current partition, select that device in the table.

Select	Name ^	Description	Physical Location Code
--------	--------	-------------	------------------------

OK Cancel

Figure 4-73 Partition Properties Optical/Tape Devices tab

Additional virtual optical devices can be created, and the media that is mounted to an existing virtual optical device can be changed in this window. Creating virtual optical media is covered in “Virtual optical devices” on page 130.

Figure 4-74 on page 160 shows an existing virtual optical device, vtopt0, being modified by changing the current media.

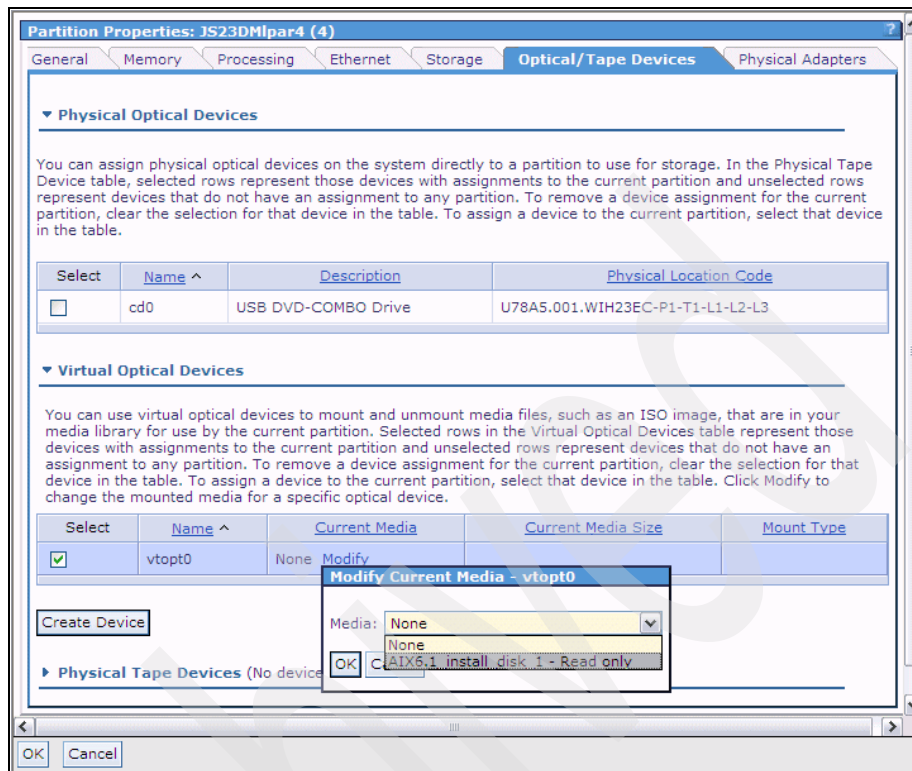


Figure 4-74 Partition Properties, changing the current mounted media

Physical adapters that are not assigned to an LPAR or any physical adapters that are already assigned to the selected LPAR are displayed when you select the **Physical Adapters** tab.

Figure 4-75 on page 161 shows a Gigabit Ethernet-SX PCI-X Adapter available for assignment to this LPAR.

Note: Partitions that are configured for shared memory or IBM i partitions cannot use physical adapters. Therefore, the Physical Adapter tab of this window is not shown when you view the properties of these types of LPARs.

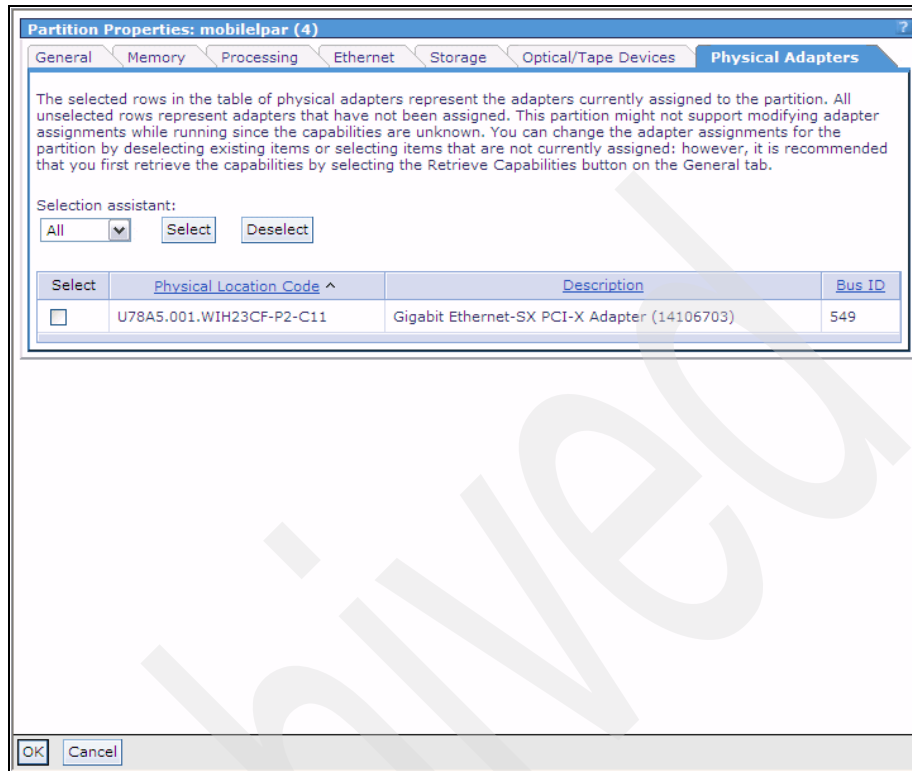


Figure 4-75 Partition Properties, Physical Adapters tab

Note: The Partition Properties window for the VIOS partition does not have the Storage and Optical Devices tabs.

4.8 Console access and activating a partition

This section discusses basic access to a partition and partition management functions.

4.8.1 Opening a virtual terminal

Accessing a partition virtual terminal from the VIOS can be done in two ways. However, only one virtual terminal to an LPAR can be open at a time.

Note: These methods are not available for IBM i. In the case of IBM i, the Operations Console (LAN) is the only supported system console.

The first method from the IVM GUI is shown in Figure 4-76. In the Partition Details section of the View/Modify Partitions view, select the check box for the desired LPAR. From the drop-down box, select **Open terminal window**.

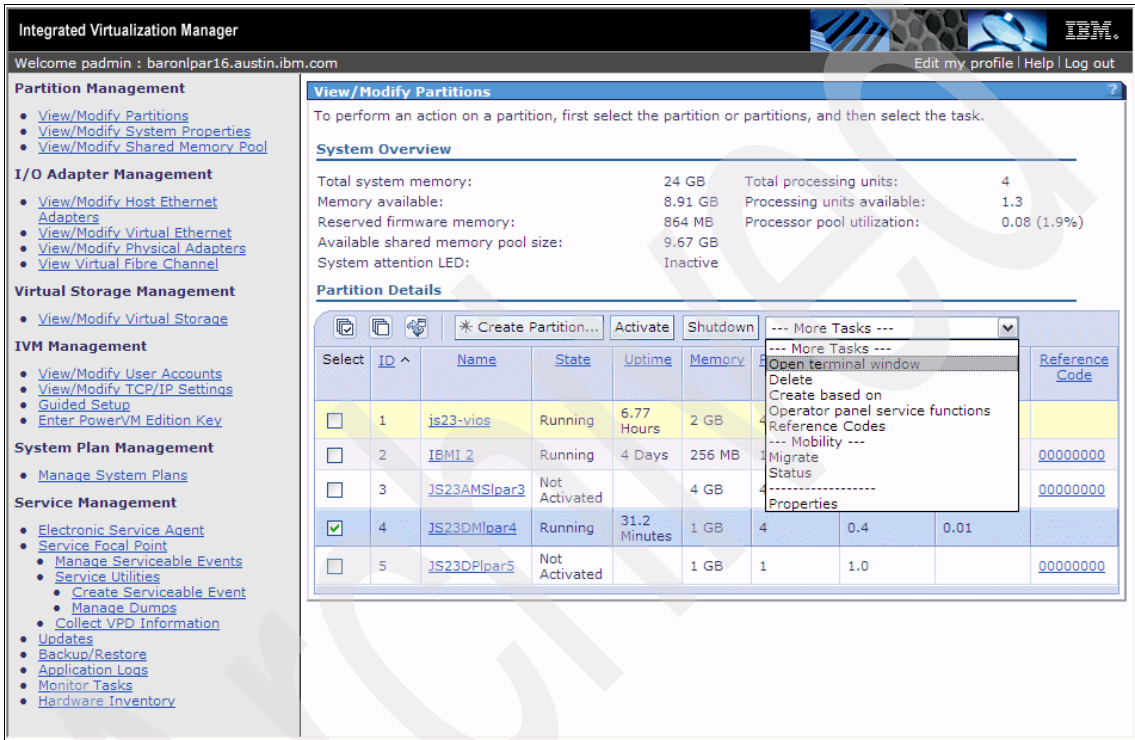


Figure 4-76 Start virtual terminal for an LPAR

A separate applet will start and open the virtual terminal window. This new window prompts for the password of the login ID of the current IVM session.

Figure 4-77 shows a successful connection to the LPAR virtual terminal.

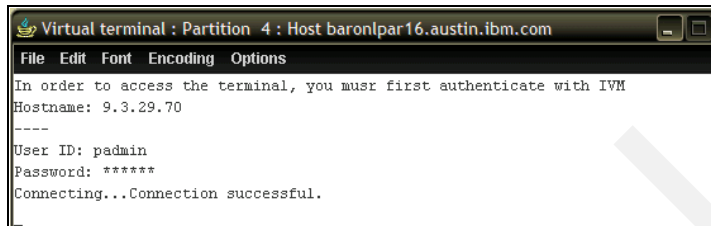


Figure 4-77 Virtual Terminal started from IVM GUI

The second method to start a virtual terminal is from the VIOS command line where you issue the **mkvt** command as shown in Example 4-23.

Example 4-23 Creating a virtual terminal from the command line

```
$ mkvt -id 4
```

Specify the partition number that you want to connect after the **-id** flag. Your command line session to the VIOS now becomes the virtual terminal for the specified partition.

Note: The following key sequence (which is the Ctrl key plus the tilde character plus period) entered from the virtual terminal enables you to break out of the **mkvt** command or close the virtual terminal applet:

Ctrl + ~.

4.8.2 Activating a partition

During initial setup and configuration of an LPAR, open a virtual terminal connection to the partition prior to activating.

Activating from the GUI

Activating a partition from the IVM GUI starts from View/Modify Partitions. In this view, select the partition to be activated and click **Activate**, as shown in Figure 4-78 on page 164.

Integrated Virtualization Manager

Welcome padmin : baronlpar16.austin.ibm.com

[Edit my profile](#) | [Help](#) | [Log out](#)

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](#)

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:

24 GB

Total processing units:

4

Memory available:

8.91 GB

Processing units available:

1.3

Reserved firmware memory:

864 MB

Processor pool utilization:

0.12 (3.1%)

Available shared memory pool size:

9.67 GB

System attention LED:

Inactive

Partition Details

Create Partition...

Activate

Shutdown

--- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js23-vios	Running	7.46 Hours	2 GB	4	0.4	0.01	
<input type="checkbox"/>	2	IBMI 2	Running	4.03 Days	256 MB	1	0.5	0.00	00000000
<input type="checkbox"/>	3	JS23AMSlpar3	Not Activated		4 GB	4	0.4		00000000
<input checked="" type="checkbox"/>	4	JS23DMlpar4	Not Activated		1 GB	4	0.4		00000000
<input type="checkbox"/>	5	JS23DPlpar5	Not Activated		1 GB	1	1.0		00000000

Figure 4-78 Activating a partition

The next window shows the current state of the partition and asks you to confirm activation by clicking **OK**, as shown in Figure 4-79 on page 165.

164 IBM BladeCenter JS23 and JS43 Implementation Guide

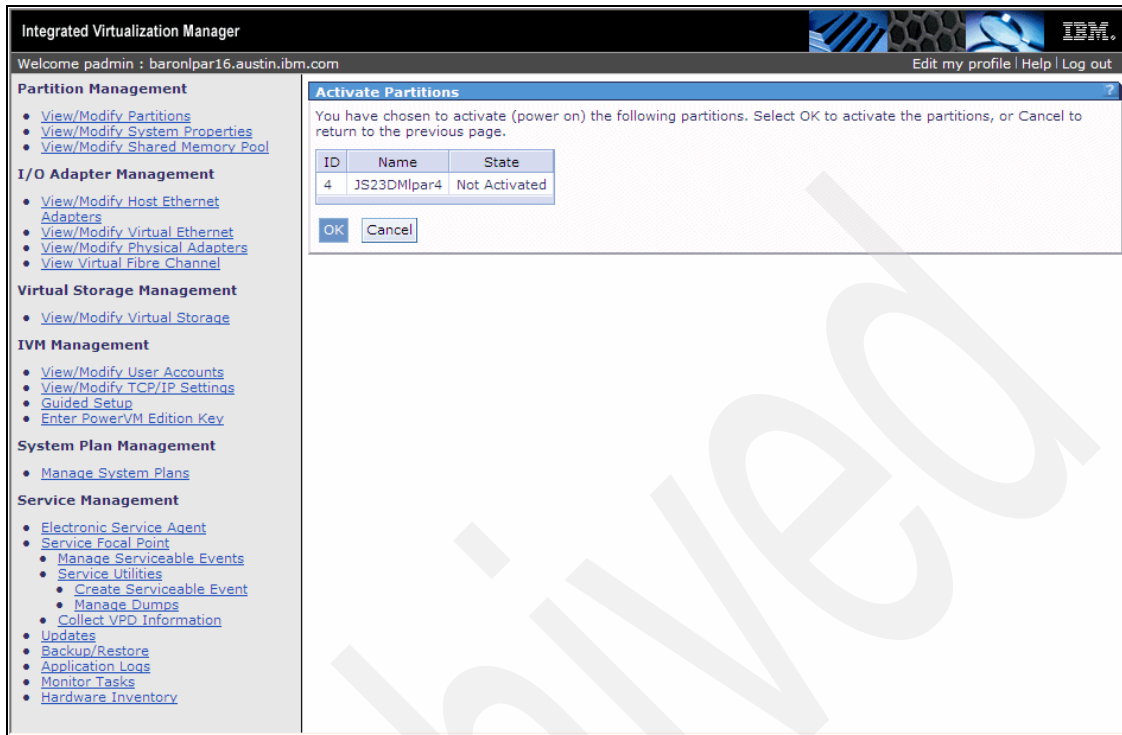


Figure 4-79 Confirm partition activation

When the LPAR activation starts, the message Performing Task - Please Wait briefly appears, then the IVM GUI returns to the View/Modify Partitions window.

Activating from the CLI

The **chsysstate** command is used to start a partition from the command-line interface (CLI) by either the LPAR number (**--id**) or name (**-n**). Example 4-24 shows LPAR 4 being activated from the CLI by using either number or name.

Example 4-24 Using the CLI to activate a partition

```
$ chsysstate -o on -r lpar --id 4
```

Or, use the name:

```
$ chsysstate -o on -r lpar -n JS23DMIpar4
```

The **lsrefcode** command can be used to monitor the status codes as the LPR becomes active. Example 4-25 on page 166 shows the **lsrefcode** command being used with both name for LPAR 4 and LPAR number.

Example 4-25 Using the CLI to monitor partition activation status codes

```
$ lsrefcode -r lpar --filter lpar_names=JS23DMlpar4 -F refcode
00000000
$ lsrefcode -r lpar --filter lpar_names=JS23DMlpar4 -F refcode
CA00E1F1
$ lsrefcode -r lpar --filter lpar_names=JS23DMlpar4 -F refcode
AA00E158
```

Or, use the number:

```
$ lsrefcode -r lpar --filter lpar_ids=4 -F refcode
00000000
$ lsrefcode -r lpar --filter lpar_ids=4 -F refcode
CA00E1F1
$ lsrefcode -r lpar --filter lpar_ids=4 -F refcode
AA00E158
```

The **lssyscfg** command can be used to display the state of the LPARS by name or the ID number of the LPAR, as shown in Example 4-26.

Example 4-26 The lssyscfg command used to display the LPAR state

```
$ lssyscfg -r lpar -F name,state
js23-vios,Running
IBMI 2,Running
JS23AMS1par3,Not Activated
JS23DMlpar4,Running
JS23DP1par5,Not Activated
```

Or, use the number:

```
$ lssyscfg -r lpar -F lpar_id,state
1,Running
2,Running
3,Not Activated
4,Running
5,Not Activated
```

4.8.3 Shut down a VIO Client partition

The shutdown of a partition can be initiated from the GUI or the CLI. The shutdown process can interact with the operating system on an LPAR, or can be immediate without notifying the operating system.

The following shutdown types, listed on the IVM GUI, are available for a partition shutdown:

- ▶ Operating System
- ▶ Delayed
- ▶ Immediate

The Operating System shutdown option is available only if the RMC connection is active. It is the preferred method.

The Delayed option is the equivalent of pushing the white control-panel power button. AIX partitions can gracefully handle this option, but Linux partitions are required to install IBM service and productivity tools for Linux on POWER.

The Immediate shutdown option should be used only as a last resort because of the potential of data loss.

IBM i partitions should be shut down by using the IBM i session commands SHTDWNSYS, ENDSBS, or PWRDWNSYS.

Shutdown from the GUI

From the View/Modify Partitions view, select the check box for the LPAR to be shut down and then click **Shutdown** as shown in Figure 4-80 on page 168.

Integrated Virtualization Manager

Welcome padmin : baronlpar16.austin.ibm.com

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

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:

24 GB

Total processing units:

4

Memory available:

8.91 GB

Processing units available:

1.3

Reserved firmware memory:

864 MB

Processor pool utilization:

0.02 (0.6%)

Available shared memory pool size:

9.67 GB

System attention LED:

Inactive

Partition Details

Create Partition...

Activate

Shutdown

--- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js23-vies	Running	7.69 Hours	2 GB	4	0.4	0.01	
<input type="checkbox"/>	2	IBMI 2	Running	4.04 Days	256 MB	1	0.5	0.00	00000000
<input type="checkbox"/>	3	JS23AMSlpar3	Not Activated		4 GB	4	0.4		00000000
<input checked="" type="checkbox"/>	4	JS23DMPpar4	Running	2.3 Minutes	1 GB	4	0.4	0.01	
<input type="checkbox"/>	5	JS23DPlpar5	Not Activated		1 GB	1	1.0		00000000

Figure 4-80 Shutdown an LPAR

The Shutdown partitions window opens, shown in Figure 4-81 on page 169.

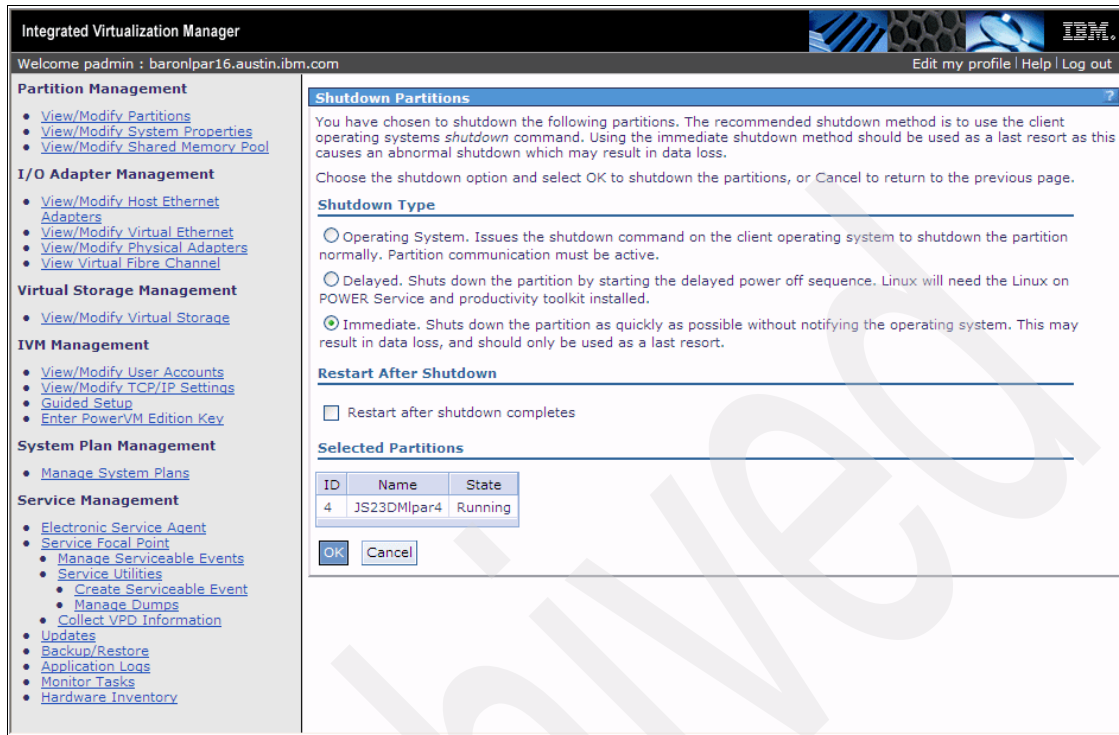


Figure 4-81 Partition shutdown options

From this window, choose the shutdown type option. The partition can also be restarted after the shutdown by selecting the restart check box option. Click **OK** and the partition shuts down.

Note: The Operating System option is disabled if RMC is not active between the LPAR and VIOS. The Delayed option is selected by default.

Shutdown from the CLI

The **chsysstate** command is used to shut down a partition from the command line by either the LPAR number or name, as shown in Example 4-27 on page 170. This example uses the operating system shutdown option on LPAR 4.

Example 4-27 CLI shutdown of a logical partition

```
$ chsysstate -r lpar -o osshutdown --id 4
```

Or, use the name:

```
$ chsysstate -r lpar -o osshutdown -n "JS23DM1par4"
```

The corresponding CLI shutdown options to use with the **-o** flag are:

- ▶ osshutdown (Operating System)
- ▶ shutdown (Delayed, white button shutdown)
- ▶ shutdown --immed (Immediate)

4.8.4 Shutdown the VIO Server

The VIO Server is shut down in a similar process to a VIO Client LPAR. Both the GUI and CLI can be used.

Shutdown from the GUI

When selecting the VIOS partition to be shut down, a warning message indicates that shutting down the IVM partition will shut down *all* partitions and the entire system. There is no shutdown option to choose only the restart option. Figure 4-82 on page 171 shows the VIOS Shutdown Partitions window.

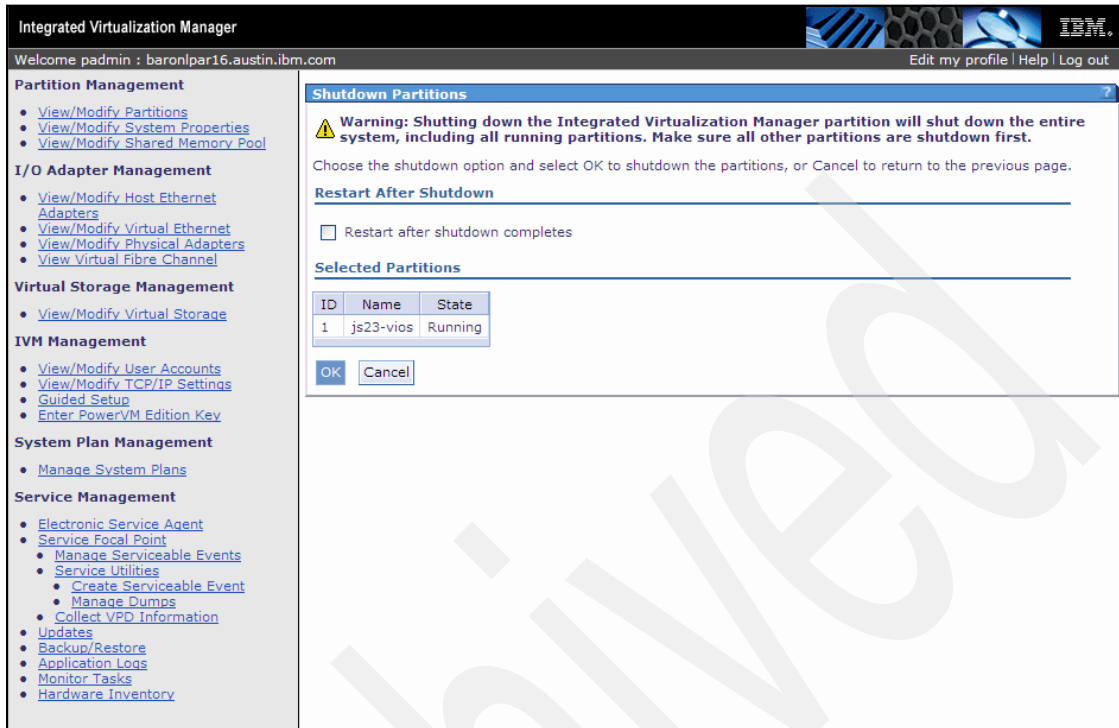


Figure 4-82 VIOS partition shutdown options

Shutdown using the CLI

The **shutdown** command to use from the CLI or console session is shown in Example 4-28 and has two options. To automatically restart after the shutdown use the **-restart** flag; to suppress the warning message and confirmation add the **-force** option.

Example 4-28 VIOS shutdown command

```
$ shutdown -restart
Shutting down the VIO Server could affect Client Partitions. Continue
[y|n]?
```

Or, use **-force**:

```
$ shutdown -restart -force
```


Active Memory Sharing configuration using IVM

Active Memory Sharing (AMS) is an IBM PowerVM advanced memory virtualization technology that provides system memory virtualization capabilities to IBM Power Systems, allowing multiple logical partitions to share a common pool of physical memory. This chapter describes how to configure AMS using the IVM GUI, and at a high level, planning considerations that should be used.

This chapter contains the following topics:

- ▶ “Planning considerations” on page 174
- ▶ “Configuration of Active Memory Sharing using IVM” on page 178
- ▶ “Active Memory Sharing summary” on page 204

For more information about AMS, see the IBM PowerVM Web site:

<http://www.ibm.com/systems/power/software/virtualization/>

For a more complete technical discussion, refer to:

- ▶ *PowerVM Virtualization Active Memory Sharing*, REDP-4470
- ▶ *IBM PowerVM Active Memory Sharing: An Overview*:
<ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/pow03026usen/POW03026USEN.PDF>

5.1 Planning considerations

Active Memory Sharing is an IBM PowerVM advanced memory virtualization technology that provides system memory virtualization capabilities to IBM Power Systems, allowing multiple logical partitions to share a common pool of physical memory.

When you use a shared memory mode, the system automatically decides the optimal distribution of the physical memory to logical partitions and adjusts the memory assignment based on demand for memory pages. The administrator just reserves physical memory for the shared memory pool and assigns logical partitions to the pool.

The PowerVM Active Memory Sharing technology (or *shared memory*) allows logical partitions to share the memory in the shared memory pool. Instead of assigning a dedicated amount of physical memory to each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*), the hypervisor constantly provisions the physical memory from the shared memory pool to the shared memory partitions as necessary. The hypervisor allows shared memory partitions that require memory to use portions of the shared memory pool that are not currently being used by other shared memory partitions. When a shared memory partition requires more memory than the current amount of unused memory in the shared memory pool, the hypervisor stores a portion of the memory that belongs to the shared memory partition in auxiliary storage. Access to the auxiliary storage is provided by a Virtual I/O Server logical partition. When the operating system attempts to access data that is located in the auxiliary storage, the hypervisor directs a Virtual I/O Server to retrieve the data from the auxiliary storage and write it to the shared memory pool so that the operating system can access the data. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition hardware feature, which also includes the license for the Virtual I/O Server software.

Successful implementation of Active Memory Sharing requires a thorough understanding of the workloads that will be placed on memory sharing partitions. Workloads that have high sustained memory requirements because of sustained peak loads or High Performance Computing applications would not be good candidates for Active Memory Sharing.

5.1.1 Prerequisites

To use Active Memory Sharing on an IBM BladeCenter JS23 or JS43, minimum software, configuration, and feature levels must be met:

- ▶ IBM PowerVM Enterprise Edition
- ▶ Virtual Input/Output Server 2.1.1
- ▶ Only virtual I/O, no physical adapters or logical ports from an HEA allowed
- ▶ Only shared processor mode, no dedicated processors
- ▶ AIX 6.1 TL3
- ▶ IBM i V6.1 plus latest cumulative PTF package + SI32798
- ▶ SUSE Linux Enterprise Server 11

5.1.2 Workload

In the Active Memory Sharing environment, three types of shared memory scenarios are available:

- ▶ Physical over-commit
- ▶ Logical over-commit
- ▶ Non over-commit

Physical over-commit occurs when the combined working set of memory, of all the shared memory partitions, is greater than the total physical memory in the shared memory pool. In this case, the actively used logical memory of the shared memory is backed by physical memory and the paging devices.

Logical over-commit occurs when actively referenced memory pages by the partitions are backed by physical memory allocated from the Active Memory Sharing Shared Memory Pool. Logical memory that is not actively referenced and that exceeds the amount of physical memory will reside on the paging device.

Non-overcommit occurs when the backing physical memory in the shared memory pool is adequate to cover all the peak requirements of the shared memory partitions.

The potential workload candidates for Active Memory Sharing should be monitored for a period of time using a dedicated memory partition to understand the size and range requirements of memory usage.

Physical over-commit would be suitable for workloads such as file and print servers and network applications that are not highly sensitive to I/O latency. Shared memory partitions in this model require paging of active memory pages.

Logical over-commit scenarios would work best with partition workloads that have peak usage at different times of the day causing utilization peaks and valleys to overlap. Other examples that would be prime candidates are test and development environments and workloads that do not have sustained load requirements.

5.1.3 Paging devices

Active Memory Sharing paging devices and operating systems paging device considerations are similar. Active Memory Sharing paging operations are typically 4 k in size. Write and read caches should be enabled. Striped disk configurations should be used when possible with a 4 k stripe size.

The sizes of the dedicated paging devices for AIX and Linux shared memory partitions must at least equal the maximum amount of logical memory assigned. IBM i shared memory partitions must have paging devices that are at least the size of the maximum logical memory assigned multiplied by 1.008.

5.1.4 Collaborative Memory Manager

Collaborative Memory Manager (CMM) running in the operating system of the shared memory partitions coordinates with the hypervisor to provide hints on logical page usage. These hints of the status of memory pages will help the hypervisor prioritize pages.

CMM is configured from the OS running in the shared memory partition. When the CMM loan policy is enabled, the OS loans pages to the hypervisor to help reduce hypervisor paging required to support the overcommitment of memory. When CMM policy is disabled the hypervisor will do the paging based on hints from the OS.

5.1.5 Processor resource sizing

This section looks at VIO Server and shared memory partition additional processor entitlements required in an Active Memory Sharing environment.

Virtual I/O Server

Additional processor resources are required for the VIO Server to properly service the Active Memory Sharing environment. The hypervisor does the

address translation and allocation of physical memory across partitions and the VIOS donates cycles to the hypervisor for these operations and also requires cycles for the paging activity.

The additional VIOS processor requirements per shared memory partition can typically range from 0.005% for light paging rates using internal storage to 0.16% for heavy paging rates using high end SAN storage as shown in Table 5-1.

Table 5-1 Estimated additional VIOS processor entitlement per shared memory LPAR

Paging rate	Storage types (value is %)			
	Internal storage	Entry level storage	Mid range storage	High end storage
Light	0.005	0.01	0.02	0.02
Moderate	0.01	0.02	0.04	0.08
Heavy	0.02	0.04	0.08	0.16

Shared memory partition

Shared memory partitions also require additional CPU entitlement compared to dedicated memory partitions running the same workload. The additional amount is dependent on the memory subscription ration (sum of logical memory divided by physical memory in the pool) and CMM loan policy settings. If memory is not over-committed, the additional CPU entitlement will be minimal.

IBM Workload Estimator (IBM WLE) is updated to produce estimated sizing for shared memory partition CPU requirements.

IBM Workload Estimator can be found at:

<http://www.ibm.com/systems/support/tools/estimator/>

5.1.6 Memory weight

The hypervisor considers several variables to allocate and manage memory across the shared memory partitions. These variables included global and partition load, global and partition memory pressure, and the one direct user adjustment memory weight. Similar workloads that run concurrently require different weights if one partition requires a priority. In an active/inactive scenario when one shared memory partition uses most of the resources, the memory weight should be equal in both partitions to enhance memory transfer when the inactive partition becomes active.

5.1.7 Consolidation factors

After the workload types in a environment have been determined and selected for consolidation, additional factors must be considered:

- ▶ The logical to physical subscription ratio must be suitable for the selected workloads.
- ▶ Assigning a memory weight must be considered. The IVM GUI allows three values, low, medium and high, with a default of medium.
- ▶ Paging device configuration must be considered. The higher the subscription ratio the higher the requirement for optimized paging devices.
- ▶ CMM configuration determines page loaning policy. Application load and loaning policy, none to aggressive, should be evaluated for acceptable performance. CMM is set at the OS level therefore a mix of loaning levels can exist in the same system.
- ▶ Resources must be rebalance to improve overall system utilization. Previous memory limits that caused under utilization of CPU and memory resources can be reviewed for new allocations or additional workloads.

5.2 Configuration of Active Memory Sharing using IVM

IVM is used to configure the main components of Active Memory Sharing, the shared memory pool, and the paging storage devices used by the hypervisor. After Active Memory Sharing is configured, IVM is used to create new, or change existing LPARS to utilize shared memory

5.2.1 Defining the shared memory pool and paging storage pool

When creating the shared memory pool for Active Memory Sharing, you also create the paging storage pool in the same step. The shared memory pool cannot exceed the current memory available minus the amount required for Reserved Firmware Memory.

Note: Reserved firmware memory requirements increase when the Active Memory Sharing shared memory pool is created. Also, as additional partitions are created and additional virtual resources are defined, the Reserved Firmware Memory value will increase. If the Active Memory Sharing shared memory pool is defined close to the maximum amount, you limit the ability to create additional resources.

Two possibilities for the paging space that is managed by the hypervisor are to:

- ▶ Create a common paging storage pool; this step is required.
- ▶ Create individual paging devices for each LPAR using shared memory; this is discussed in Figure 5.2.2 on page 187.

Because a common paging storage pool is required, the first step is to create a *storage pool* that can be assigned as the paging storage pool.

1. To create a common paging storage pool, start in the navigation area of the IVM GUI and click **View/Modify Virtual Storage** as shown in Figure 5-1.

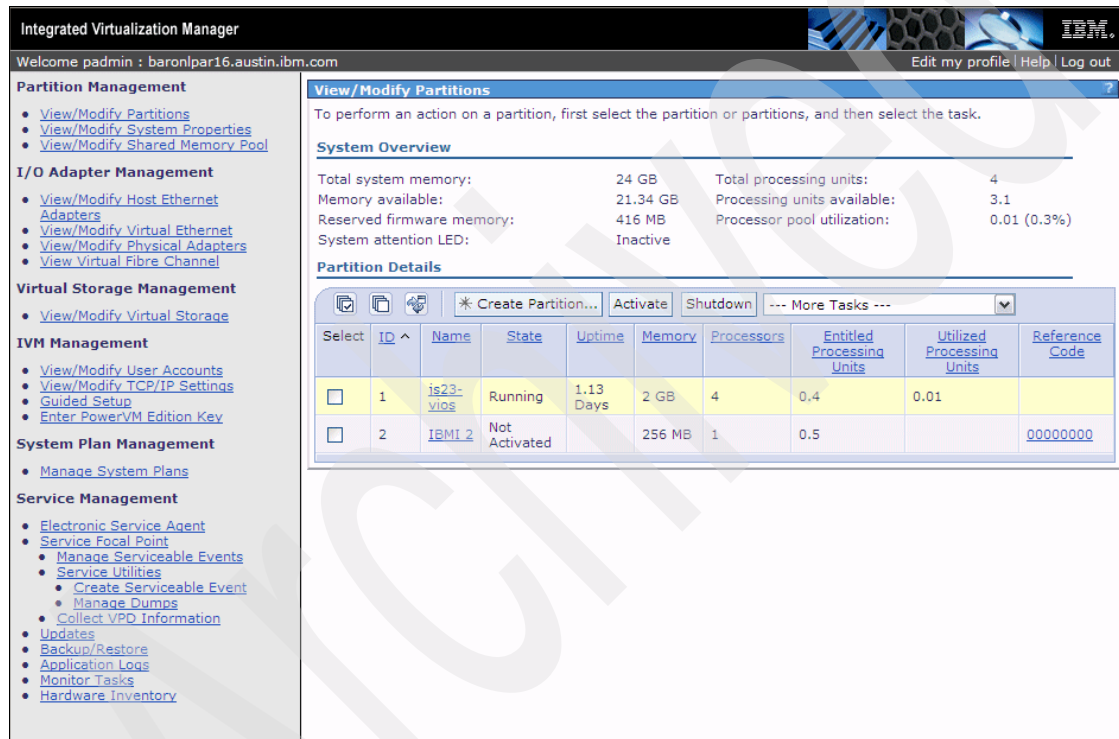


Figure 5-1 Start Active Memory Sharing configuration with View/Modify Virtual Storage

2. The next window begins the storage pool creation process. Select **Create Storage Pool** as shown in Figure 5-2 on page 180. This pool will then be used for shared paging storage.

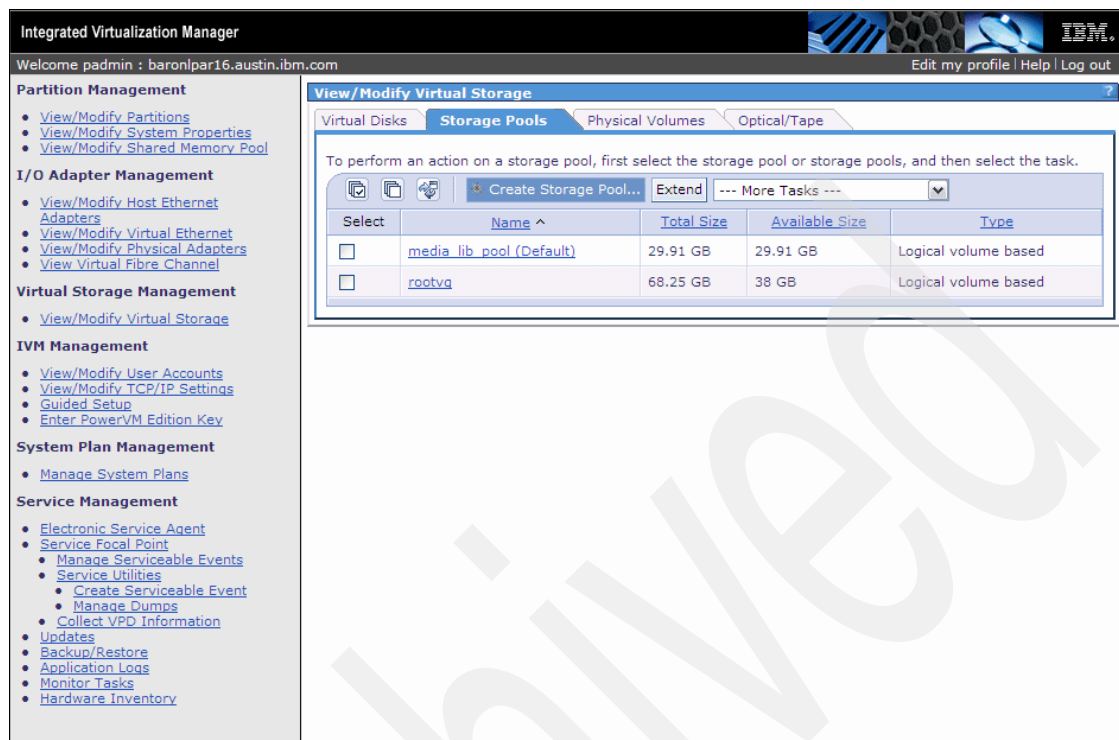


Figure 5-2 Starting the storage pool creation

- The next window, shown in Figure 5-3 on page 181, prompts for the name of the storage pool. The name must be a valid name for volume groups, for example no spaces are allowed and the name cannot exceed 15 characters. Select a storage pool type of **Logical volume based**. Next, select the physical volume or volumes to create the pool. When you are done, click **OK** to complete the storage pool creation process.

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 type="checkbox"/>	hdisk9	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010401F000000
<input type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104006000000
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T2-W500507630E85FE3F-L40104007000000
<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T2-W500507630E85FE3F-L40104008000000
<input type="checkbox"/>	hdisk5	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T2-W500507630E85FE3F-L4010400D000000
<input checked="" type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400E000000

OK Cancel

Figure 5-3 Naming the storage pool and selecting backing devices

Figure 5-4 on page 182 shows the newly created storage pool.

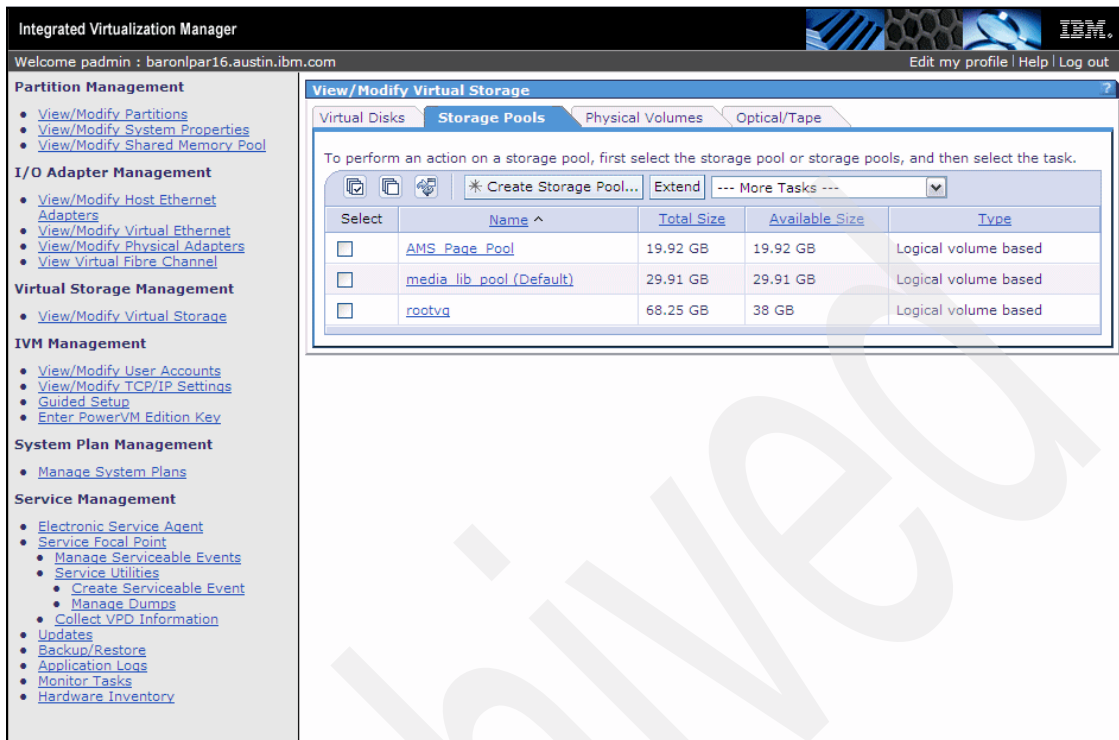


Figure 5-4 Storage pool list with new pool for paging added

Note: A new designation of Paging is added to the name field of the storage pool list when the shared memory pool is created.

With the paging storage pool created, we are ready to define the shared memory pool. From the IVM GUI, click **View/Modify Shared Memory Pool**.

Figure 5-5 on page 183 shows the shared memory pool configuration page. The first items to note are the current memory available and the reserved firmware memory values. The amount of memory to be configured for the shared memory pool cannot exceed the difference between these two values.

Note: Reserved firmware memory increases as additional LPARs, adapters, and others, are configured. If the shared memory pool is configured near to the maximum allowed value, you might have to reduce the pool size in the future to create additional LPARs or create/additional adapters.

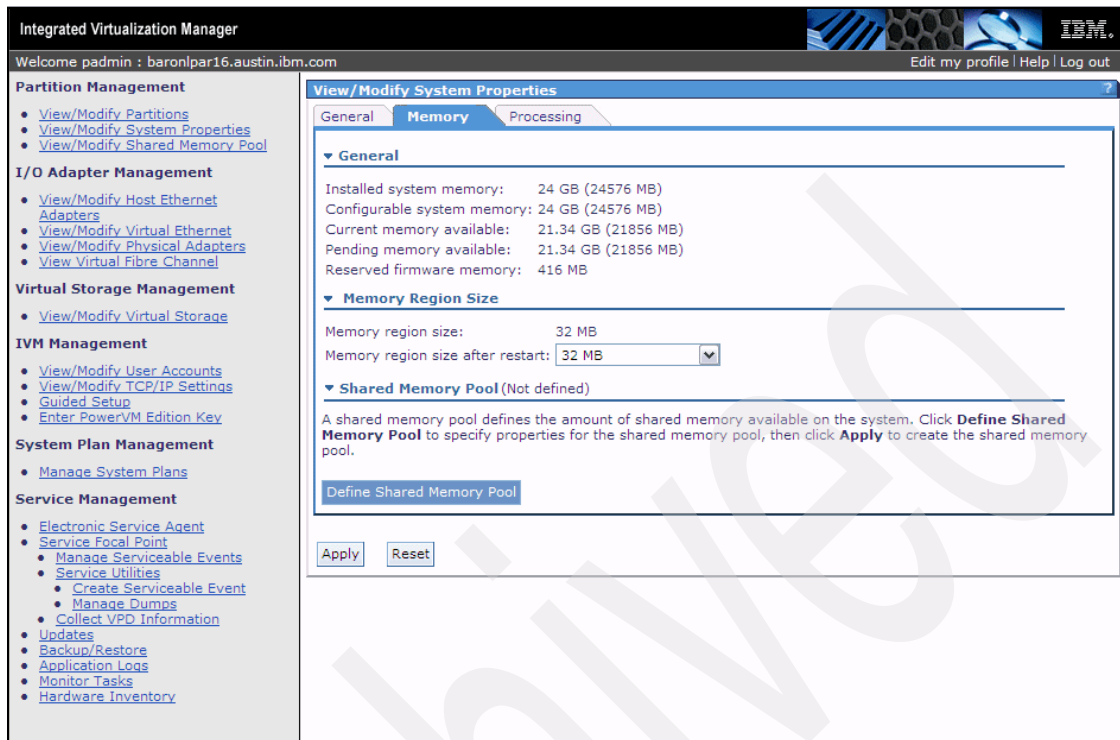


Figure 5-5 Defining a shared memory pool

- Click **Define Shared Memory Pool**. The dialog opens where you input the desired memory pool size and the storage pool to be used for the paging storage pool. After entering the information and selecting values from the drop-down box, as shown in Figure 5-6 on page 184, click **OK**.

Note: When IVM creates the shared memory pool, the value provided for the assigned memory of the pool will also be used for the maximum value of the pool.

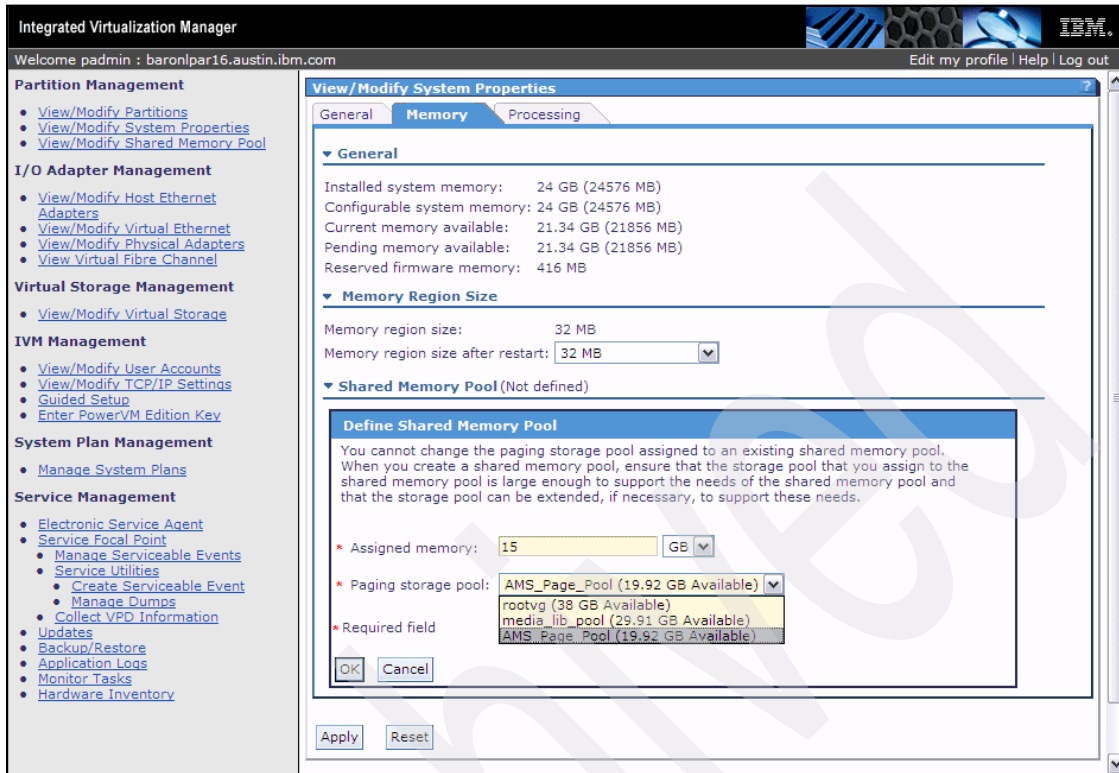


Figure 5-6 Shared memory pool configuration values

5. The window refreshes and indicates the shared memory pool has been defined as shown in Figure 5-7 on page 185. Click **OK**.

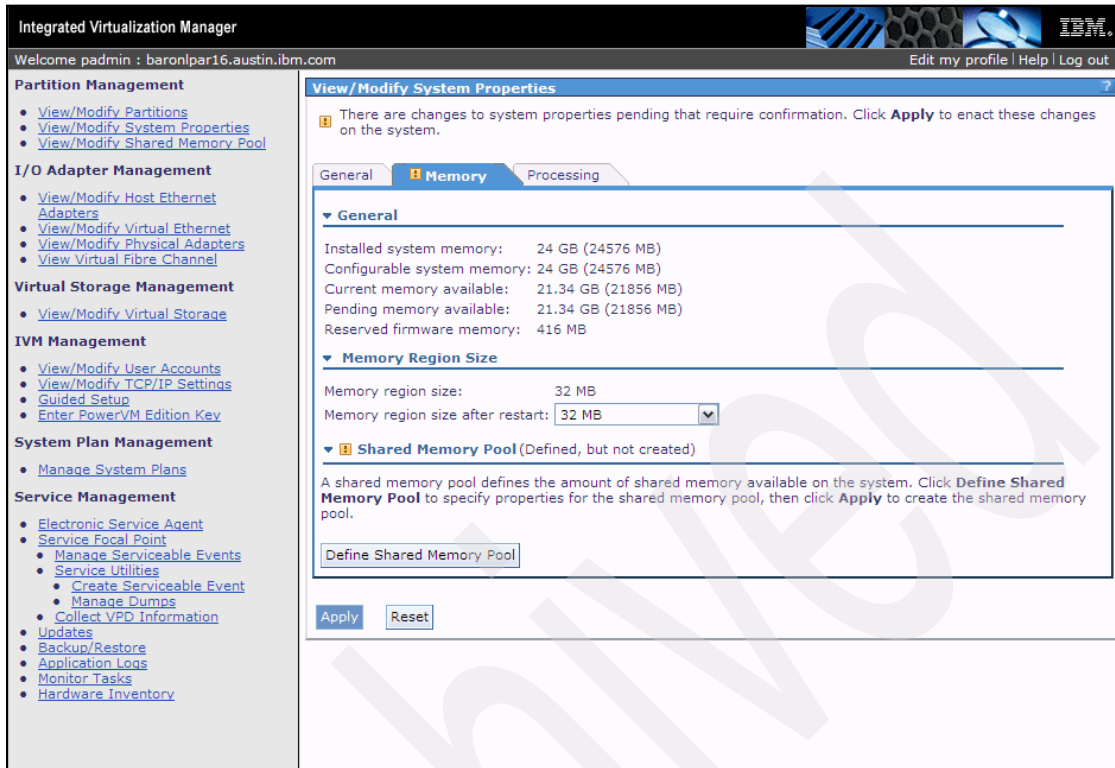


Figure 5-7 Shared memory pool defined state

- Click **Apply** to create the shared memory pool and the assignment of the paging storage pool as shown in Figure 5-8 on page 186.

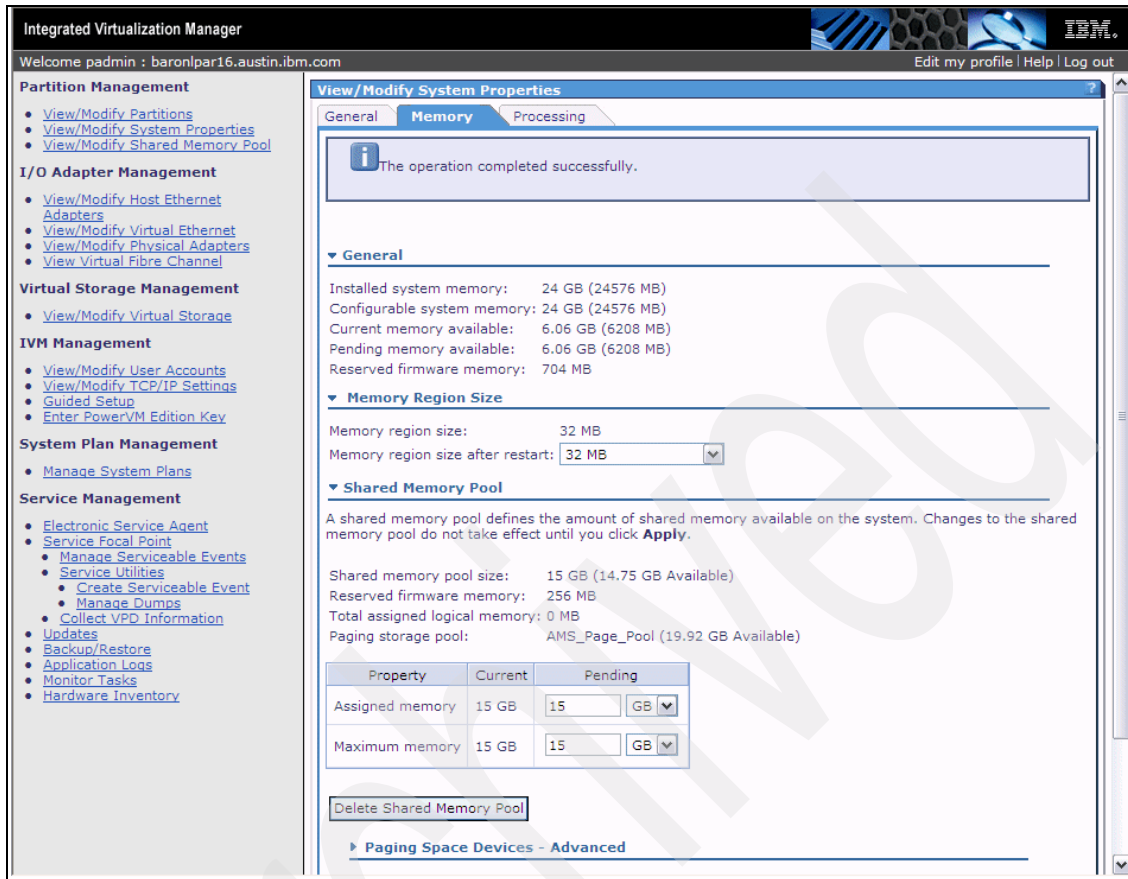


Figure 5-8 Shared memory pool information after creation

Now that we have created a shared memory pool we can create LPARs that use shared memory. As these LPARs are created, Active Memory Sharing subdivides the paging storage pool through the use of logical volumes to accommodate each LPAR.

A good practice, however, is to provide dedicated physical devices for each LPAR using shared memory as hypervisor paging devices. The next section describes how these dedicated paging devices are created.

5.2.2 Creating dedicated paging devices for partitions

To create dedicated physical devices for LPAR paging space:

1. Select **View/Modify Shared Memory Pool** from the IVM navigation area.
2. Click **Paging Space Devices - Advanced** to expand the section, as shown in Figure 5-9.

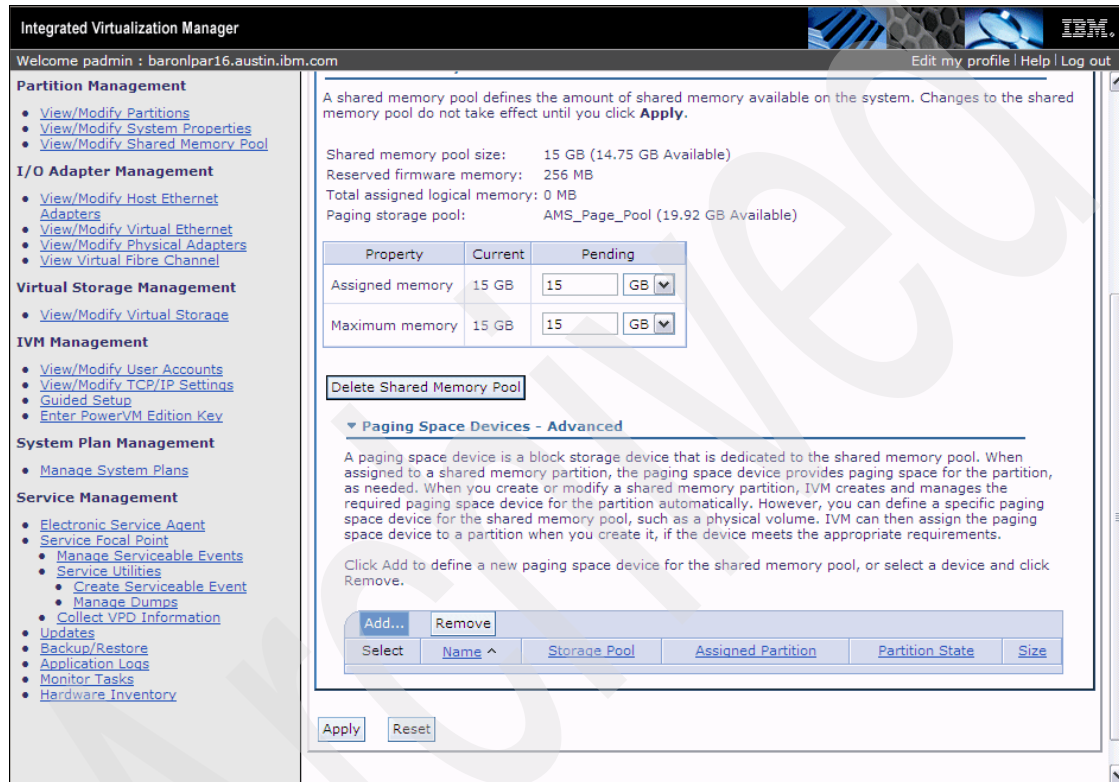


Figure 5-9 Creating dedicated paging devices for LPARS using shared memory

3. Click **Add** to display the devices that are available for selection. Select a physical volume device, and then click **OK**, as shown in Figure 5-10 on page 188.

Note: In the list of available devices, virtual disks (if available) are also displayed with physical volumes.

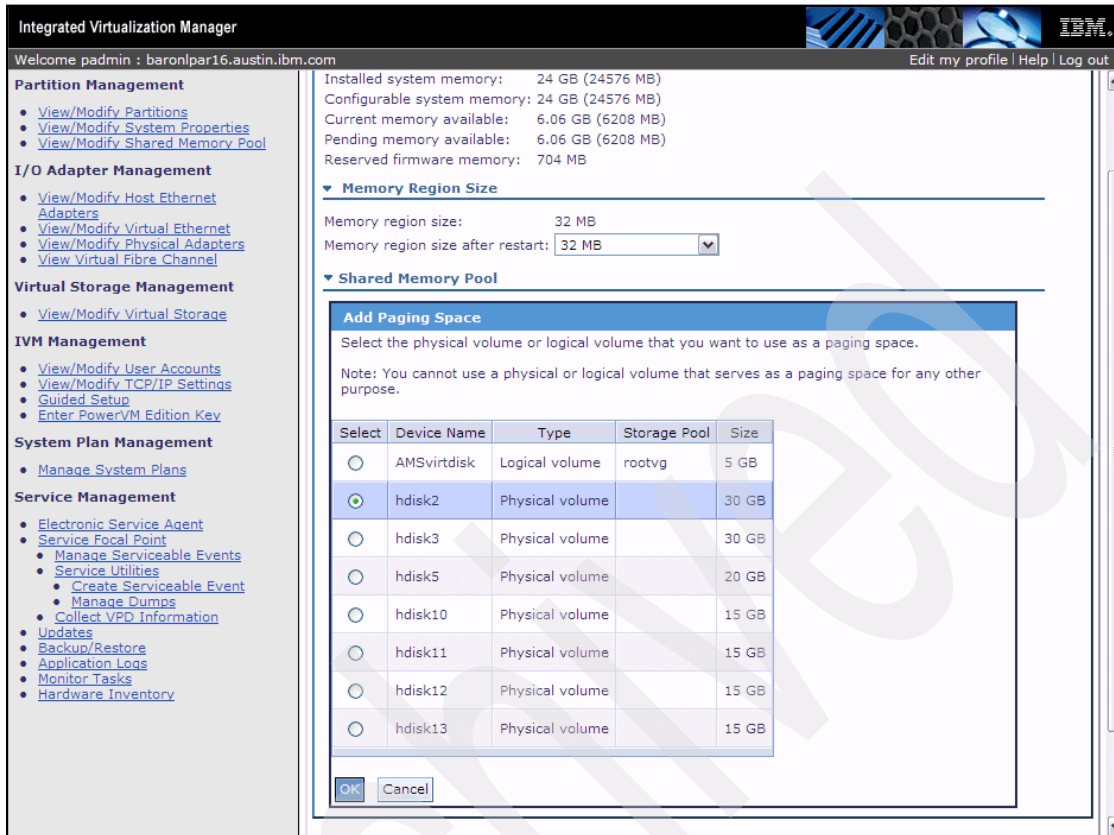


Figure 5-10 Dedicated device selection for share memory use

- Figure 5-11 on page 189 shows the selected device now defined as a paging device. Click **Apply** to complete the process.

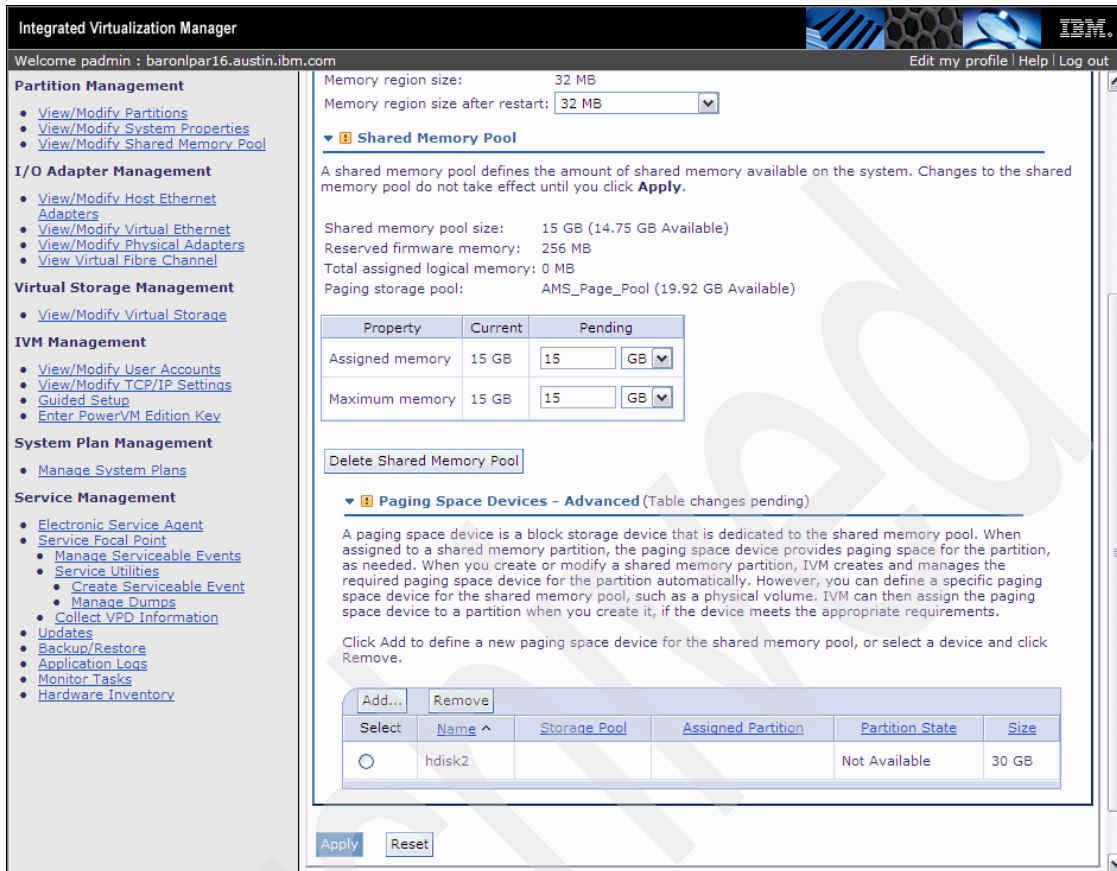


Figure 5-11 Dedicated device defined to paging space devices

Note: As LPARs are created that use shared memory, they are assigned to the smallest dedicated device available that meets memory size requirement.

5.2.3 Creating shared memory LPARs

Creating LPARs that use shared memory instead of dedicated memory uses the same wizard and process as detailed in 4.4.2, “Verify and modify VIOS partition memory and processors” on page 89.

However, several of the options that were not selectable prior to the creation of a shared memory pool can now be selected. Also, options such as selecting logical

ports from an HEA, dedicated processors, and physical adapters are no longer available.

To create shared memory LPARs:

1. Start the LPAR wizard by clicking the **View/Modify Partitions** link on the IVM GUI, and then clicking the **Create Partition** button. Figure 5-12 shows the first window of the wizard where the partition ID, partition name, and operating system environment are set. Enter the required information and click **Next**.

The screenshot shows a web-based wizard window titled "Create Partition: Name" with a sub-header "Step 1 of 8". On the left is a vertical navigation menu with links: Name, Memory, Processors, Ethernet, Storage Type, Storage, Optical/Tape, and Summary. The "Name" link is selected. The main content area is titled "Name" and contains the instruction "To create a partition complete the following information." Below this are four fields: "System name" (pre-filled with "Server-7778-62X-SN10180EA"), "Partition ID" (text box with "3"), "Partition name" (text box with "Sharedmemlp3" and a red asterisk indicating it is a required field), and "Environment" (dropdown menu showing "AIX or Linux"). At the bottom of the window are buttons for "< Back", "Next >", "Finish", "Cancel", and a "Help" link.

Figure 5-12 Creating a shared memory partition name

2. In the Create Partition Memory window shown in Figure 5-13 on page 191, select the now-available **Shared** option. The window refreshes with shared mode memory information and an area in which to enter the amount of memory. The information displayed shows the total amount of shared memory in the pool and the cumulative total of assigned logical memory. The first shared memory LPAR created indicates 0 MB in this field. Enter the desired amount of logical memory in the box and click **Next**.

Note: When creating a partition in shared or dedicated mode by using the Create Partition wizard, the memory amount you enter will be used for the assigned and maximum values.

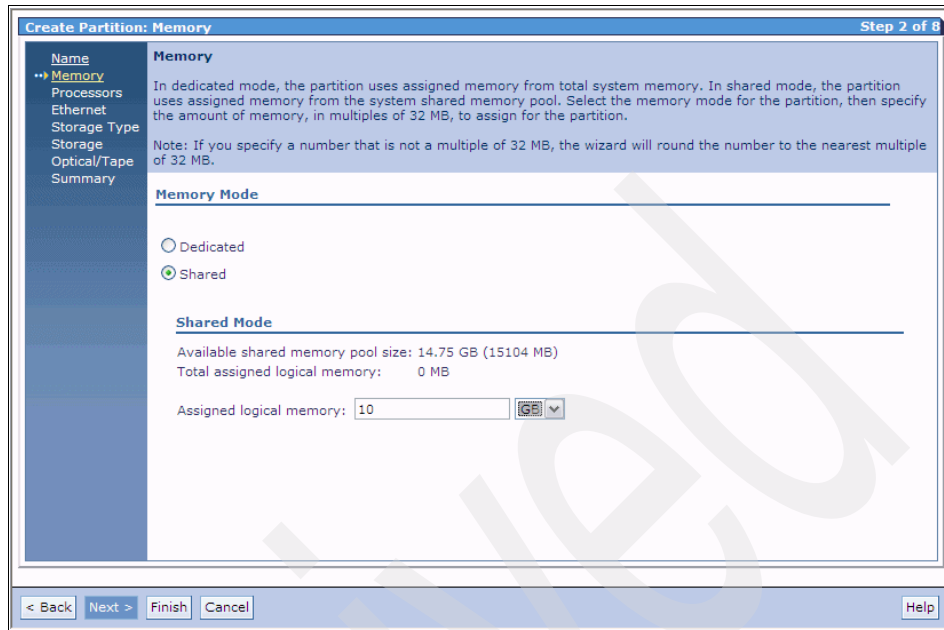


Figure 5-13 Selecting memory mode and amount for a shared memory partition

3. The next step is to select of the number of shared (virtual) processors, as shown in Figure 5-14 on page 192. Notice that the dedicated processor option cannot be selected. Use the drop-down box to select the number of assigned processors and then click **Next**.

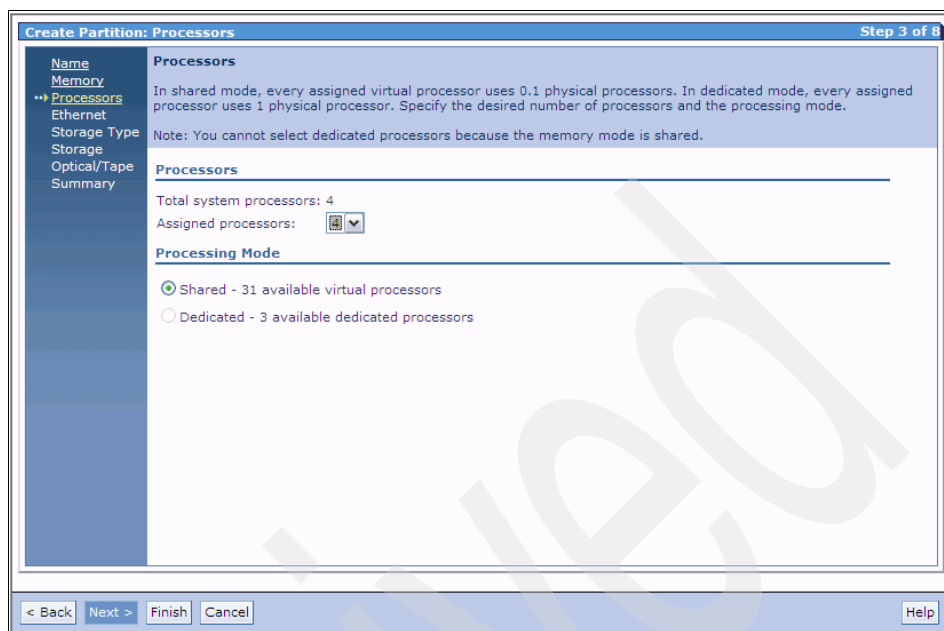


Figure 5-14 Selecting the number of processors in a shared memory partition

4. The next configuration step is Ethernet selection. As shown in Figure 5-15 on page 193, the only options are virtual Ethernet adapters. In this example we are using an existing Share Ethernet Adapter (SEA). Click **Next** to continue to the storage options.

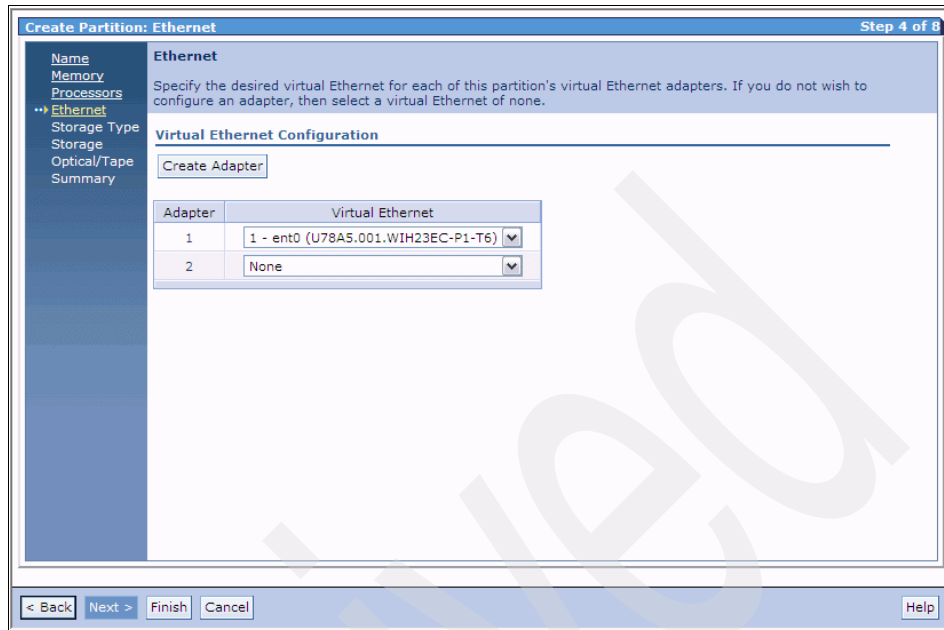


Figure 5-15 Ethernet selection for a shared memory partition

The storage selection options for a shared memory LPAR are the same as a dedicated memory LPAR. Virtual disks can be created from an existing storage pool. Existing virtual disks or physical volumes can be selected. Also, the None option is available if you do not want assigned storage at this time.

- As shown in Figure 5-16 on page 194, we chose the **Assign existing virtual disks and physical volumes** option. Click **Next** to continue to the selection window.

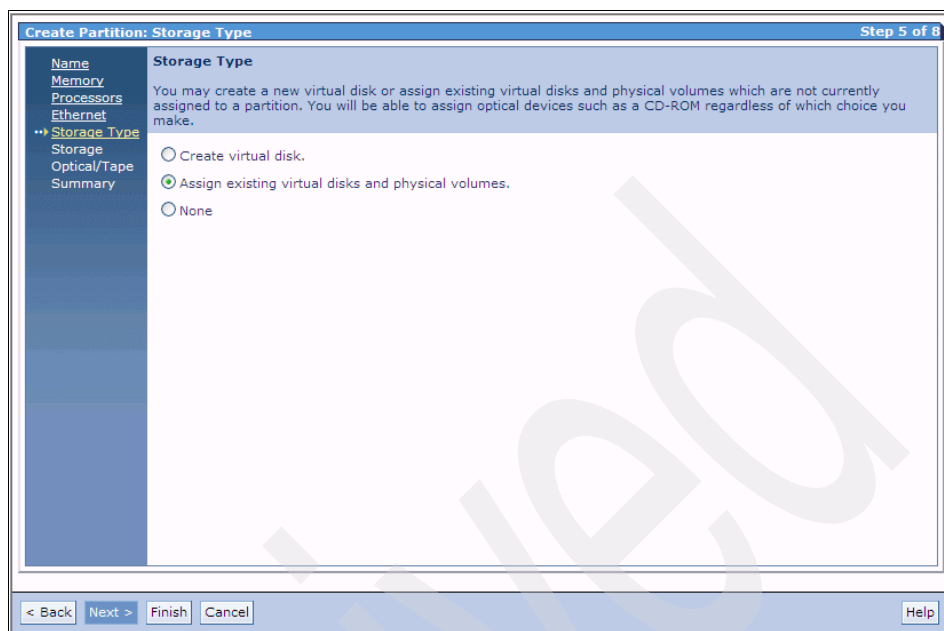


Figure 5-16 Storage selection for a shared memory partition

Figure 5-17 on page 195 shows the available selection of virtual disks (none in this example) and physical volumes that have not been assigned and are available.

In this example we selected **hdisk4**. Click **Next** to continue to the optical and tape options.

Create Partition: Storage

Step 6 of 8

Name

Memory

Processors

Ethernet

Storage Type

Storage

Optical/Tape

Summary

Storage

Select any number of physical volumes and virtual disks from the following lists of devices which are not currently assigned to a partition. You may use the Storage Management functions to change assignments at any time.

Available Virtual Disks

Select	Name ^	Storage Pool	Size

Available Physical Volumes

Select	Name ^	Size	Physical Location Code
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T2-W500507630E85FE3F-L40104007000000
<input checked="" type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T2-W500507630E85FE3F-L40104008000000
<input type="checkbox"/>	hdisk5	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T2-W500507630E85FE3F-L4010400D000000
<input type="checkbox"/>	hdisk9	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010401F000000
<input type="checkbox"/>	hdisk10	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104020000000
<input type="checkbox"/>	hdisk11	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104021000000
<input type="checkbox"/>	hdisk12	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104022000000
<input type="checkbox"/>	hdisk13	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104023000000

< Back

Next >

Finish

Cancel

Help

Figure 5-17 Storage selection for a shared memory partition

Two of the options shown in Figure 5-18 on page 196, physical optical devices and physical tape devices, will virtualize the physical device to the LPAR through the VIOS. Selecting these options does not imply a direct physical connection from the LPAR being created to the device. The virtual optical device is selected by default and can have media from the virtual media library assigned at this time.

Create Partition: Optical/Tape Step 7 of 8

Name
Memory
Processors
Ethernet
Storage Type
Storage
Optical/Tape
Summary

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

► **Physical Optical Devices**

▼ **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 type="checkbox"/>	Unknown1	None Modify		

[Create Device](#)

► **Physical Tape Devices** (No devices)

< Back Next > Finish Cancel Help

Figure 5-18 Optical and tape selections for a shared memory partition

- Click **Next** to continue to the Summary page.

The summary page as shown in Figure 5-19 on page 197 lists all of the selections made when stepping through the Create Partition wizard.

If you want to revise any of your choices, click **Back**.

After reviewing the selections, click **Finish** to complete the creation of the shared memory partition.

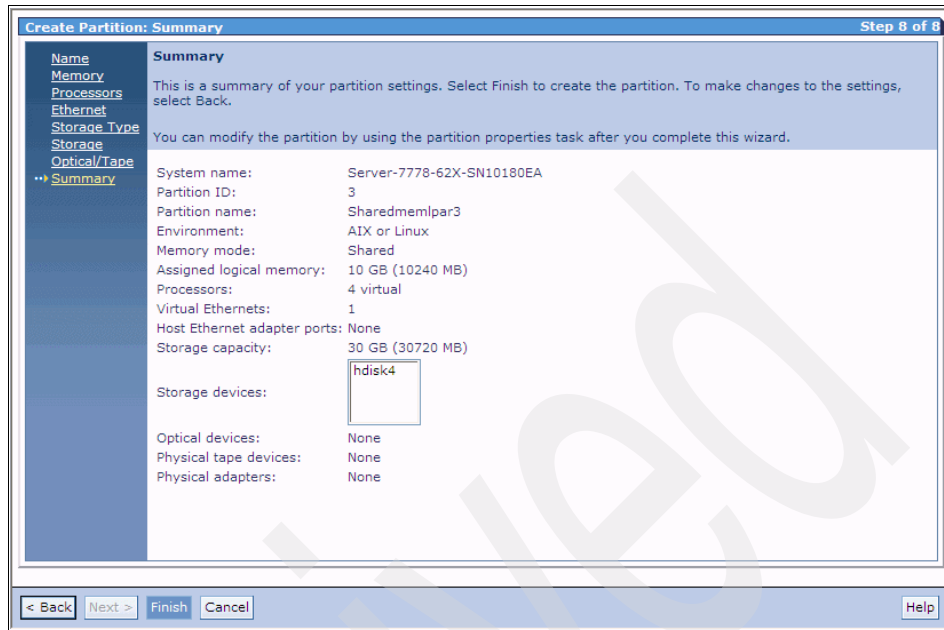


Figure 5-19 Summary of selections for a shared memory partition

Figure 5-20 on page 198 shows the View/Modify Partitions view with the new shared memory partition.

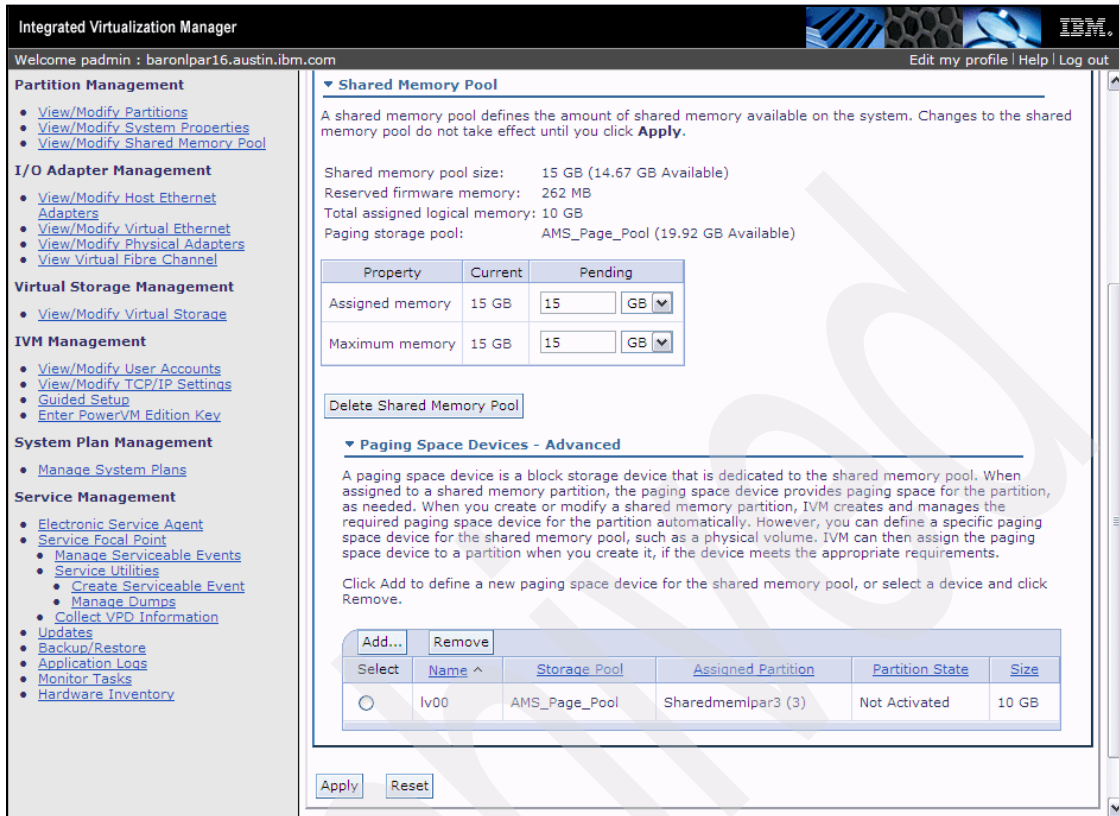


Figure 5-21 shared memory pool with paging space assignments in paging pool

5.2.4 Shared memory partitions and dedicated paging devices

During the creation of the shared memory pool, you have the option to create dedicated paging devices for shared memory partitions as detailed in Figure 5.2.2 on page 187.

By default, these dedicated devices, if available, are assigned to a shared memory partition (if of adequate size) when it is created. If the available dedicated devices have different sizes, the smallest size device that meets the requirements of the assigned logical memory will be used. If the devices are all of equal size, the first available device will be assigned.

Figure 5-22 on page 200 shows the availability of two dedicated paging devices with sizes of 20 GB and 30 GB. These dedicated devices were added after the creation of the shared memory partition Sharedmemlpar3.

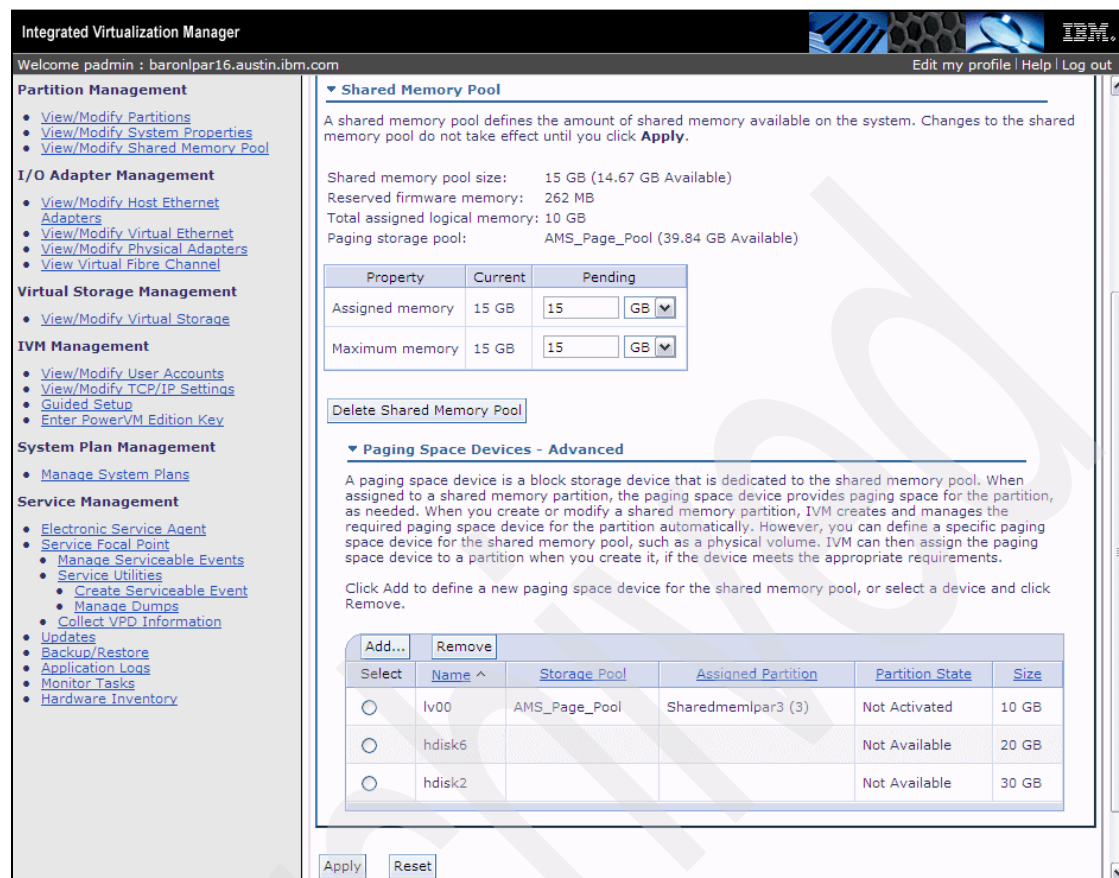


Figure 5-22 Shared memory pool view showing both types of paging devices

A new shared memory partition, Sharedmemlpar4, was created with a logical memory value of 25 GB. Figure 5-23 on page 201 shows this new partition and the assignment of hdisk2 as its dedicated paging device. Although the paging storage pool had over 39 GB available, the default is to use dedicated paging devices when available. In this case, the available hdisk2 with a size of 30 GB was assigned to the partition Sharedmemlpar4.

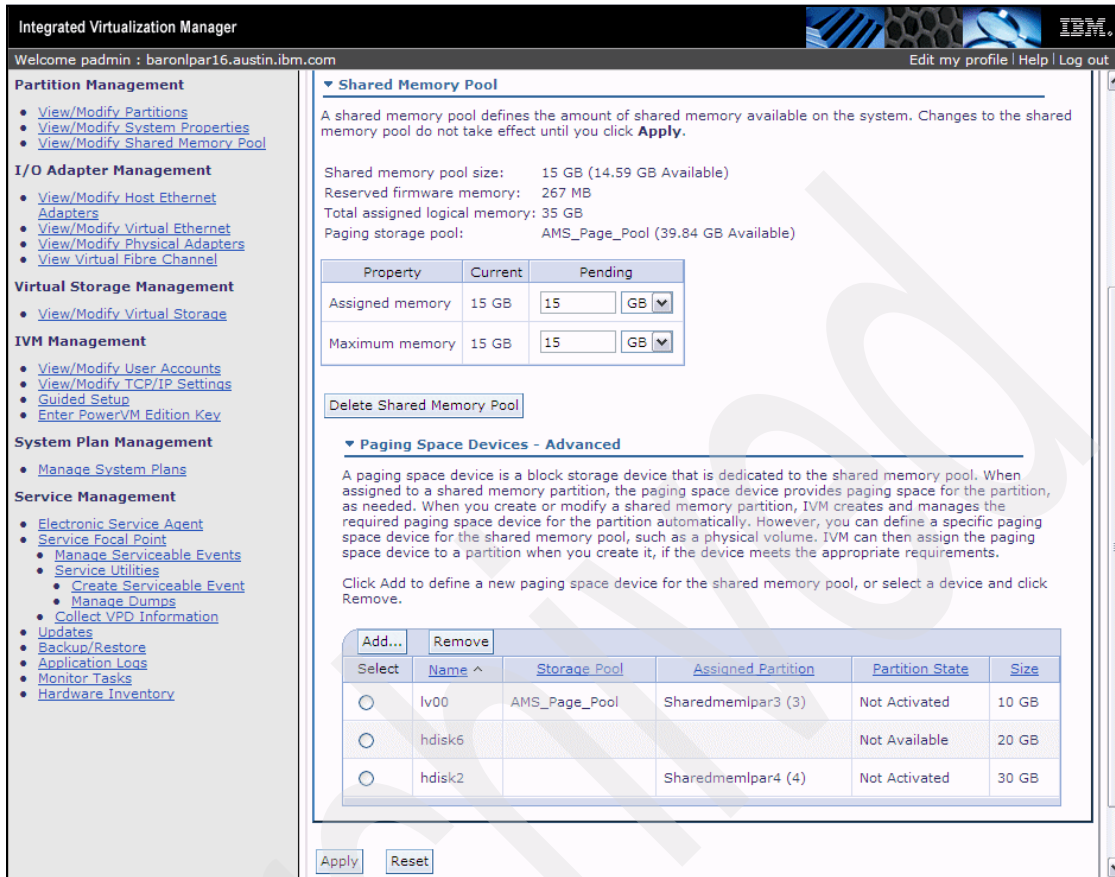


Figure 5-23 Shared memory pool view showing assigned dedicated paging device

Changing the maximum memory values of a shared memory partition can also cause a change from a paging pool logical volume to a dedicated paging device. Figure 5-24 on page 202 shows the inactive partition Sharedmemlpar3 that had its maximum memory value changed from 10 GB to 15 GB. When this change was made the paging space changed from 10 GB lv00 in the pool AMS_Page_Pool to the 30 GB hdisk6. Also note the informational message indicating current and pending values are not synchronized. A partition activation will complete the process.

Partition Properties: Sharedmem1par3 (3)

General **Memory** Processing Ethernet Storage Optical/Tape Devices

Warning: Current and pending values are not synchronized. [\[details\]](#)

Modify the settings by changing the pending values. The changes will be applied immediately; however, synchronizing the current and pending values might take some time.

Memory mode: **Shared**

Paging space: **hdisk6 (20 GB)**
Paging space storage pool: **None**
I/O entitled memory: **Auto (77 MB)**

All memory values should be in multiples of 32 MB.

Property	Current	Pending
Minimum memory	256 MB	256 MB
Assigned memory	10 GB (10240 MB)	10 GB
Maximum memory	15 GB (15360 MB)	15 GB
Memory weight	Medium - 128	Medium - 128

OK Cancel

Figure 5-24 Partition memory properties showing maximum memory and paging space changes

Figure 5-25 on page 203 shows the shared memory pool indicating the changes to the paging device used for the partition Sharedmem1par3 when the maximum memory values were changed.

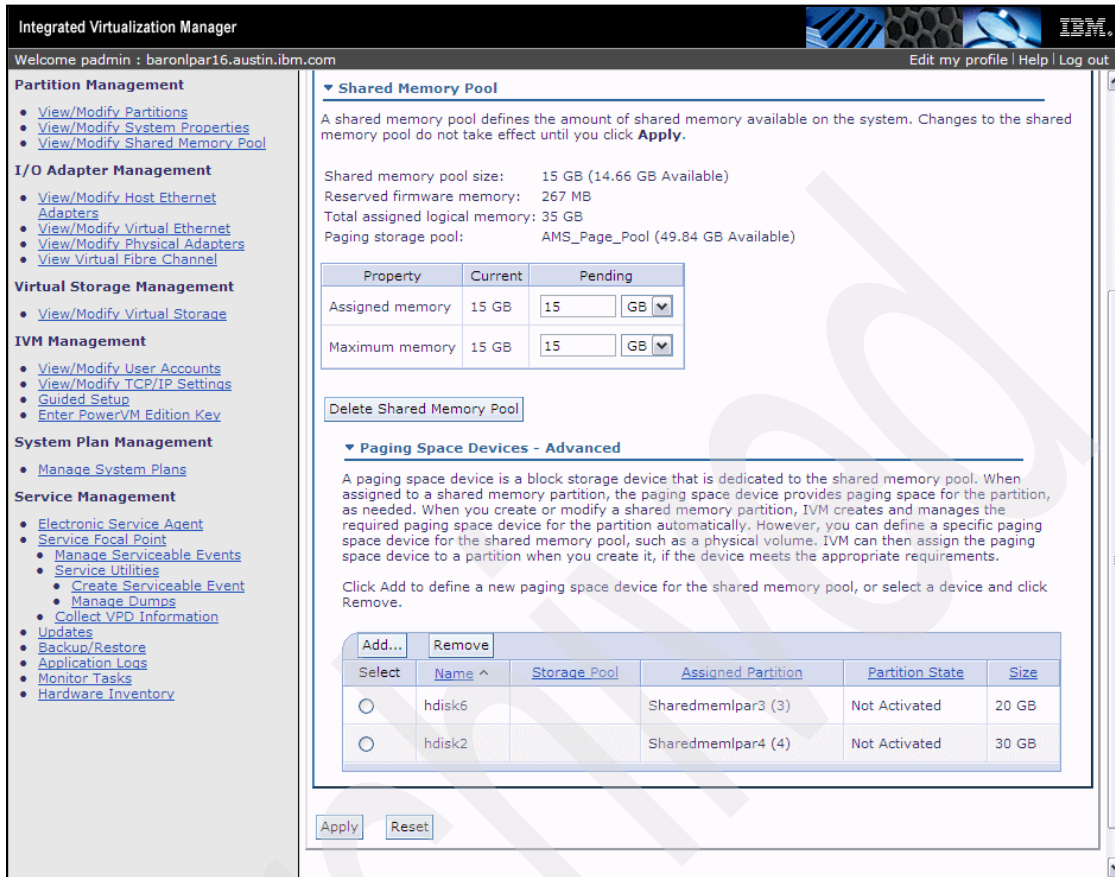


Figure 5-25 shared memory pool after partition maximum memory values changed

5.2.5 Active Memory Sharing DLPAR operations

Dynamic logical partition (DLPAR) operations can be performed on both the shared memory pool and shared memory partition logical memory assignments. The assigned memory in shared memory pool can be logically partitioned dynamically up to its maximum value, and the memory pool maximum value can be dynamically increased up to the available limits of the physical memory minus firmware requirements. The shared memory partition can be altered between the minimum and maximum values, as shown in the partition properties under the Memory tab.

5.3 Active Memory Sharing summary

Active Memory Sharing provides the ability to better utilize the memory *and* CPU resources available on a IBM BladeCenter JS23 or JS43. However, the successful implementation requires a complete understanding of current or planned workloads and the proper matching of those workloads in the right combinations. Improper matching will result in contention for memory resources and excessive paging by the VIO Server in an attempt to service the partitions memory requirements.

Using dedicated paging devices is the best method for providing paging space for the shared memory pool. If a mix of dedicated and logical volumes are used additional planning is required to determine when to add the dedicated devices in the sequence of creating shared memory partitions or changing existing dedicated memory partitions.

IBM AIX V6.1 installation

IBM AIX can be installed as a native operating system on IBM BladeCenter JS23 and JS43 Express or in a client partition of IBM PowerVM. This chapter describes details of installation on a logical partition.

This chapter contains the following topics:

- ▶ “Introduction to installing in PowerVM client partition” on page 206
- ▶ “Creating a virtual media library” on page 206
- ▶ “Preparing the PowerVM client partition” on page 212
- ▶ “Installing AIX 6.1 in a logical partition of VIOS” on page 225

6.1 Introduction to installing in PowerVM client partition

This chapter assumes that you have already installed VIOS 1.5.2.1 or any later version (latest version at the time of writing this book is V2.1.1) on the blade and performed the initial configuration. If this was not done, go to 4.2, “VIOS system management using IVM” on page 77.

To install IBM AIX 6.1 in a client partition, you must first create the client partition with the IVM before you can begin the installation of AIX. The remaining sections of this chapter describe how to accomplish the installation.

6.2 Creating a virtual media library

A virtual media library is a feature that was introduced in the Virtual IO Server version 1.5. It allows the storage of images from CDs and DVDs in a logical volume. These images can be mounted to virtual optical devices that are assigned to partitions.

A media library is created for the AIX installation DVD that is used to install the first AIX partition. This section describes how to create a storage pool that will be used to create the logical volume for the media library. After, we describe how to add a DVD image, from the AIX DVD, to the media library created.

To create a media library:

1. Click **View/Modify Virtual Storage** in the left menu under Virtual Storage Management. Then, click the **Storage Pools** register card. Create a new storage pool by clicking **Create Storage Pool**. See Figure 6-1.

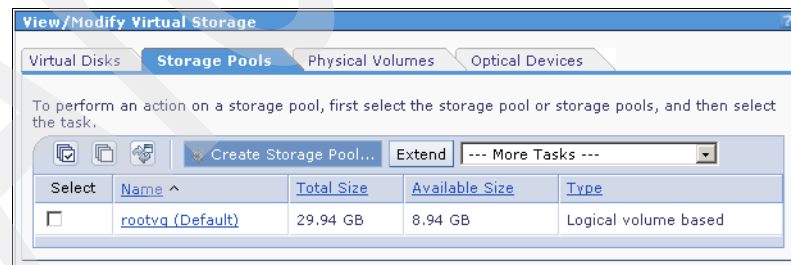


Figure 6-1 Media library - create storage pool

2. Specify the name of the storage pool and select the physical volumes that will be assigned to this storage pool. Figure 6-2 on page 207 shows that we used the name STG-Pool-Media1. The type of the volume group is *logical volume*

base. This step allows the space of the media library to be increased when necessary. Physical volume `hdisk3` is assigned to this pool. Click **OK**.

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 type="checkbox"/>	hdisk2	20 GB	U78A5.001.WIH01AA-P1-C6-T1-W204300A0B811A662-L20000000000000
<input checked="" type="checkbox"/>	hdisk3	20 GB	U78A5.001.WIH01AA-P1-C6-T1-W204300A0B811A662-L30000000000000
<input type="checkbox"/>	hdisk1	30 GB	U78A5.001.WIH01AA-P1-C6-T1-W204300A0B811A662-L10000000000000

* Required field

OK Cancel

Done

Figure 6-2 Media library - select the physical volume

3. The storage pool was created. Now, select the **Optical Devices** register card. See Figure 6-3.

View/Modify Virtual Storage

Virtual Disks **Storage Pools** Physical Volumes **Optical Devices**

To perform an action on a storage pool, first select the storage pool or storage pools, and then select the task.

Create Storage Pool... Extend --- More Tasks ---

Select	Name ^	Total Size	Available Size	Type	Parent Storage Pool
<input type="checkbox"/>	rootvg (Default)	29.94 GB	8.94 GB	Logical volume based	
<input type="checkbox"/>	STG-Pool-Media1	19.92 GB	15.03 GB	Logical volume based	

Figure 6-3 Media library - new storage pool

4. Click **Create Library**. See Figure 6-4.

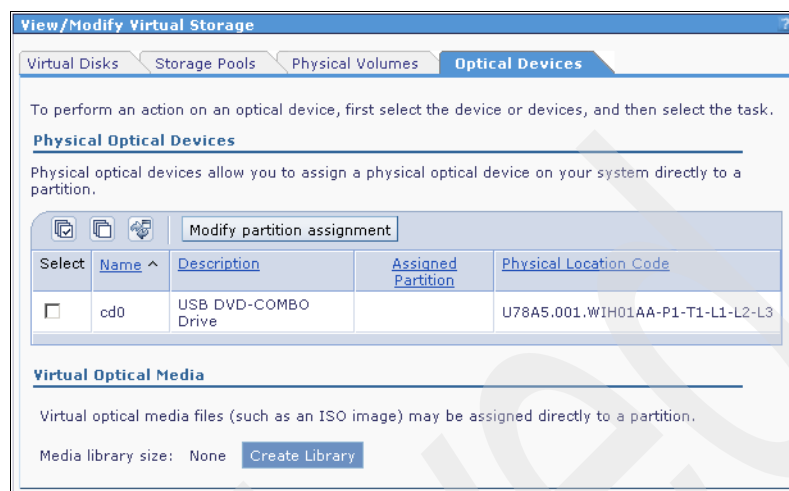


Figure 6-4 Media library - create library

5. Specify the storage pool that will contain the logical volume with the media library and the size of the media library. We used the volume group created in step 1 on page 206. The initial size was set to hold the AIX 6.1 DVD with a size of approximately 3.6 GB. See Figure 6-5. Click **OK**.




Figure 6-5 Media library - specify library size and volume group

6. Creating the library volume and file takes several moments. After that is done, return to the panel shown in Figure 6-6 on page 209. Click **Add Media** to create an image from the AIX DVD.



Figure 6-6 Media library - add media

7. The add media dialog starts and guides you through the process of adding media to the library. Click **Import from physical optical device** to get the list of available physical optical devices that you can use to import the media. Specify the media type of the new media. Choose between read-only and read/write mode of the new media. Now, enter the name of the new media. This name will be used as the file name to store the content from the CD or DVD. Click **OK** to start the copy process.

Note: Do not use a media name that contains spaces in the name. This can lead to an error message like the one shown in Figure 6-7.

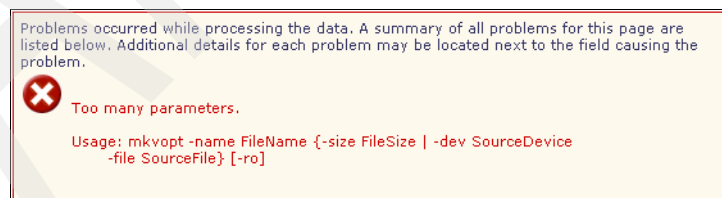


Figure 6-7 Error message using a wrong media name

You may look at existing media files in `/var/vio/VMLibrary`. The last step on this page is the specification of the optical device that contains the CD or DVD

to copy into the library. Figure 6-8 shows the optical device that is located in the media tray of the IBM BladeCenter H chassis. The remote media optical device uses the location code U78A5.001.WIH01AA-P1-T1-L1-L1.

We used the internal optical device of the BladeCenter chassis to copy the data from the IBM AIX 6.1 DVD. This took approximately two hours.

Add Media

You may upload an optical media file from your local workstation, specify an existing file in your home directory, import from a physical optical device, or create a blank media file. Depending on the size of the file, these operations may take several minutes.

☐ Upload media
☐ Add existing file
☒ Import from physical optical device
☐ Create blank media

Media type: Read only

Media name: AIX-6.1

Select an available optical device that contains the media you wish to import, and select OK. This may take some time while the media is copied into the library.

Select	Name	Description	Physical Location Code
<input checked="" type="checkbox"/>	cd0	USB DVD-COMBO Drive	U78A5.001.WIH01AA-P1-T1-L1-L2-L3

OK Cancel

Done

Figure 6-8 Media library - specify details about media creation

- The copy task takes some time. You may close the dialog browser window and proceed with other tasks in the meantime.

At any time, you can check whether the task has completed or failed by using the Monitor Task function. This function can be reached with **Monitor Task** before you close your browser window or from the main window's navigation by selecting **Service Management** → **Monitor Task**. See Figure 6-9 on page 211.

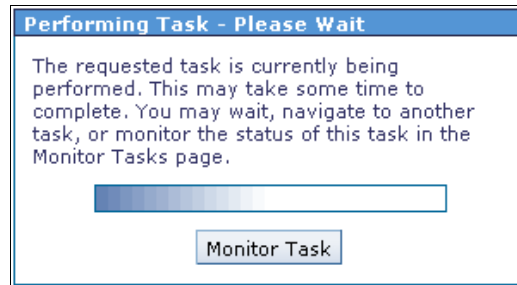


Figure 6-9 Media library - performing task

9. After closing the browser window of the Add Media dialog, you return to the view shown in Figure 6-10. The new media is already listed here. Clicking **Refresh** updates the size information during the copy operation.

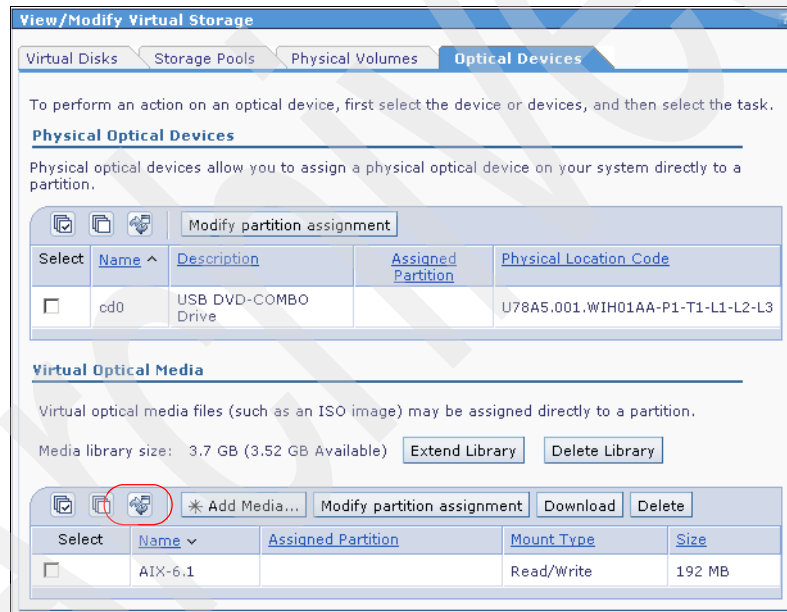


Figure 6-10 Media library - newly created media

The media library is now ready to be used. When the free space in the library is not sufficient for new CD images, expand the media library. This can be done at any time.

6.3 Preparing the PowerVM client partition

To create a client partition with the Integrated Virtualization Manager (IVM) of the Virtual I/O Server:

1. Use your preferred Web browser and enter the host name or IP address of the IVM. That is the address configured in 4.3.3, “Initial network setup” on page 81.

An IVM login window opens, shown in Figure 6-11. Use the default account that was created during setup when you had not yet created your own account. The default account is named padmin and uses the password padmin. Click **Login**.

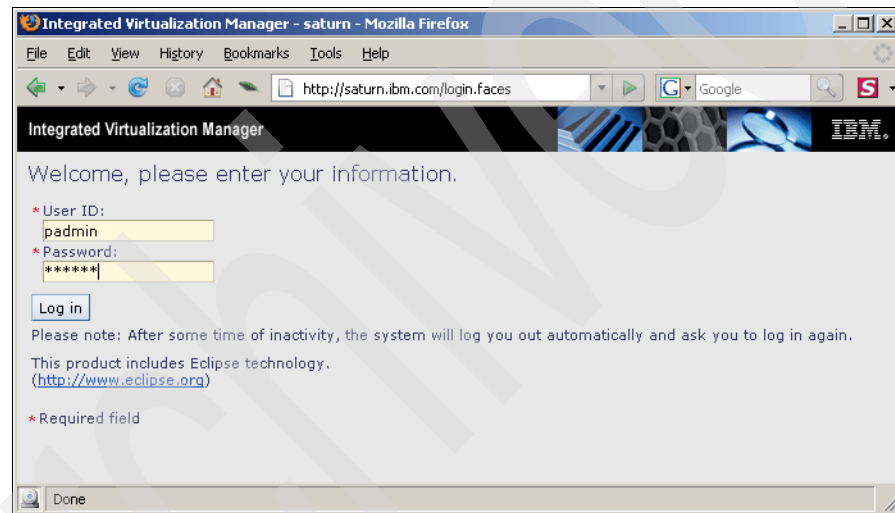


Figure 6-11 IVM login window

2. Depending on the setup of your IVM, you will be at the Guided Setup or on the View/Modify Partitions page. Figure 6-12 shows the usual page that you see after logging in when the IVM is fully configured.

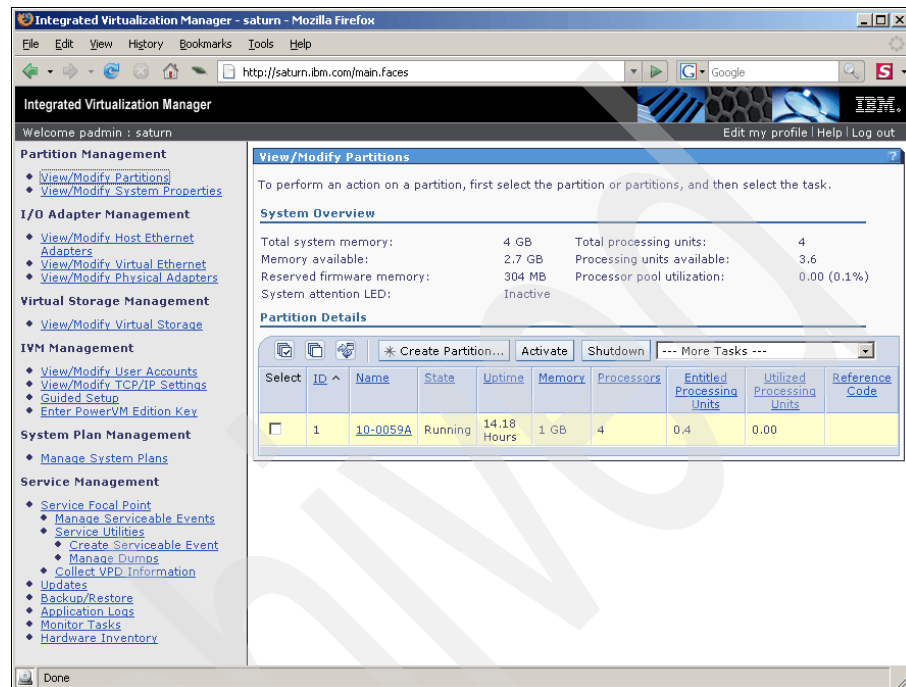


Figure 6-12 View/Modify Partitions page after login

3. Verify that your storage is available to the VIOS. Click **View/Modify Virtual Storage** in the left menu under Virtual Storage Management. See Figure 6-13.

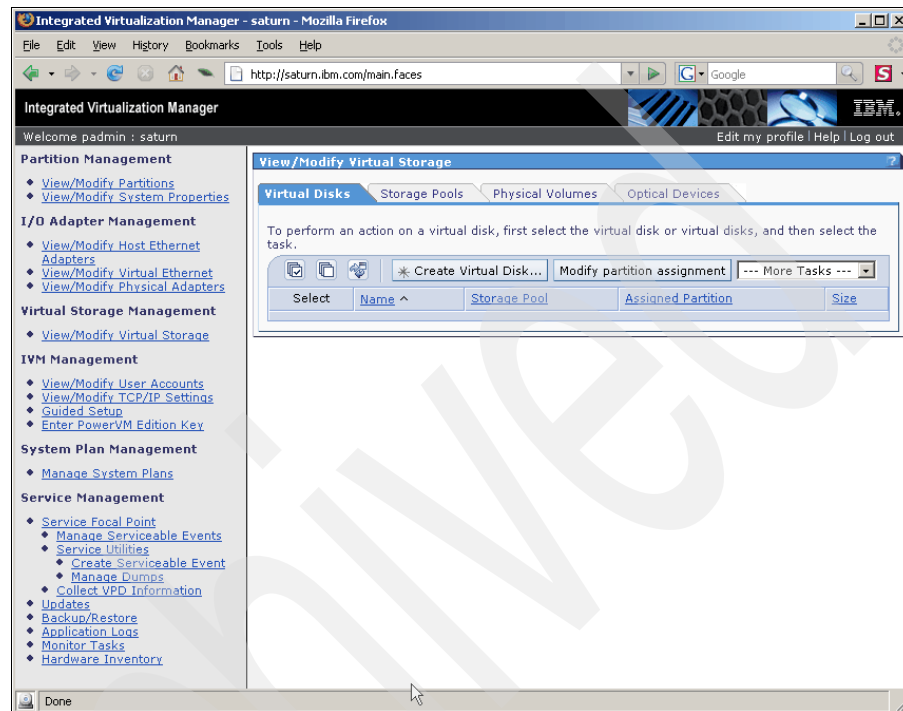


Figure 6-13 View/Modify Storage

- In the **View/Modify Virtual Storage** page, select the **Physical Volumes** tab to see a list of available hard drives to the VIOS. See Figure 6-14. Verify that the expected drives are available.

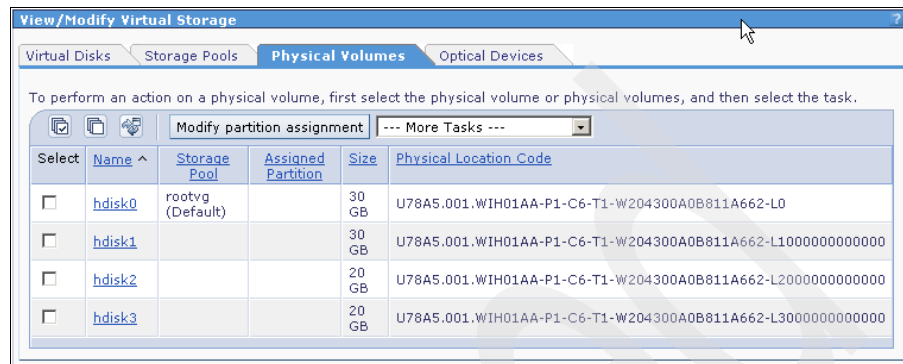


Figure 6-14 Available physical volumes

Figure 6-15 shows that four physical volumes are available. They are all located on a DS4800. The hdisk0 and hdisk1 are used for the VIOS itself. The hdisk2 will be used for AIX client partitions that will be created in the next steps.

Click **View/Modify Partitions** under Partition Management. Then click **Create Partition** as shown in Figure 6-15.

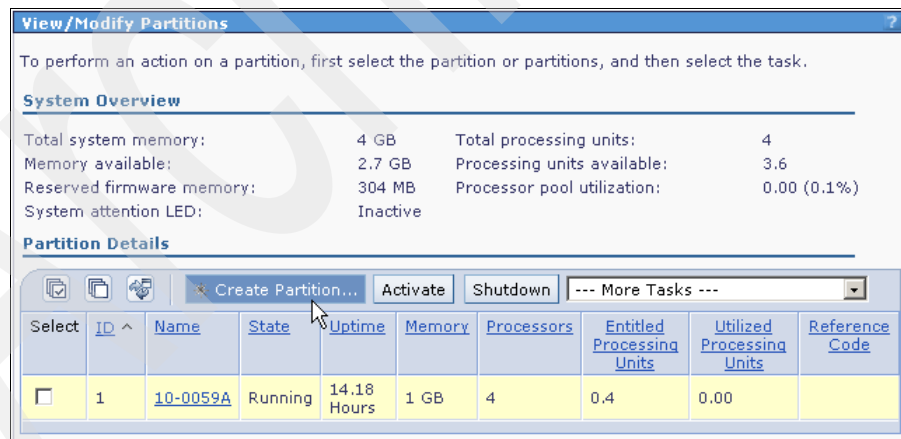


Figure 6-15 View/ Modify Partition - Create Partition

A dialog opens and guides you through the process of partition creation.

5. Specify the name and the type of the partition. The name is used to identify the partition, especially when partition mobility is later used. Using a host name might be an option here. In Figure 6-16, we chose the host name as partition name. The Environment (type) can be either **AIX or Linux**, or **i5/OS**. Choose the type according to the operating system you plan to install. We chose **AIX or Linux** for this AIX partition. Click **Next** to proceed.

http://saturn.ibm.com - Create Partition Wizard - Mozilla Firefox

Create Partition: Name Step 1 of 8

Memory
Processors
Ethernet
Storage Type
Storage
Optical
Summary

Name

To create a partition complete the following information.

System name: Server-7998-61X-SN100059A

Partition ID: 2

* Partition name: Mars

Environment: AIX or Linux

* Required field

< Back Next > Finish Cancel Help

Done

Figure 6-16 Create partition - define name

6. Define the amount of memory that will be assigned to the partition. In Figure 6-17 we chose 1 GB. Click **Next** to proceed.

http://saturn.ibm.com - Create Partition Wizard - Mozilla Firefox

Create Partition: Memory Step 2 of 8

Name

- Memory
- Processors
- Ethernet
- Storage Type
- Storage
- Optical
- Summary

Memory

Specify the amount of memory. Input should be in multiples of 16 MB.

Total system memory: 4 GB (4096 MB)
Current memory available for partition usage: 2.7 GB (2768 MB)

Assigned memory: 1 GB

< Back Next > Finish Cancel Help

Done

Figure 6-17 Create partition - configure the amount of memory

7. Choose the number of processors that will be used by the partition. You have to decide whether to use dedicated or shared processors. When a dedicated processor is used, no load can be moved to other currently free processors because this can lead to a performance issue. In Figure 6-18 you see that we configured two processors and shared processor mode. Click **Next**.

http://saturn.ibm.com - Create Partition Wizard - Mozilla Firefox

Step 3 of 8

Create Partition: Processors

Name
Memory
Processors
Ethernet
Storage Type
Storage
Optical
Summary

Processors

In shared mode, every assigned virtual processor uses 0.1 physical processors. In dedicated mode, every assigned processor uses 1 physical processor. Specify the desired number of processors and the processing mode.

Processors

Total system processors: 4
Assigned processors: 2

Processing Mode

☒ Shared - 36 available virtual processors
☐ Dedicated - 3 available dedicated processors

< Back Next > Finish Cancel Help

Done

Figure 6-18 Create partition - CPU configuration

8. Depending on the setup of the network in the BladeCenter chassis and in the Virtual IO Server, you might have different settings. When you define host Ethernet adapters to the partition, no Ethernet bridge in the Virtual IO Server is required. The disadvantage is that you are not able to use partition mobility. For partition mobility, the partition is required to be configured with virtual Ethernet adapters. The Ethernet bridge is configured in the Virtual IO Server. For more details, see 4.5.2, “Virtual Ethernet Adapters and SEA” on page 97.

As shown in Figure 6-19, we chose one virtual Ethernet adapter. Click **Next**.

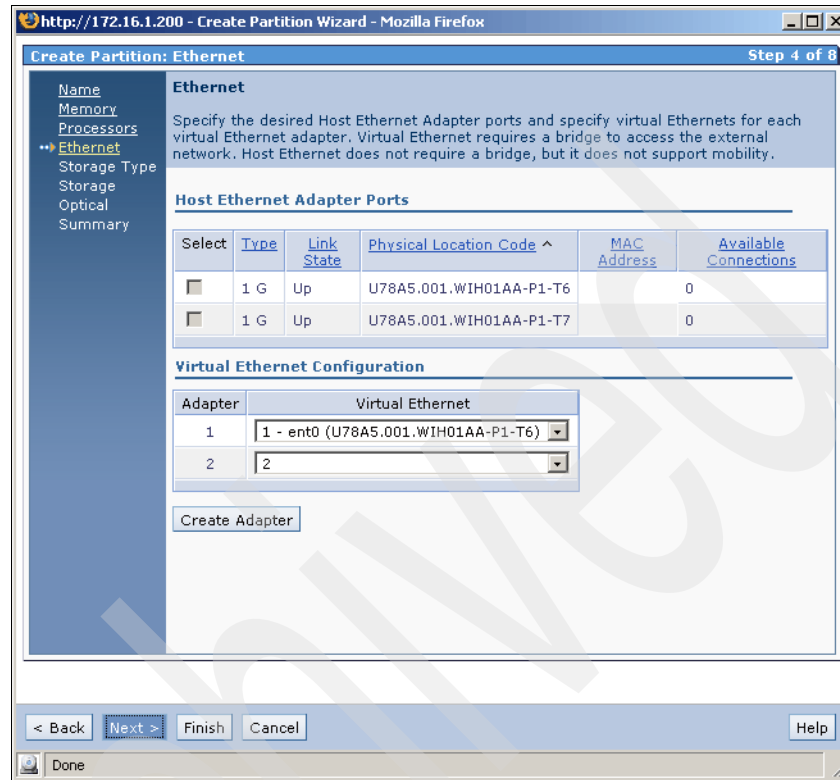


Figure 6-19 Create partition - Ethernet network

9. Set up the storage type you plan to use. Three options are available:
 - Volume group
 - File-based storage
 - Dedicated physical volume for the partition

Depending on the type of storage subsystem, there are limitations about the maximum number of physical volumes you can have per host adapter or host adapter group. In the case of the DS4800, you can have up to 32 logical drives from the DS4000 assigned to a host or host group. In our scenario, we used a dedicated physical volume for each partition.

As shown in Figure 6-20, select **Assign existing virtual disks and physical volumes**. Click **Next**.

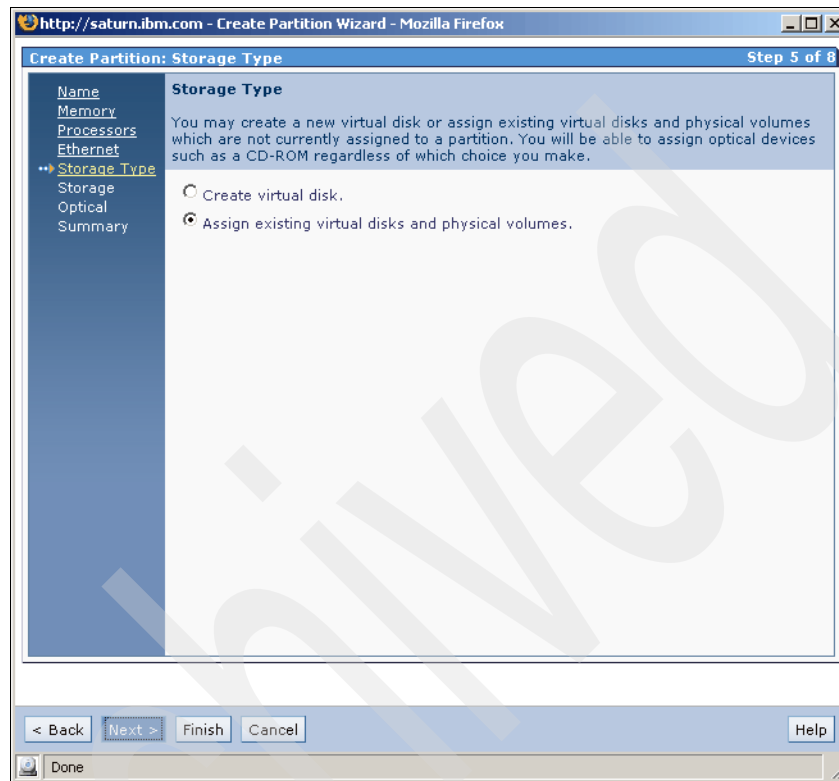


Figure 6-20 Create partition - storage type

10. Select the physical volume or volumes that must be available to the partition.
Figure 6-21 shows the section of hdisk1. Click **Next**.

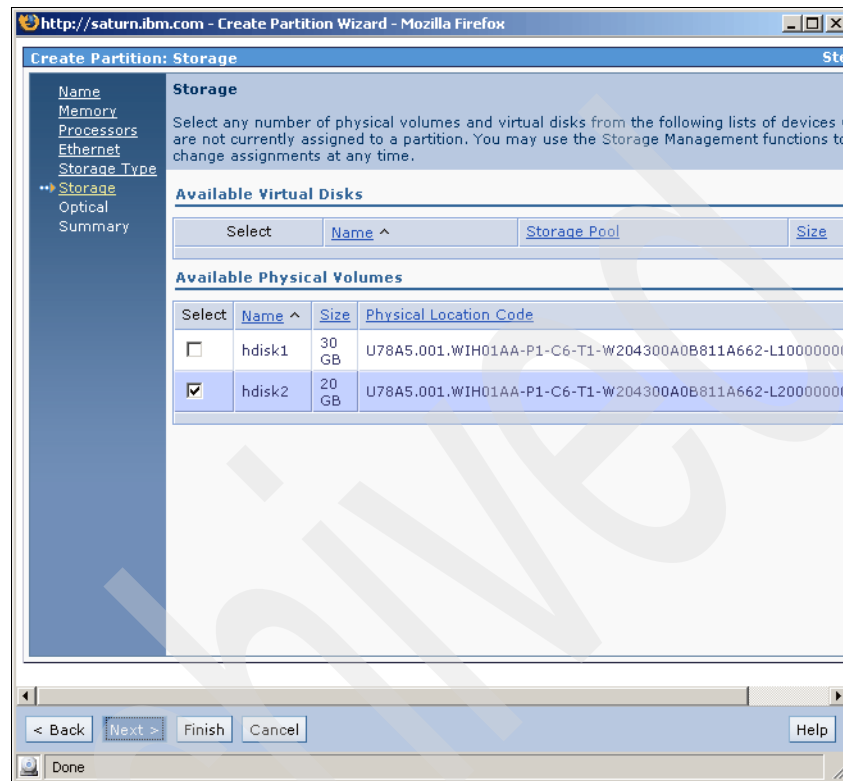


Figure 6-21 Create partition - select physical volumes

11. In the optical section of the partition creation process, you can define the CD-ROM drives that will be used by the partition. Two options are possible:

- Physical drive attached to the partition
- Virtual drive attached to the partition

Multiple physical CD-ROM drives might be available. Use the location code to differentiate between the CD-ROM drives:

U78A5.001.WIH01AA-P1-T1-L1-L2-L3 CD-ROM drive in the media tray
U78A5.001.WIH01AA-P1-T1-L1-L1 Remote media CD-ROM drive

The local drive installed in the media tray of the IBM BladeCenter chassis is identified by the location code U78A5.001.WIH01AA-P1-T1-L1-L2-L3. The CD-ROM drive that has the location code U78A5.001.WIH01AA-P1-T1-L1-L1 is the CD-ROM drive that is provided with the Remote Control Web interface of the Advanced Management Module.

Note: When you attach the media tray of the BladeCenter chassis to a blade that is already up and running you might have to issue **cfgdev** on the command line of the Virtual IO Server to have it be recognized by VIOS.

Virtual CD-ROM drives are used to mount CDs that are placed in the media library. See 4.6.2, “Storage pools” on page 119 and 4.6.4, “Optical and tape devices” on page 127.

The current setup uses an AIX CD that was placed in the media library. The first virtual optical device is checked by default. Click the **Modify** link to select an image from the media library. See Figure 6-22.

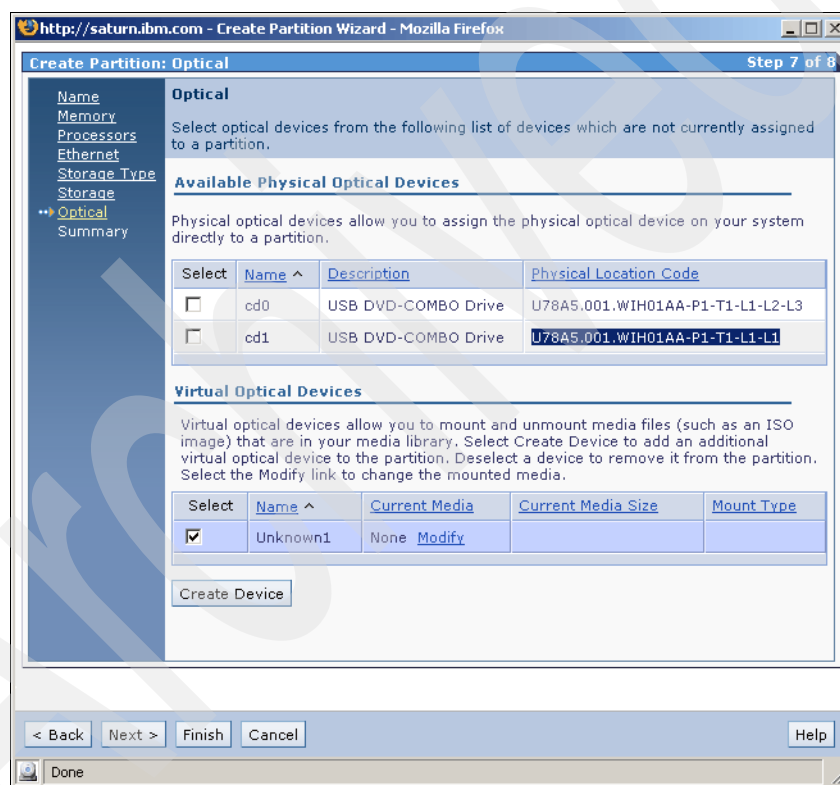


Figure 6-22 Create partition - modify virtual optical device

12. Change the selected media from None to **AIX-6.1 - Read only** and click **OK**. See Figure 6-23.




Figure 6-23 Create partition - modify current media of virtual optical device

13. Click **Next** to see an overview of the setting of the new partition. See Figure 6-24.

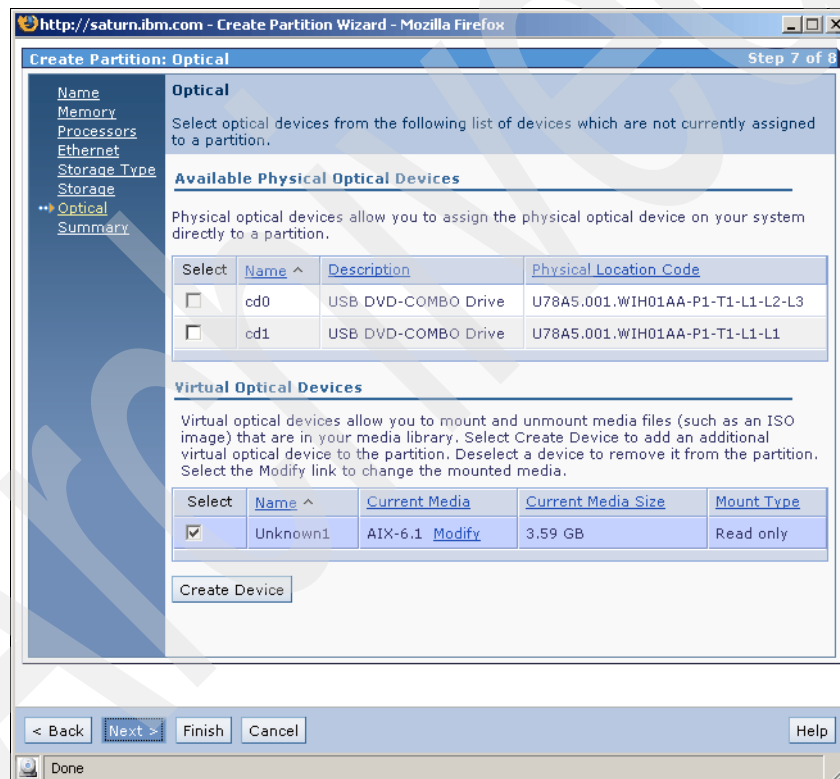


Figure 6-24 Create partition - virtual optical device

14. Verify your setting and click **Finish** to create a partition with the settings you defined. See Figure 6-25.

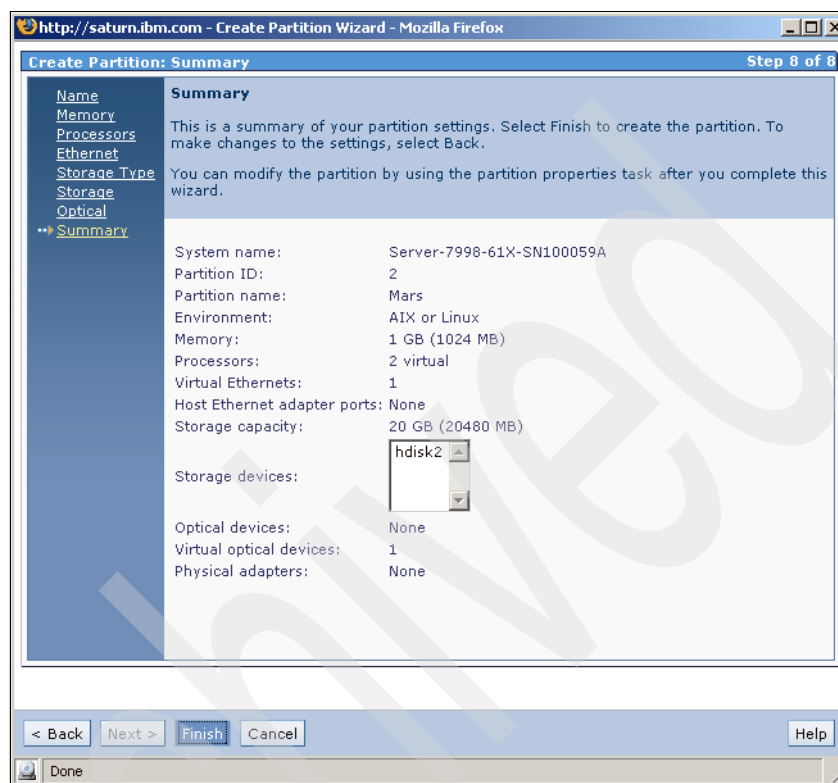
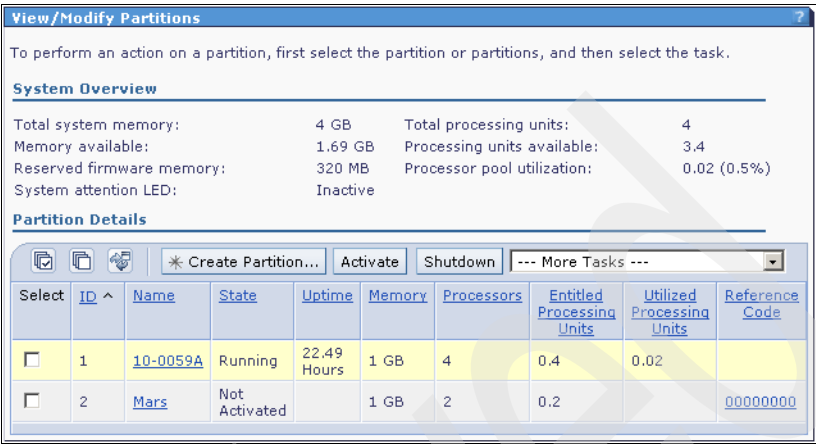


Figure 6-25 Create partition - summary

15. The new partition will be listed under View/Modify Partitions, as shown in Figure 6-26.



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:	4 GB	Total processing units:	4
Memory available:	1.69 GB	Processing units available:	3.4
Reserved firmware memory:	320 MB	Processor pool utilization:	0.02 (0.5%)
System attention LED:	Inactive		

Partition Details

Buttons: ☐ ☐ ☐ * Create Partition... Activate Shutdown --- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	10-0059A	Running	22.49 Hours	1 GB	4	0.4	0.02	
<input type="checkbox"/>	2	Mars	Not Activated		1 GB	2	0.2		00000000

Figure 6-26 Newly created AIX/Linux partition

The preparation of the partition is done. Proceed with the installation of AIX in the newly created partition.

6.4 Installing AIX 6.1 in a logical partition of VIOS

The previous sections described how to prepare the media library that contains the AIX 6.1 DVD image that will be used to install the first logical partition, and how to create a logical partition. This section describes the activation of the logical partition and the installation of AIX 6.1 from a virtual optical device.

To install:

1. To activate the partition, click the check box of the partition and click **Activate**. See Figure 6-27.

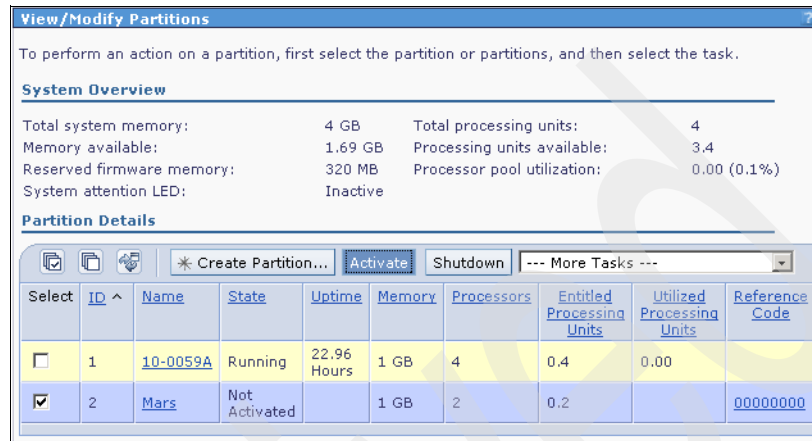


Figure 6-27 Activate a partition

2. Confirm the activation of the partition by clicking **OK** as shown in Figure 6-28.

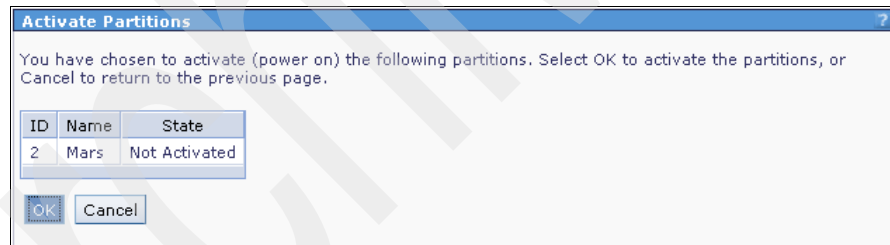


Figure 6-28 Confirm partition activation

3. The state of the partition has changed to Running. Select **Open terminal window** from the More Tasks drop-down list box to open a terminal connected to the selected partition. See Figure 6-29.

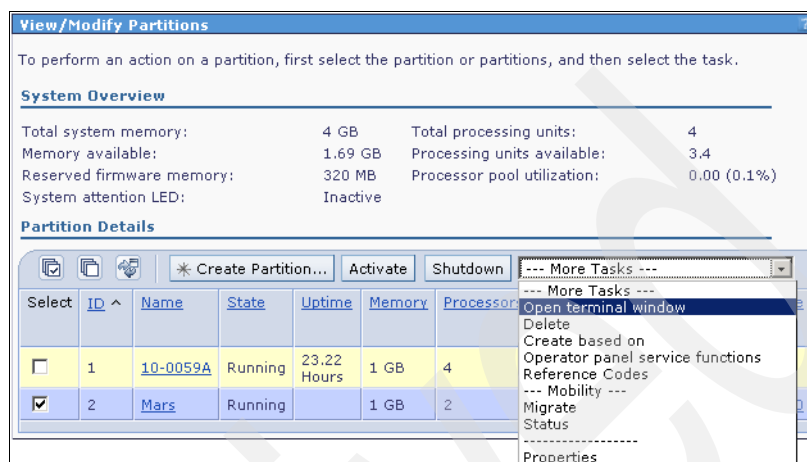


Figure 6-29 Open a virtual terminal to the partition

4. Authenticate on the Virtual IO Server to get the virtual terminal connected. You may use the account `padmin` with the default password `padmin` here if you have not yet created your own account. After the authentication is done, a message will be shown that the terminal has connected, as shown in Figure 6-30.




Figure 6-30 Virtual terminal connection

On the virtual terminal you see the POST of the partition with the possibility to enter the SMS menu. No change is required in this stage. The partition will boot from the assigned virtual optical device.

5. When the partition has booted from the virtual optical device, you see several messages, as shown in Example 6-1 on page 229.

Example 6-1 AIX install kernel load

Elapsed time since release of system processors: 1665 mins 26 secs

```
-----  
Welcome to AIX.  
boot image timestamp: 23:19 10/31  
The current time and date: 01:23:30 06/11/2008  
processor count: 2; memory size: 1024MB; kernel size: 26145029  
boot device:  
/vdevice/v-ccsi@30000002/disk@8200000000000000:\ppc\chrp\bootfile.exe  
kernel debugger setting: enabled  
-----
```

```
AIX Version 6.1  
Starting NODE#000 physical CPU#001 as logical CPU#001... done.  
Starting NODE#000 physical CPU#002 as logical CPU#002... done.  
Starting NODE#000 physical CPU#003 as logical CPU#003... done.  
Preserving 126407 bytes of symbol table [/usr/lib/drivers/hd_pin]  
Preserving 199549 bytes of symbol table [/usr/lib/drivers/hd_pin_bot]
```

6. Define the current virtual terminal as system console by entering **1**. Click Enter to proceed; see Example 6-2. Depending on the console you are using, you might have to also enter **F1** or **2**.

Example 6-2 Select the system console

Preserving 199549 bytes of symbol table [/usr/lib/drivers/hd_pin_bot]

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

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

7. Select the number of the language that you want to use during the installation of IBM AIX 6.1. You can define the language of the operating system that will

be installed in a later step. Press Enter to proceed. See Example 6-3 on page 230.

Example 6-3 Select the language used during installation

```
>>> 1 Type 1 and press Enter to have English during install.
      2 Entreu 2 i premeu Intro per veure la instal·lació en català.
      3 Entrez 3 pour effectuer l'installation en français.
      4 Für Installation in deutscher Sprache 4 eingeben
        und die Eingabetaste drücken.
      5 Immettere 5 e premere Invio per l'installazione in Italiano.
      6 Digite 6 e pressione Enter para usar Português na instalação.
      7 Escriba 7 y pulse Intro para la instalación en español.
```

88 Help ?

>>> Choice [1]:

8. Modify required settings such as language or time zone, and then continue the installation by entering 1 and pressing Enter, as shown in Example 6-4.

Example 6-4 AIX installation summary

Overwrite Installation Summary

```
Disks: hdisk0
Cultural Convention: en_US
Language: en_US
Keyboard: en_US
JFS2 File Systems Created: Yes
Graphics Software: Yes
System Management Client Software: Yes
Enable System Backups to install any system: Yes
```

Optional Software being installed:

>>> 1 Continue with Install

88 Help ?		WARNING: Base Operating System Installation will
99 Previous Menu		destroy or impair recovery of ALL data on the
		destination disk hdisk0.

>>> Choice [1]:

The copy process starts after you select 1 followed by the Enter key.

9. After the installation is done, the partition is automatically rebooted.

Then, select your terminal type, as shown in Example 6-5 on page 231.

Example 6-5 Select the terminal type you are using

```

                                Set Terminal Type
The terminal is not properly initialized. Please enter a terminal type
and press Enter. Some terminal types are not supported in
non-English languages.

    ibm3101      tvi912      vt330      aixterm
    ibm3151      tvi920      vt340      dtterm
    ibm3161      tvi925      wyse30      xterm
    ibm3162      tvi950      wyse50      lft
    ibm3163      vs100       wyse60      sun
    ibm3164      vt100       wyse100
    ibmpc        vt320       wyse350

                                +-----Messages-----
                                | If the next screen is unreadable, press Break (Ctrl-c)
                                | to return to this screen.
88 Help ?
>>> Choice []:
```

10. Select **Show Installed License Agreements** and press Enter to read the license agreement; see Example 6-6.

Example 6-6 License agreement menu

Software License Agreements

Move cursor to desired item and press Enter.

```

Show Installed License Agreements
Accept License Agreements
```

F1=Help	F2=Refresh	F3=Cancel	Esc+8=Image
Esc+9=Shell	Esc+0=Exit	Enter=Do	

11. Select the software package from which you would like to read the license agreements. The default is to show all license agreements. Press Enter to start showing the license text; see Example 6-7 on page 232.

Example 6-7 Select a software license

Show Installed License Agreements

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

			[Entry Fields]	
* SOFTWARE name			[all]	+
SHOW license agreement text?			yes	+
F1=Help	F2=Refresh	F3=Cancel	F4=List	
Esc+5=Reset	Esc+6=Command	Esc+7=Edit	Esc+8=Image	
Esc+9=Shell	Esc+0=Exit	Enter=Do		

12. Navigate through the licenses. When you have finished reading, press F3 twice. You are returned to the Software License Agreements panel. Select **Accept License Agreements** and press Enter; see Example 6-8.

Example 6-8 License agreement menu

Software License Agreements

Move cursor to desired item and press Enter.

Show Installed License Agreements
Accept License Agreements

F1=Help	F2=Refresh	F3=Cancel	Esc+8=Image
Esc+9=Shell	Esc+0=Exit	Enter=Do	

13. Press Tab and then Enter to accept the license and change the no to a **yes**, as shown in Example 6-9.

Example 6-9 Accept license agreements

Accept License Agreements

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

ACCEPT Installed License Agreements			[Entry Fields]
yes			+
F1=Help	F2=Refresh	F3=Cancel	F4=List
Esc+5=Reset	Esc+6=Command	Esc+7=Edit	Esc+8=Image
Esc+9=Shell	Esc+0=Exit	Enter=Do	

14. After the status of the command has changed to OK, press F10 or ESC+0 to leave the license agreement. The keys you require depend on the terminal you are using; see Example 6-10.

Example 6-10 Command status

COMMAND STATUS

Command: OK stdout: no stderr: no

Before command completion, additional instructions may appear below.

F1=Help	F2=Refresh	F3=Cancel	Esc+6=Command
Esc+8=Image	Esc+9=Shell	Esc+0=Exit	/=Find
n=Find Next			

15. The installation assistant guides you through the first administrative tasks, such as setting a root password or configuring the network connection. Proceed with the setup as described in the AIX documentation. To complete this task and open a login prompt, press ESC+0 or F10. You may start this installation assistant at any time again by using the command **install_assist** after logging in as *root*. The installation assistant is shown in Example 6-11.

Example 6-11 AIX installation assistant

Installation Assistant

Move cursor to desired item and press Enter.

Set Date and Time
Set root Password
Configure Network Communications
Install Software Applications
Using SMIT (information only)
Tasks Completed - Exit to Login

F1=Help	F2=Refresh	F3=Cancel	Esc+8=Image
Esc+9=Shell	Esc+0=Exit	Enter=Do	

The installation of AIX 6.1 in the logical partition has completed. You may now start with configuring your AIX or installing your applications.

IBM i V6.1 installation

This chapter explains the installation process of the IBM i V6.1 operating system on an IBM BladeCenter JS23/JS43 Express server installed in a BladeCenter S chassis using the disks provided in the disk storage modules.

For the IBM BladeCenter JS23/JS43 in a BladeCenter H chassis, the installation process is similar to the information provided here, except that the storage is provided from a SAN environment.

This chapter contains the following topics:

- ▶ “Preparing for installation” on page 236
- ▶ “IBM System Access for Windows V6R1” on page 253
- ▶ “Creating an IBM i V6.1 partition” on page 263
- ▶ “Installing and setting up IBM i V6.1” on page 296
- ▶ “IBM i V6.1 backup and restore” on page 307

For a technical overview and complete information, as well as the latest updates for IBM i on Power blades, refer to the readme file available at:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

7.1 Preparing for installation

Important considerations exist for setting up and using IBM i V6.1 client logical partitions on IBM Power servers or the IBM BladeCenter JS23 or JS43 Express server. On Power blades, you use the Integrated Virtualization Manager (IVM) to manage partitions.

A *client logical partition* is a partition that uses some of the I/O resources of another partition.

When the IBM i V6.1 client logical partition is managed by Integrated Virtualization Manager (IVM), you can assign only virtual resources to the 6.1 partition. Disk units, optical devices, and Ethernet are accessed using virtual I/O adapters. The Virtual I/O Server (VIOS) logical partition provides the disk, optical, and network resources to the client logical partition. This configuration is sometimes referred to as a *pure virtual partition*.

A major benefit of using virtual I/O is that you can share the hardware among the client logical partitions through the server logical partition. This benefit allows you to optimize the amount of hardware used by the host partition.

7.1.1 Software installation process

The IBM i V6.1 installation process involves three phases:

1. Pre-installation activities
2. Main installation
3. Post-installation activities

Figure 7-1 on page 237 illustrates one type of software installation of the IBM i 6.1 operating system. The example this process illustrates is one of performing an upgrade to a new release of the operating system. However, many of the steps are similar when performing a new installation.

Note: IBM i V6R1 can be ordered pre-installed on the JS23/JS43. Use feature code FC 8141 (IBM i pre-load) and FC 8147 (VIOS Pre-load pre-req) when ordering.

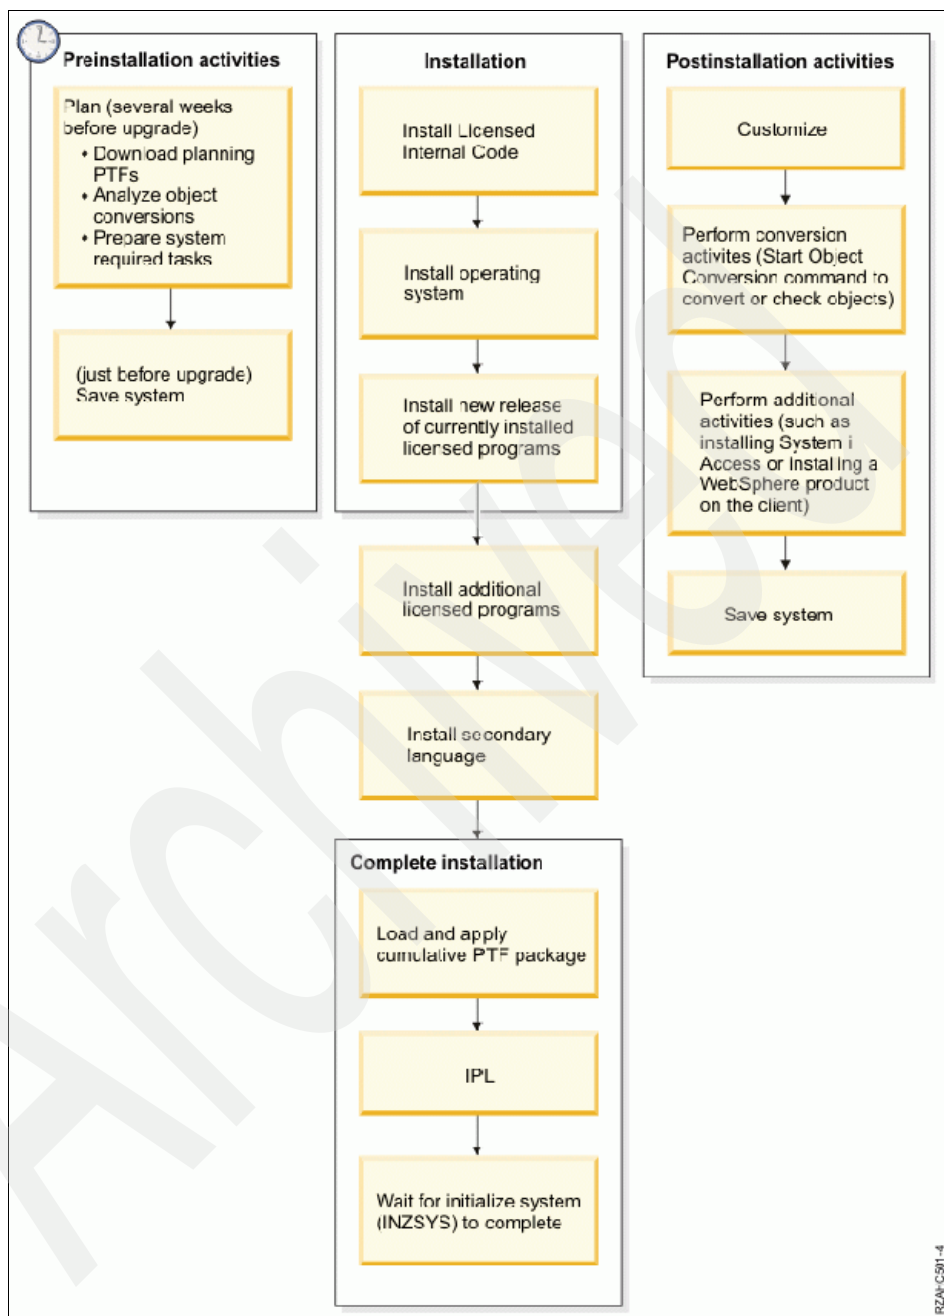


Figure 7-1 IBM i V6.1 installation process

7.1.2 Hardware environments

This section describes an example IBM BladeCenter chassis and IBM BladeCenter JS23/JS43 Express server configuration with recommended firmware levels.

Note: The disk configurations are dependent on the I/O requirements. For example, two SAS disk drives are not enough with mirroring and backup to the media library. For performance reasons, the best practice is to install IBM i to disk units other than the internal disks of the JS23/JS43.

A combination of BladeCenter S chassis and IBM BladeCenter JS23 or IBM BladeCenter JS43 Express server are included in the following list:

- ▶ BladeCenter S chassis
 - 1 Advanced Management Module
 - 1 supported Ethernet I/O Module
 - 1 IBM BladeCenter S 6-Disk Storage Module
 - 2 SAS hot-swap disks
 - 1 SAS Connectivity Module
- ▶ IBM BladeCenter JS23 Express
 - 1 JS23 Express server
 - 4 GB memory
 - SAS Expansion Card (CFFv) for IBM BladeCenter
 - 1 SAS disk drive
- ▶ IBM BladeCenter JS43 Express
 - 1 JS43 Express server
 - 4 GB memory
 - SAS Expansion Card (CFFv) for IBM BladeCenter
 - 1 SAS disk drive

A combination of BladeCenter H chassis and IBM BladeCenter JS23 or IBM BladeCenter JS43 Express server are included in the following list:

- ▶ BladeCenter H chassis
 - 1 Advanced Management Module
 - 1 supported SAN I/O Module
 - 1 supported Ethernet I/O Module
 - 1 Multi-Switch Interconnect Module
- ▶ IBM BladeCenter JS23 Express
 - 1 JS23 Express server
 - 4 GB memory

- 1 QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)
- 1 SAS disk drive
- BM BladeCenter JS43 Express
 - 1 JS43 Express server
 - 4 GB memory
 - 1 QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)
 - 1 SAS disk drive

Table 7-1 lists the minimum and required features necessary to manage an IBM BladeCenter JS23 Express system with the IBM i V6.1 operating system.

Table 7-1 BladeCenter JS23 feature list

Feature	Description	Notes
7778-23X	IBM BladeCenter JS23	4-Way single-wide blade
8395	Processor Entitlement List Priced	Four-processor Entitlement required
8393	Processor Entitlement Express Priced	
8229	4 GB DDE2 667 MHz VLP RDIMMs	4 GB minimum required, up to 32 GB maximum
8239	8 GB DDE2 667 MHz VLP RDIMMs	
8245	16 GB DDE2 400 MHz VLP RDIMMs	
8237	73 GB SAS 10k SFF disk	Disk is not required on base offering
8236	146 GB SAS 10k SFF disk	
8274	300 GB SAS 10K SFF disk	
8273	69 GB SDD SFF disk	
8252	QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)	<ul style="list-style-type: none"> ► Option for SAN connection and Ethernet in BladeCenter H - Connects to MSIM in high speed switch bays 7-8 or 9-10 ► Not supported in BladeCenter S
8240	Emulex 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in bay 3 or 4 of an H or S chassis
8241	Qlogic 4 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in bay 3 or 4 of an H or S chassis

Feature	Description	Notes
8242	Qlogic 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in bay 3 or 4 of an H or S chassis
8271	Qlogic 8 GB Fibre Channel Expansion card (CFFh)	-

Table 7-2 lists the minimum and required features necessary to manage an IBM BladeCenter JS43 Express system with the IBM i V6.1 operating system.

Table 7-2 BladeCenter JS43 feature list

Feature	Description	Notes
7778-23X Plus 8446	IBM BladeCenter JS43 Express Double wide 8-way blade	Adding feature 8446 to the 7778-23X equals the JS43 double wide blade package
8395	Processor Entitlement List Priced	Eight processor Entitlement required
8393	Processor Entitlement Express Priced	
8229	4 GB DDE2 667 MHz VLP RDIMMs	4 GB minimum required, up to 32 GB maximum
8239	8 GB DDE2 667 MHz VLP RDIMMs	
8245	16 GB DDE2 400 MHz VLP RDIMMs	
8237	73 GB SAS 10k SFF disk	Disk is not required on base offering
8236	146 GB SAS 10k SFF disk	
8274	300 GB SAS 10K SFF disk	
8273	69 GB SDD SFF disk	
8252	QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)	<ul style="list-style-type: none"> ► Option for SAN connection and Ethernet in BladeCenter H - Connects to MSIM in high speed switch bays 7-8 or 9-10 ► Not supported in BladeCenter S
8240	Emulex 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in bay 3 or 4 of an H or S chassis
8241	Qlogic 4 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in bay 3 or 4 of an H or S chassis

Feature	Description	Notes
8242	Qlogic 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in bay 3 or 4 of an H or S chassis
8271	Qlogic 8 GB Fibre Channel Expansion card (CFFh)	-

For more information about supported devices on a BladeCenter JS23/JS43 server, refer to the following Web site:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

This site contains links to a specifications and a readme file that has excellent detail for implementing this type of configuration. The site also has information relating to the different BladeCenter chassis configurations and requirements.

7.1.3 BladeCenter hardware preparation

For a smooth installation process of IBM i V6.1, all installed hardware components should be at the latest firmware levels.

Check the following components and update them if necessary:

- ▶ IBM JS23/JS43 Express server firmware (located on the service processor, and updated through the VIOS)
- ▶ Advanced Management Module (AMM) firmware
- ▶ Fibre Channel I/O module firmware
- ▶ Ethernet I/O module firmware
- ▶ SAS Connectivity Module firmware
- ▶ Expansion card firmware (examples: QLogic Ethernet and 4 GB Fibre

For firmware updates for the IBM BladeCenter JS23 or IBM BladeCenter JS43 Express, BladeCenter S, and BladeCenter H chassis and integrated components:

1. Go to:
<http://www.ibm.com/support/us/en/>
2. When you reach that window, select **BladeCenter** from the product drop-down box, and then click the blue arrow.
3. In the next window, select a product from the product family box: IBM BladeCenter JS23, BladeCenter JS43, BladeCenter S, or BladeCenter H.

4. Click the down arrow button in the Operating system box and select **IBM i V6.1**, as shown in Figure 7-2. Then click the **Go** button to activate the search.
(Figure 7-2 is an example of the search options when using the support Web site to locate updates.)



The screenshot shows the 'Support for IBM BladeCenter' page. At the top, there is a breadcrumb trail: 'IBM Systems > Systems support > BladeCenter >'. Below this is the title 'Support for IBM BladeCenter' and an image of three BladeCenter server units. The main section is titled 'Select your product' and contains a note: 'Fields marked with an asterisk (*) are required.' There are four dropdown menus: 'Product family: *' with 'BladeCenter JS22' selected, 'Type:' with 'All types' selected, 'Model:' with 'All' selected, and 'Operating system:' with 'IBM i 6.1' selected. A blue 'Go' button is located at the bottom right of the form.

Figure 7-2 Firmware information and download

5. On the next window, scroll down to view all the available updates and drivers related to the selected product and operating system for downloading.

Figure 7-3 on page 243 shows an example of the available firmware and BIOS updates. Scroll through the list to find the update or tailor the results by using the Refine results option.

➔ Alternate download resources view

Refine results:

Observe all prerequisites and co-requisites listed in each download.
Click the category below to quickly jump to the code you need, then click the version to access the download page.

- Critical update
- Advanced Management Module
- BIOS
- Fibre Channel driver
- Fibre Channel firmware
- Fibre Channel solution
- Fibre Channel utility
- Firmware update
- Hard drive
- Infiniband
- Management Module
- RAID
- Security
- Serial attached SCSI (SAS)
- Utility

➔ Having trouble downloading a file from our site?

Critical update	
Frequently Asked Questions (FAQ) Seagate OEM hard drive issue - IBM BladeCenter and System x	13 Mar 2009
Advanced Management Module	
Advanced Management Module (AMM) firmware release matrix - IBM BladeCenter	30 Mar 2009
(BOFM enabled) (Functional release) Advanced Management Module firmware v2.48D (BPET48D) - IBM BladeCenter	30 Mar 2009 v2.48D
Enhanced Role Based Security Snap-in (LDAP snap-in) - IBM BladeCenter	25 Jun 2008 v2.01
BIOS	
BIOS/Firmware Inband update v3.00, 01EA340_039_039 - IBM BladeCenter JS12, JS22	21 Nov 2008 v01EA340_039_039
Firmware v01EA320_046_030 - IBM BladeCenter JS22	29 May 2008 v01EA320_046_030
Fibre Channel driver	
Boot Code for Emulex 4Gb fibre channel expansion card v6.02a7 - IBM BladeCenter	29 May 2008 v6.02a7

Figure 7-3 Example: Partial list of available downloads by type

7.1.4 VIO Server software environments

VIO Server is part of IBM PowerVM Editions (formerly Advanced POWER Virtualization). It is required in the IBM i V6.1 for IBM BladeCenter JS23/JS43 Express environment. At minimum, VIOS level 1.5 is required for IBM i. Using version 2.1 or later works best.

Work with your local sales channel to ensure that PowerVM (Standard or Enterprise Edition) and the latest fix pack are part of the BladeCenter JS23/JS43 order. Click the **Supported Environments** link to open the PDF file and verify that you have the minimum supported release of VIOS:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

For detailed information about using the Integrated Virtualization Manager and VIO Server, refer to Chapter 4, “System planning and configuration using VIOS with IVM” on page 65.

7.1.5 Network considerations

Before the initial installation process starts, reserve a set of IP addresses for setting up the minimum environment:

- ▶ AMM

The Advanced Management Module (AMM) address is a physical LAN IP address. It is used to manage the BladeCenter chassis and IBM BladeCenter JS23/JS43 Express server remotely.

Note: Do not place the Advanced Management Module (AMM) IP address in the same subnet as the other addresses to prevent I/O traffic between AMM and integrated Ethernet Switch modules (Proxy-Loop).

- ▶ Ethernet I/O module

This IP address connects the Ethernet I/O Module to the physical LAN, thus allowing any blades in the BladeCenter chassis to access the LAN.

- ▶ VIOS/IVM

This IP address connects the Virtual I/O Server (VIOS) and Integrated Virtualization Manager (IVM) to the LAN.

- ▶ SAS modules

This IP address is used to communicate with the SAS modules.

- ▶ 6.1 LAN console

This IP address on the LAN enables the 5250 console to connect to the VIOS using IBM System i Access for Windows software.

- ▶ 6.1 production interface

This IP address on the external LAN provides 5250 production network access. This address will be configured after 6.1 is installed using LAN console. A good practice is for the 6.1 LAN console and production network interface to use two separate Virtual Ethernet adapters in the 6.1 partition.

- ▶ PC for LAN Console

When the System i for Windows LAN console is first established, the PC console must be on the same subnet as the 6.1 partition. After the console is established, this restriction is removed.

7.1.6 Storage consideration BladeCenter H

An IBM BladeCenter JS23/JS43 Express server installed in a BladeCenter H chassis has no access to any physical devices in the BladeCenter H chassis. Storage must be provided by attaching LUNs on a SAN to the VIO Server. The VIO Server provides a virtualized access from IBM i on the JS23/JS43 to the logical unit numbers (LUNs).

Note: When you configure LUNs for IBM i, configure them as 512-byte AIX LUNs, not as 520-bytes IBM i LUNs.

For detailed information about how to define a SAN environment for an IBM BladeCenter JS23 or JS43 blade installed in a BladeCenter H chassis using LUN attached disks to install IBM i V6.1, refer to:

http://www.ibm.com/systems/power/hardware/blades/i_on_blade_readme.pdf

7.1.7 Disk consideration in BladeCenter S

This section describes the configuration environment (used in this chapter) to install IBM i V6.1 on an IBM BladeCenter JS23/JS43 Express server in a BladeCenter S chassis.

IBM i V6.1 running on an IBM BladeCenter JS23/JS43 Express server does not have physical access to storage or other devices; instead, the hardware does. Every hardware device is provided to the IBM BladeCenter JS23/JS43 server as a virtual device using the Integrated Virtualization Manager (IVM) functionality.

To provide access to a SAS drive in the BladeCenter S chassis to the partition, at least one SAS I/O module must be installed in the BladeCenter S chassis. An SAS expansion adapter (CIOv) also must be installed in each IBM BladeCenter JS23 or IBM BladeCenter JS43 Express server. A single SAS I/O module provides access to both Disk Storage Modules (DSM) and all 12 disks.

The physical connection to tape drives is owned and managed by VIOS. The IBM i does not have direct access to the tape. The SAS I/O expansion module also provides a connection of an LTO tape connection for backup and restore.

After SAS disk drives are assigned to an IBM BladeCenter JS23/JS43, they become available in an hdiskX hard disk drive. Each virtualized SAS disk drive will be recognized as a DDxx physical disk drive in IBM i V6.1 after it has been installed.

Note: SATA drives installed in the Disk Storage Modules (DSM) in the BladeCenter S chassis are not supported for IBM i V6.1.

The concept of a chassis with a self-contained disk subsystem or storage modules is unique to the BladeCenter S chassis. The module is a collection of disk drives that are made accessible through a SAS module and a SAS card in the BladeCenter JS23/JS43 server. The SAS module is responsible for both the provisioning of physical disk drives through zoning and for failover redundancy when two SAS modules are present. Each SAS drive in the DSM is assigned individually.

Zoning allows you to map hard drives in storage module 1 and storage module 2 to the blade servers, and to map the blade servers to the external ports on the connectivity module.

Note: To provide better reliability, availability, and serviceability (RAS), use the IBM i V6.1 disk mirroring functionality for the SAS disks provided from the VIO Server; this will protect the IBM i 6.1 installation.

To ensure higher availability, the configuration of assigned SAS disk drives in the Disk Storage Modules (DSM) can be individually configured. Information provided in 7.1.8, “Disk configuration in BladeCenter S” on page 247, helps you to visualize or select a predefined DSM configuration.

7.1.8 Disk configuration in BladeCenter S

To use a predefined configuration to a BladeCenter JS23/JS43 server, you must establish a connection to the SAS module, as shown in Figure 7-4, using a browser window directly connected to the SAS module.

An alternative, that is more intuitive for clients is the SCM GUI.

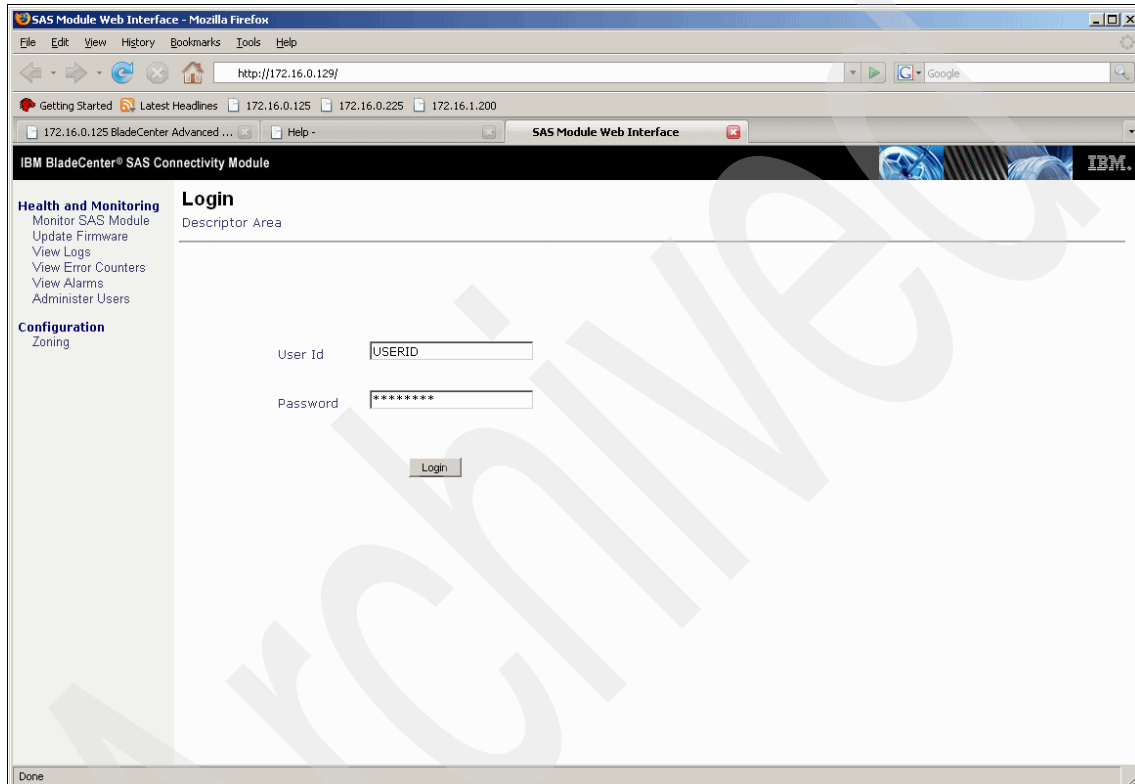


Figure 7-4 SAS Connection module login

To use the SAS module:

1. Enter **User ID** and **Password** of the account that has access to the SAS module and click **Login**, as shown in Figure 7-4.

Figure 7-5 on page 248 shows the next window that appears is the SAS module welcome window.

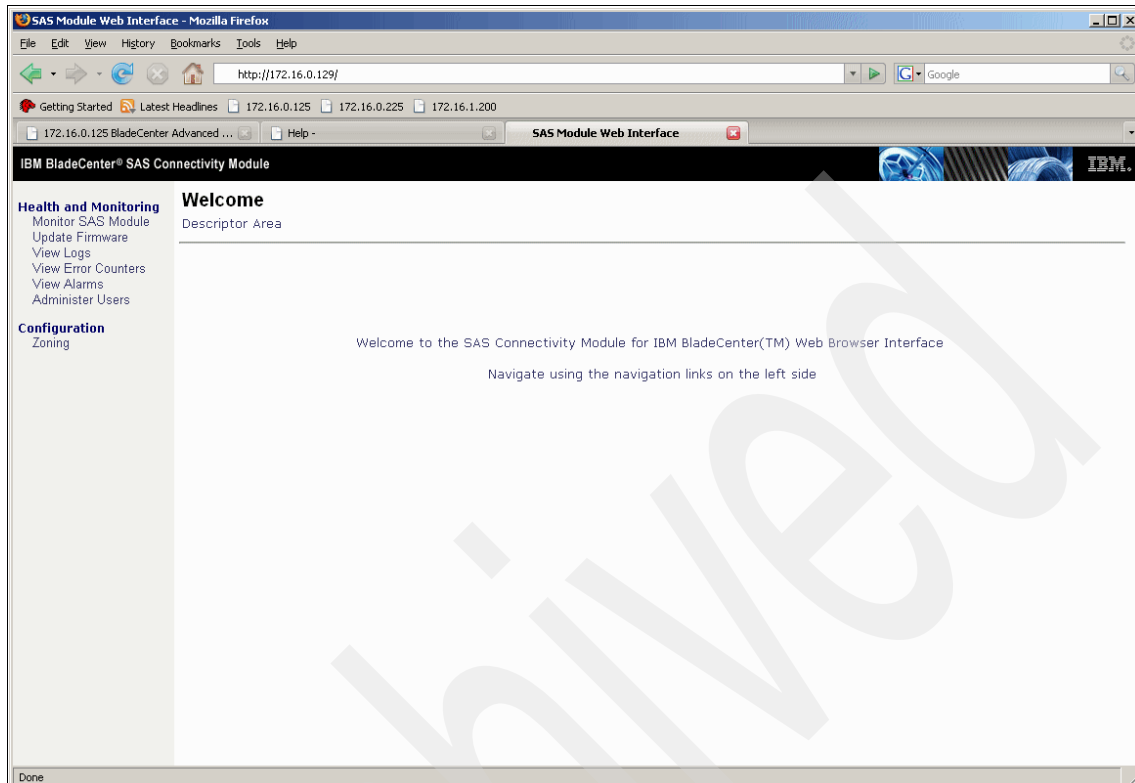


Figure 7-5 SAS connection module welcome

2. Select **Zoning**. In the Manage Fabric panel, shown in Figure 7-6 on page 249, **Predefined Config 09** is selected and active. Notice that our BladeCenter JS23/JS43 installed in slot 4, Zone Group ID 37 is configured. Remember the Zone Group ID for the following window to examine the corresponding hard disk drives.
3. Click **Basic Zone Permission Table**.

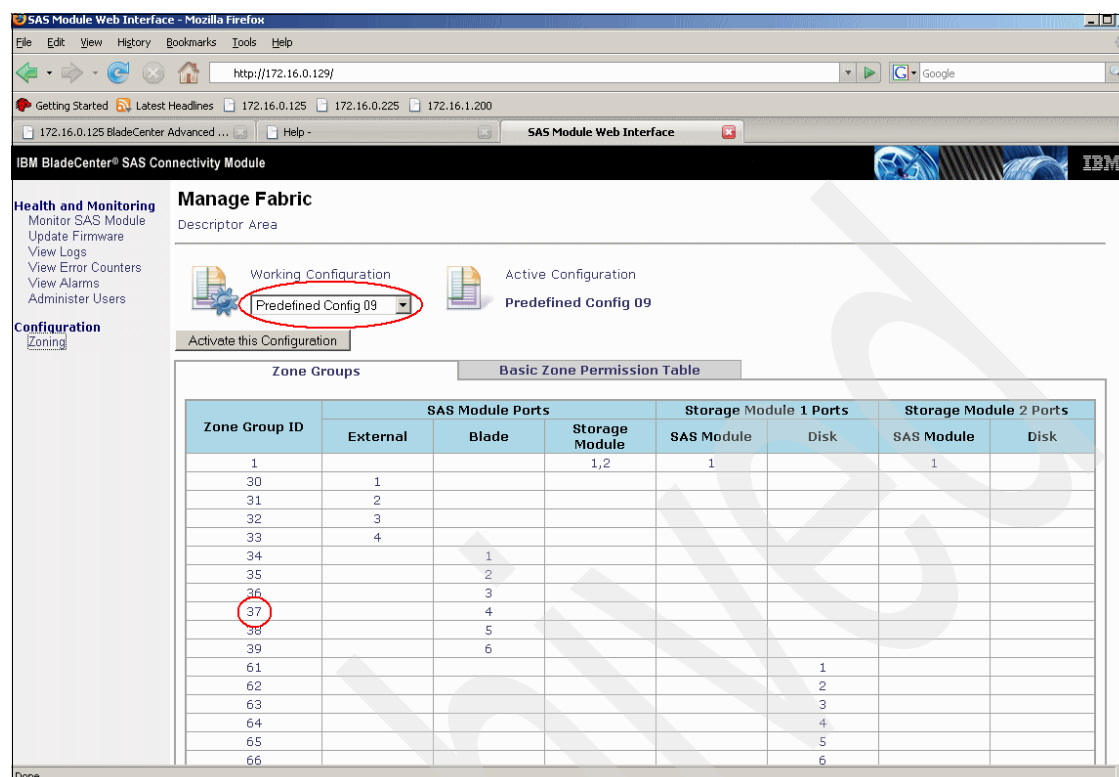


Figure 7-6 SAS connection module zone groups

Figure 7-7 on page 250 shows the definition and setup window for the actual configuration. In this configuration three disks from SAS module 1 and three disks from SAS module 2 are defined for Predefined Config 09. Individual User Defined Configs are provided for specific configurations.

For more information about this topic, refer to: *Implementing the IBM BladeCenter S Chassis*, SG24-7682.

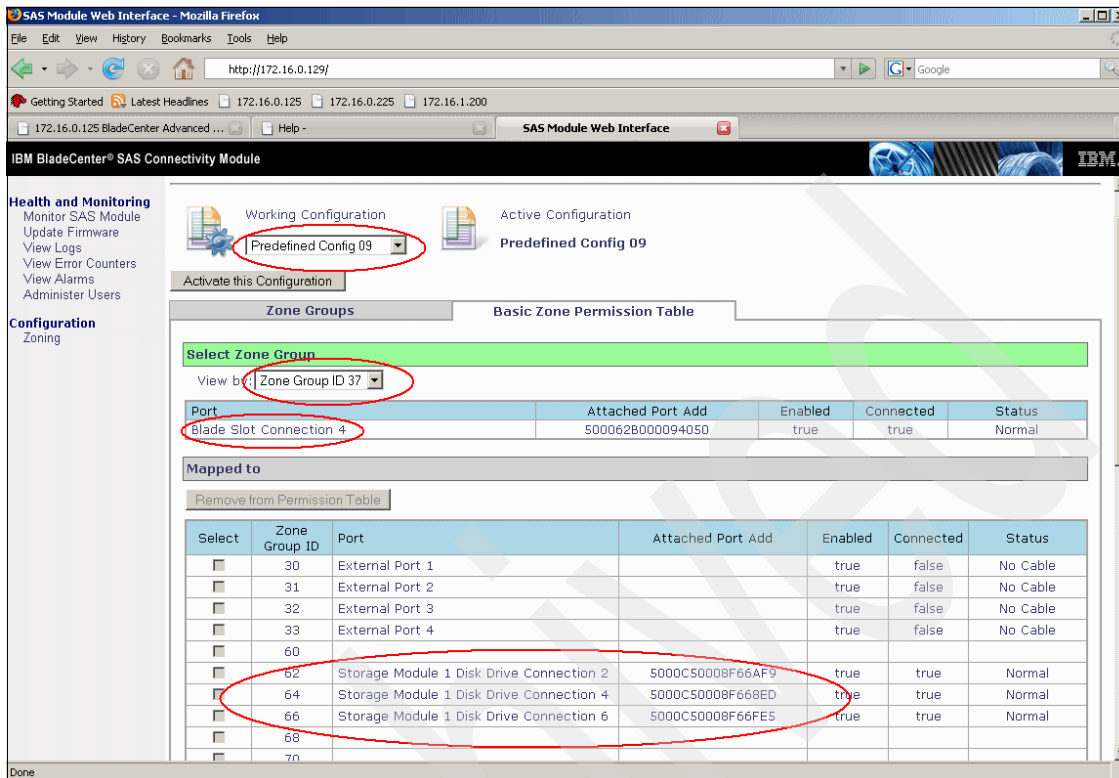


Figure 7-7 SAS connection module zoning

- To verify the configuration in the SAS module configuration menus, logon to the IBM BladeCenter Advanced Management Module. Then, under Storage Tasks, select **Configuration** as shown in Figure 7-8 on page 251.




Figure 7-8 AMM SAS configuration zone

- Click **Predefined Config 09** to proceed.

Figure 7-9 on page 252 shows the current configuration. Select the blade in the upper rectangle to highlight the assigned disks to that blade.

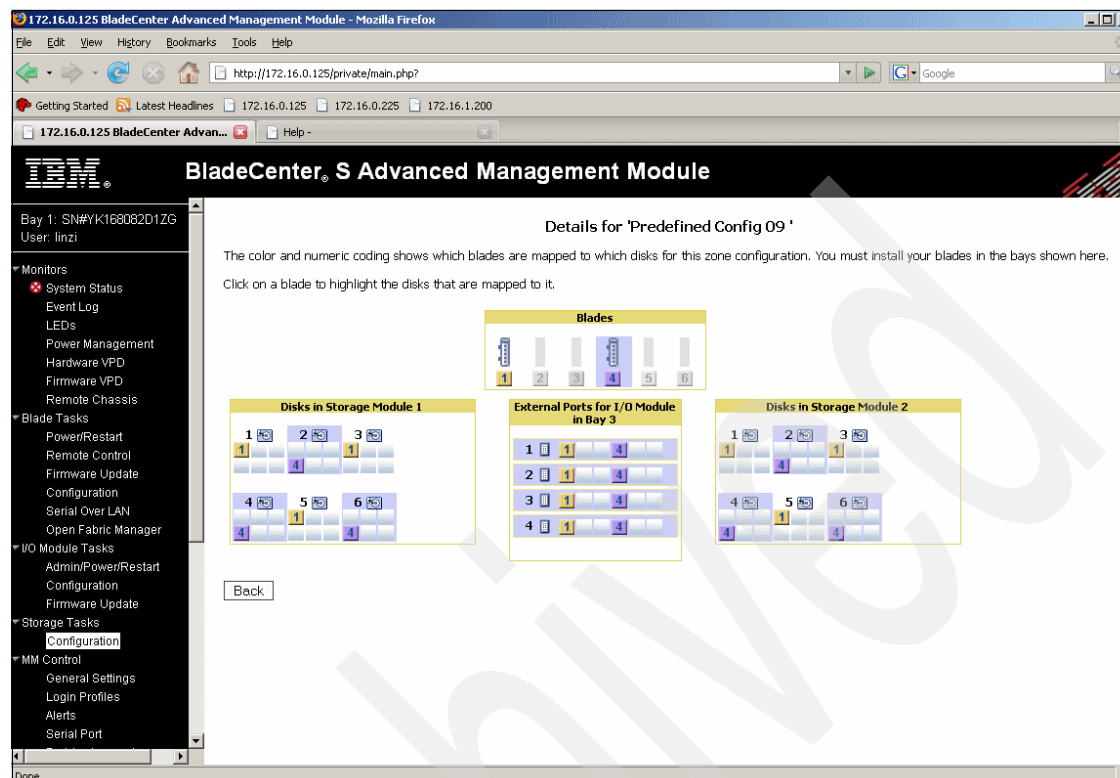


Figure 7-9 AMM SAS configuration zone 9

For detailed information, refer to *Implementing the IBM BladeCenter S Chassis*, SG24-7682 and *IBM BladeCenter Products and Technology*, SG24-7523.

7.1.9 Individual BladeCenter S disk configuration

If one of the eight predefined SAS I/O module disk configurations does not match the target configuration, four user-predefined configurations are available for individual use.

To create an individual customer configuration, one of the four user predefined configuration templates must be changed. Two access methods are available to create a custom configuration:

- ▶ Using Telnet functionality to access the SAS I/O module command-line interface
- ▶ Using the Storage Configuration Manager (SCM) graphical interface

The IBM Storage Configuration Manager (SCM) might be used to create an individual configuration if you are not familiar using the SAS I/O module command-line interface. The SCM software can be downloaded from:

<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5502070&brandind=5000008>

7.2 IBM System Access for Windows V6R1

System i Access for Windows fully integrates the power of the IBM i V6.1 operating system into the desktop so that users have a simple and optimized interface to use the systems capabilities in business environments. The product provides a single solution to work with databases and other data on the system, to run 5250 applications, and to administer the IBM i environment.

The main reason for installing this software is to access the console session on your IBM i partition. System i Access provides two different connection options for the system console. One option is a direct cable from an adapter to a PC comm port and the other method is to connect using the LAN. When you install IBM i in the blade server environment you can only use the LAN Console option for the partition. No physical adapters that are compatible with the blade environment can be used for a direct console connection.

Prior to version 6.1, the IBM System i Access for Windows product was known as the IBM iSeries® Access for Windows products.

With System i Access for Windows, a full suite of client/server capabilities to support a PC and system interaction:

- ▶ System i Access for Windows required programs
- ▶ System i Access for Windows optional features
- ▶ System i Navigator graphical interface and functions
- ▶ Data access functions
- ▶ AFP Workbench Viewer
- ▶ Secure Sockets Layer (SSL)
- ▶ Printer driver
- ▶ Display and printer emulation
- ▶ Operations console

7.2.1 Preparation of System i Access for Windows

To meet all the System i Access for Windows hardware requirements, follow the instructions described in:

<http://www.ibm.com/systems/i/software/access/windows/v6r1pcreq.html>

For more information about the IBM System i Access for Windows V6R1, see:

<http://www.ibm.com/systems/i/software/access/index.html>

To obtain the IBM System i Access for Windows software, go to:

<http://www.ibm.com/systems/i/software/access/caorder.html>

Note: When the IBM i Access for Windows connection is first established, the console PC must be on the same subnet as the 6.1 partition. After the console is established, this restriction is removed.

7.2.2 Installing System i Access for Windows

To install System i Access for Windows:

1. After you insert the System i Access for Windows V6R1 DVD, the dialog shown in Figure 7-10 appears. Select the language you require, and click **OK**.

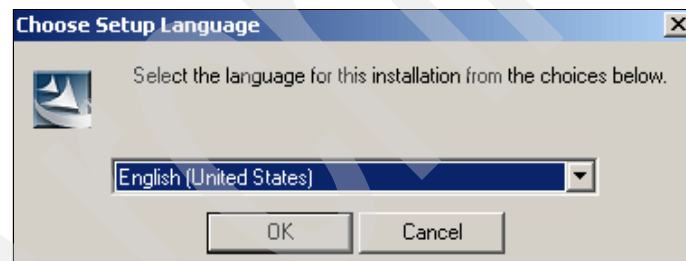


Figure 7-10 Setup LAN Console

2. The InstallShield Wizard Welcome window, shown in Figure 7-11 on page 255, opens. Click **Next** to continue.




Figure 7-11 IBM System i Access for Windows welcome window

3. The License Agreement shown in Figure 7-12 on page 256 appears. If you agree with the terms, select **I accept the terms in the license agreement**. Click **Next** to continue.

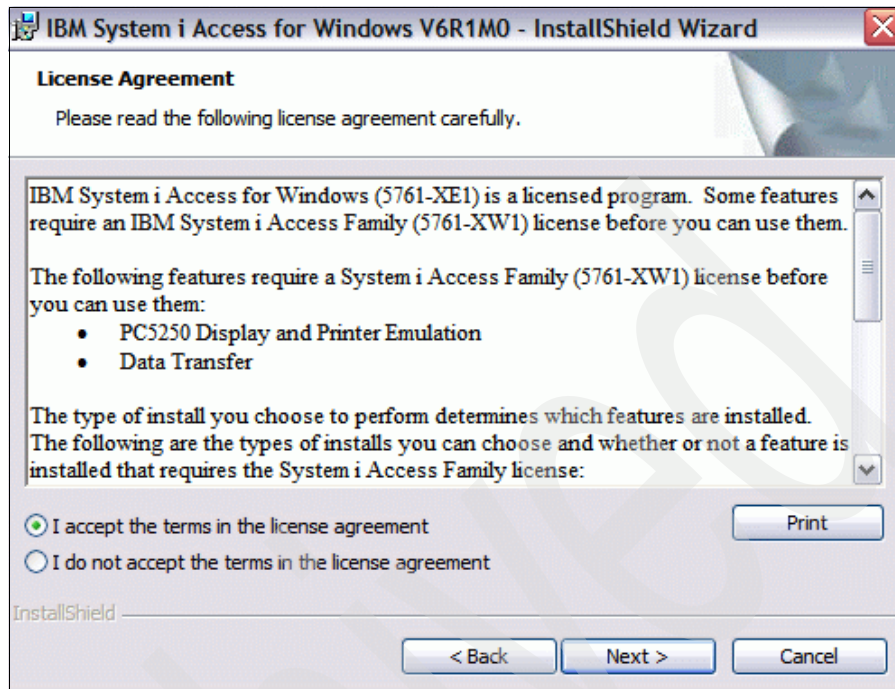


Figure 7-12 IBM System i Access for Windows License Agreement

4. IBM System i Access for Windows can be installed at a different location, as shown in Figure 7-13 on page 257. Do one of the following steps:
 - To store the software at a different location, click **Change** and choose a new location.
 - To accept the predefined path and click **Next** to continue.




Figure 7-13 IBM System i Access for Windows install location

5. Depending on the native language, a selection can be made in the following window as shown in Figure 7-14 on page 258. Normally, you would select the same language as the language for the IBM i V6.1 operating system.

Click **Next** to continue.

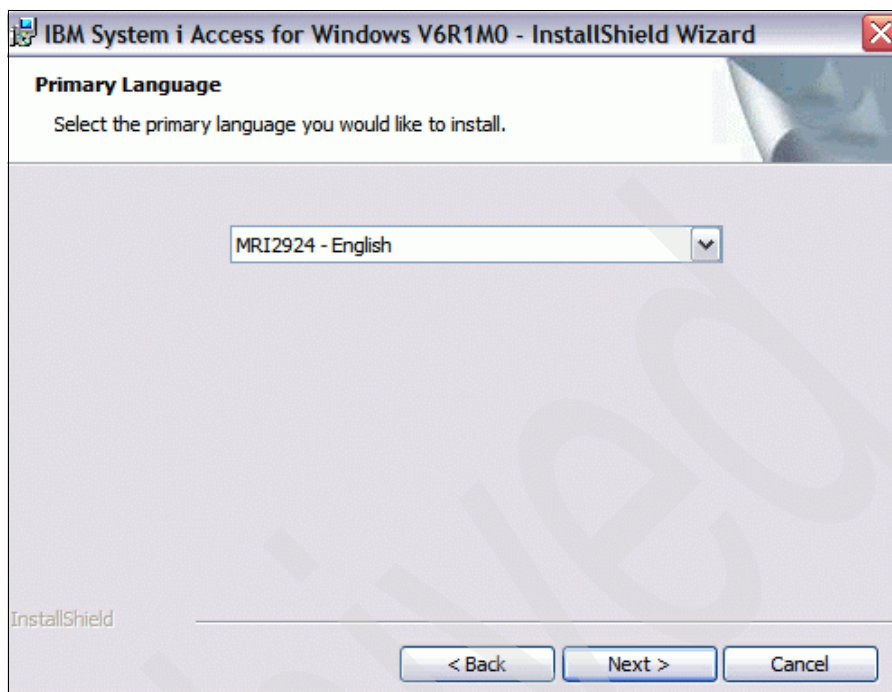


Figure 7-14 IBM System i Access for Windows Primary language

6. Depending on the complexity of functions, several choices are available as shown in Figure 7-15 on page 259. The normal case is a complete installation. Experienced administrators can select the custom installation to save disk space, or install determined functions only.
Click **Next** to continue.




Figure 7-15 IBM System i Access for Windows Setup Type

7. Select **Complete** and click **Next**
8. Certain features require a license agreement to use their functionality, as shown in Figure 7-16 on page 260. Ask your service representative for a valid license key.
Click **Next** to continue.




Figure 7-16 IBM System i Access for Windows Restricted Features

The installation starts automatically. Figure 7-17 on page 261 shows the progress of the installation process.




Figure 7-17 IBM System i Access for Windows installation progress

9. Figure 7-18 on page 262 indicates the installation process was successful.
Click **Finish** to continue.




Figure 7-18 IBM System i Access for Windows installation completed

To finalize the IBM System i Access for Windows installation a reboot is required, as indicated in Figure 7-19.

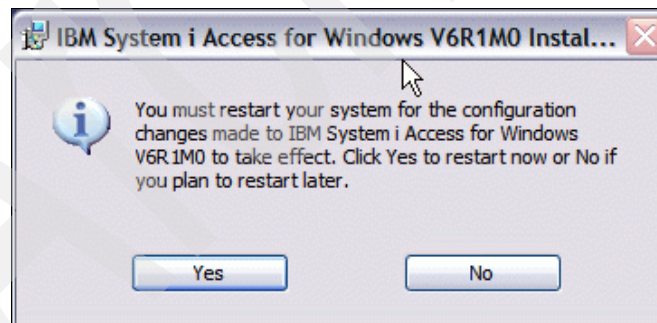


Figure 7-19 IBM System i Access for Windows Reboot

10. Click **Yes** to reboot the system.

After the console PC is successfully rebooted, the information window shown in Figure 7-20 on page 263 is displayed. The Welcome window provides additional

information about the software just installed. (For some information, the administration PC requires a connection to the Internet.)

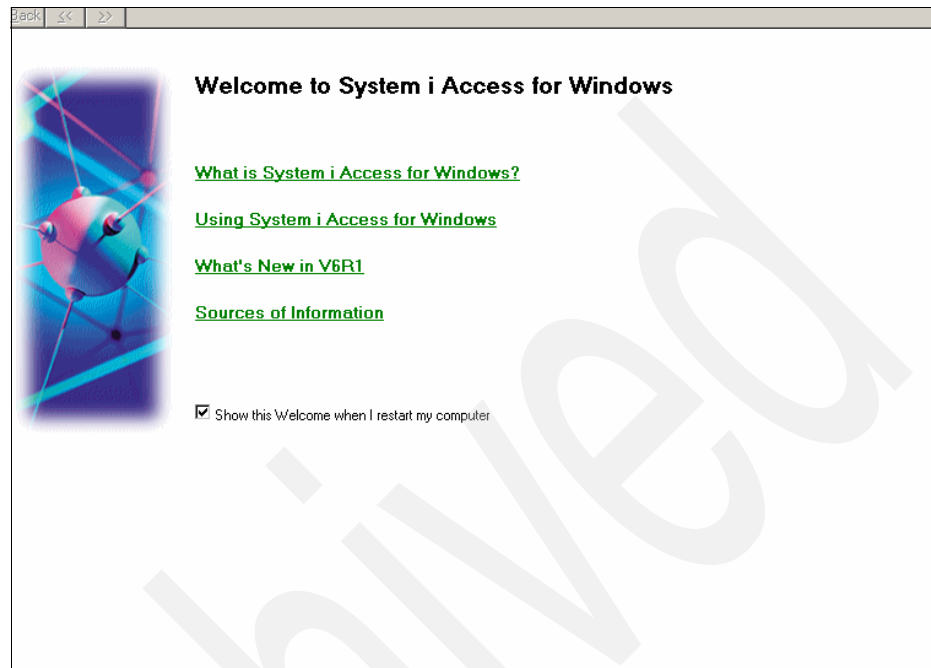


Figure 7-20 IBM System i Access for Windows Welcome

After the System i Access for Windows product is installed, two conditions must be met before you can run a System i Access for Windows application:

- ▶ Your PC and the system must be on a TCP/IP network.
- ▶ A user ID must exist on the IBM i V6.1 operating system.

7.3 Creating an IBM i V6.1 partition

Using Integrated Virtualization Manager (IVM) to create an IBM i V6.1 partition is similar to using the HMC. IVM uses a number of defaults that simplify partition creation. For example, because IBM i V6.1 partitions cannot own physical hardware on an IVM-managed system such as a BladeCenter JS23/JS43, those windows are omitted from the creation wizard. Other windows that relate to shared processor pool settings and memory settings are simplified as well. Typically, you will have to tailor the partition properties to change the settings so they are more granular.

One other difference between the HMC method and using IVM, if you are familiar at all with the HMC you know that the partition has properties and the partition profile contains the definition of your selections for hardware, load source, console etc. With the HMC you might even have multiple partition profiles for one partition. In IVM, only one entity describes both the partition properties and the partition profile. IVM does not have separate partition profiles.

Because your IBM i partition is virtual, IVM defaults the load source and alternate initial program load (IPL) adapters to the Virtual SCSI client adapter in the IBM i V6.1 partition, and the console adapter to the first Virtual Ethernet adapter. If you plan to use separate Virtual Ethernet adapters for LAN console and production traffic, and you want to use the second Virtual Ethernet adapter for the LAN console, you can make the change in the partition properties.

7.3.1 IBM i V6.1 minimum requirements

The minimum amount of memory for an IBM i V6.1 client partition on the BladeCenter JS23/JS43 blade should be 1 GB. If you are running two or more IBM i 6.1 partitions on the same blade, ensure enough memory is on your blade.

The actual memory and CPU values should be sized individually for each IBM i V6.1 workload using the IBM Systems Workload Estimator, available at:

<http://www.ibm.com/systems/support/tools/estimator/index.html>

Note: Also consider that the system firmware or POWER Hypervisor for the BladeCenter JS23/JS43 blade will allocate some memory for its usage. This amount can vary but is generally less than 1 GB. This value can be found on the main panel of IVM listed by *Reserved firmware memory*.

7.3.2 VIO Server configuration

For a detailed explanation of how to set up and configure the VIOS partition to use the Integrated Virtualization Manager (IVM), refer to Chapter 4, “System planning and configuration using VIOS with IVM” on page 65.

7.3.3 Creating an IBM i partition

This section provides a brief explanation of how to create an IBM i V6.1 partition.

It assumes you have previously configured disk space (LUNs) for this partition’s usage.

To create an IBM i V6.1 partition using the IVM:

1. Click **View/Modify Partitions** to display the panel, as shown in Figure 7-21.

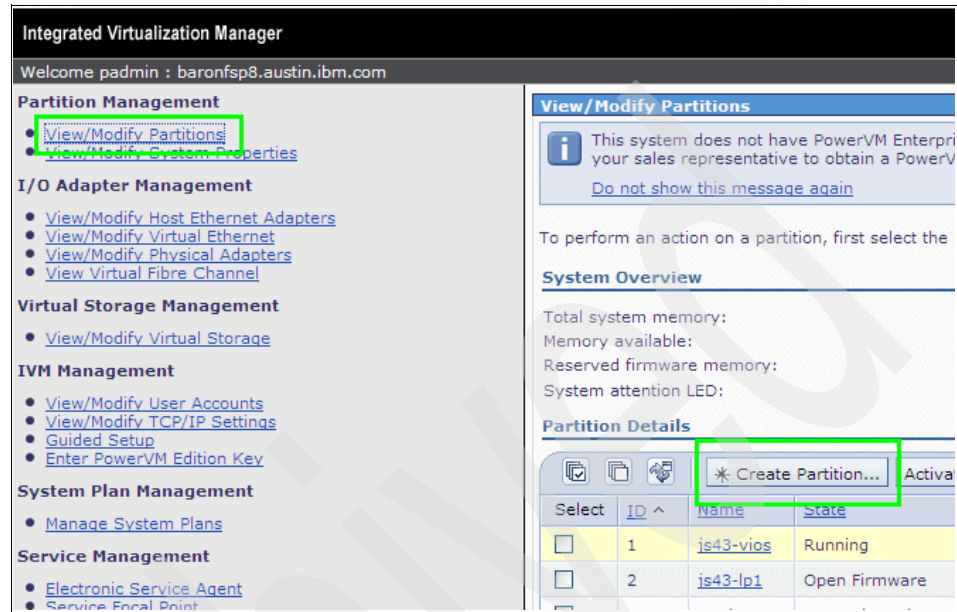


Figure 7-21 View Modify Partitions - Create Partition option

2. Click **Create Partition**.
3. In the next window that opens, shown in Figure 7-22 on page 266:
 - a. The available Partition ID is preselected. It is filled in by the wizard, using the next available number.
If you want, you may change this value to an unused number.
 - b. Enter a name for the partition in the Partition name field.
 - c. Select **IBM i** or **i5/OS** in the Environment field.
 - d. Click **Next** to continue.

Create Partition: Name

Navigation: Name, Memory, Processors, Ethernet, Storage Type, Storage, Optical/Tape, Summary

Name

To create a partition complete the following information.

System name: Server-7778-63X-SN10181CA

Partition ID: 3

* Partition name: IBMi

Environment: IBM i

Figure 7-22 Partition id, name and environment options

- Figure 7-23 shows an example of the memory definitions. You can see the total system memory and what memory is available for partition usage. The available memory is the value of the total minus allocated partition memory and hypervisor allocated memory.

Define the Assigned memory value (Assigned equates to Desired on the HMC). Available memory for the partition is listed above the Assigned memory field. Click **Next** to continue.

Create Partition: Memory Step 2 of 8

Navigation: Name, Memory, Processors, Ethernet, Storage Type, Storage, Optical/Tape, Summary

Memory

Specify the amount of memory, in multiples of 32 MB, to assign for the partition.

Note: If you specify a number that is not a multiple of 32 MB, the wizard will round the number to the nearest multiple of 32 MB.

Total system memory: 8 GB (8192 MB)

Current memory available for partition usage: 4.44 GB (4544 MB)

Assigned memory: 1 GB

GB, MB, GB, TB

Figure 7-23 Partition memory definition panel

- Figure 7-24 on page 267 is an example of the processor selection panel. In this example, the blade server had 8 processors total. Other partitions are created that also use some processor capacity. In the Assigned processors drop-down box, you select how many processing units to assign to this partition. For example, using shared if you select 1 as shown in the figure, you will have .10 units configured in the partition. The value that shows 63 available virtual processors indicates that out of the 80 virtual processors

available (8.0 physical processors), the blade has 17 virtual processors in use by other partitions. That also equates to 1.7 processors used.

Dedicated processing is another option where one or more whole processors can be assigned to the partition.

Select the desired processor configuration. Click **Next** to continue.

Figure 7-24 Partition processor selection

- Figure 7-25 provides an example of the Ethernet options available to select for this partition. Because our console for the IBM i partition is LAN based, we must have an Ethernet selected that is bridged. In our example, we bridged one of the HEA ports prior to create this partition. For more information about bridging the HEA ports, see 4.5.1, “Host Ethernet Adapters” on page 95.

Specify the VLAN used for bridging. Click **Next** to continue.

Figure 7-25 Partition Ethernet selection

- Figure 7-26 on page 268 shows an example of the selection for disk units to use in the partition. You can use virtual disks or physical disks. For an IBM i partition, use physical volumes. Selecting virtual disks will work, but you might have performance issues in the partition. Also assumed at this point is that you have previously configured LUNs or have physical units available to assign to this partition.

- a. Select **Assign existing virtual disks and physical volumes**. Click **Next** to continue.

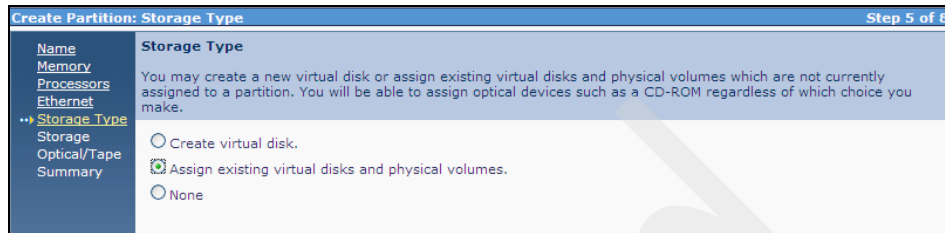


Figure 7-26 Select storage type

8. Figure 7-27 on page 269 shows an example of available disk units. For this partition we selected to use `hdisk8` and `hdisk9`, which are LUNs created in a storage subsystem that have been assigned to this JS43. Depending on your configuration, you might also have virtual disks available. If so, they would be listed under the Available Virtual Disks section.

Depending on the BladeCenter chassis configuration in the storage window, you can select either:

- A LUN that is already configured for the IBM i V6.1 partition from the list of **Available Physical Volumes**.
- A disk provided from the VIO Server to the IBM i V6.1 partition from the list **Available Virtual Disks**. This disk might have been created from a storage pool.

9. Click **Next** to continue.

Storage

Storage

Select any number of physical volumes and virtual disks from the following lists of devices which are assigned to a partition. You may use the Storage Management functions to change assignments at any time.

Available Virtual Disks

Select

Name ^

Storage Pool

Available Physical Volumes

Select	Name ^	Size	Physical Location Code
<input type="checkbox"/>	hdisk1	68.37 GB	U78A5.001.WIH23CF-P1-T5-L60000-L0
<input type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400
<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400
<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401
<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401
<input checked="" type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401
<input checked="" type="checkbox"/>	hdisk9	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401

Figure 7-27 Disk selection

10. Also depending on the installation preparation in the Optical devices menu, you can select either:
 - a. The USB DVD drive to the IBM i V6.1 partition (the media tray is assigned to this blade) in the **Available Physical Optical Devices** section.
 - b. Or a predefined Virtual Optical Device in the **Virtual Optical Devices** section. This shortens the installation time using Virtual Optical Devices with predefined images of the LIC, BOSS_1, and BOSS_2 DVDs installing a second or third IBM i V6.1 partition on the same BladeCenter JS23/JS43 server, and a manual exchange of the install media is omitted.
 - c. Select the first row in the Virtual Optical Devices section.
 - d. Click **Modify** underneath the Current Media field and select the desired Virtual Optical Media.
 - e. Click **Next** to continue.
11. Figure 7-28 on page 270 provides an example of the optical selection panel. If the media tray for the BladeCenter has been assigned to the blade server you are creating the partition on, the device is available. Under the Physical Optical Devices section is the check box to select **cd0**.

Optical/Tape

Select optical or tape devices from the following list of devices which are not currently assigned to a p

▼ Physical Optical Devices

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

Select	Name ^	Description	Physical Location Code
<input checked="" type="checkbox"/>	cd0	USB DVD-COMBO Drive	U78A5.001.WIH23CF-P1-T1-L1-L2-L3

▼ Virtual Optical Devices

You can use virtual optical devices to mount and unmount media files, such as an ISO image, that ar
media library for use by the partition. Select a virtual optical device in the table to assign it to the ne
Clear the selection for a device if you do not want to assign it to the partition. Click Modify to change
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	Mour
<input checked="" type="checkbox"/>	Unknown1	None Modify		

Create Device

► Physical Tape Devices (No devices)

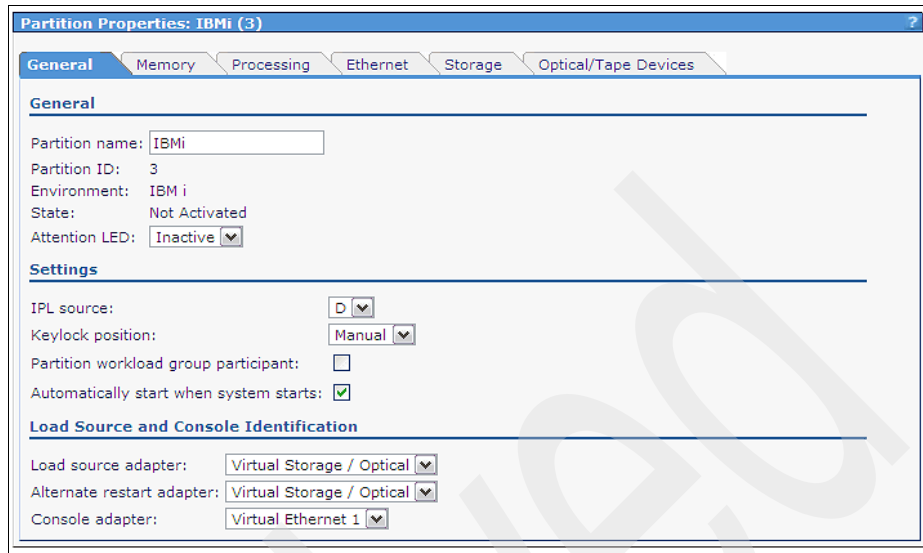
Figure 7-28 Partition optical selections

- Review the summary of your definition and click **Finish** to create the IBM i V6.1 partition.

7.3.4 Partition properties

Earlier in the process of creating the IBM i partition, selections were made for processor and memory values. With the partition create wizard there are not options to tailor the memory and processor values like there are when using an HMC. From the Memory and Processor tabs you can change the selections to allow a minimum and maximum value. This approach allows some movement of memory and processor resources using dynamic allocations.

Figure 7-29 on page 271 shows the fields for Load Source and Console Identification. This information can be found by selecting the partition and using the Properties task. The first tab of the properties box is the General tab. Here, you can view the fields for the load-source adapter and the console adapter. The selections should be the virtual adapters when in the blade environment. Also note that the IPL source is set to D, which uses the *Alternate restart adapter*.



Partition Properties: IBMi (3)

General | Memory | Processing | Ethernet | Storage | Optical/Tape Devices

General

Partition name:

Partition ID:

Environment:

State:

Attention LED:

Settings

IPL source:

Keylock position:

Partition workload group participant: ☐

Automatically start when system starts: ☒

Load Source and Console Identification

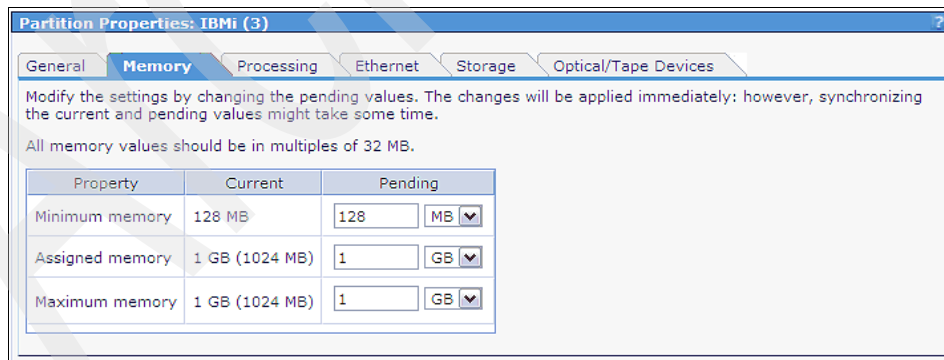
Load source adapter:

Alternate restart adapter:

Console adapter:

Figure 7-29 Load Source and Console Identification fields

Figure 7-30 provides an example of the Memory tab. You may adjust the partition memory allocation by changing the values and clicking **OK**. By setting the minimum and maximum values, you can create a range of memory to stay within when performing dynamic allocation. Changing the Assigned value and clicking on OK would dynamically adjust the memory for the partition provided you stay within the minimum and maximum values. You may change the minimum and maximum values also, but it will cost an IPL of the partition to set the new values.



Partition Properties: IBMi (3)

General | **Memory** | Processing | Ethernet | Storage | Optical/Tape Devices

Modify the settings by changing the pending values. The changes will be applied immediately; however, synchronizing the current and pending values might take some time.

All memory values should be in multiples of 32 MB.

Property	Current	Pending
Minimum memory	128 MB	<input type="text" value="128"/> <input type="text" value="MB"/>
Assigned memory	1 GB (1024 MB)	<input type="text" value="1"/> <input type="text" value="GB"/>
Maximum memory	1 GB (1024 MB)	<input type="text" value="1"/> <input type="text" value="GB"/>

Figure 7-30 Partition memory allocation

Figure 7-31 on page 272 shows the Processing tab. You may adjust the partition processor allocations by changing the values and clicking **OK**. As with the

memory, you may adjust the minimum and maximum values to create a range of processing units to stay within when performing dynamic allocation.

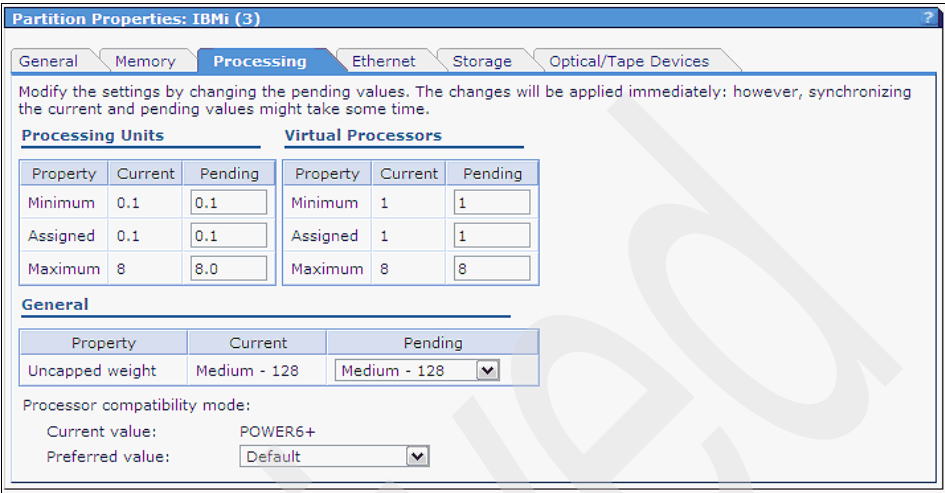


Figure 7-31 Partition processing properties tab

Figure 7-32 shows an example of a modified set of values for processing units. In this example, the values were modified to have 1.2 Assigned units and a maximum value of 4.0 units. Also notice that Virtual Processors values were modified. With these settings, we can have anywhere from .1 processing units to 4.0 processing units. Anytime a change is desired within the minimum and maximum values, that change can be performed dynamically by modifying the Assigned value.




Figure 7-32 Processing units value change

7.3.5 IBM i V6.1 installation of media preparation

Two general methods are available for installing IBM i Licensed Internal Code (LIC) and the 6.1 operating system on a BladeCenter JS32/JS43 blade in an IBM BladeCenter chassis. You may use the CD/DVD drive in the IBM BladeCenter

chassis Media module attached to the IBM i V6.1 partition, or you may create virtual optical media devices.

Using virtual optical media devices is very practical during the installation process of LIC, operating system, or related software because the entire contents of the required CD/DVDs is dumped to a virtual optical image. Therefore, during the installation process you do not have to replace the CD/DVD media physically if a second CD/DVD media is required.

For further information about how to create a virtual optical media device, refer to 7.5.2, “Creating virtual media library using IVM” on page 315.

One other option is to use a remote PC's CD/DVD device and configure it to the AMM in the BladeCenter. This option would allow you to remotely install your operating system from your PC. See 7.5.4, “Attaching a remote PC file or media device” on page 320 for more information.

7.3.6 Connect the System i LAN console

Connecting a local console on a network (LAN) to a system enables you to have an active console and a functional remote control panel. As we previously mentioned, to establish a connection to the IBM i V6.1 partition in the BladeCenter S chassis, the PC must be in the same TCP/IP subnet as the BladeCenter JS23/JS43 blade.

If you experience problems establishing a connection, check the settings of the installed firewall. The ports used by the System i LAN console might have been disabled by the firewall.

To establish a connection to an IBM i V6.1 created partition:

1. In the active Windows session, select **Start → All Programs → IBM System i Access for Windows → Operations Console**. The System i Operations Console window shown in Figure 7-33 on page 274 opens.

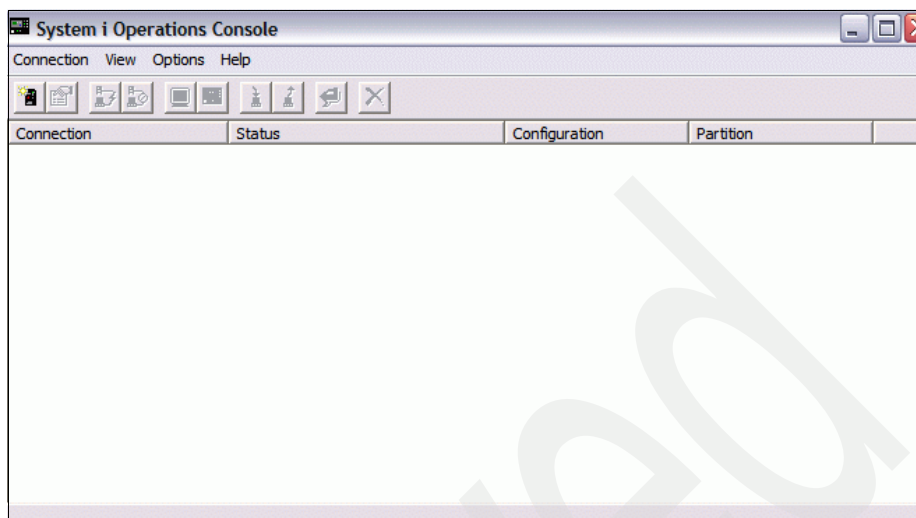


Figure 7-33 IBM System i Operations Console

2. Select **Connection** menu and then select **New Connection** to continue as shown in Figure 7-34.

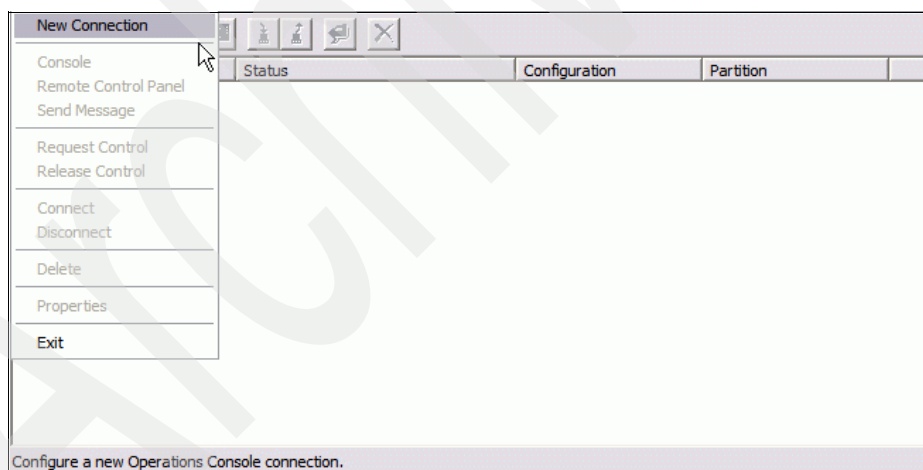


Figure 7-34 IBM System i operator console

3. You reach the Operations Console Configuration wizard welcome window, as shown in Figure 7-35 on page 275.

A connection to the Internet is required to reach the InfoCenter services.

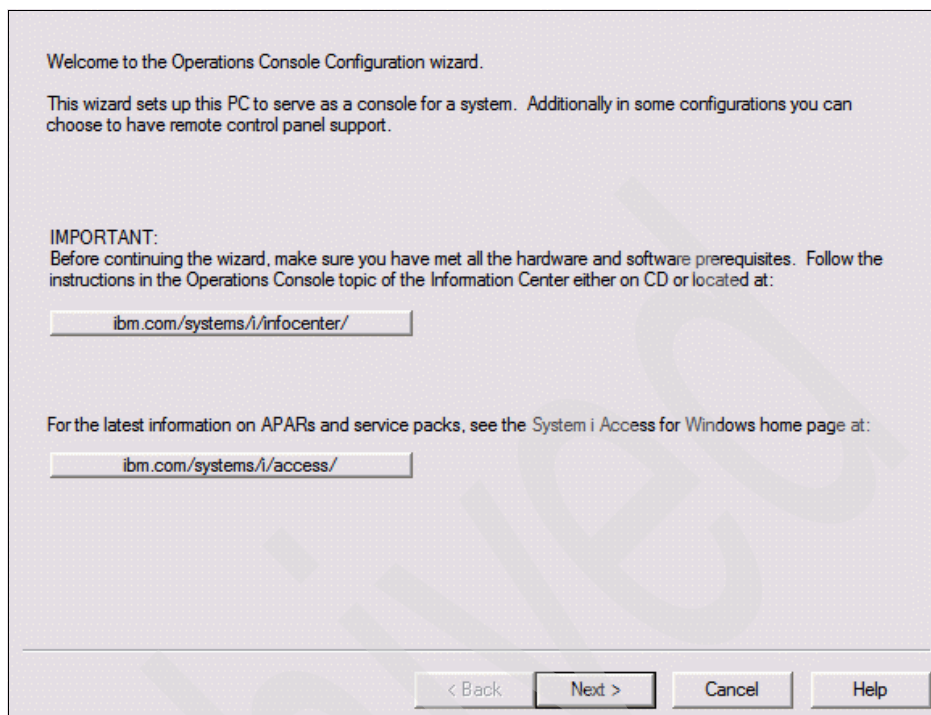


Figure 7-35 IBM System i Operations Console Welcome

Click **Next** to continue.

You might also see a dialog box appear asking you to confirm that the prerequisites for Operations Console have been met. Clicking on the **Help** button will provide the necessary information. If the prerequisites have been met then click **Yes** to continue.

4. Depending on the actual infrastructure or existing systems, a connection can be made by using a direct connection or a LAN connection, as shown in Figure 7-36 on page 276.

In our case, we selected the method **Local console on a network (LAN)**.




Figure 7-36 IBM System i Operations Console - choose a configuration

Click **Next** to continue.

5. The System i service host name must be defined first to establish a connection to the BladeCenter JS23/JS43 blade; see Figure 7-37 on page 277.

The System i service host name (interface name) is the name that identifies the service connection on your network that is used for service tools, which includes an Operations Console local console on a network (LAN) configuration. It is assigned by your system or network administrator and must be resolved through DNS. Because this is the first connection, the service host name information is passed to the service tools LAN adapter.

You must have a service host name (interface name) any time a console or remote control panel is being connected using a network connection. For example, if a server is logically partitioned, then even though the primary might have a non-networked console, having a remote control panel to a secondary partition might be desirable.

Enter the service host name and click **Next**.




Figure 7-37 IBM System i Operations Console - enter the Service host name

6. If this is a first-time connection, the console responds as shown in Figure 7-38 on page 278.

If the Service host name is already defined on the host file of the console PC, then the window will resolve the IP address automatically and fill in the field as shown in Figure 7-39 on page 278.

Note: Choose a service host name that is related to the IBM i V6.1 partition name created in Integrated Virtualization Management (IVM) so that you can more easily remember which partition is meant.

The service host name and service TCP/IP address are stored automatically on the host file of the IBM System i Access for Windows console PC.

For Windows XP environments you will find the hosts file at the following path:

C:\WINDOWS\system32\drivers\etc\hosts




Figure 7-38 IBM System i operator console

Click **OK** to continue.

7. In the next window, shown in Figure 7-39, enter the service TCP/IP address, and then, click **Next** to continue.

Important: Do not use the IP address from the VIO Server here. The System i Access for Windows software establishes a connection to the IBM i partition through the VIOS.

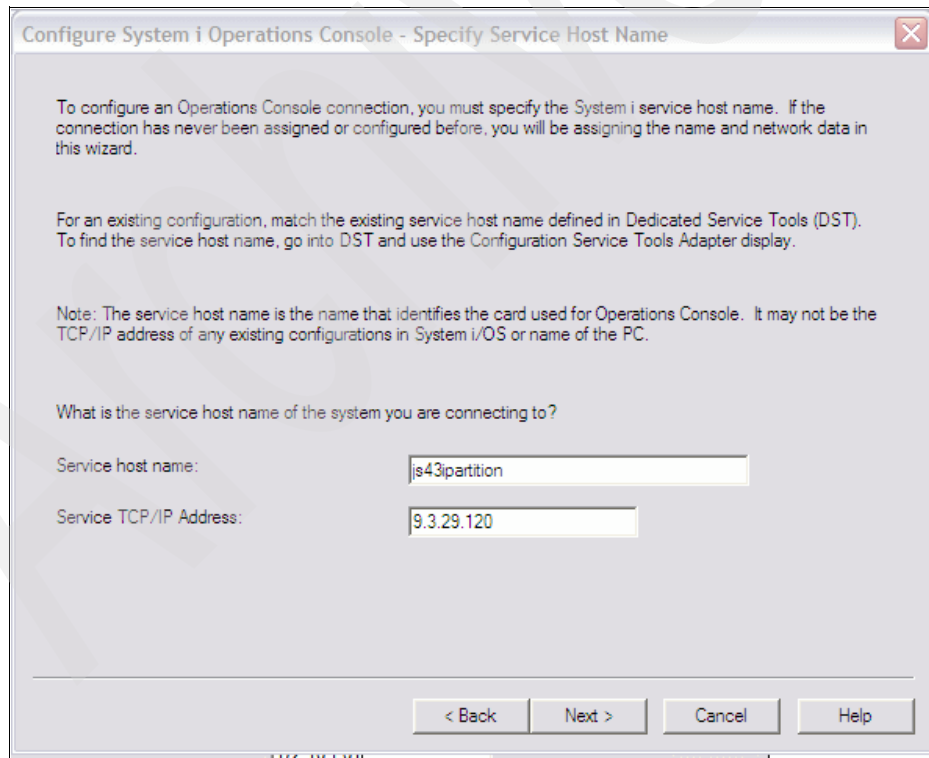


Figure 7-39 IBM System i Operations Console - enter the Service TCP/IP Address

8. After providing an IP address, the wizard asks for the subnet mask, gateway, serial number and target partition ID, as shown in Figure 7-40.

Configure System i Operations Console - Specify Interface Information - JS43IPARTITION

What is the TCP/IP information and serial number of the system to which you are making a service connection?

Service host name: JS43IPARTITION

Service TCP/IP Address: 9.3.29.120

Service subnet mask: 255.255.

Service gateway address 1: 9.3.29.

System serial number:

Target partition:

< Back Next > Cancel Help

Figure 7-40 Specify Interface Information window

9. Modify the required fields to the actual implementation. In our hardware scenario, a gateway was implemented. Two important fields are System serial number and Target partition, also shown in Figure 7-43 on page 281.

- System serial number

This is the BladeCenter JS23/JS43 unique system number. To find the system serial number, use the IVM console and look under System Properties. An example is shown in Figure 7-41 on page 280.

View/Modify System Properties

General Memory Processing

General

System name: Server-7778-63X-SN10181CA

Type/Model: 7778-63X

Serial number: 10181CA

Figure 7-41 System Properties - Serial number

– Target partition

This is the ID of the IBM i V6.1 partition. To see if partition ID 1 is predefined to VIOS, use IVM. If no other partition is created at this time, the IBM i V6.1 partition ID is 2. The partition ID can be found by looking at the View/Modify partition panel. Next to the partition name is the ID field as shown in Figure 7-42. In our example the partition ID is 3.

Partition Details

☑ ☐ ⚙ * Create Partition... Activat

Select	ID ^	Name	State
<input type="checkbox"/>	1	js43-vios	Running
<input type="checkbox"/>	2	js43-lp1	Not Activated
<input type="checkbox"/>	3	IBMi	Not Activated
<input type="checkbox"/>	4	testlpar	Open Firmware

Figure 7-42 Partition ID

As an example in Figure 7-43 on page 281, enter the appropriate values and click **Next** to continue.

Configure System i Operations Console - Specify Interface Information - JS43IPARTITION

What is the TCP/IP information and serial number of the system to which you are making a service connection?

Service host name: JS43IPARTITION

Service TCP/IP Address: 9.3.29.120

Service subnet mask: 255.255.255.0

Service gateway address 1: 9.3.29.1

System serial number: 10181CA

Target partition: 3

< Back Next > Cancel Help

Figure 7-43 IBM System i Operations Console - enter System serial number

10. The next window that appears requests a Service tool device ID to authenticate the communication between the LAN console PC and the IBM i partition, as shown in Figure 7-44 on page 282.

Service tool user IDs are user IDs that are required for accessing service functions through dedicated service tools (DST), system service tools (SST), the System i Navigator (for logical partitions and disk unit management), and the Operations Console. QCONSOLE is the default service tool user ID to be used with console connections. Use ID 11111111 because it is the only service tools ID that is not set as expired.

Service tools user IDs are created through DST or SST and are separate from IBM i V6.1 user profiles. A good practice is to create additional backup service tool IDs after you have installed your partition. For more information about this topic, refer to Service tools user IDs at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzamh/rzamhwhatuserids.htm>

In the window, enter the **Service tool device ID** and click **Next** to continue.

To use restricted system functions securely using a LAN connection, you must specify a service tools device ID. This protects access to system service functions such as the console and remote control panel.

For the first LAN console, QCONSOLE can be used as the default. For subsequent LAN configurations a new service tools device ID must be created on the system and specified here.

Service tools device ID for this PC:

< Back Next > Cancel Help

Figure 7-44 IBM System i Operations Console - enter Service tools device ID

11. Figure 7-45 on page 283 shows the final window that is displayed after you define the information for an IBM System i Operations Console.

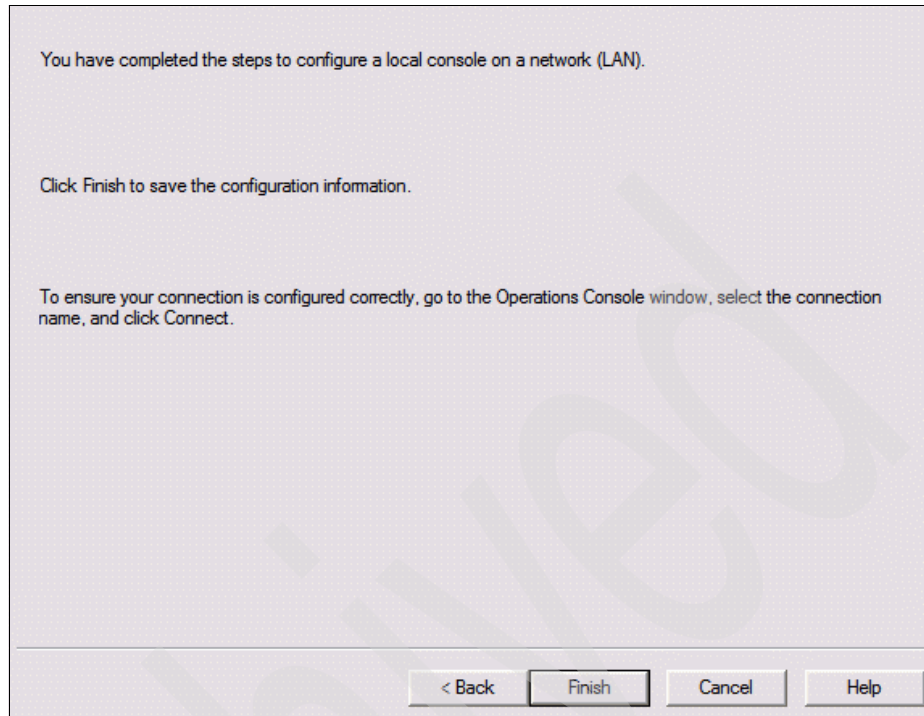


Figure 7-45 IBM System i Operations Console - finalizing the setup

Click **Finish** to save the configuration information.

The configuration window closes immediately and you are returned to the initial window with the predefined console definitions for a BladeCenter JS23/JS43 blade, as shown in Figure 7-46.




Figure 7-46 IBM System i Operations Console

To connect the IBM System i Operations Console to the IBM i V6.1 partition, click the connection name, then click the connect icon or select the **Task Connection** → **Connect**. Figure 7-47 on page 284 shows an example of the

connection icon. After the session starts the connection, the partition can be activated. Partition activation is discussed in the next section.

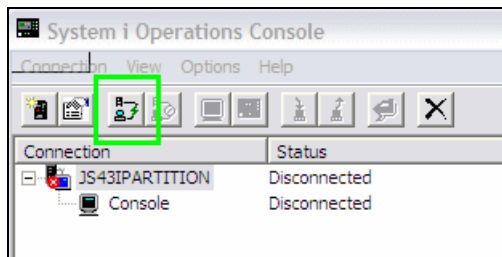


Figure 7-47 Connect console session

7.3.7 IBM i V6.1 IPL types

The IPL type determines which copy of programs your system uses during the initial program load (IPL).

- IPL type A** Use IPL type A when directed for special work, such as applying fixes (PTFs) and diagnostic work.
- IPL type B** Use the B copy of Licensed Internal Code during and after the IPL. This copy resides in System Storage Area B. This copy contains temporarily applied fixes.
- IPL type C** Development support reserves this type of IPL for hardware service representatives.
Note: Do *not* use this function. Data loss can occur with improper use of this function.
- IPL type D** Use IPL type D when directed for special work, such as installing and reloading programs. IPL type D loads the system programs from an alternate IPL load source, such as a tape drive or CD-ROM.
 Typically, an IPL uses programs that are stored on the primary IPL load source (typically a disk drive). However, sometimes it is necessary to perform an IPL from another source, such as programs that are stored on tape. To do this, you must use IPL type D to perform an IPL from the alternate IPL load source.

When configuring the IBM i V6.1 partition, use IPL type D to install and set up the environment. After the Licensed Internal Code (LIC) is successfully installed, the installation process automatically changes the IPL type to A.

Note: Typically after installation of PTFs, you run the partition on the B side. This value is changed on the General tab of the partition properties.

7.3.8 Completing the partition

After the prerequisites are completed, the steps required to install 6.1 on a BladeCenter JS23/JS43 are essentially the same as on any other supported system:

1. Place the IBM i V6.1 installation media in the DVD drive in the BladeCenter media tray, which at this point should be assigned to your BladeCenter JS23/JS43. Or use the previous created Virtual Optical Media device described in 7.5.3, “Adding image files to media library” on page 317.
2. In IVM, select the View/Manage partitions task.
3. Select the IBM i V6.1 partition and click **Activate**.
4. After connecting the IBM System i for Windows LAN console, the panel shown in Figure 7-48 on page 286 appears. The default language feature for English is 2924 but can be changed on the next panel.

Press Enter to continue.




Figure 7-48 LIC initial installation panel

5. Depending on the native language, a selection can be made in the next panel, as shown in Figure 7-49 on page 287.

Normally, you select the same language as the language for the IBM i V6.1 operating system. Language feature 2924 enables the English environment.




Figure 7-49 Confirm Language setup

Press Enter to continue.

6. The next panel displays several options, as shown in Figure 7-50 on page 288. To install the Licensed Internal Code, type 1 and press Enter.





Figure 7-50 Install LIC

7. Select the target installation device. Move the cursor to the target device, type 1 and press Enter; see Figure 7-51 on page 289.



8. Confirm the previous selection of the Load Source Device by pressing F10; see Figure 7-52 on page 290.




Figure 7-52 Confirm source device

- The install Licensed Code (LIC) menu appears on the console as shown in Figure 7-53 on page 291. Type 2 for Install Licensed Internal Code and Initialize system, then press Enter to continue.

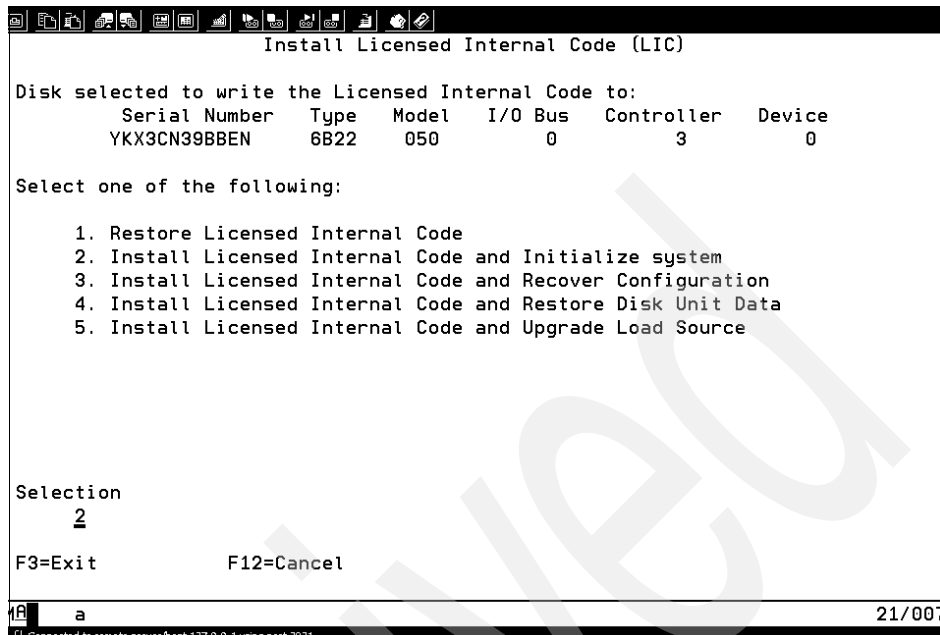


Figure 7-53 Select options

10. The Confirmation panel opens, as shown in Figure 7-54 on page 292. This procedure causes existing data on the disk assigned to this logical partition to be lost. Press F10 to continue or press F12 to cancel and return to the previous panel.




Figure 7-54 Confirm definition

After you confirm the definition, you reach the Initialize the Disk status panel, as shown in Figure 7-55 on page 293. Depending on the predefined size of the virtual disk, this procedure can take 60 minutes or more.

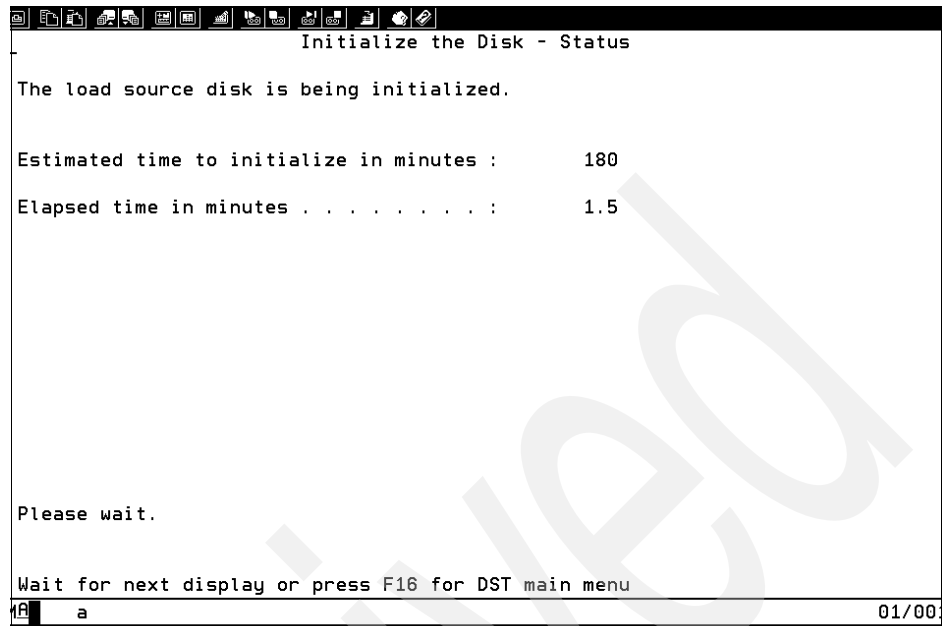


Figure 7-55 Initialize disk

11. Next, the Install Licensed Internal Code status display appears on the console as shown in Figure 7-56 on page 294. It remains on the console for approximately 30 minutes.

After the LIC has completed installing, the logical partition is automatically restarted to IPL to DST at this time to complete the Licensed Internal Code installation.

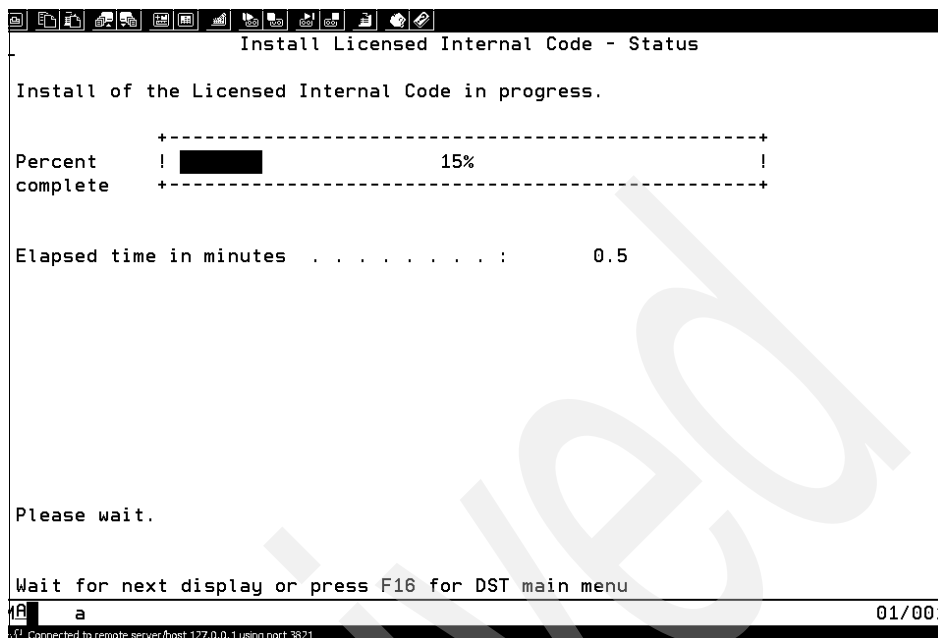


Figure 7-56 Install LIC status

12. The Disk Configuration Attention Report display might appear on the console.

Figure 7-57 on page 295 shows the report for a new disk configuration.

Press F10 to accept the action to define a new disk configuration.

Note: If the Disk Unit Not Formatted For Optimal Performance Attention Report appears on the console, then further actions should be performed as described in the InfoCenter Web site:

<http://publib.boulder.ibm.com/infocenter/iserics/v5r4/index.jsp?topic=/rzahc/rzahcdiskw.htm>

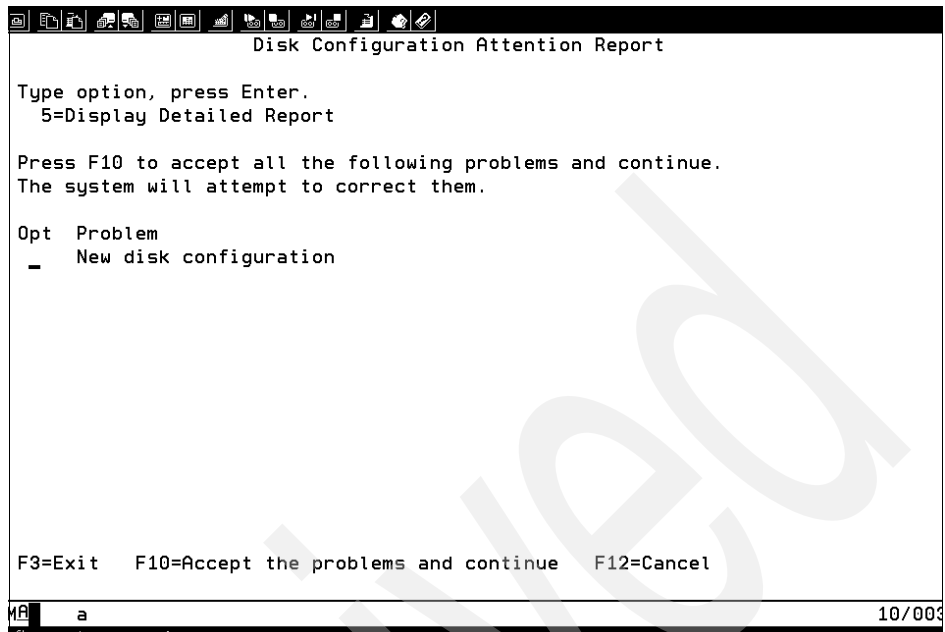


Figure 7-57 Attention Report

After the Licensed internal Code installation is complete, you see the panel shown in Figure 7-58 on page 296.

At this time, a good practice is to complete disk-unit configuration before installing the operating system. When completing disk configuration, you add more units and possibly start mirroring on the disk units.

Go to the following Web address to assist with performing disk configuration. Not all steps have to be performed.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzarm/rzarmrcvaft.htm&tocNode=toc:rzahg/i5os/17/0/5/5/4/>

Note: Disk configuration is not a required action at this time but should be used in the case of a failure that might cause a reload. It is much quicker to install LIC than it would be to install LIC and the operating system again.

After completing disk configuration you may continue the installation process for IBM i. Select option 2 to install the IBM i operating system after disk configuration steps are completed.

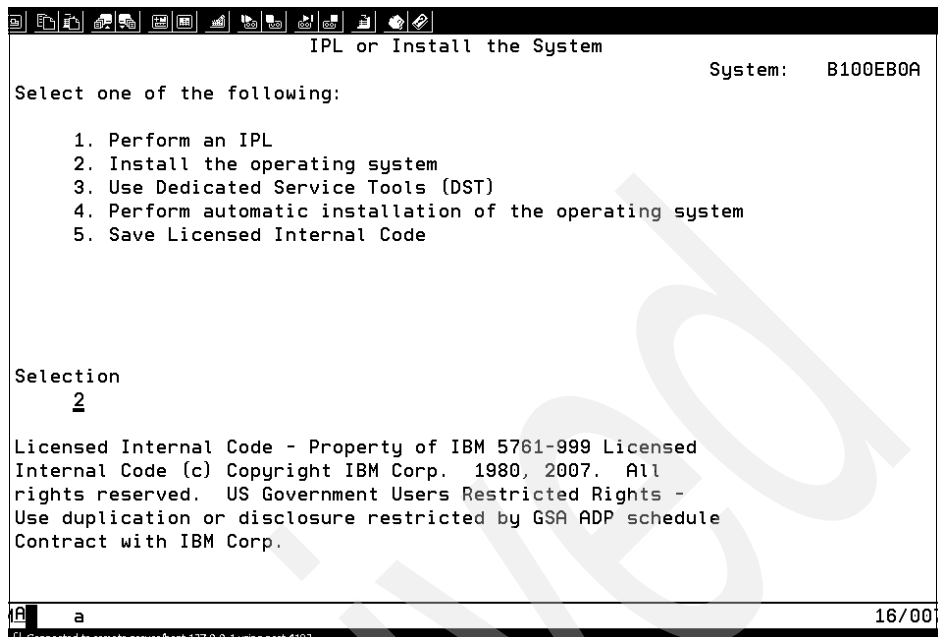


Figure 7-58 Install the operating system

7.4 Installing and setting up IBM i V6.1

From the IPL or Install the System window, the installation process of the operating system can be continued without an interruption. If you use the virtual optical device method of having the two IBM i V6.1 DVDs previously unloaded to virtual optical devices, the only action necessary is to assign the virtual optical device with the IBM i DVD 1 content to the IBM i partition.

To install the IBM i V6.1 operating system:

1. On the IPL or Install the System panel on the console, type the number that corresponds to the correct device type as shown in Figure 7-59 on page 297. In our case we used option 2 for the assigned virtual optical devices.

If a virtual optical device is not defined, then replace the Licensed Internal Code DVD with the IBM i DVD 1 in the CD/DVD drive in the media tray.

Type 2 and press Enter to continue.




Figure 7-59 Select install device

2. The Confirm Install of the Operating System panel is displayed on the console screen, as shown in Figure 7-60 on page 298.

Press Enter to continue the installation process.




Figure 7-60 Confirm installation

3. The Select a Language Group panel displays the primary language preselection, as shown in Figure 7-61 on page 299. This value should match the Language feature number that is printed in the installation media.

The following Web address provides the language feature codes:

<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/index.jsp?topic=/rzahc/rzahcnlvfeaturecodes.htm>




Figure 7-61 Select language feature

4. Type your choice and press Enter to continue.

The Confirm Language Feature Selection panel appears on the console, as shown in Figure 7-62 on page 300.

Press Enter to confirm and continue.

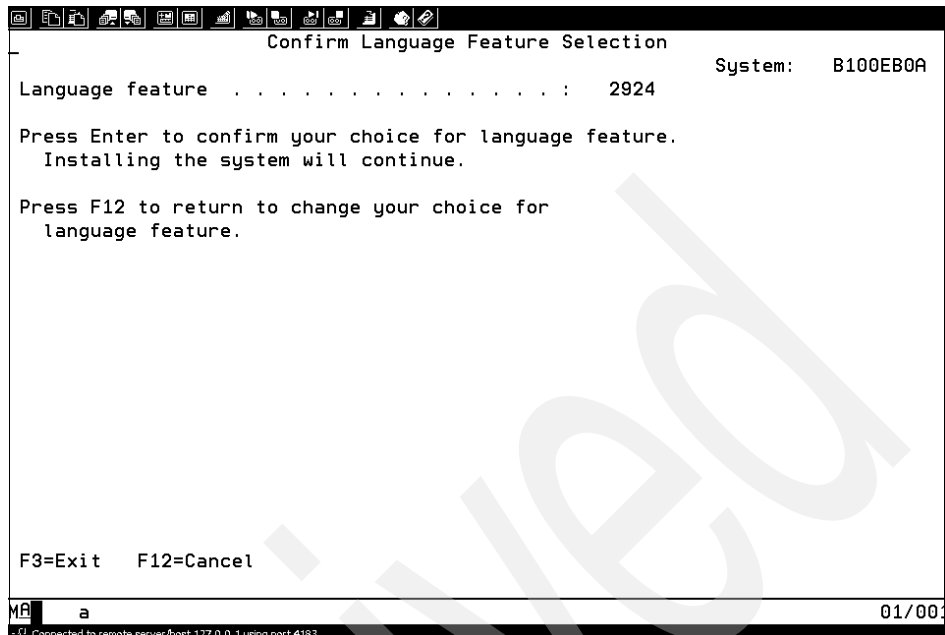


Figure 7-62 LIC install confirm language

5. The Licensed Internal Code IPL in Progress panel appears on the console, as shown in Figure 7-63 on page 301. No administrator action is required.

```

Licensed Internal Code IPL in Progress                                06/17/08  00:03:57

IPL:
  Type . . . . . : Attended
  Start date and time . . . . . : 06/17/08  00:03:23
  Previous system end . . . . . : Abnormal
  Current step / total . . . . . : 12      16
  Reference code detail . . . . . : C6004059

IPL step                                Time Elapsed   Time Remaining
Data Base Recovery                      00:00:13      00:00:00
Journal Synchronization                  00:00:01      00:00:00
>Commit Recovery
Data Base Initialization
Journal IPL Clean up

Item:
  Current / Total . . . . . :

Sub Item:
  Identifier . . . . . :
  Current / Total . . . . . :

01/001
Connected to remote server/host 127.0.0.1 using port 4183
```

Figure 7-63 IPL in progress

6. The Install the Operating System panel appears on the console, as shown in Figure 7-64 on page 302.

Change the date and time values to the appropriate settings. You must use the 24-hour clock format to set the current time.

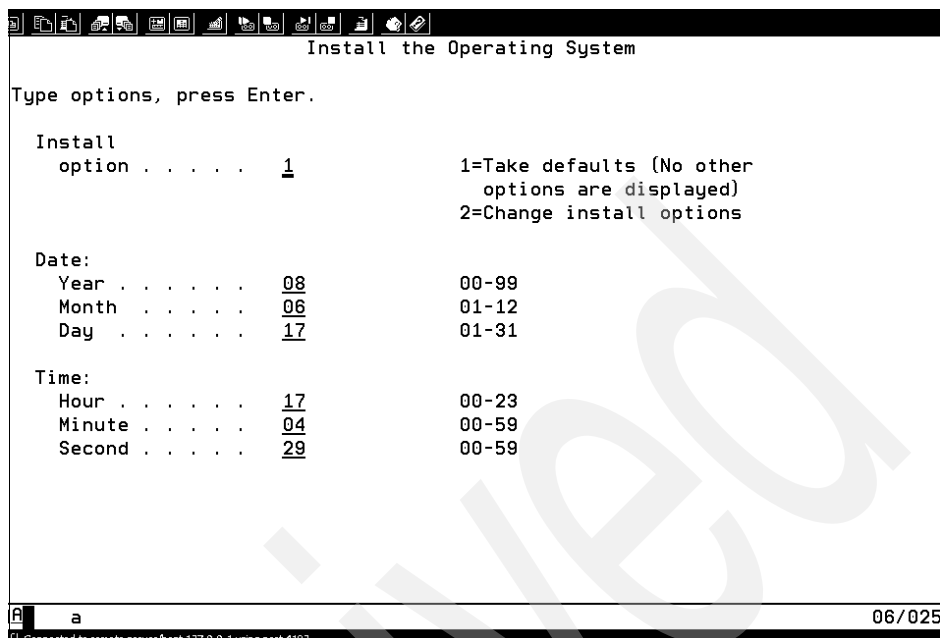


Figure 7-64 Set date and time

7. Figure 7-65 on page 303 shows an example of a status display in the operator console during the installation process. No further action required. Note that the display will be blank for awhile between Installation Stage 4 and 5.




Figure 7-65 Installation status

8. When the Sign On panel is displayed, as shown in Figure 7-66 on page 304, the base installation of the IBM i 6.1 Operating System is completed.




Figure 7-66 Sign On panel

At this stage, the IBM i V6.1 system is ready to use. Information about installing libraries or Licensed Program Products and system configuration is beyond the scope of this book. For detailed software installation information, refer to the following Web site:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzam8/rzam81.htm>

7.4.1 IBM i V6.1 installing PTFs

For detailed explanations about the following tasks, refer to the corresponding sources on the Web.

Use the IBM i recommended fixes Web site to obtain a list of the latest fixes:

http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixes

Refer to the primary Web site for downloading fixes for all operating systems and applications, Fix Central:

<http://www-933.ibm.com/support/fixcentral/>

For IBM i V6.1, IBM i5/OS, or OS/400® Operating Systems, fixes are available. To obtain an IBM i V6.1 fix overview for downloading:

1. Select **System i** in the Product family field.
2. Select **IBM i, i5/OS, and OS/400** in the Product field.
3. Select one the following options in the Ordering option field:
 - Groups, Hyper, Cumulative fixes
 - Individual fixes
 - Search for fixes
4. Select, for example, **V6R1** in the OS level field for fixes for the actual IBM i Operating System version.

Note: To download fixes or obtain information about fixes from the Fix Central Web site, you must have a valid IBM ID and password.

7.4.2 IBM i V6.1 TCP/IP setup

Ethernet on a BladeCenter JS23/JS43 server supports TCP/IP, Advanced Peer-to-Peer Networking (APPN), Advanced Program-to Program Communication (APPC), and remote workstation.

You can use this information to configure and manage TCP/IP on the BladeCenter JS23/JS43 server. To configure the IBM i V6.1 communication, refer to the Ethernet on System i V6.1 topic in the IBM Information Center:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzai2/rzai2configuring.htm&tocNode=toc:rzahg/i5os/10/5/5/>

Note: IBM i V6.1 installed on a IBM BladeCenter JS23/JS43 server will communicate to the external LAN using Virtual Ethernet Adapter only, and only those with the Virtual Ethernet bridge enabled.

7.4.3 IBM i V6.1 Navigator for i

The Navigator for i or IBM Systems Director Navigator for i is a Web-based console that consolidates all Navigator for i functions available on the Web. IBM i Systems Director Navigator for i is the tool for IBM i V6.1 management and can work with IBM Systems Director products that enable heterogeneous and cross-system management.

The IBM Systems Navigator for i provides a graphical interface to manage a BladeCenter JS23/JS43 server or Power Systems, as shown in Figure 7-67.

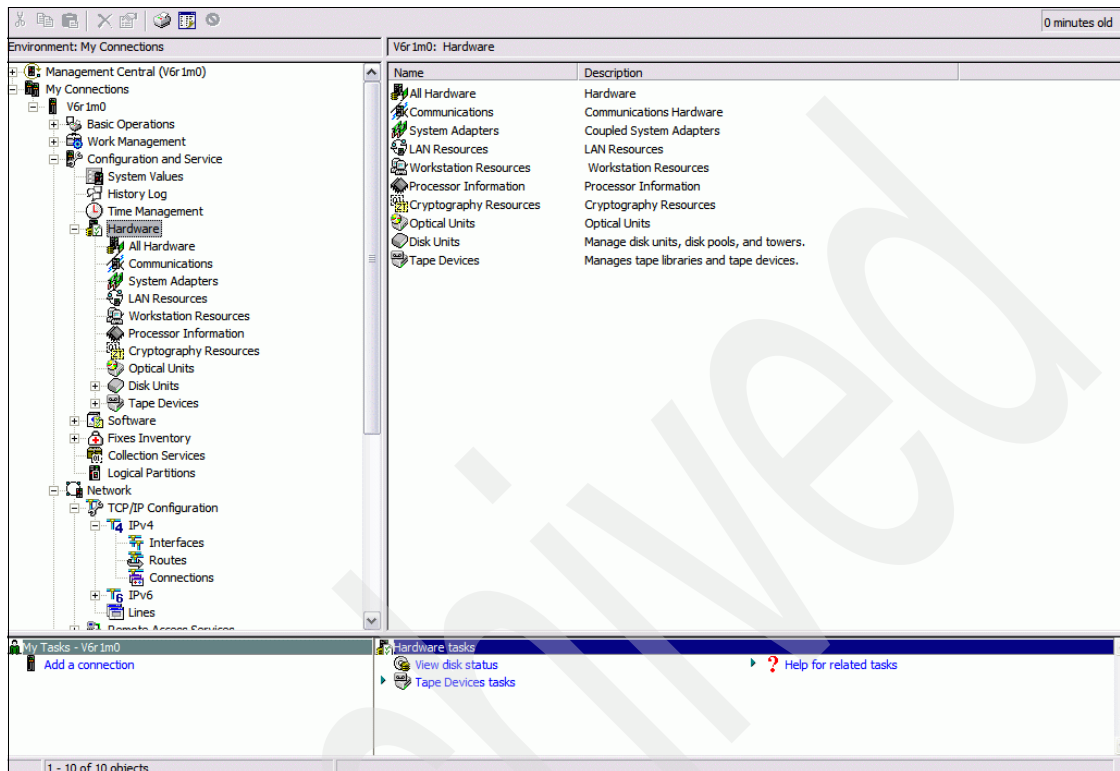


Figure 7-67 IBM Systems Navigator for i

For more information about IBM Systems Director Navigator for i functionality refer to:

- Navigator for i Web site:
<http://www.ibm.com/systems/i/software/navigator/index.html>
- *Managing OS/400 with Operations Navigator V5R1 Volume 1: Overview and More*, SG24-6226

7.5 IBM i V6.1 backup and restore

Two methods are available to backup or restore an IBM i partition:

Important: The virtualized DVD-ROM drive in the chassis *cannot* be used for IBM i V6.1 backups, because it is not writable.

- ▶ One method is to use file-backed space provided as a virtual optical device. After the file has been created it can be written to any BSH or BCS supported SAS tape device.
- ▶ Another method is to use a virtual tape device backed by a SAS tape drive that is virtualized by VIOS. The virtual tape drive method is the newer method. This option is only supported using the TS2240 LTO 4 SAS tape drive attached to a SAS switch module; any other tape drive will not work. When the tape drive is virtualized to the IBM i partition a second virtual SCSI adapter is created for the tape drive connection.

Note: As of this writing, the TS2240 LTO 4 SAS tape drive is supported for IBM i in the BCH and BCS. This tape drive can be virtualized to the IBM i partition and will appear as a 3580 Model 004.

For further information about using the IBM SAS LTO tape library attached to a BladeCenter S chassis, refer to the readme guide:

http://www.ibm.com/systems/resources/systems_power_hardware_blades_i_on_blade_readme.pdf

IBM i V6.1 backup to virtual optical device

Performing an IBM i V6.1 operating system backup is a three-stage process on the IBM BladeCenter JS23/JS43 blade:

1. Create a virtual media library. See section 7.5.1, “Creating a virtual media library for backup” on page 309 for more information.
2. A standard IBM i V6.1 **save** command or Backup, Recovery, and Media Services (BRMS) is used to perform a save on a writable optical device that contains an optical volume. The writable optical device is a file-backed virtual optical drive created in VIOS. The optical volume is a virtual media image, which is a file on disk in VIOS. The file location is `/var/vio/VMLibrary/filename`.
3. The virtual media image, containing the 6.1 save, is written out to a SAS-attached tape drive using the VIOS command **backup**.

IBM i V6.1 restore of virtual optical device

Restoring follows the same two-stage process in reverse:

1. The virtual media image file is restored from the SAS tape drive onto VIOS disk by using the VIOS command **restore**. The image file is then mounted on the correct virtual optical drive assigned to the IBM i V6.1 partition and becomes available as a volume from which to restore.
2. A standard IBM i V6.1 restore is performed from the volume using a **restore** command or BRMS. A full system restore can be executed by first performing a D-mode IPL from the virtual optical image, provided the image contains a full system save performed previously.

IBM i V6.1 backup to virtual tape device

Performing an IBM i V6.1 operating system backup to virtual tape is a two-stage process on the IBM BladeCenter JS23/JS43 blade:

1. Ensure that the virtual tape device is assigned to the partition you are performing the backup on. To change or view the assignment use the **View/Modify Virtual Storage** task, then select the **Optical/Tape** tab. Under the **Physical Tape Devices** section, select the tape drive and click **Modify Partition Assignment**.

Figure 7-68 shows an example of the Physical Tape Devices option.

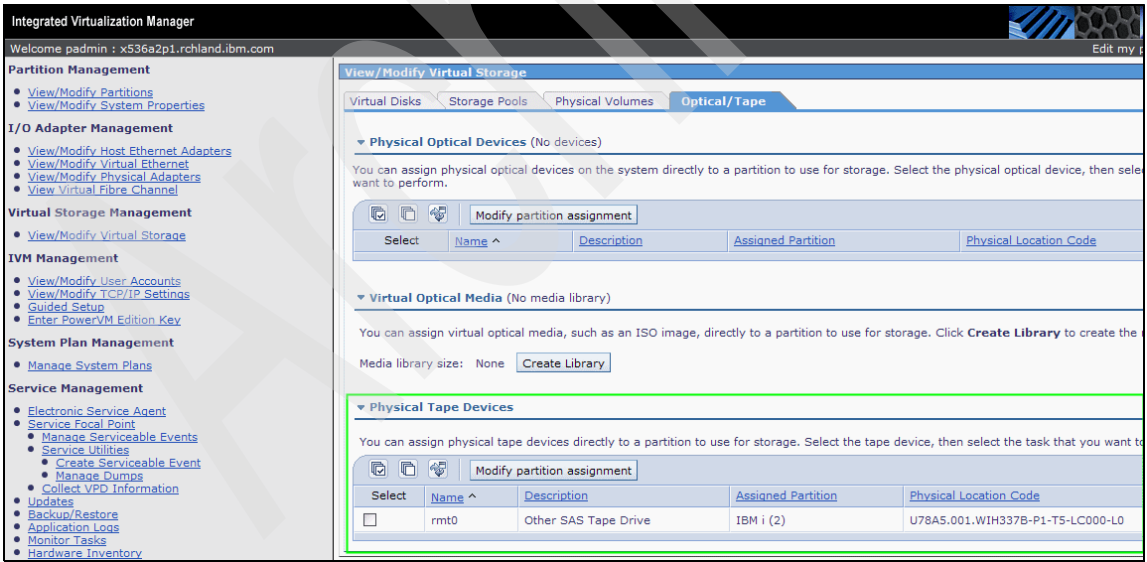


Figure 7-68 Physical Tape Devices

2. A standard 6.1 **save** command or BRMS is used to perform a save on the tape device (tap0x). If autocfg is set to *on*, the tape device will configure as an 3580 model 004.

IBM i V6.1 restore of virtual tape device

Restoring follows the same two-stage process.

1. Ensure the virtual tape device is assigned to the partition you are performing the backup on. To change or view the assignment use the **View/Modify Virtual Storage** task, then select the **Optical/tape** tab. Under the **Physical Tape Devices** section select the tape drive and use the **Modify Partition Assignment** button.
2. Use the standard 6.1 **restore** command and restore the image from the tape device. The tape device can also be used to perform a D-IPL and full system restore.

7.5.1 Creating a virtual media library for backup

This section describes how to create a virtual media library for storing the IBM i 6.1 operating system backup, as follows:

1. Log on to the IVM on the BladeCenter JS23/JS43.
2. Select **View/Modify Virtual Storage**, the **Storage Pool** tab, and then click **Create Storage Pool**. The window displayed in Figure 7-69 on page 310 is opens.

Note: A best practice is to use a separate storage pool. However, rootvg can be used if enough free space exists.

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"/>	hdisk5	279.4 GB	U4545.001.07CM0ST-P2-D2

* Required field

OK Cancel

Done

Figure 7-69 IVM Create Storage Pool

3. Enter a name for the storage pool (in our case the internal disk in the BladeCenter S disk module was used), or in a SAN environment, a predefined LUN. Click **OK** to continue.
4. To create the virtual media library, select the **Optical Devices** tab and click **Create Library**.
5. Select the name of the new storage pool and enter an appropriate size for the media library. Click **OK** to continue.
6. To add a new virtual optical device to the media library, select **Add Media** in the Virtual Optical Media section. The window, shown in Figure 7-70 on page 311, opens.

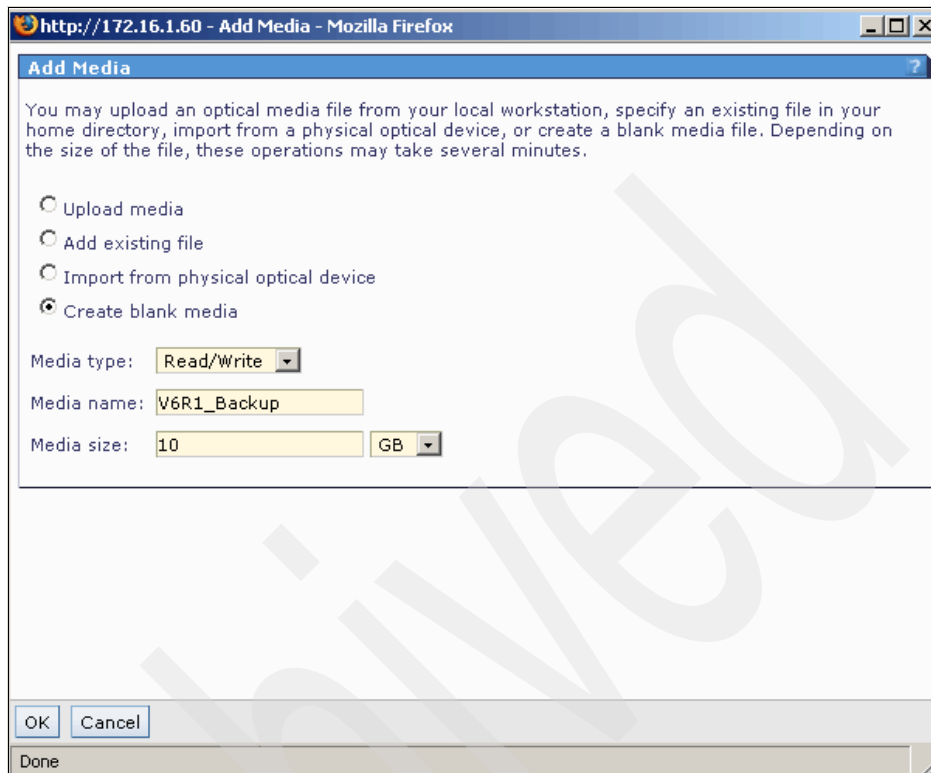


Figure 7-70 IVM Create blank media

7. Select **Create blank media** and enter a meaningful media device name and an appropriate size for the new volume. Ensure the media type is set for read/write. Click **OK** to continue.
8. The new virtual optical device should be listed in the Virtual Optical device list, as shown in Figure 7-71 on page 312.




Figure 7-71 iVMM Virtual optical device created

To assign the new created virtual optical device to the IBM i V6.1 partition, select the virtual optical device and click **Modify partition assignment** as shown in Figure 7-72 on page 313.

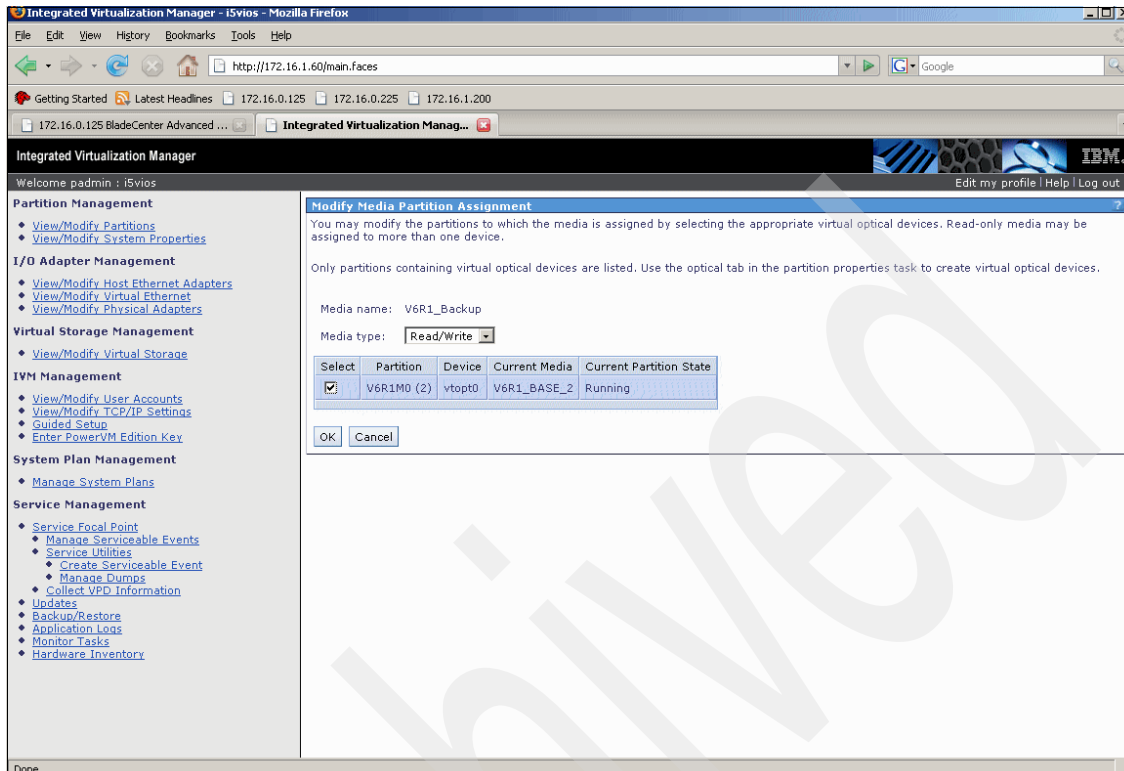


Figure 7-72 Virtual optical device assign to partition

9. Select the IBM i V6.1 partition and click **OK** to continue.

Figure 7-73 on page 314 shows the IVM Virtual Storage Management window with the current assignment of the virtual optical device to the partition.

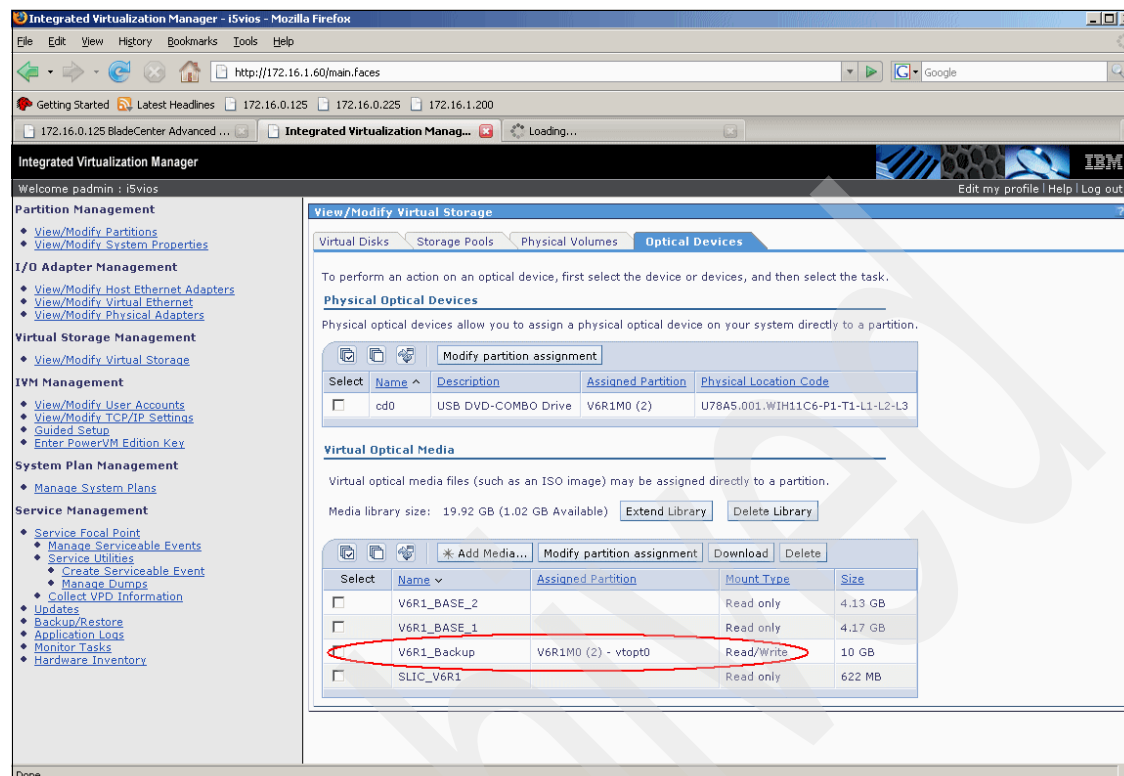


Figure 7-73 iVMM Virtual optical device assignment done

After the virtual optical device is mounted to the correct virtual optical device, it becomes available in the IBM i V6.1 partition. The IBM i V6.1 operating system does not use the device name of the virtual optical device given in Integrated Virtualization Manager.

Execute the IBM i V6.1 command WRKOPTVOL; the panel shown in Figure 7-74 on page 315 opens. The virtual optical device is identified with a time stamp volume ID.

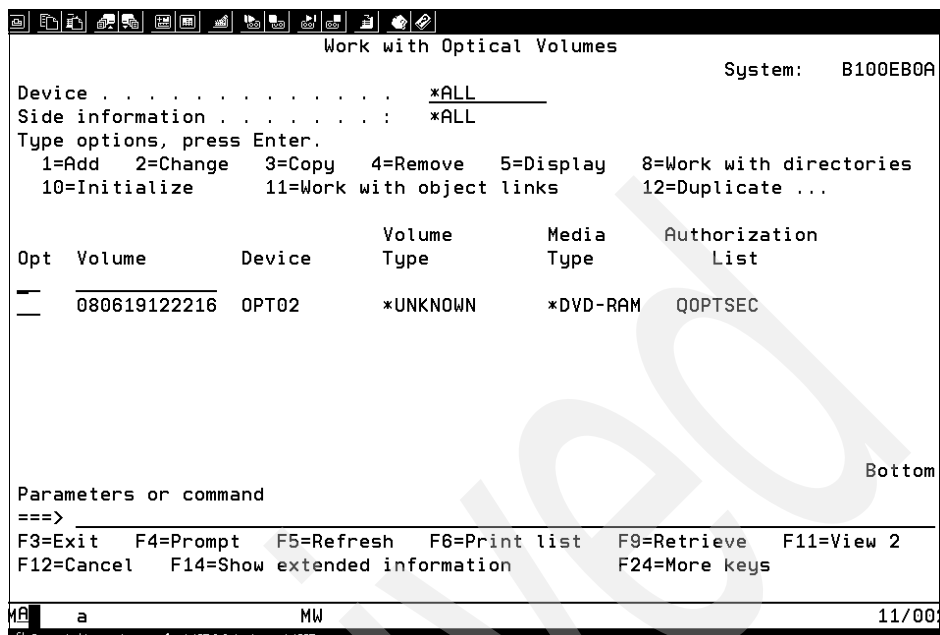


Figure 7-74 Virtual optical device check device

7.5.2 Creating virtual media library using IVM

This section describes the process to create a virtual media library using IVM. This library is created using the IVM options and is located in the directory /var/vio/VMLibrary. After the library has been created you can add files such as ISO images to perform installations of partitions.

To create a virtual media library:

1. First create a storage pool to contain the virtual optical library, as follows (if a usable storage pool already exists, skip to step 2 on page 316):
 - a. Use the IVM options to select **View/Modify virtual storage**.
 - b. Select the **Storage Pools** tab.
 - c. Click on **Create Storage Pool**.

Figure 7-75 on page 316 shows an example of the Storage Pools tab highlighting the Create Storage Pool option.




Figure 7-75 Create Storage Pool option

- d. Provide a storage pool name.
- e. Select the option **Logical Volume Based** for storage pool type.
- f. Select one of the available hdisk resources to create the storage pool on.

Figure 7-76 provides an example of the storage pool name, size and hdisk selection.

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 storage.

* 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 type="checkbox"/>	hdisk9	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010401F000000
<input checked="" type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104006000000
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104007000000
<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104008000000
<input type="checkbox"/>	hdisk5	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400D000000
<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400E000000
<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400F000000
<input type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104010000000
<input type="checkbox"/>	hdisk10	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104020000000
<input type="checkbox"/>	hdisk11	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104021000000
<input type="checkbox"/>	hdisk12	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104022000000
<input type="checkbox"/>	hdisk13	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104023000000

* Required field

Figure 7-76 Selecting storage pool name, size and resource

2. Now that the storage pool has been created, the virtual media library can be created using the new storage pool, as follows:
 - a. Use the IVM options to select **View/Modify virtual storage**.
 - b. Click on the **Optical/Tape** tab.
 - c. Expand the section **Virtual Optical Media**.

d. Click on **Create Library**.

Figure 7-77 shows an example of the Create Library option.



Figure 7-77 Create Media Library

e. Define the media library size.

Figure 7-78 shows an example of the storage pool name field. Select the correct storage pool to contain your virtual media library.




Figure 7-78 Select storage pool name

f. Click **OK** to finish.

7.5.3 Adding image files to media library

One other method of adding media files to the library can be accomplished by creating an ISO image of the media and using FTP to copy the file to the JS23/JS43. If your media library has not been created yet, refer to 7.5.1, “Creating a virtual media library for backup” on page 309.

For example, this process could be used when installing the IBM i OS to a new IBM i partition that has not been loaded, as follows:

1. Create ISO image files of the installation media:
 - a. Load the IBM i SLIC media in your PC CD-ROM.
 - b. Using Record Now or another burning program, create an ISO image of the CD. Usually this is performed using a backup function.

The next several figures show examples of using Record Now to create an ISO image of your media.

Figure 7-79 shows the Save Image option. The option is found under the Backup Projects task.



Figure 7-79 Save Image - Create ISO file

Figure 7-80 on page 319 shows an example of selecting the source and destination folders. Click **Browse** to define the output folder.

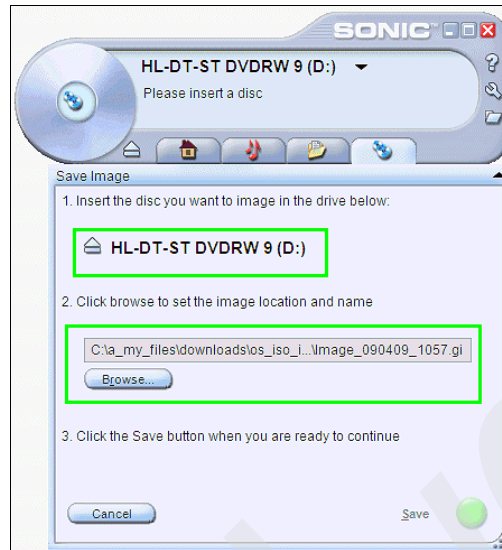


Figure 7-80 Select the output destination folder

Figure 7-81 shows an example of the destination folder. Select the **Save as type** option and ensure the save is set for .iso type. This type is not the default so you have to change it in the Save as type box.




Figure 7-81 Change file type to .iso

2. Copy the .iso image file to the JS23/JS43 using FTP. The file is copied to the /home/padmin directory. Make sure to use image mode when copying the file with FTP so that the files is transferred in binary format.
3. Move (**mv** command) the .iso file from /home/padmin to /var/vio/VMLibrary. You have to use oem_setup_env to escape the VIOS restricted shell environment to be able to use the **mv** command. You should also change the file name so it is easier to identify the files. For example, the file name `ibmi_slic.iso` indicates that the image is the Licensed Internal Code CD necessary for a D-IPL (stand-alone IPL).

7.5.4 Attaching a remote PC file or media device

Using the AMM you can attach a remote PC file or the remote PC media device to the AMM for use as an input device for files. This feature creates a CD device under the AMM for assignment to one of the blade servers. The assignment is performed through the media tray of the BladeCenter. After creating the link between the AMM and the remote PC you must assign the BladeCenter media tray to the blade server you are intending to use the file on.

After the media tray has been assigned to the blade you can use the View/Modify Virtual Storage task and the Optical/Tape tab to assign the CD device to the partition you are working with. The remote PC file or CD-ROM device (or both the PC file and CD-ROM device) appear as an additional CD device assignable to a partition.

The figures in this section demonstrate how to attach the remote PC file and CD-ROM. After the file or CD has been assigned, the device or file can be used by one of the blade partitions by modifying the partition assignment.

To attach a remote PC file or media device:

1. Open a session to the AMM you want to work with.
2. Select the **Blade Tasks** option.
3. Select the **Remote Control** option.

Figure 7-82 on page 321 shows an example of the AMM interface and the blade tasks selected.

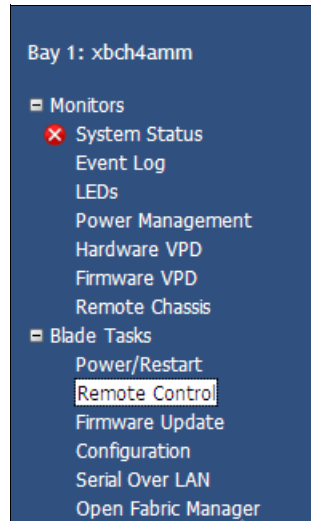


Figure 7-82 Blade Task - Remote Control

4. After the task Remote Control has been selected, use the **Start Remote Control** button, which invokes a Java™ window.

Figure 7-83 shows an example of the Start Remote Control button.

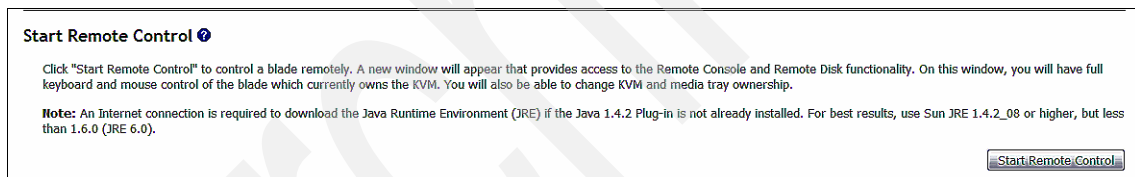


Figure 7-83 Start Remote Control

5. After the java interface has started, select the **Remote Drive** option.

Figure 7-84 shows an example of the java interface for remote control.

Note: A best practice is to leave the Java interface window running until you have completed the media installation action.




Figure 7-84 Remote Control Java window

6. The Remote Disk window opens. Select either or both of the following options: **CD ROM** or **Select Image**. Figure 7-85 provides an example of the Select Image option.

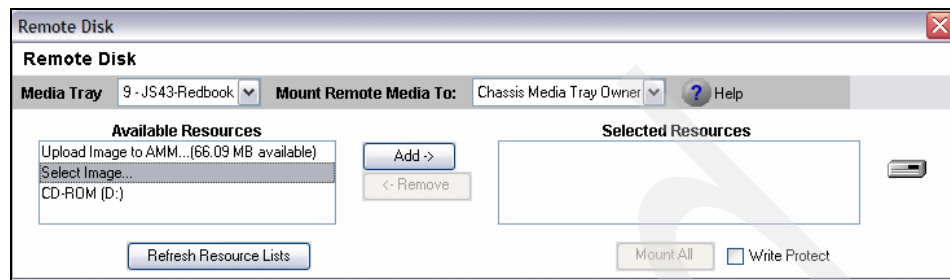


Figure 7-85 Select image option

7. Click **Add**. You may now browse for the specific file you want to add, as shown in Figure 7-86.

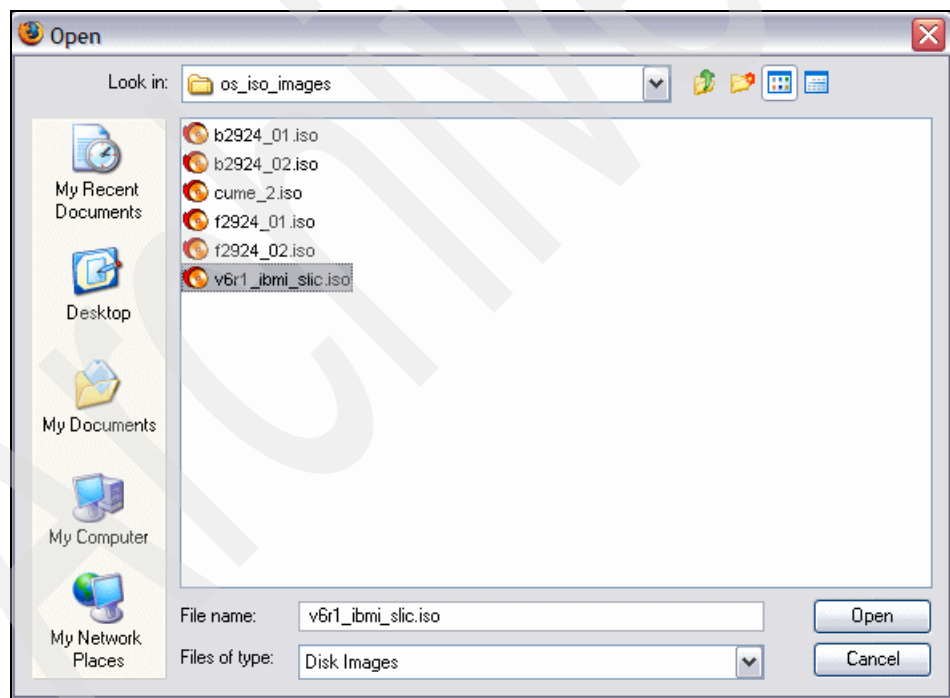


Figure 7-86 Browse and select file

After the file has been added it is listed in the Selected Resources list. Figure 7-87 on page 323 provides an example of this view.

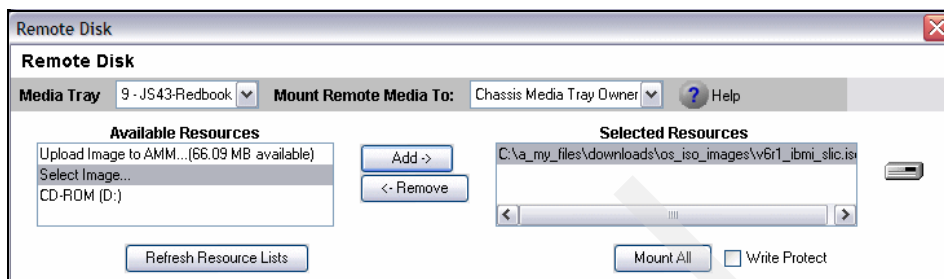


Figure 7-87 File added to Selected Resources list

8. To add the CD-ROM, select the CD-ROM listed and click **Add**. It is then listed under the Selected Resources list.
9. After all selections have been made, click **Mount all** to add your resources to the AMM and make them available to the blade that has the media tray selected.

Figure 7-88 shows an example of the panel after the devices have been mounted. Using the *Unmount All* option would remove the devices from the media tray, however be sure the partition assignment is clear before you remove the resource using the *Unmount* option.

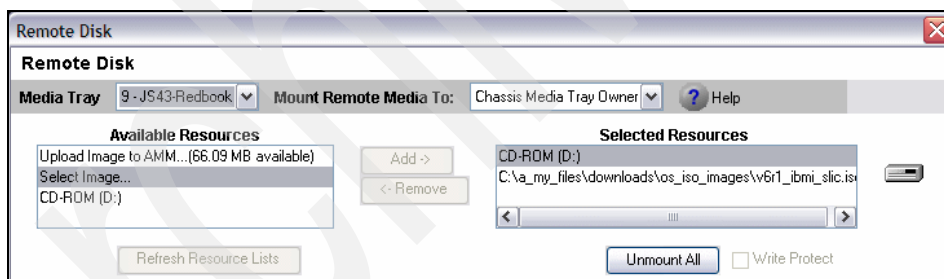


Figure 7-88 Selected Resources mounted

10. After mounting the resources and assigning the media tray to the blade, run the **cfgdev** command to complete the assignment of the media to the blade. This command is run from the command line of the blade that has the media tray assignment.
11. After **cfgdev** has completed, use the **View/Modify Virtual storage** task. Then, select the **Optical/Tape** tab. The new resources will appear as physical optical devices. Modify the partition assignment for the device to add it into the partition configuration for use.

Figure 7-89 on page 324 shows an example of the resources added using the above process.

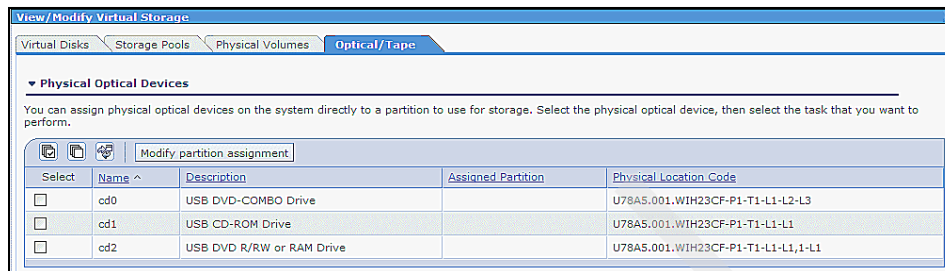


Figure 7-89 New physical optical devices

7.5.5 IBM Tivoli Storage Manager

Starting with Integrated Virtualization Manager V1.4, you can install and configure the IBM Tivoli® Storage Manager (TSM) client on the Virtual I/O Server (VIOS). With IBM Tivoli Storage Manager, you can protect your data from failures and other errors by storing backup and disaster recovery data in a hierarchy of offline storage.

IBM Tivoli Storage Manager can help to protect computers running a variety of operating environments, including the VIO Server, on a variety of hardware. Configuring the IBM Tivoli Storage Manager client on the Virtual I/O Server enables you to include the Virtual I/O Server in your standard backup.

The TSM client software is included in the VIO Server installation images by default. To ensure the IBM Tivoli Storage Manager client is installed on the VIO Server, run the command shown in Example 7-1. This command is performed outside of the VIOS restricted shell environment. The command `lssw` can be used inside of the shell to see similar results.

Example 7-1 TSM client check

```
ls1pp -L| grep TSM
5.4.0.0    C    F    TSM Client - Application
5.4.0.0    C    F    TSM Client - Backup/Archive
5.4.0.0    C    F    TSM Client - Backup/Archive
5.4.0.0    C    F    TSM Client - IMAGE Backup
....
```

Using the IBM Tivoli Storage Manager does not eliminate having to perform an IBM i V6.1 save operation on a virtual optical media device in IVM.

Providing details of configuring and using the IBM Tivoli Storage Manager client and server is beyond the scope of this book. For detailed information about how to configure and manage the VIO Server as a IBM TSM client, refer to:

[http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphb1/iphb1tivagents.htm?resultof="tivoli"](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphb1/iphb1tivagents.htm?resultof=)

For more technical information about integrating IBM Tivoli Storage Manager, refer to *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

7.5.6 IBM i V6.1 shutdown and restart

This section describes the shutdown and restart procedure on an IBM i V6.1 partition.

Shut down an IBM i V6.1 partition

Using the View/Modify Partitions - Shutdown option to shutdown an IBM i partition is not advisable. Although you may use this option, it is similar to using the white power-off button on a true (physical) system. If you decide to use this option, do not select Immediate for the shutdown type.

Before you shut down an IBM i V6.1 logical partition, ensure that all jobs are completed and all applications are ended. The proper way to do this is to end all the subsystems.

To access the IBM i OS command line, perform one of the following options:

- ▶ Open a Console connection using the System i Access for Windows software.
- ▶ If TCP/IP is already configured, you can use a Telnet 5250 session; however if you are using Telnet, the session will be dropped when the TCP/IP subsystem and running jobs end. You would have to establish the LAN console session to complete the proper shutdown sequence.

To shut down the IBM i 6.1 partition:

1. Open a console connection to the IBM i partition.
2. Sign in with a user ID that has the privilege to shut down the partition.
3. End the running subsystems by issuing the ENDSBS *ALL command.
4. Monitor the message queue by using the command DSPMSG QSYSOPR.
5. Wait for the message System ended to restricted condition, which might take several minutes to appear while various jobs and subsystems end. You might have to refresh the display by pressing F5.

6. Execute the command PWRDWNSYS, then press F4 to prompt for options as shown in Figure 7-90. Change the Controlled end delay time to **300**. Press Enter when ready to power down the partition.

```

Power Down System (PWRDWNSYS)
Type choices, press Enter.

How to end . . . . . *CNTRLD      *CNTRLD, *IMMED
Controlled end delay time . . . 300      Seconds, *NOLIMIT
Restart options:
  Restart after power down . . . *NO      *NO, *YES
  Restart type . . . . . *IPLA      *IPLA, *SYS, *FULL
  IPL source . . . . . *PANEL      *PANEL, A, B, D, *IMGCLG
  
```

Figure 7-90 IBM i power down partition

7. Confirm the shutdown action by pressing F16.
8. This process can take awhile. Check the Integrated Virtualization Manager (IVM) window for the message Not Activated in the State column of the IBM i partition.

Start (activate) an IBM i V6.1 partition

Before activating an IBM i partition, verify the IPL type you want to perform is set correctly. Use View/Modify partitions, then select the partition and view the properties. On the General tab you will see the IPL type.

To activate an IBM i V6.1 partition, follow these steps:

1. In Integrated Virtualization Manager (IVM), select the IBM i partition and click **Activate**.
2. Click **OK** in the next window to activate the IBM i partition.
3. After the partition starts its IPL sequence you can activate the Operations Console session to connect to the partition.

Red Hat Enterprise V5.3 Linux installation

This chapter describes the procedures to install Red Hat Enterprise Linux V5.3 on a JS23 BladeCenter.

This chapter contains the following topics:

- ▶ “Red Hat operating systems and prerequisites” on page 328
- ▶ “Linux LPAR installation using DVD” on page 329
- ▶ “Linux network installation” on page 333
- ▶ “Native Red Hat Enterprise Linux 5.3 installation” on page 346
- ▶ “Red Hat Enterprise Linux 5.3 automated installation” on page 346
- ▶ “IBM service and productivity tools” on page 358

8.1 Red Hat operating systems and prerequisites

This section lists supported Red Hat operating system versions and describes considerations for Red Hat Enterprise Linux.

8.1.1 Supported Red Hat operating system versions

Supported Red Hat operating systems for installation on a JS23 are:

- ▶ Red Hat Enterprise Linux for POWER Version 4.6 or later
- ▶ Red Hat Enterprise Linux for POWER Version 5.1 or later
- ▶ Red Hat Enterprise Linux for POWER Version 5.3

This chapter specifically covers installing Red Hat Enterprise Linux for POWER Version 5.3 with a DVD and over the network on a PowerVM logical partition (LPAR).

8.1.2 Considerations and prerequisites

This section discusses system configuration considerations and prerequisites that you should understand prior to installing Red Hat Enterprise Linux 5.3 on a JS23 partition.

PowerVM LPAR considerations and prerequisites

By the characteristics of the virtualization features in the System p JS23 BladeCenter, the operating system and applications do not know they are running in a micro-partitioned or virtualized I/O environment. This approach allows applications to run unmodified in a partition that takes advantage of virtualization features.

Because the virtual partition handles the transition of the virtual adapter's I/O operation to the physical adapter, it is important to guarantee that the partition is properly sized to handle the I/O requirements. A good source for processor and memory considerations for PowerVM partitions based on I/O considerations is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

Additional information about virtualization and LPAR setup is available in Chapter 4, "System planning and configuration using VIOS with IVM" on page 65.

Red Hat Enterprise Linux 5.3 installation considerations

Although the minimum RAM required to install Red Hat Enterprise Linux 5.3 is 1 GB, the best practice is to have 2 GB RAM.

Note: Make sure you have 1G of RAM or more allocated to your LPAR. Systems with less than 1 GB of RAM can experience hanging of the installation process or other unexpected failures.

In addition, ensure enough unpartitioned disk space exists or that you have one or more partitions that can be deleted to free up disk space for the Linux installation. The Red Hat Recommended Partitioning Scheme is available at:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/ch11s03.html

8.2 Linux LPAR installation using DVD

With PowerVM installed and the system partitioned into LPARs using the PowerVM LPAR considerations and Red Hat Enterprise Linux 5.3 prerequisites, we are ready to install Linux on the JS23 BladeCenter LPAR.

Starting an LPAR installation of Red Hat Enterprise Linux 5.3 using a DVD

To start the installation of Red Hat Enterprise Linux 5.3 from a DVD:

1. On the Advanced Management Module (AMM) select **Blade Tasks** → **Remote Control**.
2. In the right panel, select your blade bay location from the Media tray owner drop-down menu, which lists the various blade bays, as shown in Figure 8-1.

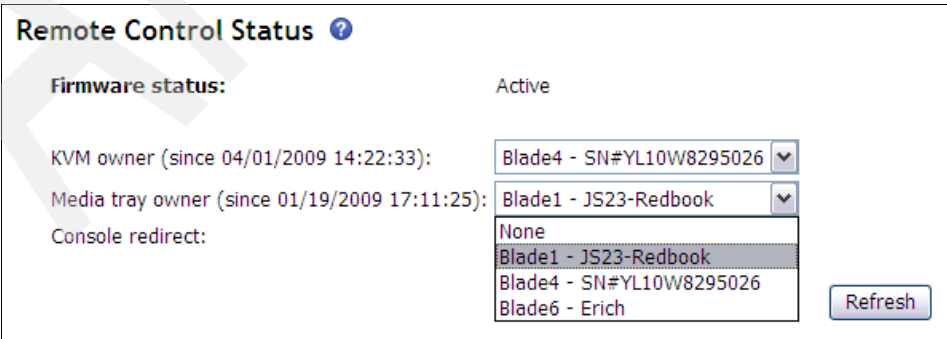


Figure 8-1 Remote Control window - assign media tray

3. Click **Refresh**.

Important: The other option is to press the MT (media tray) button on the blade to assign the media tray to the blade.

Make sure no other blade in the BladeCenter is using the media tray before you click this button. If the media tray is assigned to another blade, the MT light is on.

4. Verify that your blade bay owns the media tray by opening the AMM window and selecting **Monitors** → **System Status**. The window on the right shows a check mark in the MT column of your blade bay location. Figure 8-2 gives one example of this assignment.

Blades ?

Click the icon in the Status column to view detailed information about each blade.


Bay	Status	Name	Pwr	Owner**		cKVM*	I/O Compatibility
				KVM	MT*		
1		JS23-Redbook	On	✓	✓		OK
2		No blade present					

Figure 8-2 BladeCenter System status.

5. Place the DVD into the BladeCenter media tray.
6. Log in to the Integrated Virtualization Manager (IVM) with a Web browser.

Important: Make sure the latest Java Runtime Environment (JRE™) is installed on the native system to run the IVM terminal. At the time of this publication, the JRE is Sun's JRE 1.4.2_19, or later.

Perform the following steps:

- Select one of the available PowerVM LPARs to install with Red Hat Enterprise Linux 5.3 by placing a check mark in the Select box.
- Click **Activate**, as shown in Figure 8-3 on page 331.

Partition Details									
<div> <div> <div></div> <div></div> <div></div> </div> <div> <div>* Create Partition...</div> <div>Activate</div> <div>Shutdown</div> <div>--- More Tasks ---</div> </div> </div>									
Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	js23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 8-3 Activating an IVM partition

- c. From the More Tasks drop-down list, select **Open terminal window**. See Figure 8-4.

Partition Details									
<div> <div> <div></div> <div></div> <div></div> </div> <div> <div>* Create Partition...</div> <div>Activate</div> <div>Shutdown</div> <div>--- More Tasks ---</div> </div> </div>									
Select	ID ^	Name	State	Uptime	Memory	Processors	<div> <div>--- More Tasks ---</div> <div> <div>Open terminal window</div> <div>Delete</div> <div>Create based on</div> <div>Operator panel service functions</div> <div>Reference Codes</div> <div>-----</div> <div>Properties</div> </div> </div>		
<input type="checkbox"/>	1	js23-vios	Running	20.33 Hours	1 GB	4			
<input checked="" type="checkbox"/>	2	js23-lp1	Not Activated		1 GB	1			

Figure 8-4 Opening a terminal window from the IVM

7. The console is a pop-up and it asks you to authenticate with the PowerVM User ID and password.

Note: Although this section covers installation with the Integrated Virtualization Manager (IVM) console, other console options are available on the JS23. They are covered in Appendix A, “Consoles, SMS, and Open Firmware” on page 487.

The SMS menu appears in the IVM terminal; see Figure 8-5 on page 332.

Tip: The SMS menu opens almost instantly, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. If you happen to miss the SMS option section window, just shut down and reactivate the LPAR again. The IVM terminal will hold the connection to the PowerVM LPAR even if the LPAR is shut down.

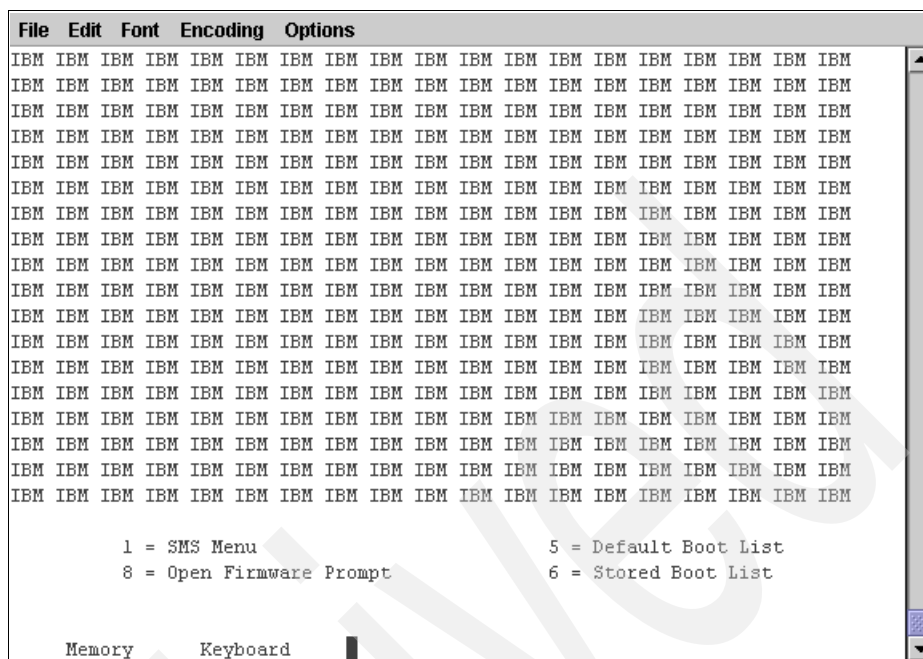


Figure 8-5 SMS menu

Perform the following steps:

- a. Select **1 = SMS Menu** by pressing the number 1 on the keyboard.

Tip: To select and navigate through the SMS menu, press the number next to the desired system function.

- b. Select option **5. Select Boot Options.**
- c. Select option **1. Select Install/Boot Device.**
- d. Select **3. CD/DVD.**
- e. Select **6. USB.**
- f. Finally, select **1. USB CD-ROM.**
- g. Select **2. Normal Mode Boot.**
- h. Select **1.Yes** to exit the SMS menu.
- i. At the boot prompt press the Enter key.

The system begins reading from the DVD; it can take several minutes.

8. Complete the installation; go to “Completing the installation” on page 333.

Completing the installation

After successfully booting from the DVD, the Red Hat Anaconda Installer asks if you want to perform a media check. If this is your first time installing with this media, run the media check. After the media check is complete, Anaconda assists with the completion of the installation. Detailed installation instructions are available at:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/pt-install-info-ppc.html

8.3 Linux network installation

This section describes a Network File System (NFS) installation on an external Storage Area Network (SAN) device, which is part of a PowerVM LPAR. Although this example uses a specific storage device and network option to complete the installation, this process is generic enough that it can work for all supported internal or external storage devices on JS23; and only a few slight modifications are necessary in Anaconda's setup if a different network installation option is selected. The goal of this section is to show the external SAN disk installation in detail so the steps can be referenced later.

This installation is using an NFS installation method, but NFS is one of many supported network installation types on Linux. The NFS server configuration and setup are described in detail in Appendix C, "Additional Linux installation configuration options" on page 529.

This section assumes that you have already set up the NFS server properly and have read "PowerVM LPAR considerations and prerequisites" on page 328 and followed the installation prerequisites in "Red Hat Enterprise Linux 5.3 installation considerations" on page 329.

LPAR installation of Red Hat Enterprise Linux 5.3 over the network

To perform the installation over the network:

1. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.

Perform the following steps:

- a. Select one of the available PowerVM LPARs to install with Red Hat Enterprise Linux 5.3 by placing a check mark in the Select box.
- b. Click **Activate**, as shown in Figure 8-6 on page 334.

Partition Details									
<div> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> * Create Partition... Activate Shutdown --- More Tasks --- </div>									
Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	js23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 8-6 Activating an IVM partition

- c. From the More Tasks drop-down box, (Figure 8-7) select **Open terminal window**.

Partition Details									
<div> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> * Create Partition... Activate Shutdown --- More Tasks --- </div>									
Select	ID ^	Name	State	Uptime	Memory	Processors	<div> <div>--- More Tasks ---</div> <div> Open terminal window Delete Create based on Operator panel service functions Reference Codes ----- Properties </div> </div>		
<input type="checkbox"/>	1	js23-vios	Running	20.33 Hours	1 GB	4			
<input checked="" type="checkbox"/>	2	js23-lp1	Not Activated		1 GB	1			

Figure 8-7 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

2. The SMS menu appears in the IVM terminal (Figure 8-8 on page 335).

Tip: The SMS menu appears very quickly after activating the LPAR, therefore, have your hand ready on the 1 key so you can immediately press it when you see the SMS menu. If you happen to miss the SMS selection window, just shut down and reactivate the LPAR again. The IVM terminal holds the connection to the PowerVM LPAR even if the LPAR is shut down.




Figure 8-8 SMS menu

- a. Select **1 = SMS Menu** by pressing the 1 key on the keyboard.

Tip: Press the number next to the desired system function to select and navigate through the SMS menu.

- b. Select **5. Select Boot Options.**
- c. Select **1. Select Install/Boot Device.**
- d. Select **6. Network.**
- e. Select a specific network port.
- f. Select **2. Normal Mode Boot**
- g. Select **1.Yes** to exit the SMS menu.

If everything is set up correctly, the system will receive an IP address from the TFTP server and start loading the boot image for the first stage of the installation.

Note: Review Appendix C, “Additional Linux installation configuration options” on page 529 prior to starting an NFS installation.

You will notice the packet count value increasing. After the complete boot image is uploaded, the system boots from it to show the Red Hat Enterprise Linux 5.3 welcome panel, shown in Figure 8-9.

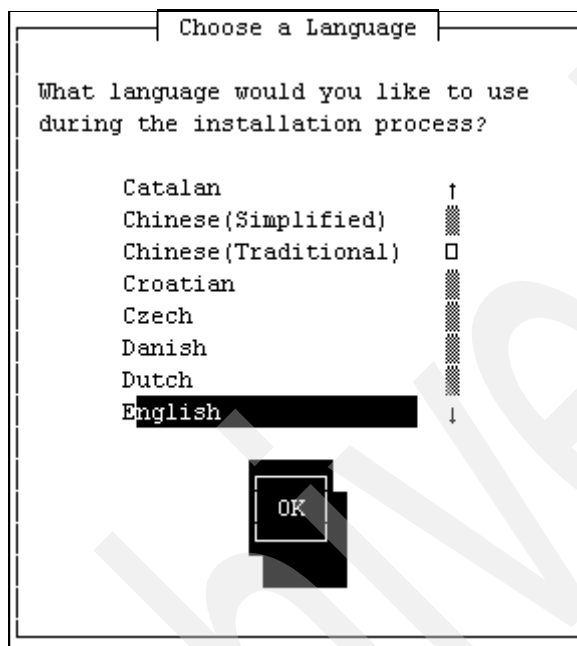


Figure 8-9 Red Hat Enterprise Linux 5.3 welcome panel

3. Select the language to use during the installation process. In this example, we are using English. Then, press the Tab key to move to and select the **OK** button, and then press Spacebar to confirm.

Tip: Use the Tab key to navigate between selections, the Up/Down Arrow keys to move within a list, and Spacebar to select a specific option inside the Anaconda installer.

4. The installation method is NFS. Use the Arrow key to scroll down to NFS image and then press the Tab key to move to and select the **OK** button, and then press Spacebar to confirm.
5. Select the networking device to install the image from, as shown in Figure 8-10 on page 337




Figure 8-10 Select network device

Note: This step appears only when you run Anaconda on machines with more than one network card. The **Identify** option can be used to find the physical port for the selected interface, by flashing the LED lights of the correspondent physical port for a number of seconds.

6. Perform one of the following steps:

- If you are *not* configuring DHCP, go to step 7 on page 338.
- If you are configuring DHCP, use the Configure TCP/IP window (shown in Figure 8-11 on page 338) to:
 - i. Select either **IPv4** or **IPv6** support.
 - ii. Select **Dynamic IP configuration (DHCP)**.
 - iii. Select **OK**.
 - iv. Go to step 9 on page 339.

As an example, if you want to configure static IP parameters with IPv4 support, configure the TCP/IP parameters and **Enable IPv4** support, then select **Manual configuration** and disable **Enable IPv6 support** as shown in Figure 8-11 on page 338, and configure the device as shown in Figure 8-12 on page 338.

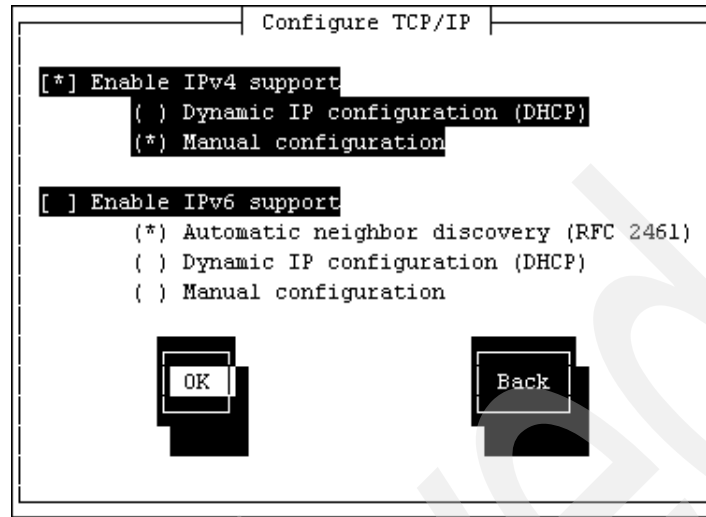


Figure 8-11 TCP/IP configuration panel

7. In the next panel, configure the LPAR's IPv4 address, subnet mask, gateway, and name server. An example configuration is shown in Figure 8-12

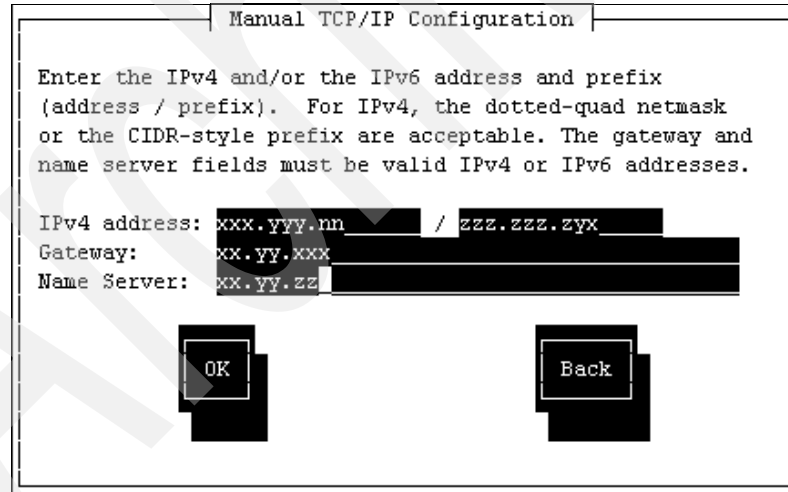


Figure 8-12 TCP/IP configuration of IP address, gateway, and name server

8. In the NFS Setup window, shown in Figure 8-13 on page 339, enter the IP address of the NFS server, and enter the NFS directory that contains the Red Hat Enterprise Linux 5.3 installation image.

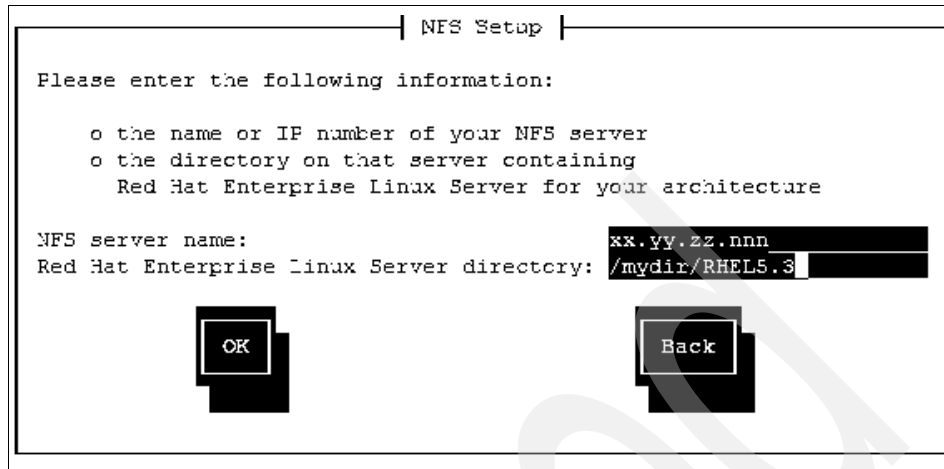


Figure 8-13 NFS server configuration window panel

In this step it is possible to start a *Virtual Network Computing* (VNC) server and continue the installation from Anaconda's graphical interface, but for this example, we continue with the text mode interface, as shown in Figure 8-14.

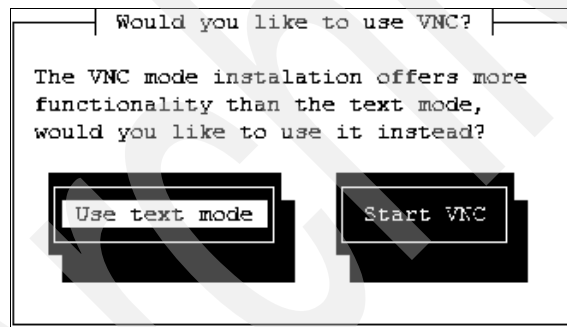


Figure 8-14 Select between VNC or text installation modes panel

9. Approximately one minute later, the Welcome to Red Hat Enterprise Linux Server message panel appears. Select **OK**.
10. Enter the installation number (Figure 8-15 on page 340); select **OK** to confirm.

Note: The installation number is specific to systems running Red Hat Enterprise Linux 5 or later. This number comes from Red Hat subscription management or the poster included in the media kit, and it determines the packages available to the installer.

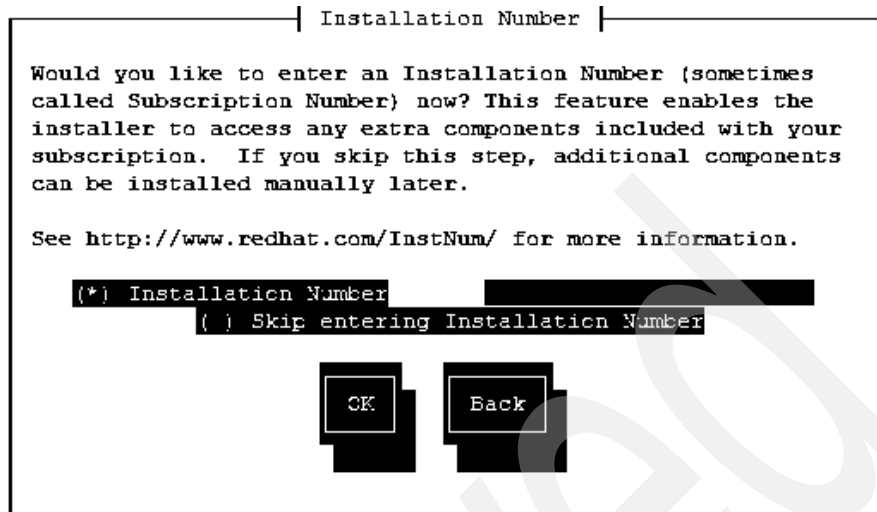


Figure 8-15 Installation number panel

Note: If you do not enter the Installation number, then later on, you will only have the basic packages to select from. In this case, a warning indicates that you have to select **Skip** to proceed.

11. Select the disk partitioning type for this installation. In this scenario, we have selected the option **Remove all partitions on selected drives and create a default layout**. The drive we have selected is /dev/sda, which is the externally connected storage SAN disk that was partitioned using the IVM tools. See Figure 8-16 on page 341 as an example.

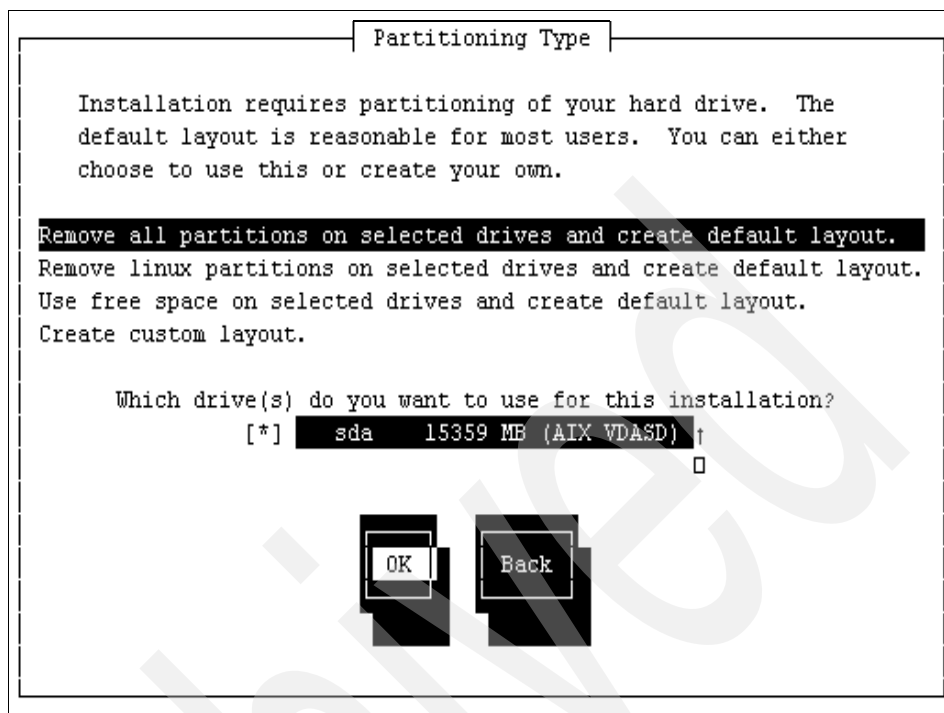


Figure 8-16 Select Partitioning Type panel

12. A warning message asks if the selection is OK. Press **Yes** to confirm.
13. Select **Yes** to review the suggested disk partition layout.
14. Review the allocated size for swap, ext3 file system, and /boot, as shown in Figure 8-17 on page 342. Press **OK** to confirm.




Figure 8-17 Review Partitioning panel

Note: This configuration can only be edited by a graphical installer such as Virtual Network Connection (VNC). This cannot be done from the IVM terminal, so only the default values selected by the Anaconda Installer are allowed.

15. Select **OK** on the Network Configuration panel. The default is fine because this was already set up in Figure 8-12 on page 338.
16. Select **OK** for the Miscellaneous Network Setting window. The gateway and primary DNS are already configured.
17. Select **OK** after editing or confirming the host name.
18. Select the time zone for the LPAR using the Tab key to enter the time zone options field and use the Arrow keys to move up or down inside of it until you find your locale. Press the Tab key and then Spacebar on the **OK** button to confirm.
19. Enter the root password.

20. Select any additional software applications with the Spacebar and select **OK** when complete, as shown in Figure 8-18.

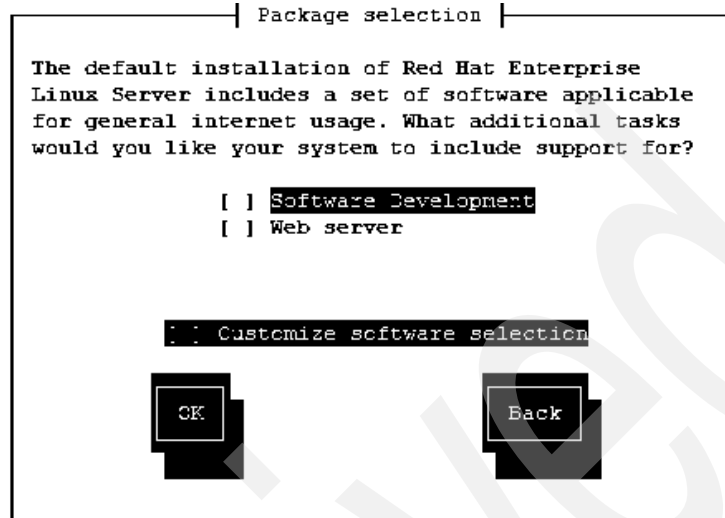


Figure 8-18 Select additional packages panel

Note: If, during the installation, you skip this step, these packages can be installed later by using `yum` command from the command line.

21. Select **OK** to allow the installation to begin.

The next window has two progress bars: One for the package currently being installed and another detailing the overall progress of the installation (Figure 8-19 on page 344).

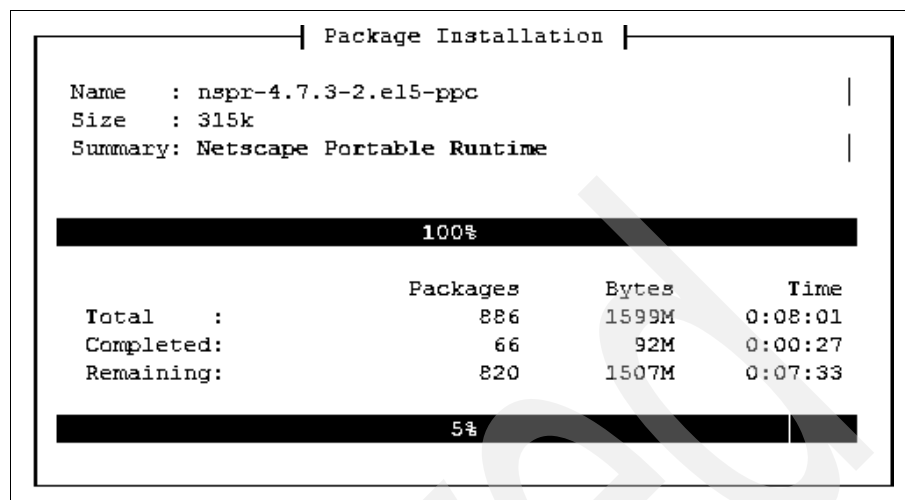


Figure 8-19 Installation progress window

22. Select **Reboot** after the Install Complete window appears, as shown in Figure 8-20

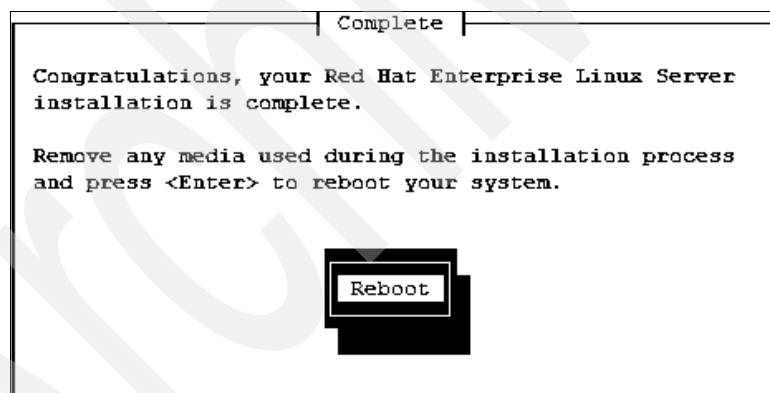


Figure 8-20 Installation complete panel

Note: If the LPAR does not automatically boot from the intended hard disk (boot device) after reboot, try this:

- a. Shut down and reactivate the LPAR from the IVM.
- b. Enter the SMS Menu.
- c. Select **5. Select Boot Options** → **1. Select Install/Boot Device** → **5. Hard Drive** → **9. List All Devices**.
- d. Choose the appropriate hard disk with the Linux image from the given list.
- e. Select **2. Normal Mode Boot** → **1. Yes**.

23. During boot, the Setup Agent panel opens (Figure 8-21). You may modify any of the fields or select **Exit** to finish booting the LPAR.

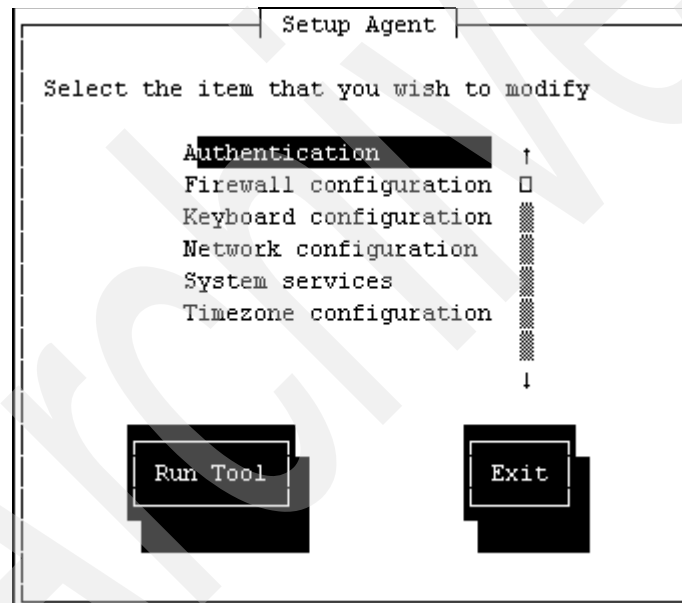


Figure 8-21 Setup Agent panel

The Red Hat Enterprise Linux 5.3 login prompt appears, as shown in Figure 8-22 on page 346. Installation is complete.

```
[ OK ]
Starting smartd: [ OK ]

Red Hat Enterprise Linux Server release 5.3 (Tikanga)
Kernel 2.6.18-128.el5 on an ppc64

localhost.localdomain login: █
```

Figure 8-22 Finished Red Hat Enterprise Linux 5.3 installation

8.4 Native Red Hat Enterprise Linux 5.3 installation

A native Red Hat Enterprise Linux 5.3 installation on a IBM BladeCenter JS23 Express follows a similar process to the one given in the previous VIOS LPAR installation sections. However, several key differences do exist:

- ▶ In a native installation, the IVM terminal is no longer available to complete the Linux installation, but you can use the Serial Over LAN (SOL) console, or the graphical display (with BladeCenter's KVM), as an alternative. See Appendix A, “Consoles, SMS, and Open Firmware” on page 487 for more information. Use the SOL console to display the SMS menu and the Anaconda options during the installation.
- ▶ The resource allocation of processors, I/O adapters, memory, and storage devices in a native environment is *fixed*.
- ▶ Virtualization functions and features are not available.

8.5 Red Hat Enterprise Linux 5.3 automated installation

Red Hat provides an automated installation functionality known as Kickstart. The system administrator performs a Kickstart automated installation by creating a single file containing answers to all the questions normally asked during a Red Hat installation. This file can reside on a single server system, and multiple clients can read it during installation.

This section explains how to create a Kickstart file and use it to perform an unattended installation of Red Hat Enterprise Linux 5.3

8.5.1 Methods for creating a Kickstart file

The three methods to create a Kickstart file are:

- ▶ Create the file using the Red Hat Kickstart Configurator tool.
- ▶ Have Red Hat Anaconda Installer generate this file for you during a manual installation. After the installation completes, the Kickstart file resides at `/root/anaconda-ks.conf` and you can copy it to an installation server. The basic file created during a manual NFS installation is given in Example 8-1.
- ▶ Generate the file yourself using a text editor.

Example 8-1 Basic Kickstart configuration file created during manual installation

```
# Kickstart file automatically generated by anaconda.

install
nfs --server=9.3.80.16 --dir=/install/linuxRHEL5.3
key 2515-dd4e-2152-25dd
lang en_US.UTF-8
network --device eth1 --bootproto static --ip 9.3.233.177 --netmask
255.255.254.0 --gateway 9.3.232.1 --nameserver 9.3.192.21 --hostname
JS23-turkey-3-lpl.austin.ibm.com
rootpw --iscrypted $1$0VjuAHA9$6cqrT/Nfnh6pRYFBIBECy/
firewall --enabled --port=22:tcp
authconfig --enablesshadow --enablemd5
selinux --enforcing
timezone --utc America/New_York
bootloader --location=partition --driveorder=sda --append="console=hvc0
rhgb quiet"
# The following is the partition information you requested
# Note that any partitions you deleted are not expressed
# here so unless you clear all partitions first, this is
# not guaranteed to work
#clearpart --all --drives=sda
#part prepboot --fstype "PPC PReP Boot" --size=4 --ondisk=sda
#part /boot --fstype ext3 --size=100 --ondisk=sda
#part pv.14 --size=0 --grow --ondisk=sda
#volgroup VolGroup00 --pesize=32768 pv.14
#logvol / --fstype ext3 --name=LogVol100 --vgname=VolGroup00 --size=1024
--grow
#logvol swap --fstype swap --name=LogVol101 --vgname=VolGroup00
--size=1000 --grow --maxsize=1984

%packages
@office
```

```
@editors
@text-internet
@gnome-desktop
@ dialup
@core
@base
@games
@java
@legacy-software-support
@base-x
@graphics
@printing
@sound-and-video
@admin-tools
@graphical-internet
emacs
kexec-tools
device-mapper-multipath
hfsutils
xorg-x11-utils
xorg-x11-server-Xnest
libsane-hpaio
-sysreport
```

Notice that all partition information is commented out with a number sign (#) symbol. This section has to be uncommented and edited to support the partition schemes of systems that will use the automated Kickstart installation process. The automated Kickstart process does not work without these edits.

8.5.2 Creating Kickstart file with Kickstart Configurator tool

In this section, we use the Kickstart Configurator tool with a graphical interface to demonstrate how to create a basic Kickstart text file.

Although this file has many optional settings, some are mandatory settings or dependencies. Although describing every configuration option here is not possible, we do provide a general overview so you can become familiar enough with the Kickstart Configurator tool to navigate on your own.

Restriction: It is necessary to have a running X Server to use the administration tools provided by Red Hat, because in most cases the tools use a graphical user interface (GUI).

You can install a *Virtual Network Computer* (VNC) server in your system and export the graphical interface to some other host in the network. RHEL 5.3 distributes the *vnc-server* rpm package, which you can install to get access to VNC client/server functionalities.

Red Hat provides a utility called `system-config-kickstart` to assist with the creation of the configuration file for an unattended installation. However, after you create the Kickstart file with the tool, you might have to make several manual changes to it.

Note: Because Kickstart is not included in the default software installation, install the following rpm packages from the installation media to get Kickstart working:

```
pykickstart
system-config-kickstart
```

After you install the required Kickstart packages, perform the following steps:

1. issue the `system-config-kickstart` command from the command line to launch this utility.
2. A window opens showing the Basic Configuration panel. The most important configuration setting is the Target Architecture setting highlighted with a rectangle in Figure 8-23 on page 350. It is also important to define a root password to enable SSH login after installation. This password is encrypted in the configuration file.

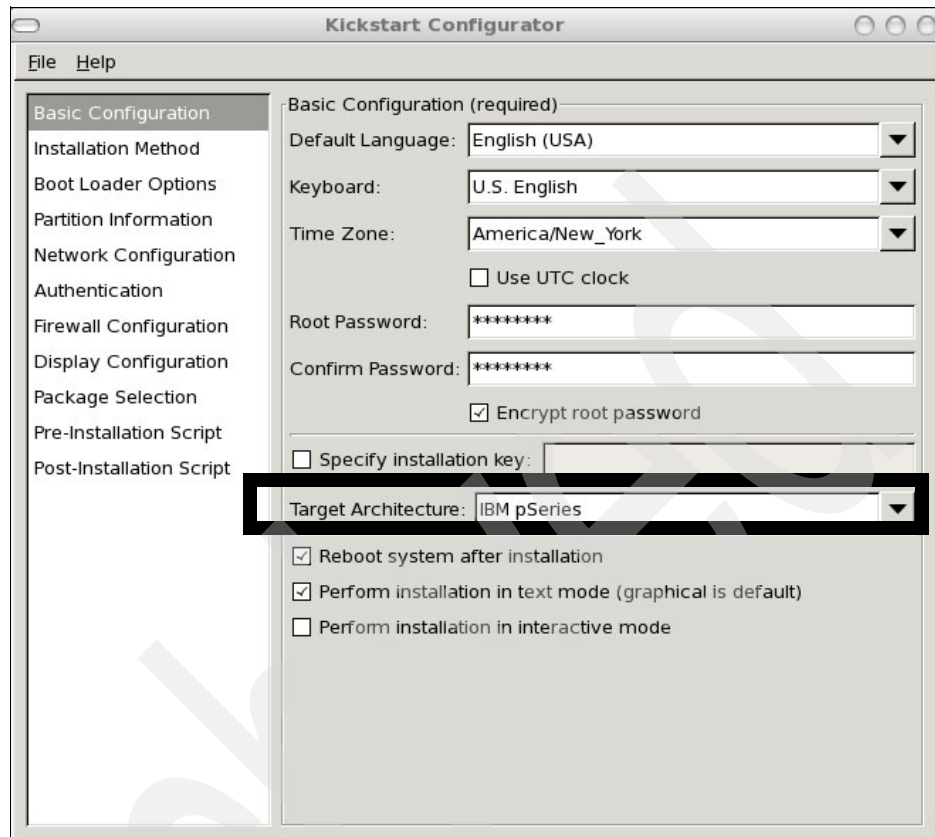


Figure 8-23 Kickstart main window with Basic Configuration panel (©2008 Red Hat, Inc.)

3. In the Installation Method panel (shown in Figure 8-24), all the basic parameters for a network installation using NFS are shown.




Figure 8-24 Installation Method panel (©2008 Red Hat, Inc.)

4. The next editable panel is the Partition Information panel, shown in Figure 8-25. Click **Add** to create a partition. The tool helps you select the mount point, file system type, and partition size.

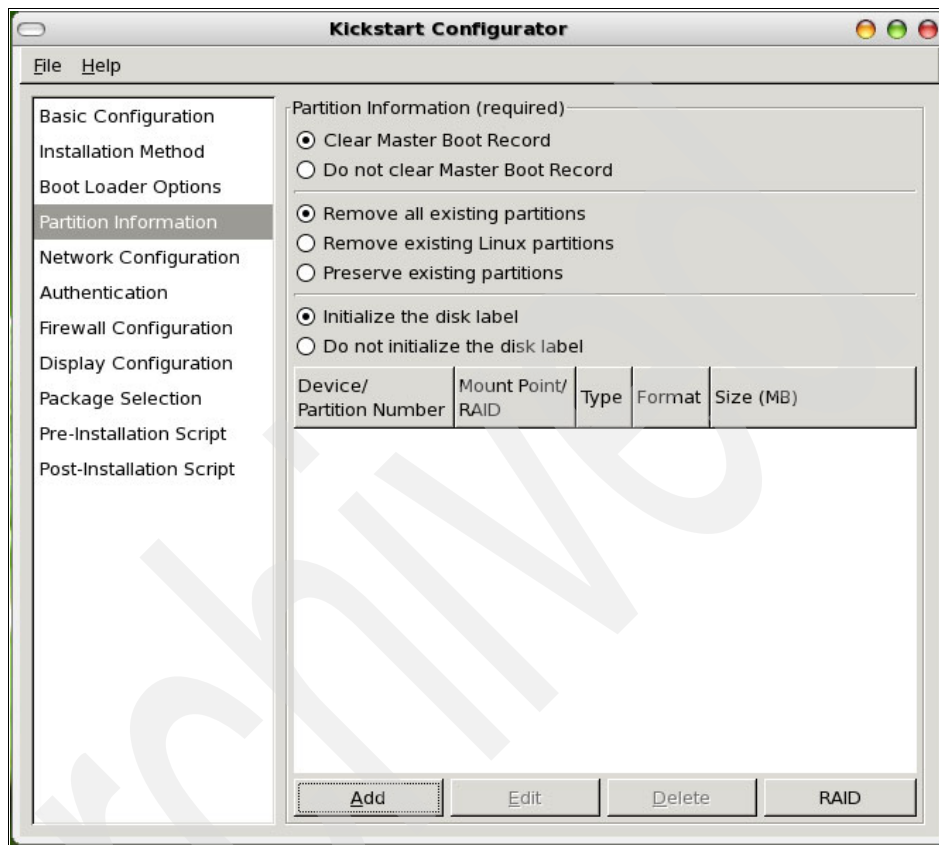


Figure 8-25 Partition Information panel (©2008 Red Hat, Inc.)

5. In the Network Configuration panel, click **Add Network Device** to add the devices you are installing from. If you have to go back and make changes to this setup, click **Edit Network Device** (see Figure 8-26).



Figure 8-26 Kickstart Configurator Network Configuration panel (©2008 Red Hat, Inc.)

6. The next panel is the Authentication panel. In this configuration, we use the default settings.
7. Figure 8-27 shows the Firewall Configuration panel. As an example, a good practice is to enable SSH and to trust interface eth1, at the very minimum to access the system later using the network.




Figure 8-27 Firewall Configuration panel (©2008 Red Hat, Inc.)

8. Figure 8-28 on page 355 shows the Package Selection panel. Selecting individual packages from this panel is not possible. However, you can add individual packages to the %packages section of the Kickstart file after saving it.

Note: If you see the following message in the Package Selection panel, you have no repositories defined:

Package selection is disabled due to problems downloading package information

You can either register the system with Red Hat Network (RHN) to a local internal RHN Satellite through the yum-rhn-plugin, or copy all the files from the Red Hat Enterprise Linux 5.3 installation disks to a directory and use them as a reference. Refer to the Red Hat knowledge base article at:

http://kbase.redhat.com/faq/FAQ_103_12748.shtm



Figure 8-28 Package Selection panel (©2008 Red Hat, Inc.)

9. After package selection, save the configuration using the **File** → **Save**.

The basic Kickstart configuration file created with the Kickstart Configurator is shown in Example 8-2 on page 356.

Important: The order of the main sections in the Kickstart configuration file is important for the functionality.

Example 8-2 Basic Kickstart configuration file

```
#platform=IBM pSeries
# System authorization information
auth --useshadow --enablemd5
# System bootloader configuration
bootloader --location=mbr
# Clear the Master Boot Record
zerombr
# Partition clearing information
clearpart --all --initlabel
# Use text mode install
text
# Firewall configuration
firewall --enabled --trust=eth0,eth1
# Run the Setup Agent on first boot
firstboot --disable
# System keyboard
keyboard us
# System language
lang en_US
# Installation logging level
logging --level=info
# Use NFS installation media
nfs --server=9.3.29.121 --dir=/nsf/RHEL
# Network information
network --bootproto=dhcp --device=eth0 --onboot=on
network --bootproto=dhcp --device=eth1 --onboot=on
# Reboot after installation
reboot
#Root password
rootpw --iscrypted $1$ug.Uehug$WGBeIYgPLXt8A3QcsnDa90

# SELinux configuration
selinux --enforcing
# Do not configure the X Window System
skipx
# System timezone
timezone America/New_York
# Install OS instead of upgrade
install

%packages
@sound-and-video
@gnome-desktop
@office
@graphical-internet
```

10.If you want, manually adjust the Kickstart configuration file that you have created with a text editor.

Note: If you have not defined any disk partition options or you were unsure of your disk partition layout, manually edit the Kickstart file to include the following information after the #Partition clearing information section:

```
#Disk partitioning information
autopart
```

This option automatically creates disk partitions.

Red Hat Enterprise Linux 5.3 Installation Guide is a good source of information regarding all Kickstart file options. It is also available at the following Web site:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/ch-kickstart2.html

8.5.3 Performing a Kickstart installation

It is possible to run unattended installations with the Kickstart file stored on a floppy disk, on a CD/DVD, on a USB key, on a partition in the system, or on some NFS/HTTP/FTP server. In this section we provide instructions for how to run an automatic installation with the kickstart file stored on an NFS-exported directory, accessible through the network.

Pass the location of the Kickstart file during boot. BOOTP does not have the ability to provide anything more than the location to the bootable image and the server IP address hosting the Kickstart file. Use the Open Firmware prompt to pass the required parameters:

1. Type 8 during the LPAR boot process to go to the Open Firmware prompt, as shown in Figure 8-29.

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

Memory  Keyboard  Network  SCSI  Speaker  ok
0 > _
```

Figure 8-29 Open Firmware prompt

2. At the Open Firmware prompt, enter the command to start automated installation. For example, if the configuration file is served using NFS, enter the following command:

```
boot net ks=nfs://192.168.1.254/ks.cfg ksdevice=eth1 ip=dhcp
```

Press Enter. The process begins and the automated Red Hat Enterprise Linux installation is now complete.

Tip: For a good source of information regarding all Kickstart file options, refer to the Kickstart Installations at:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/ch-kickstart2.html

8.6 IBM service and productivity tools

The IBM service and productivity tools are packages that are installed after a successful Red Hat Enterprise Linux installation.

Important: These packages are *not* shipped with the Red Hat installation CDs. IBM owns and distributes them.

These packages enable features such as:

- ▶ Reliability, availability, and serviceability (RAS) functionality
- ▶ I/O hotplug
- ▶ Dynamic Logical Partitioning (DLPAR) capabilities
- ▶ Live partition migration capabilities

See Appendix D, “Service and productivity tools for Linux” on page 539 for more information about installing the service and productivity tools specific to your system’s configuration.

SUSE Linux Enterprise Server V11 installation

This chapter describes the procedures to install SUSE Linux Enterprise Server (SLES) V11 on a JS43 BladeCenter.

This chapter contains the following topics:

- ▶ “Operating systems and prerequisites” on page 360
- ▶ “Linux LPAR installation using DVD” on page 361
- ▶ “Linux network installation (detailed)” on page 366
- ▶ “Native SLES 11 installation” on page 387
- ▶ “SLES 11 automated installation” on page 387
- ▶ “IBM service and productivity tools” on page 389

9.1 Operating systems and prerequisites

This section lists supported operating systems and describes considerations for SUSE Linux Enterprise Server.

9.1.1 Supported operating systems

SUSE Linux Enterprise Server 10 Service Pack 1 (SLES 10 SP1) for POWER or later supports installation on a JS43.

This chapter specifically covers installing SUSE Linux Enterprise Server 11 for POWER with a DVD and over the network on a PowerVM LPAR.

9.1.2 Considerations and prerequisites

This section discusses system configuration considerations and prerequisites that you should understand prior to installing SLES 11 on a JS43 partition.

PowerVM LPAR considerations and prerequisites

By the characteristics of the virtualization features in the System p JS43 BladeCenter, the operating system and applications do not know they are running in a micro-partitioned or virtualized I/O environment. This approach allows applications to run unmodified in a partition that takes advantage of virtualization features.

Because the PowerVM partition handles the transition of the virtual adapter's I/O operation to the physical adapter, it is important to guarantee that the partition is properly sized to handle the I/O requirements. A good source for processor and memory requirements for PowerVM partitions based on I/O requirements is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

Additional information about virtualization and LPAR setup is available in Chapter 4, "System planning and configuration using VIOS with IVM" on page 65.

SUSE Linux Enterprise Server 11 installation considerations

The *SUSE Linux Enterprise Server 11 Deployment Guide* recommends at least 512 MB of RAM to install SUSE Linux Enterprise Server 11. Refer to:

http://www.novell.com/documentation/sles11/pdfdoc/book_sle_deployment/book_sle_deployment.pdf

Tip: We use at least 1 - 2 GB per PowerVM LPAR and 4 GB on LPARs running processor-intensive loads.

In addition, the SLES 11 deployment guide suggests to have at least 1.5 GB of hard disk space or have one or more hard disk partitions that can be deleted to free up the miNFSum disk space for the Linux installation.

Tip: We use 10 GB or more *total* hard disk space for each PowerVM LPAR.

The Novell Web site has additional installation preparation information for SLES 11 available at:

<http://www.novell.com/documentation/sles11/index.html>

9.2 Linux LPAR installation using DVD

With the PowerVM installed and the system partitioned into LPARs using the PowerVM LPAR considerations and the SUSE Linux Enterprise Server 11 prerequisites, we are ready to install Linux on the JS43 BladeCenter LPAR.

Note: Although this section covers installation with the Integrated Virtualization Manager (IVM) console, other console options are available on the JS43; these are covered in Appendix A, “Consoles, SMS, and Open Firmware” on page 487.

Start an LPAR installation of SLES 11 using a DVD

To start the installation of SLES 11 from a DVD:

1. On the Advanced Management Module (AMM) select **Blade Tasks** → **Remote Control**.
2. In the right panel, select your blade bay location from the Media tray owner drop-down menu, which lists the various blade bays, as shown in Figure 9-1 on page 362.

Remote Control Status ?

Firmware status:

Active

KVM owner (since 04/01/2009 14:22:33):

Blade4 - SN#YL10W8295026

Media tray owner (since 01/19/2009 17:11:25):

Blade1 - JS23-Redbook

Console redirect:

None

Blade1 - JS23-Redbook

Blade4 - SN#YL10W8295026

Blade6 - Erich

Refresh

Figure 9-1 Start Remote Console panel

- Click **Refresh**.

Important: The other option is to press the **MT** button on the blade to assign the media tray to the blade.

Make sure no other blade in the BladeCenter is using the media tray before you click this button. If the media tray is assigned to another blade, the MT light is on.

- Double-check that your blade bay owns the media tray by opening the AMM panel and selecting **Monitors** → **System Status**. The left side will show an X in the MT column of your blade bay location. Figure 9-2 gives one example of this assignment.

Blades ?

Click the icon in the Status column to view detailed information about each blade.

Bay	Status	Name	Pwr	Owner**		cKVM*	I/O Compatibility
				KVM	MT*		
1		JS23-Redbook	On	✓	✓		OK

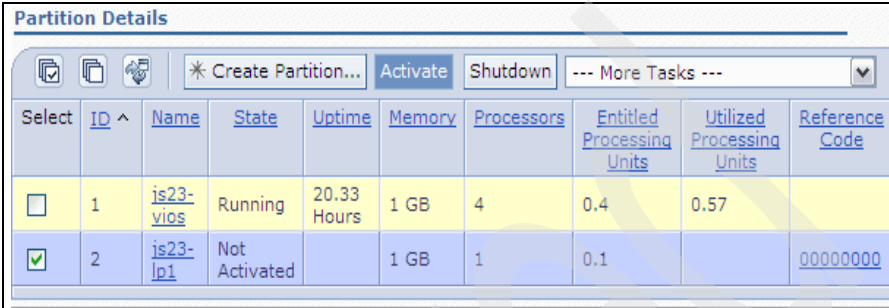
Figure 9-2 Checkmark indicates media tray (MT) is owned by the blade in bay 1

- Place the DVD into the BladeCenter media tray.
- Log in to the Integrated Virtualization Manager (IVM) using a Web browser.

Important: Make sure the latest Java Runtime Environment (JRE) is installed on the native system to run the IVM terminal. At the time of this publication, the recommended JRE is Sun's JRE 1.6.0_13, or higher.

Perform the following steps:

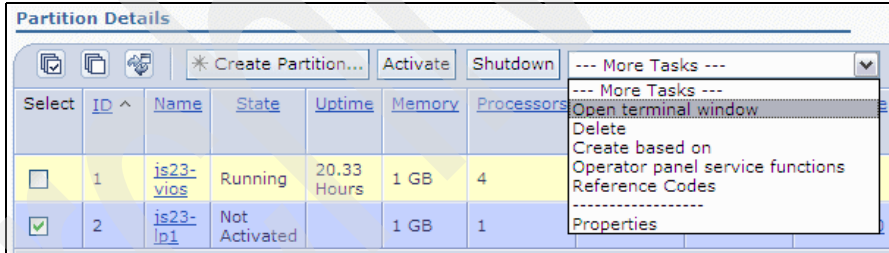
- a. Select one of the available PowerVM LPARs to install with SLES 11 by placing a checkmark in the Select box.
- b. Click **Activate** as shown in Figure 9-3



Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 9-3 Activating an IVM partition

- c. From the More Tasks drop-down list, select **Open terminal window**. See Figure 9-4.



Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 9-4 Opening a terminal window from the IVM

7. The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

The System Maintenance Services (SMS) menu appears in the IVM terminal. See Figure 9-5 on page 364.

Tip: The SMS menu opens almost instantly, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. If you happen to miss the SMS option section window, just shut down and reactivate the LPAR again. The IVM terminal holds the connection to the PowerVM LPAR even if the LPAR is shut down.

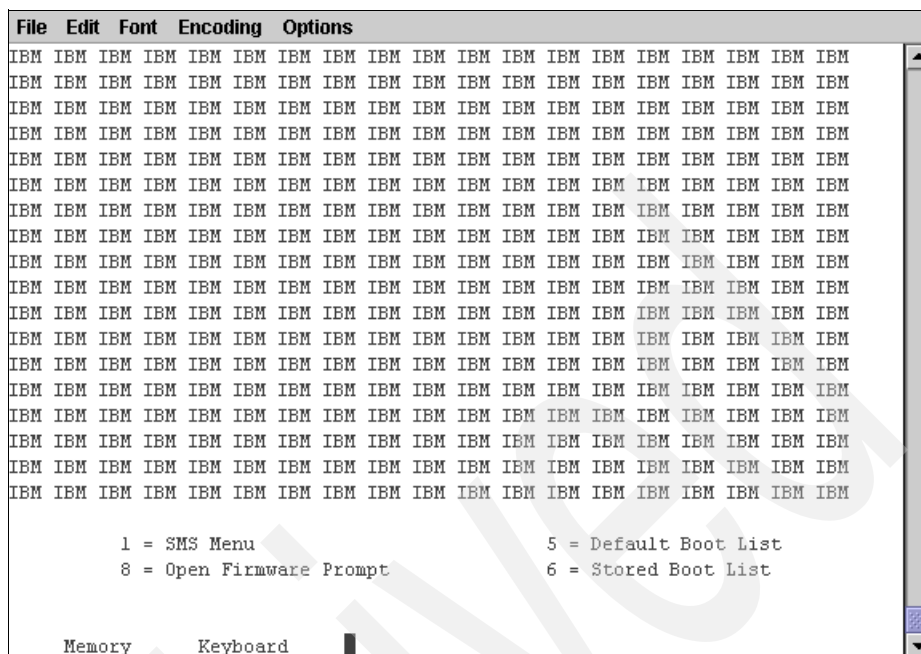


Figure 9-5 SMS Menu

Perform the following steps:

- Select **1 = SMS Menu** by pressing the number 1 on the keyboard.

Note: Press the number next to the desired system function to navigate through the SMS menu.

- Select option **5. Select Boot Options**.
- Select option **1. Select Install/Boot Device**.
- Select **3. CD/DVD**.
- Select **6. USB**.
- Finally, select **1. USB CD-ROM**. See Figure 9-6 on page 365.

```
SMS 1.5 (c) Copyright IBM Corp. 2000,2003 All rights reserved.¶
```

```
¶  
Select Task¶  
USB CD-ROM¶  
( loc=U787A.001.DNZ00XG-P4-D3 )¶
```

Figure 9-6 Select CD-ROM

- g. Select **2. Normal Mode Boot**. See Figure 9-7.

```
1. Information¶  
2. Normal Mode Boot¶  
3. Service Mode Boot¶  
-----  
¶  
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:¶
```

Figure 9-7 Select Mode Boot

- h. Select **1.Yes** to exit the SMS menu.
- i. At the Linux boot prompt, type **install**. See Figure 9-8 then press Enter to confirm. The LPAR starts reading from the DVD, which can take several minutes.

```
Welcome to SuSE:SLE-11:GA!¶  
Type "install" to start the YaST installer on this CD/DVD¶  
Type "slp" to start the YaST install via network¶  
Type "rescue" to start the rescue system on this CD/DVD¶  
Welcome to yaboot version 1.3.11.SuSE¶  
Enter "help" to get some basic usage information¶  
boot:¶
```

Figure 9-8 Select installation type

Complete the installation

After successfully booting off the DVD, the SUSE language selection panel appears on the IVM terminal and Your awesome Setup Tool (YaST) assists with the completion of the installation.

More detailed installation instructions are available at:

http://www.novell.com/documentation/sles11/book_sle_admin/index.html?page=/documentation/sles11/book_sle_admin/data/book_sle_admin_pre.html

9.3 Linux network installation (detailed)

This section describes a Network File System (NFS) installation on a PowerVM LPAR using an external Storage Area Network (SAN) device. Even though this example uses a specific storage device and network option to complete the installation, this network installation process is generic enough that it works for all supported internal or external storage devices on JS43, and only a few slight modifications are required in the YaST installer setup if a different network installation option is selected. The goal of this section is to show the SAN disk installation in detail so the steps can be referenced later.

This installation is using an NFS installation method, but NFS is one of many supported network installation types on Linux. The NFS server configuration and setup are described in detail in “Services required for installing Linux using the network” on page 530 of Appendix C, “Additional Linux installation configuration options” on page 529.

This section assumes that you have already:

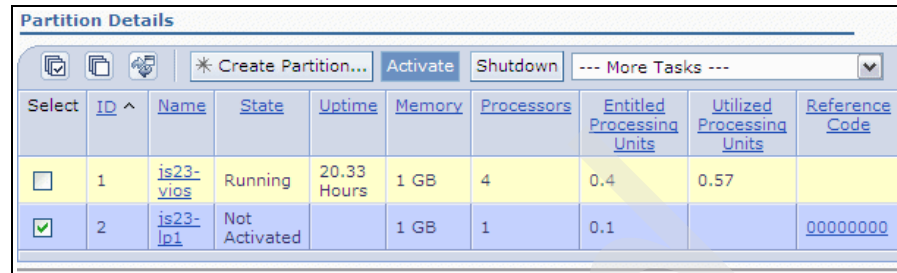
1. Set up the NFS server properly.
2. Read through “PowerVM LPAR considerations and prerequisites” on page 360
3. Followed the installation prerequisites in “SUSE Linux Enterprise Server 11 installation considerations” on page 360.

Start an LPAR installation of SLES 11 over the network

To start an LPAR installation of SLES 11 over a network.

1. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
Perform the following steps:
 - a. Select one of the available PowerVM LPARs to install with SLES 11 by placing a checkmark in the Select box.

- b. Click **Activate** as shown in Figure 9-9.

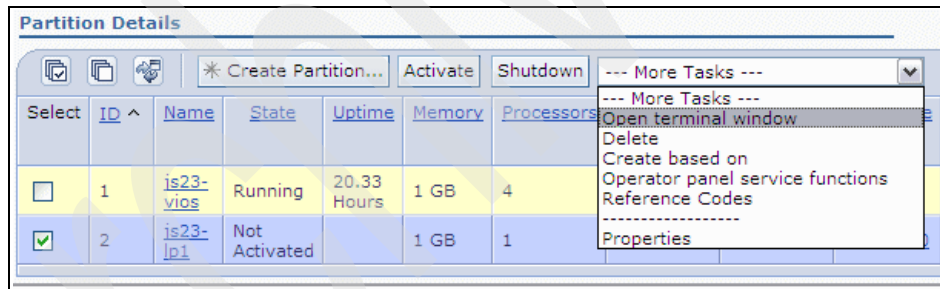


Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 9-9 Activating an IVM partition

- c. From the More Tasks drop-down box, select **Open terminal window**, as shown in Figure 9-10.

Important: Make sure the latest Java Runtime Environment (JRE) is installed on the native system to run the IVM terminal. At the time of this publication, the recommended JRE is Sun's JRE 1.6.0_13, or higher.



Select	ID ^	Name	State	Uptime	Memory	Processors
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1

Figure 9-10 Opening a terminal window from the IVM

2. The console is a pop-up and asks you to authenticate with the PowerVM User ID and password.

The firmware boot panel appears in the IVM terminal.

Note: The SMS menu option opens almost instantly after the LPAR is activated, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. See Figure 9-11. If you happen to miss the SMS selection window, just shut down and reactivate the LPAR again. The IVM terminal will hold the connection to the PowerVM LPAR even if the LPAR is shutdown.

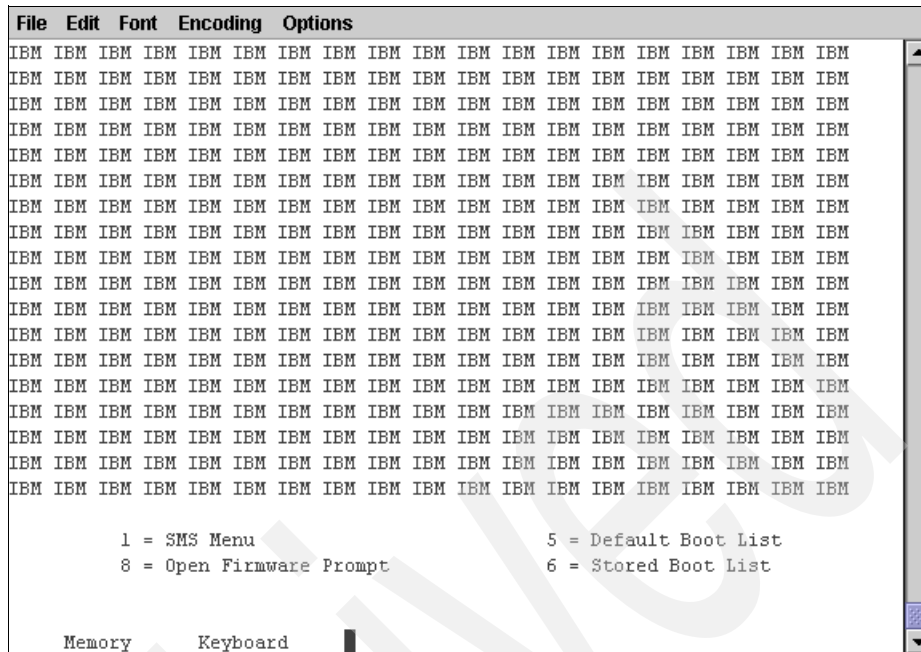


Figure 9-11 SMS Menu

Perform the following steps:

- a. Select **1 = SMS Menu** by pressing the number 1 on the keyboard.

Note: Press the number next to the desired system function to select and navigate through the SMS menu.

- b. Select **5. Select Boot Options**.
- c. Select **1. Select Install/Boot Device**.
- d. Select **6. Network**.
- e. Select **1. BOOTP**.
- f. Select a network port.
- g. Select **2. Normal Mode Boot**.
- h. Select **1.Yes** to exit the SMS menu.

Continue with the installation

If everything is set up correctly, the system connects using DHCP to the TFTP service to start loading the boot image for the first stage of the install. In several minutes, the Main Menu panel opens, as shown in Figure 9-12, and you can continue with the installation, as follows:

1. Select **3) Expert**, as shown in Figure 9-12.

```
Main Menu

1) Start Installation
2) Settings
3) Expert
4) Exit or Reboot

>
```

Figure 9-12 Main Menu

Tip: Press the number next to the desired configuration option and then the Enter key to select it in the Main Menu window.

The Enter key alone moves you back to the previous option window.

2. Select **2) Kernel Modules (Hardware Drives)**, as shown in Figure 9-13.

```
Expert -- Time: 04:34

1) System Information
2) Kernel Modules (Hardware Drivers)
3) Verify Installation CD-ROM/DVD
4) Eject CD
5) Show config
6) Change config
7) Start shell

>
```

Figure 9-13 Expert

3. Select **1) Load ppc Modules**, as shown in Figure 9-14 on page 370.

```

Kernel Modules (Hardware Drivers)

1) Load ppc Modules
2) Load USB Modules
3) Load FireWire Modules
4) Load File System Modules
5) Load Unknown Modules
6) Show Loaded Modules
7) Unload Modules
8) Add Driver Update
9) Show Driver Updates

> 1

```

Figure 9-14 Load ppc Modules

4. Select each individual module to pre-install based on your LPAR's network configuration. Press the number next to the module name and then the Enter key, then press the Enter key again to confirm.

Tip: Use the vertical scroll bar on the IVM terminal to navigate the module list. The most commonly used modules are:

- ▶ 5) e1000 : Intel PRO/1000
- ▶ 15) ehea : EHEA and IBMVETH

5. Press the Enter key after you finish loading the modules to go back to the Main Menu.
6. Select **1) Start Installation or Update**, as shown in Figure 9-15.

```

Start Installation or System

1) Start Installation or Update
2) Boot Installed System
3) Start Rescue System

>

```

Figure 9-15 Start installation or update option

7. Select **2) Network**.
8. Select **3) NFS** as the network protocol.
9. Select **1) Yes** or **2) No** for network configuration through DHCP.

If you selected **2) No**, enter the static IP information (Figure 9-16 on page 372 shows a sample configuration):

- LPAR's IP address
- LPAR's netmask
- LPAR's gateway
- LPAR's name server
- The NFS server's IP address
- The directory on the NFS server which contains the SLES 11 image

```

Choose the network protocol.

1) FTP
2) HTTP
3) NFS
4) SMB / CIFS (Windows Share)
5) TFTP

> 3
Detecting and loading network drivers

Automatic configuration via DHCP?

1) Yes
2) No

> 2

Enter your IPv4 address.
Example: 192.168.5.77/24
> 9.8.234.221

Enter your netmask. For a normal class C network, this is usually
255.255.255.0
[255.255.255.0]> 255.255.255.128

Enter the IP address of the gateway. Leave empty if you don't need one
> 9.8.234.138

Enter your search domains, separated by a space:
> 9.8.234.180

Enter the IP address of your name server. Leave empty or enter "+++" if you
don't need one
> 9.8.234.180

Enter the IP address of the NFS server
> 9.8.234.180

Enter the directory on the server
[/]> /srv/repository/suse/install/sles-11-ppc

```

Figure 9-16 Static network configuration example

The LPAR begins reading from the SLES 11 image directory and then displays the YaST¹ welcome panel, as shown in Figure 9-17 on page 373.

¹ Your awesome Setup Tool (YaST)

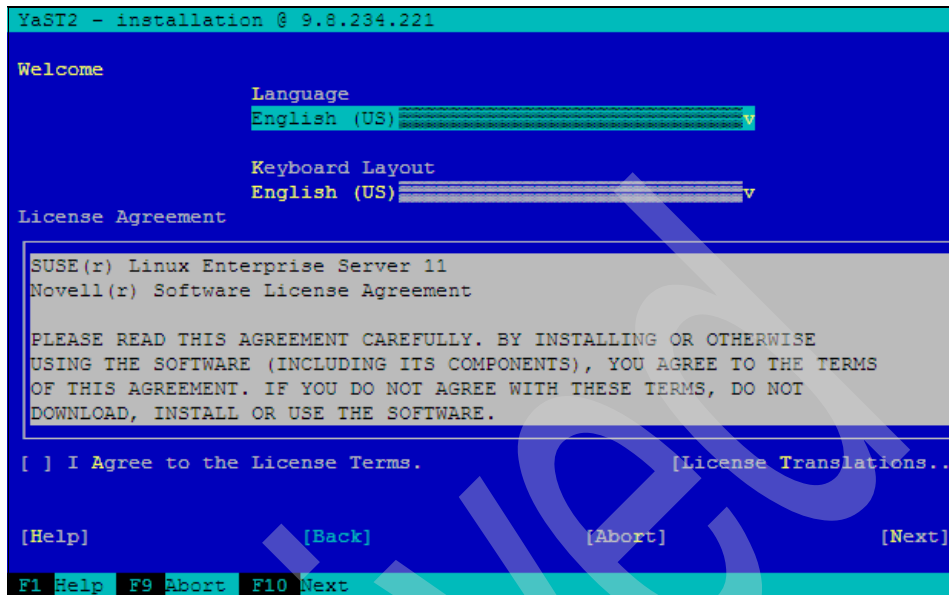


Figure 9-17 YaST Welcome panel

Tip: Navigate the YaST tool by using the:

- ▶ Tab key to move between sections
- ▶ Up and Down Arrow keys to move within a specific window section
- ▶ Spacebar to check for any entries that are indicated by (x)
- ▶ Enter key to confirm a selection with square brackets “[]” around it
- ▶ Delete key to erase entries

10. Highlight **English** in the Language selection panel. Press the Tab key to **[Next]** and then press Enter.
11. On the next panel, if you agree with the terms, select **Yes, I Agree to the License Agreement** and then press the Tab key to **[Next]** and press Enter to confirm.
12. Select **New Installation** for the Installation Mode, as shown in Figure 9-18 on page 374.




Figure 9-18 Installation Mode

13. Configure your clock and time zone information, as shown in Figure 9-19.




Figure 9-19 Clock and Time Zone

14. The Installation Settings window provides the Keyboard layout, Partitioning information, Software installation options, and the install Language configuration. Select the **[Change...]** option to edit any of these fields. Select **[Accept]** when these settings are complete, as shown in Figure 9-20.

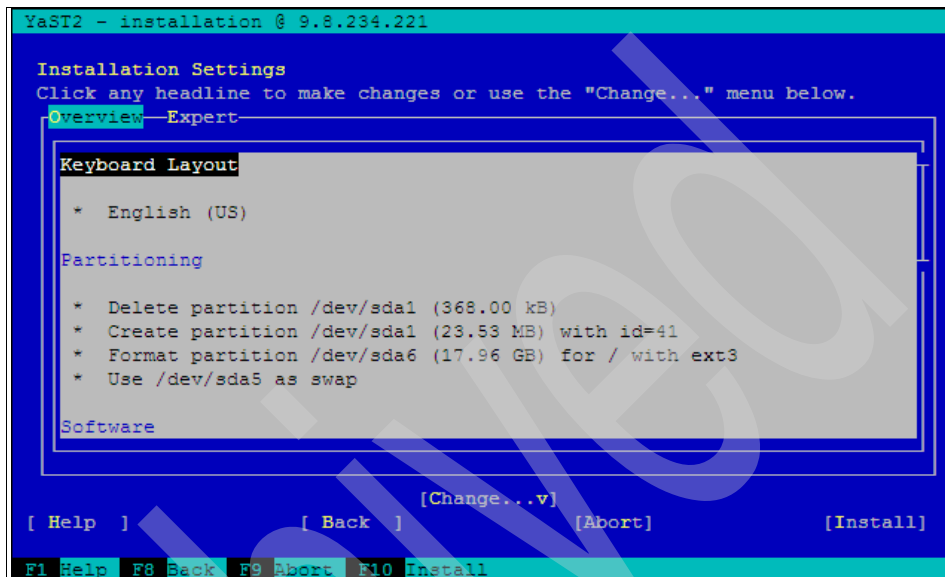


Figure 9-20 Installation Settings

15. If you accept the terms, select **[I Agree]** to the AGFA Monotype Corporation License Agreement, as shown in Figure 9-21 on page 376.




Figure 9-21 AGFA License Agreement

16. Select **[Install]** to start the installation, as shown in Figure 9-22.




Figure 9-22 Confirm Installation

The YaST window refreshes to the installation progress bars, as shown in Figure 9-23. The top status bar shows the progress YaST has made installing a specific package and the bottom is the progress of the entire installation. The system will reboot after the installation completes.

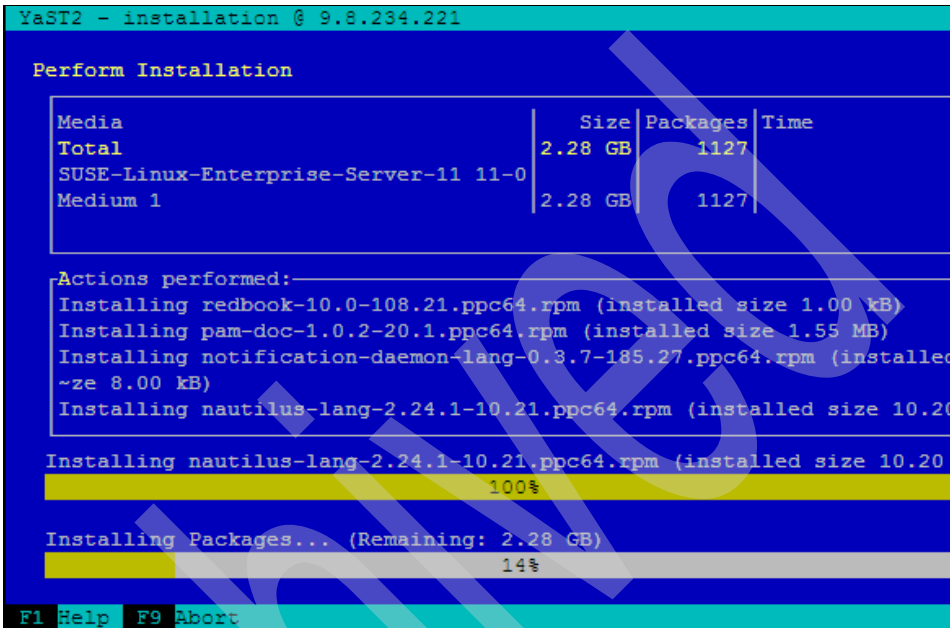


Figure 9-23 YaST installation progress window

Note: If the LPAR does not automatically boot from the intended hard disk (boot device) after reboot, try this:

- Shut down and reactivate the LPAR from the IVM.
- Enter the SMS Menu.
- Select **5. Select Boot Options** → **1. Select Install/Boot Device** → **5. Hard Drive** → **9. List All Devices**.
- Select the appropriate hard disk with the Linux image from the given list.
- Select **2. Normal Mode Boot** → **1. Yes**.

17. After booting from the appropriate storage device YaST will start again. Select **[Continue]** for each hardware device YaST detects, as shown in Figure 9-24 on page 378.

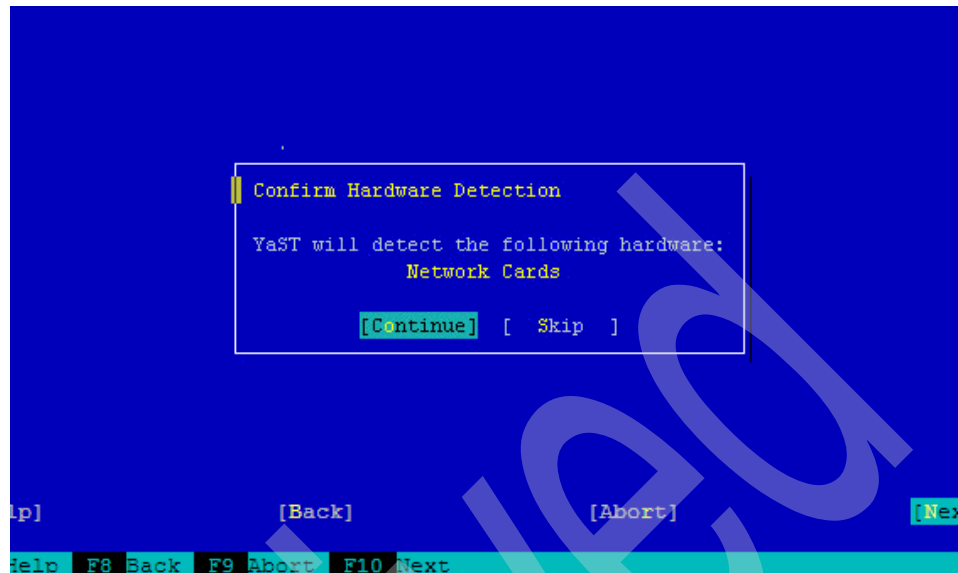


Figure 9-24 Confirm hardware detection window

18. Boot the system. See Figure 9-25.




Figure 9-25 Reboot now

19. Enter the root user's password. Select **[Next]** to confirm, as shown in Figure 9-26.




Figure 9-26 root User Password

20. Provide the host name and the domain. Select **[Next]** to confirm. See Figure 9-27 on page 380.




Figure 9-27 Hostname and Domain Name

21. Select **Use Following Configuration** in the Network Configuration window (Figure 9-28 on page 381) and verify that the Firewall is marked as **enabled**. Press Tab to **[Change]** to change the Secure Shell (SSH) port settings to open, as follows:

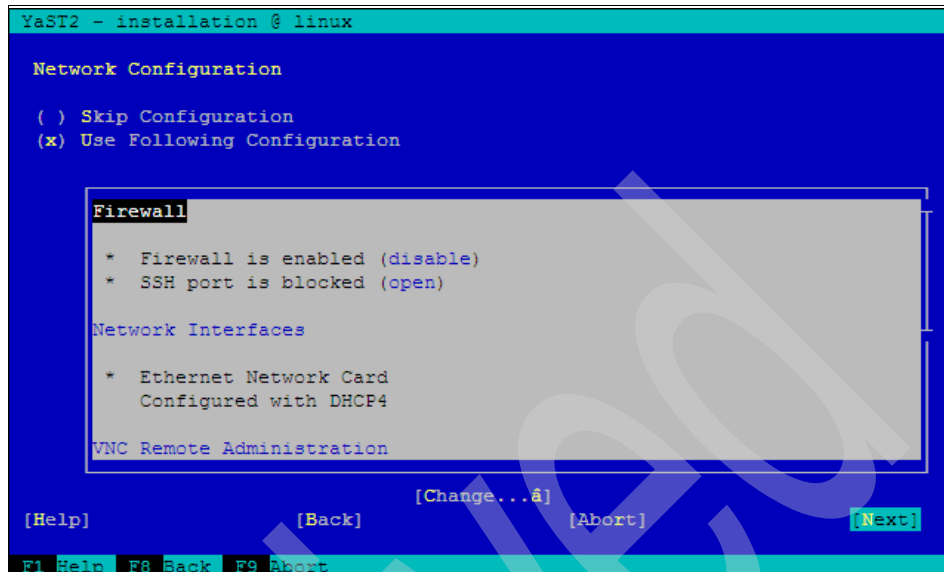


Figure 9-28 Change network configuration

- Select **Firewall** as shown in Figure 9-29.




Figure 9-29 Select firewall

- b. Scroll to **Allowed Services**.
- c. Find and highlight **SSH** in the new window, as shown in Figure 9-30. Finally, press Enter to confirm.

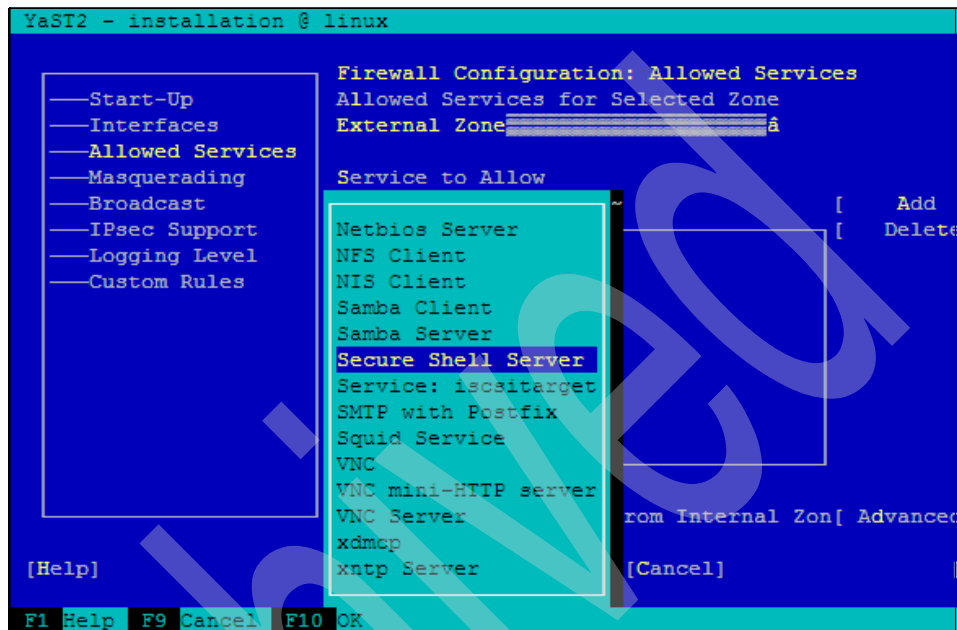


Figure 9-30 Services to allow list and selecting SSH service

- d. Press the Tab key to highlight **[Add]** and the press Enter to confirm.
- e. SSH will appear in the **Allowed Service** list, as shown in Figure 9-31 on page 383. Press **[Next]** to confirm.




Figure 9-31 Allowed Service Secure Shell Server (SSH)

- f. Now the Firewall section of the Network Configuration window (Figure 9-32) shows SSH port is open.

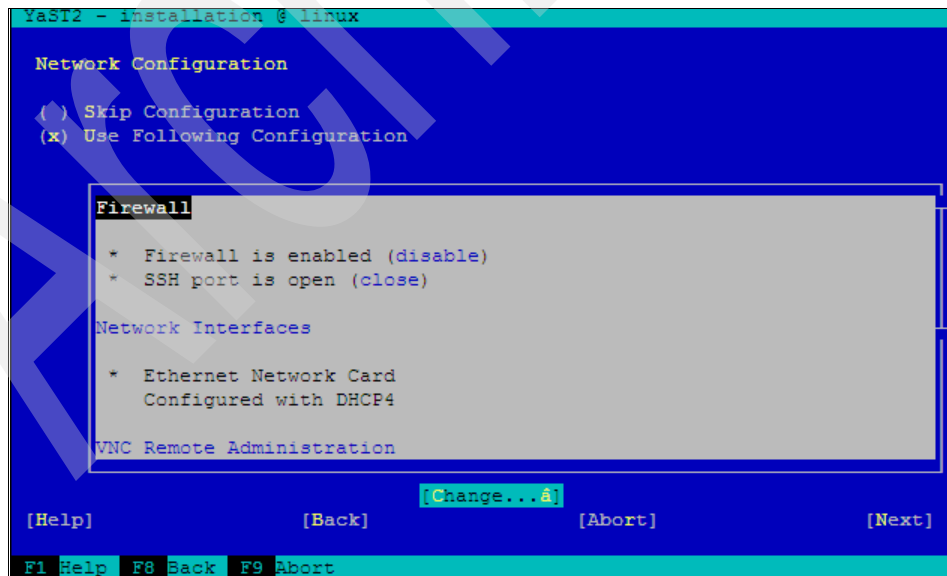


Figure 9-32 SSH port is open

22. Test the Internet connection, if desired.
23. If you want, change the certificate authority (CA) installation setting. Select **[Next]** to confirm the changes.
24. Select the user authentication method appropriate for this LPAR and select **[Next]**. See Figure 9-33.




Figure 9-33 User Authentication Method

25. Create a local user and select **[Next]**. See Figure 9-34 on page 385.




Figure 9-34 New Local User

26. YaST writes the configuration settings and then displays the release notes. Select **[Next]** after reading the release notes.
27. If you want, configure Hardware (Printers), then confirm the described configuration with **[Next]**.
28. YaST displays the Installation Completed window (Figure 9-35 on page 386). If you want to, select **Clone This System for AutoYaST** (see 9.5, “SLES 11 automated installation” on page 387 for more information) and then select **[Finish]**.

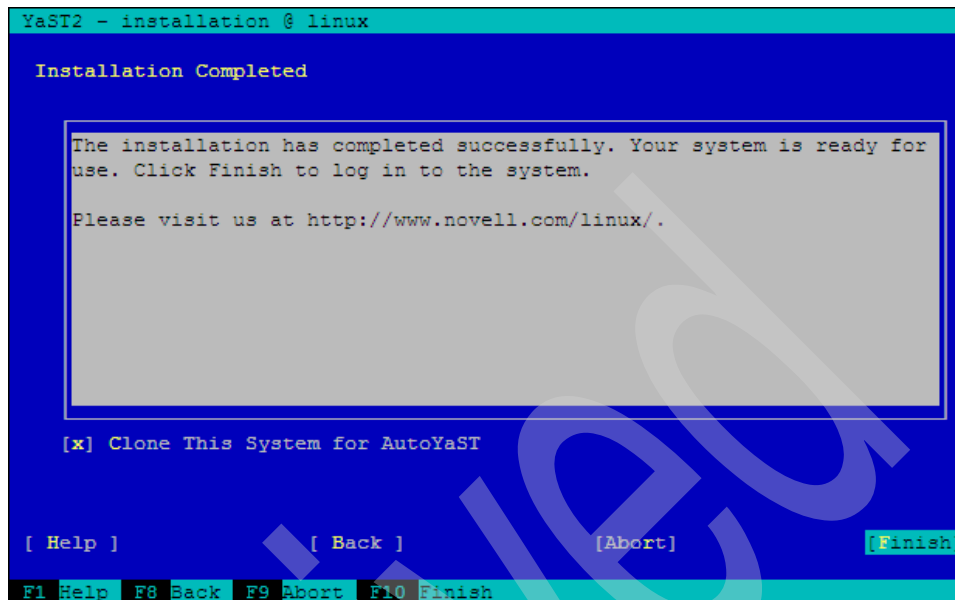


Figure 9-35 Installation completed window

29. Log in to the system with the new user, as shown in Figure 9-36.

```
Method:          3 (Parameter/Value Pairs via stdin)
Match:           serviceable=1 and (type=2 or type=3)

Servicelog ID:   4
Log Timestamp:   Sun Mar  8 14:21:21 2009
Update Timestamp: Sun Mar  8 13:21:21 2009
Notify:          0 (EVENT)
Command:         /etc/ppc64-diag/ppc64_diag_migrate
Method:          3 (Parameter/Value Pairs via stdin)
Match:           refcode="#MIGRATE" and serviceable=0

Starting rtas_errd (platform error handling) daemon: done
Starting Firewall Initialization (phase 2 of 2) done
Master Resource Control: runlevel 3 has been read
Skipped services in runlevel 3: smbfs nfs smartd splash

Welcome to SUSE Linux Enterprise Server 11 (ppc64) - Kernel 2.6.27.19-5-ppc64
onsole).

sles11 login: █
```

Figure 9-36 Login screen

9.4 Native SLES 11 installation

A native SLES 11 installation of a JS43 blade follows a similar process to those given in the VIOS LPAR installation sections. However, several key differences exist, as follows:

- ▶ In a native installation, the IVM terminal is no longer available to complete the Linux installation, but you may use the Serial Over LAN (SOL) console as an alternative. See Appendix A, “Consoles, SMS, and Open Firmware” on page 487 for more information. Use the SOL console to display the SMS menu and the yaboot/lilo/grub or yast options during the installation.
- ▶ The resource allocation of processors, I/O adapters, memory, and storage devices in a native environment is *fixed*.
- ▶ Virtualization functions and features are not available.

9.5 SLES 11 automated installation

SUSE has an automated installation functionality known as AutoYaST to install multiple systems in parallel. The system administrator performs an AutoYaST automated installation by creating a single file containing answers to all the questions normally asked during a SUSE installation. This file resides on a single server system and multiple clients can read it during installation. Several methods are available for creating an AutoYaST profile:

- ▶ Clone the install configuration information from a reference machine.
- ▶ Use the AutoYaST GUI to create and modify the AutoYaST profile. See Appendix B, “SUSE Linux Enterprise Server AutoYaST” on page 515 for more information about this method.
- ▶ Use an XML editor to create an AutoYaST profile from scratch.

This section describes a method to clone installations to identical machines from a reference machine:

1. Perform a CD/DVD or network installation.
2. In the Installation Completed YaST window (see Figure 9-35 on page 386), select **Clone This System for AutoYaST** and then select **[Finish]**. This step creates an AutoYaST profile in `/root/autoinst.xml` that is ready for immediate use.

Tip: *SUSE Linux Enterprise Server* guide is a good source of information regarding all AutoYaST profile options. Read the section about automated installation. It is available at:

http://www.novell.com/documentation/sles11/pdfdoc/book_sle_deployment/book_sle_deployment.pdf

3. Place the `autoinst.xml` file on a server that is accessible for all the systems with *identical* hardware configurations to use.

Performing an AutoYaST installation

Pass the location of the AutoYaST profile and install images during boot. BOOTP does not have the ability to provide anything more than the location to the bootable image and the server IP address hosting the AutoYaST profile. Use the Open Firmware prompt to pass the required parameters.

To perform an AutoYaST installation

1. Type 8 during the LPAR boot process to go to the Open Firmware prompt, as shown in Figure 9-37.

1 = SMS Menu	5 = Default Boot List
8 = Open Firmware Prompt	6 = Stored Boot List
Memory Keyboard Network SCSI Speaker ok	
0 > _	

Figure 9-37 Open Firmware prompt

2. Type the following command in the Open Firmware prompt to start automated installation. For example, if the profile is served using NFS:

```
boot net autoyast=nfs://193.200.1.80/home/autoinst.xml
install=nfs://192.168.1.254/SuSE/SLES11
```

Note: This automated installation assumes that a DHCP server will provide the client with the correct IP address information. A system using static IP requires additional parameters for `sshpassword`, `hostip`, `netmask`, `gateway`, and `nameserver`.

3. Press Enter to start the process. The automated SUSE Linux Enterprise Linux installation is now complete.

9.6 IBM service and productivity tools

The IBM service and productivity tools are packages that are installed after a successful SLES installation.

Important: These packages are *not* shipped with the SUSE installation CDs. IBM owns and distributes them.

These packages enable features such as:

- ▶ Reliability, availability, and serviceability (RAS) functionality
- ▶ I/O hotplug
- ▶ Dynamic Logical Partitioning (DLPAR) capabilities
- ▶ Live partition migration capabilities

See Appendix D, “Service and productivity tools for Linux” on page 539 for more information about installing the service and productivity tools specific to your system’s configuration.

JS23 and JS43 power management using EnergyScale technology

The EnergyScale technology described in 3.4, “IBM EnergyScale technology” on page 42 can be used by the BladeCenter Advanced Management Module (AMM) and IBM Systems Director Active Energy Manager (AEM) to monitor and control power usage of the IBM BladeCenter JS23 and JS43 blades. This chapter describes how to use the BladeCenter AMM and the AEM extension of IBM Systems Director to manage power usage.

This chapter contains the following topics:

- ▶ “Power management through the AMM” on page 392
- ▶ “Power management through AEM” on page 401

An in-depth discussion of installing, configuring, and using Active Energy Manager can be found in *Going Green with IBM Systems Director Active Energy Manager*, REDP-4361.

EnergyScale options can be performed on the AMM and the blade servers JS23/JS43 using AMM options or AEM.

10.1 Power management through the AMM

The IBM BladeCenter Advanced Management Module (AMM) provides a Web-based and command line user interface to monitor and control individual blades and switch modules installed in the BladeCenter. The AMM also collects historical or trend data for individual components in the IBM BladeCenter. This data can be reviewed from the Web-based interface or CLI. The information can also be collected by the Active Energy Manager extension for IBM Systems Director. This section describes how to use both interfaces of the AMM to make changes in the Power Saving and Power Capping modes. The examples shown here are from a BladeCenter H (BCH). The BladeCenter S (BCS) screens look alike except the BCS has only one power domain.

10.1.1 Using the AMM Web GUI for blade power management

A detailed description of how to access the AMM is not covered in this section but detailed access and use information can be found in *IBM eServer BladeCenter Systems Management*, REDP-3582. The Web GUI can be used from a Web browser pointed to the IP address of the AMM.

AMM Power Management view

From the main menu panel in the Monitors section, click **Power Management**; see Figure 10-1 on page 393 for an example of the Power Management option. The BladeCenter H (BCH) has two power domains. The BladeCenter S (BCS) only has one power domain.

The AMM Power Management view has the following sections:

- ▶ BladeCenter Power Domain Summary
- ▶ BladeCenter Power Domain Planning
- ▶ BladeCenter Chassis Power Summary
- ▶ BladeCenter Chassis Configuration Setting
- ▶ BladeCenter Chassis Power Consumption

Figure 10-1 on page 393 shows highlighted boxes. Number 1 links to the AMM Power Management Policy. Number 2 links to the Power Domain, which contains blade server slots and chassis components.

In each section, a complete description of the fields can be obtained by clicking the blue question mark (?).

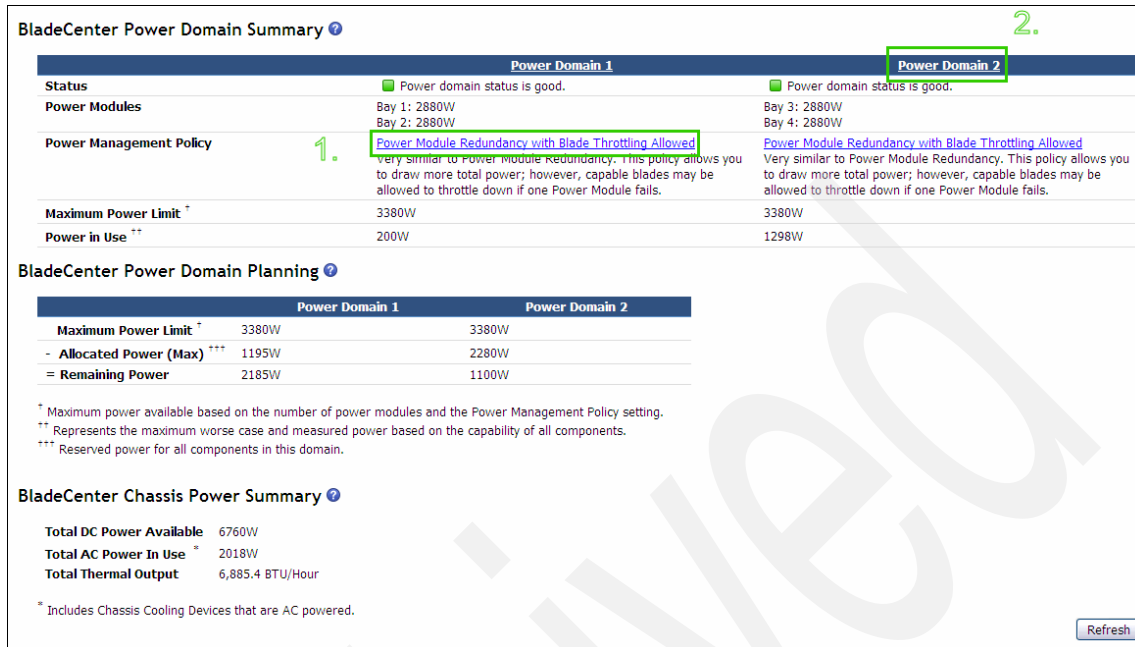


Figure 10-1 BladeCenter Power Domain Summary

Scrolling the page down below the Blade Chassis Power Summary provides access to the acoustical settings for the chassis, power consumption history, and links to view the thermal and power trending history for several of the chassis components. An example of the options is shown in Figure 10-2 on page 394 and Figure 10-3 on page 394.




Figure 10-2 Additional power settings




Figure 10-3 Chassis thermal and trending options

Choose the **Power Management Policy** link (number 1 as shown in Figure 10-1 on page 393), which enables you to select three different management policies. Figure 10-4 on page 395 shows an example of this option.

Three different selections can be applied to manage the power domain. As mentioned previously, the BCH has two power domains. Each domain can set this policy separately and they do not have to match.

The Basic Power Management option is the least conservative of the three available options. This policy does not guarantee redundancy so if a power module fails, the BladeCenter chassis could be affected. When this policy is in effect the value for Maximum Total Power Limit is used to determine if a blade can power on. As long as the power used is below this maximum value the blade can continue the power on procedure. If one of the power modules fails, blades capable of throttling can be throttled back to keep the domain usage below the maximum limit. If the limit is reached, the domain can be affected by power loss.

The Power Module Redundancy option is used when only one AC source is present. One AC source in this case means the electrical grid. For example, the BCH has two line cord inputs. Each is capable of connecting to its own AC power source. If the two line cords attach to the same power grid, it is considered a single AC source. Another possibility is to have a data center wired so that each AC line cord of the BCH can be plugged into a separate power grid or AC source. If an outage occurs on one power module, this option allows some redundancy and can keep the chassis and blades running. A second power module failure can result in a domain outage.

The third option, Power Module Redundancy with Blade Throttling is basically the same as the Power Module Redundancy option except in case of a failure of a module, the blades can be throttled back to allow the chassis and blades to continue to run. After an outage has been restored, the blades return to running at the un-throttled power level.

BladeCenter Domain 1 Power Management Policies ⓘ

Links [Power Summary](#) | [Domain 2 Power Management Policies](#)

This table lists the power management policies ordered from most conservative to least conservative.

Select	Option Name	Power Supply Failure Limit*	Maximum Power Limit (Watts)	Estimated Utilization**
<input type="radio"/>	Power Module Redundancy Intended for a single AC power source into the chassis where each Power Module is on its own dedicated circuit. Total allowed power draw is limited to one less than the number of Power Modules when more than one Power Module is present. One Power Module can fail without affecting blade operation. Multiple Power Module failures can cause the chassis to power off. Note that some blades may not be allowed to power on if doing so would exceed the policy power limit. More...	1	2880	6%
<input checked="" type="radio"/>	Power Module Redundancy with Blade Throttling Allowed Very similar to Power Module Redundancy. This policy allows you to draw more total power; however, capable blades may be allowed to throttle down if one Power Module fails. More...	1	3380	5%
<input type="radio"/>	Basic Power Management Total allowed power is higher than other policies and is limited only by the total power capacity of all the Power Modules up to the maximum of chassis power rating. This is the least conservative approach, since it does not provide any protection for AC power source or Power Module failure. If any single power supply fails, blade and/or chassis operation may be affected. More...	0	3520	5%

* This is the maximum number of power supplies that can fail while still guaranteeing the operation of the domain in the selected policy. ** The estimated utilization is based on the maximum power limit allowed in this policy and the current aggregated power in use of all components in the domain.

Save Refresh

Figure 10-4 Power Management Policies

Selecting the Power Domain link (referenced by number 2 in Figure 10-1 on page 393) enables you to select different components of the chassis to modify the power settings. This is where you can select the blade server JS23/JS43 to modify the power policies.

AMM Power domain details view

The BladeCenter Power Domain Details view shows each module that is serviced by the power domain. The chassis location (bay), status, module type, power state, current power in use, allocated maximums and minimums in watts, and CPU duty cycles are listed for each module. Modules that have specific

capabilities or collect power trend data appear as a link to a module-specific detail view. Figure 10-5 provides an example of this selection.

BladeCenter Power Domain 2 Details ?

Links [Power Summary](#) | [Power Domain 1](#)

Bay(s)	Status	Module	State	Power In Use	Allocated Power		CPU Duty Cycles
					Maximum	Minimum	
Chassis Components							
		Midplane	On	5W	10W	10W	n/a
1		Media Module	On	5W	10W	10W	n/a
Power Module Cooling Devices							
1		Power Module	On	15W	30W	30W	n/a
2		Power Module	On	15W	30W	30W	n/a
3		Power Module	On	15W	30W	30W	n/a
4		Power Module	On	15W	30W	30W	n/a
Management Modules							
1		xbch4mm	On	12W	25W	25W	n/a
2		Advanced Management Module Bay 2 (not present)		0W	15W	15W	n/a
I/O Modules							
1		Ethernet SM	On	22W	45W	45W	n/a
2		Ethernet SM	On	22W	45W	45W	n/a
7		Fibre Channel SM	On	15W	30W	30W	n/a
9		Fibre Channel SM	On	32W	65W	65W	n/a
Blades							
[8]		JS23-Redbook	On	213W	406W	320W	n/a **
[9-10]		JS43-Redbook	On	462W	745W	573W	n/a **
[13-14]		Japan_demo	On	456W	764W	607W	n/a **

* This blade may throttle if redundancy is lost in this power domain.

** Click on the module name to view CPU speeds.

* Cannot communicate with the blade. The power values for this blade are assumed.

Figure 10-5 Power Domain Details

Selecting the components, such as a blade enables you to set several power management options. Figure 10-6 on page 397 lists the options available for a blade that is capable of power management.

In this panel you can see what the blade power capabilities are. In this example the blade supports:

- ▶ Power metering
- ▶ Power capping
- ▶ Static low power saving
- ▶ Dynamic power saving

BladeCenter JS23-Redbook (Bay 8) Power Summary

[Links](#) [Power Summary](#) | [Power Domain 1](#) | [Power Domain 2](#)

The following capabilities are supported:

- Power metering
- Power capping
- Static low power saving
- Dynamic power saving

Processors	4
Effective CPU Speed	3621 MHz
Maximum CPU Speed	4204 MHz

BladeCenter JS23-Redbook (Bay 8) Configuration Setting

Power Capping Options

Power Capping

Enabled

Maximum Power Limit (range 180-406, guaranteed range 320-406)

406

Power Savings

☐ Static High Performance

☐ Static Low Power Saver

☒ Dynamic Power Saver

☐ Favor Performance over Power

Save

Figure 10-6 Blade power configuration settings

Power capping allows you to allocate less power and cooling to a system. This can help save on data center infrastructure costs, and then potentially allow more servers to be put into an existing infrastructure.

To enable the Power Capping option, select **Enable** from the pull-down menu. Then, set a cap level using the Maximum Power Limit range box. This value will limit the power usage to the value specified. When the limit has been reached the blade is throttled back to limit the consumption to the maximum value.

Power Savings setting is used to put the server into a mode that consumes less energy.

Power savings can be enabled using one of the various options. If you select the **Dynamic Power Saver** option, you may also select the **Favor Performance over Power** usage option.

Scroll down the page down to view trend data for the particular blade server selected. Figure 10-7 on page 398 shows an example of the trend data.




Figure 10-7 Blade server trend data

10.1.2 Using the AMM CLI for blade power management

Similar to the Web interface, the command-line interface (CLI) can be used to display power domain and specific module information. The AMM CLI can be accessed by either a Telnet or SSH to the IP address of the AMM. Complete the login by using the same user ID and password that is used for the Web interface.

The `fuelg` command

The `fuelg` command used from the AMM CLI controls and displays BladeCenter power functions and the EnergyScale features of an IBM BladeCenter JS23 or 43. This command with all options is shown in Example 10-1 on page 399.

Example 10-1 *fuelg command with help flag*

```
system> fuelg -h
fuelg {<pdx> {-pm}}|{-pt}}|{-am|-int|-pcap|-pme|-ps}}|{-pt}}|{-tt}
Power management settings and values by domain
pdx:   power domain (pd1, pd2)
-pm:   power management settings (redwoperf, redwperf, nonred)
-am:   acoustic mode setting (on, off)
-int:  polling interval (between 10 and 60 minutes, in increments of 5)
-pcap: power cap for blades. Note: setting a pcap will automatically enable
pme.
-pme:  power management and capping enabling for blades (off, on). Note: the
blade must be powered on before enabling capping.
-ps:   power saver mode for blades (off, on). Note: the blade must be powered
on before enabling power saver mode.
-pt:   power trending data (1, 6, 12, or 24 hours)
-tt:   thermal trending data (1, 6, 12, or 24 hours)
```

Example 10-2 shows the **fuelg** command used from the `system>` prompt with no flags to display the BladeCenter Power Domain information.

Example 10-2 *The fuelg command used to display BladeCenter Power Domain*

```
system> fuelg
Note: All power values are displayed in Watts.

Total DC Power Available: 2900
Total AC Power In Use:    807
Total Thermal Output:     2753 BTU/hour

Power Domain
-----
Status: Power domain status is good.
Modules:
  Bay 1: 1450
  Bay 2: 1450
  Bay 3: 1450
  Bay 4: 1450
Power Management Policy: AC Power Source Redundancy
Power in Use:           769
Total Power:            2900
Allocated Power (Max):  1346
Remaining Power:        1554
Power Service: 220 VAC

-am off
-int 10
```

From the AMM system> prompt use the **env** command to set the specified blade as the persistent target for commands during the current session. Example 10-3 shows the environment being set to blade slot 4. Note that the prompt changes.

Example 10-3 The env command used to set a persistent target

```
system> env -T blade[4]
OK
system:blade[4]>
```

Example 10-4 shows the **fuelg** command with no other parameters being used to display the capabilities, current settings, and power consumption values of the blade in BladeCenter slot 4.

Example 10-4 The fuelg command used to display blade slot power information

```
system:blade[4]> fuelg
-pme off
-ps off
PM Capability: Dynamic Power Measurement with capping and power saver
mode
Effective CPU Speed: 3800 MHz
Maximum CPU Speed: 3800 MHz
-pcap 256 (min: 256, max: 282)
Maximum Power: 150
Minimum Power: 150
Average Power: 150
```

The following **fuelg** command flags change blade-specific EnergyScale parameters:

- | | |
|-----------------------------|---|
| -int <i>interval</i> | Polling interval is between 10 and 60 minutes in increments of 5. |
| -pcap <i>value</i> | Power cap for blades where the value is between the minimum and maximum amounts. Setting pcap also enables pme. |
| -pme <i>on/off</i> | Power management and capping enabled/disabled uses minimum value unless a different pcap value has been used. |
| -pt <i>interval</i> | Power trending data is in intervals of 1, 6, 12, or 24 hours. |

Power Saver Mode enabled from the AMM CLI

The baseline values and power consumption are shown in Example 10-4. Power Saver Mode was enabled using the **fuelg -ps** command and then the changes reviewed using the **fuelg** command shown in Example 10-5 on page 401.

Example 10-5 Power Saver Mode enabled

```
system:blade[4]> fuelg -ps on
OK
system:blade[4]> fuelg
-pme off
-ps on
PM Capability: Dynamic Power Measurement with capping and power saver
mode
Effective CPU Speed: 3440 MHz
Maximum CPU Speed: 3800 MHz
-pcap 256 (min: 256, max: 282)
Maximum Power: 139
Minimum Power: 139
Average Power: 139
```

Power trend data for the last hour was reviewed using the **fuelg -pt 1** command shown in Example 10-6.

Example 10-6 Power trend data in Power Saver Mode

```
system:blade[4]> fuelg -pt 1
Date      Time      Max  Min  Avg
-----
07/02/08 13:20:20 162  150  150
07/02/08 13:30:20 154  150  150
07/02/08 13:40:20 154  138  147
07/02/08 13:50:19 150  138  139
07/02/08 14:00:20 142  138  139
07/02/08 14:10:20 142  138  139
```

10.2 Power management through AEM

The Active Energy Manager (AEM) plug-in for IBM Systems Director provides the ability to monitor and manage energy capabilities of a resource. In addition, AEM provides management functions through integration with IBM Systems Director such as scheduling options to apply power saver mode and power capping to support systems or groups of systems at predetermined times.

AEM also supports the application of power policies to supported systems or groups of systems. As a result, managing energy across multiple systems can be deployed, maintained, and modified with minimal effort. IBM Systems Director

Version 6.1.1 and Active Energy Manager version 4.1.1 are required for energy management of the JS23/JS43 blade servers.

The intent of this book is *not* to explain and demonstrate all the different options available using AEM. The references in this publication only intend to show some of the options available and what can be configured. In most instances, multiple paths exist to the same options in AEM. These options will be described in greater detail in a subsequent publication.

The following information and examples assume that IBM Systems Director and the Active Energy Manager extension have been installed and configured. Refer to the following publications for more information:

- ▶ Complete planning, installation, configuring, and usage information of IBM Systems Director is in *Implementing IBM Systems Director 6.1*, SG24-7694.
- ▶ IBM Active Energy Manager installation and use is in *Going Green with IBM Systems Director Active Energy Manager*, REDP-4361.

10.2.1 Active Energy Manager console

After IBM Systems Director has been installed and your target systems have been configured, you may use Active Energy Manager to manage the power capabilities of the chassis and blade servers.

To access the Active Energy Manager option, use the menu option as shown in Figure 10-8 on page 403.



Figure 10-8 Director menu options

After you select Active Energy Manager, the options shown in Figure 10-9 on page 404 become available. In this example, we have four resources that can be managed by Active Energy Manager. One of the resources is the BCH chassis and the other three are blade servers within the chassis.

Note: When a JS43 is present in the chassis, the AMM might have problems reporting the JS43 blade server to Active Energy Manager. To correct this issue, be sure that the AMM firmware level is at BPET48F or later. Otherwise, the JS43 might not appear as a resource that can be managed by Active Energy Manager.

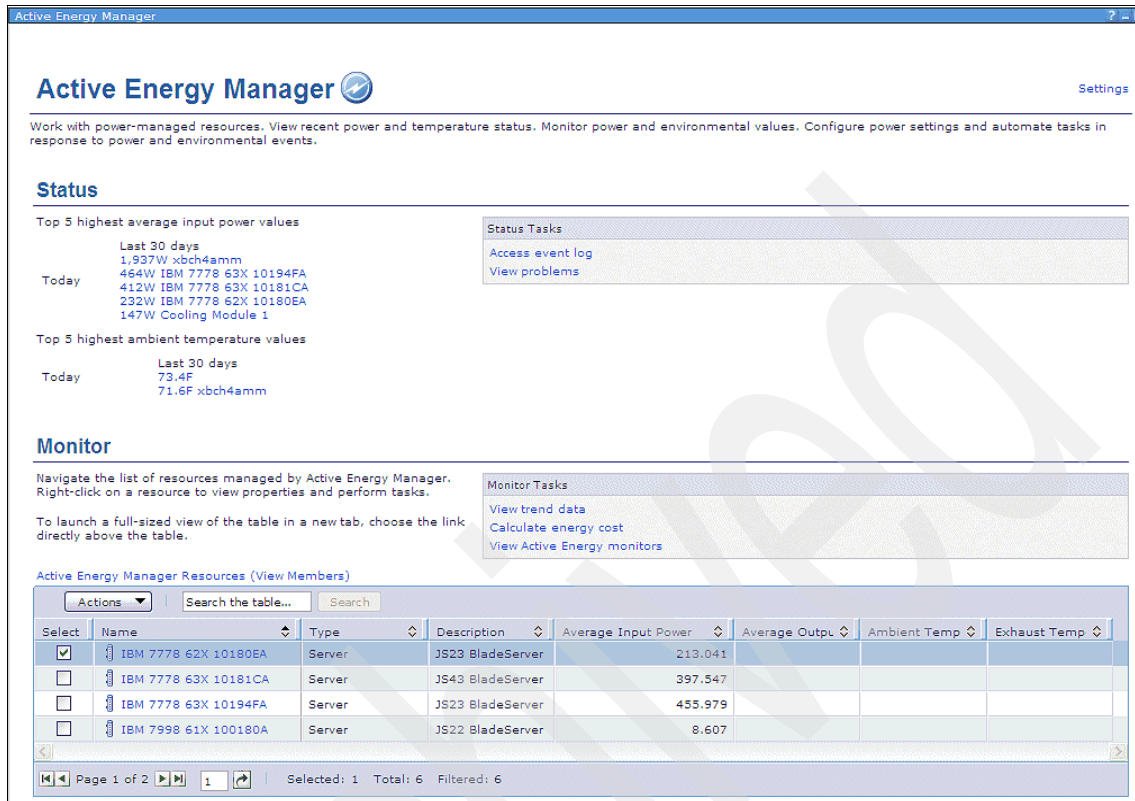


Figure 10-9 Active Energy Manager options

10.2.2 Active Energy Manager energy properties

In Active Energy Manager window, you may use the check boxes to select the resource to work with. Figure 10-10 on page 405 shows an example of selecting the **BladeCenter Chassis** and then clicking the **Actions** button to select the **Properties** option, as displayed in Figure 10-11 on page 405.

Active Energy Manager Resources (View Members)

Actions | Search the table... Search

Select	Name	Type	Description
<input type="checkbox"/>	IBM 7778 62X 10180EA	Server	
<input type="checkbox"/>	IBM 7998 61X 100180A	Server	
<input type="checkbox"/>	IBM 7998 61X 1001C8A	Server	
<input checked="" type="checkbox"/>	xbch4amm	BladeCenter Chass	

Figure 10-10 Select resource

Related Resources	▶
Topology Perspectives	▶
Create Group	
Change Password	
Manage MIBs...	
Remove...	
Rename...	
SNMP Browser	
Add to	▶
Energy	▶
Inventory	▶
Release Management	▶
Security	▶
System Configuration	▶
System Status and Health	▶
Properties	
Active Energy Manager Resources	▶
Import Groups	
Columns...	
Export	
Select All	
Deselect All	
Show Filter Row	
Clear All Filters	
Edit Sort	
Clear All Sorts	

Figure 10-11 Actions options

Using the various tabs properties view you can see information about the resource selected. Select the Active Energy tab to view data available about the chassis, as shown in Figure 10-12 on page 406.

xbch4amm (Properties)

Name: Actions

Access: ☒ OK

Status: ☒ Critical

General Active Status Applied Activities Configuration Event Log Inventory **Active Energy**

Active Energy Manager support level: Full

Last time metered: Friday, April 17, 2009 11:57:27 AM CDT

Nameplate power: 8,800 watts

Average input power: 1,948 watts

Average output power: 1,736 watts

Ambient temperature: 23 Celsius

Exhaust temperature: 36 Celsius

Energy price: 0.12

Currency type: USD (\$)

Cooling rate multiplier: 1.5

Metering active: True

Metering interval: 5 minutes

Edit...

Figure 10-12 Properties - Active Energy tab

Use the Edit tab to modify the energy price and metering values. This data can then be used for cost estimating of the power used for the chassis. Figure 10-13 shows an example of the values available to edit.

Edit Properties

Active Energy Manager support level: Full

Last time metered: Friday, April 17, 2009 11:57:27 AM CDT

Nameplate power: 8,800 watts

Average input power: 1,948 watts

Average output power: 1,736 watts

Ambient temperature: 23 Celsius

Exhaust temperature: 36 Celsius

Energy price:

Currency type: ▼

Cooling rate multiplier:

Metering active: ▼

Metering interval: minutes

OK Cancel

Figure 10-13 Edit values

10.2.3 BladeCenter Energy Properties

In this section, we look at the energy management options available on the JS23/JS43. Using Active Energy Manager you can configure power capping, power saver and view trend data on the blade server.

Enabling power capping

To enable power capping on the blade server, use Active Energy Manager and select the desired blade resource. Click the **Actions** button and then, as shown in Figure 10-14, select **Energy** → **Manage Power** → **Power Capping**.

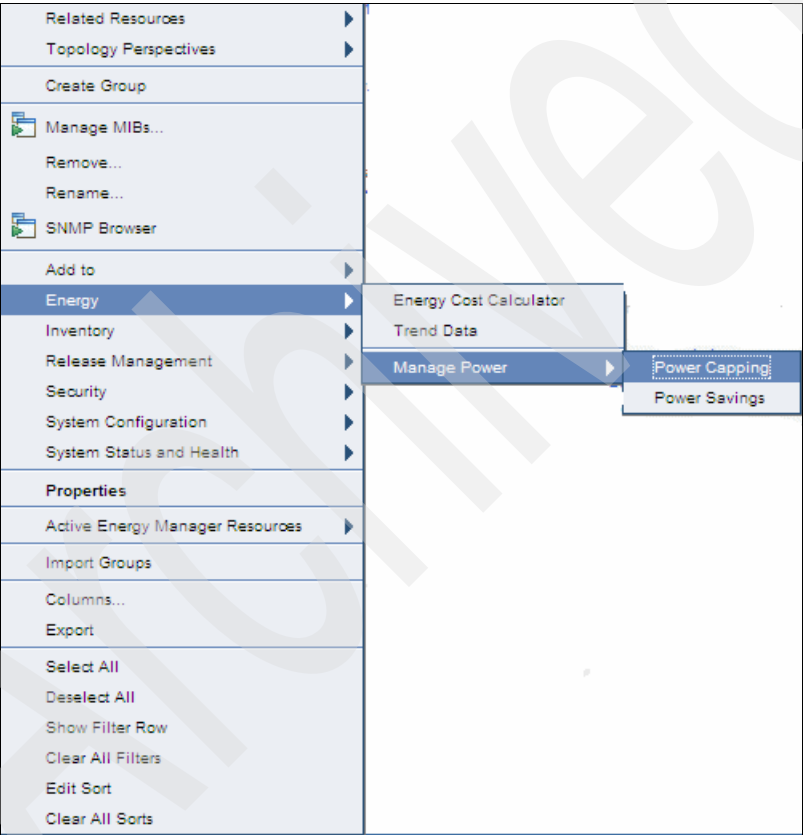


Figure 10-14 Power Capping option

To activate power capping, select the **Activate Power Capping** option. Then, you may modify the parameters for Power cap type field and set the value. Two options are available for the power cap type selection. One option is based on a percentage, the other is based on wattage. Choose the preferred value and set it by using the slider bar or input a number in the entry box.

Click **Save** to save your settings. An example of the power capping options are shown in Figure 10-15.

Choose either an absolute power cap, or a percentage of the available power cap.

☐ Activate Power Capping ☒ Deactivate Power Capping

Power cap type:
Absolute value (Watts)

Power cap value:
180W 406W W
Values between 180W and 320W are not guaranteed

Targets:

Name	Current power cap	Power Capping
IBM 7778 62X 10180EA	None	Inactive

Page 1 of 1 1 Total: 1

Save Close

Figure 10-15 Power Capping options

Figure 10-16 shows an example of the power capping features enabled for the blade server.

Choose either an absolute power cap, or a percentage of the available power cap.

☒ Activate Power Capping ☐ Deactivate Power Capping

Power cap type:
Absolute value (Watts)

Power cap value:
180W 338W 406W W
Values between 180W and 320W are not guaranteed

Targets:

Name	Current power cap	Power Capping
IBM 7778 62X 10180EA	338W (69.91%)	Active

Page 1 of 1 1 Total: 1

Save Close

Figure 10-16 power capping enabled

Enabling power savings

To enable power savings on the blade server, use Active Energy Manager and select a blade resource. Click **Actions** and then, as shown in Figure 10-17 on page 409, select **Energy** → **Manage Power** → **Power Savings**.

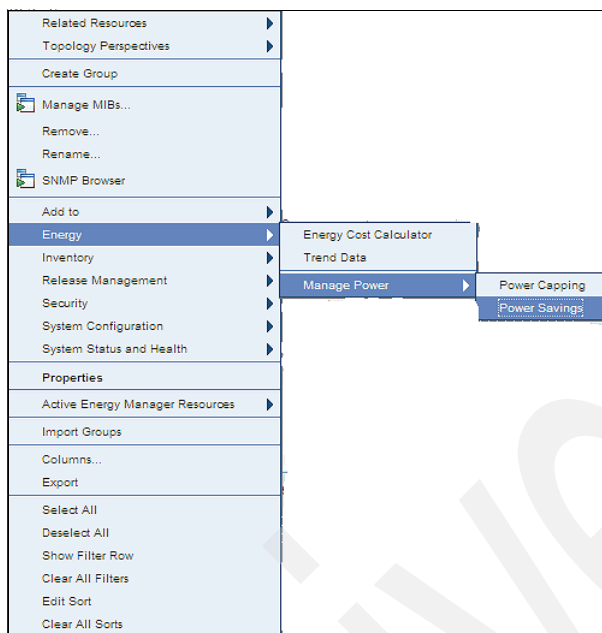


Figure 10-17 Power Savings option

The power savings options are as follows:

- ▶ No power savings
Choose this option to have no power savings. The processor runs at high speed.
- ▶ Static power savings
Choose this option to reduce power usage by lowering processor speed. This option saves energy while maintaining a reasonable processor performance.
- ▶ Dynamic power savings
Choose this option to automatically balance power usage and processor performance. This option saves energy while minimizing performance impact. When dynamic power savings is enabled, you can also select to **Favor power** or **Favor performance**.

Figure 10-18 on page 410 shows an example of the power savings options. Click **Save** to save your modified preferences.

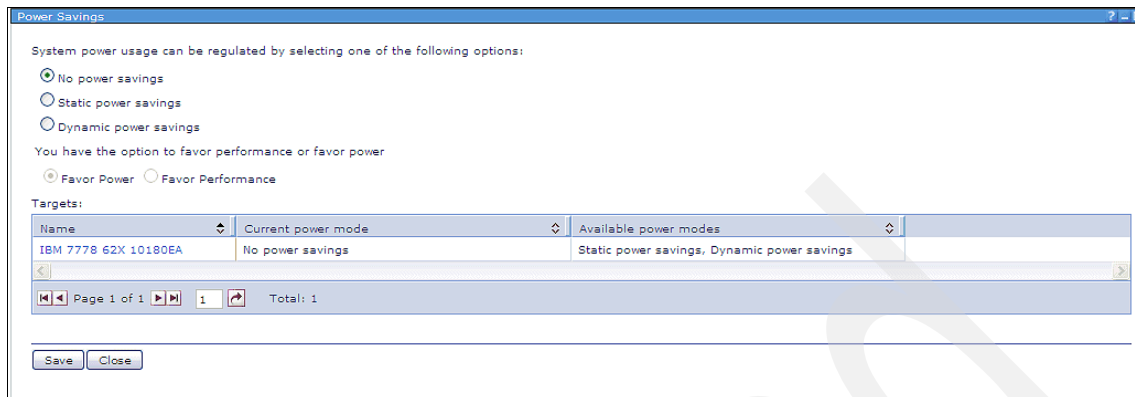


Figure 10-18 Power Savings options

Viewing blade server JS23/JS43 trend data

Using the Active Energy Manager you can view trend data for the JS23/JS43. Trend data provides information usable for viewing details relating to power usage, capping values, and informational events. This data can be charted for the last hour up to the last year in different intervals. Click **Actions** and then, as shown in Figure 10-19, select **Energy** → **Trend Data**.

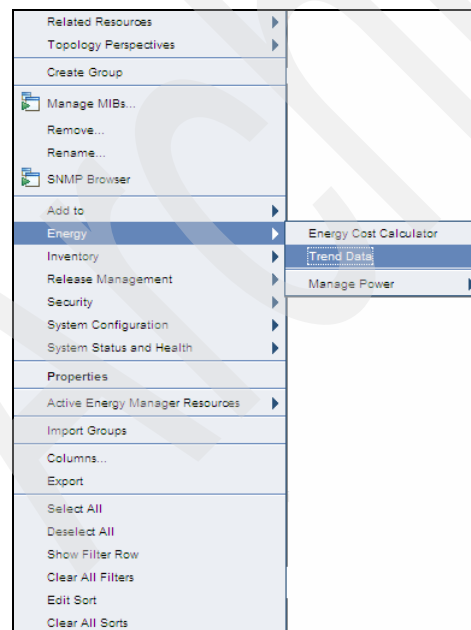


Figure 10-19 Trend Data option

The trend data is displayed, as shown in Figure 10-20.

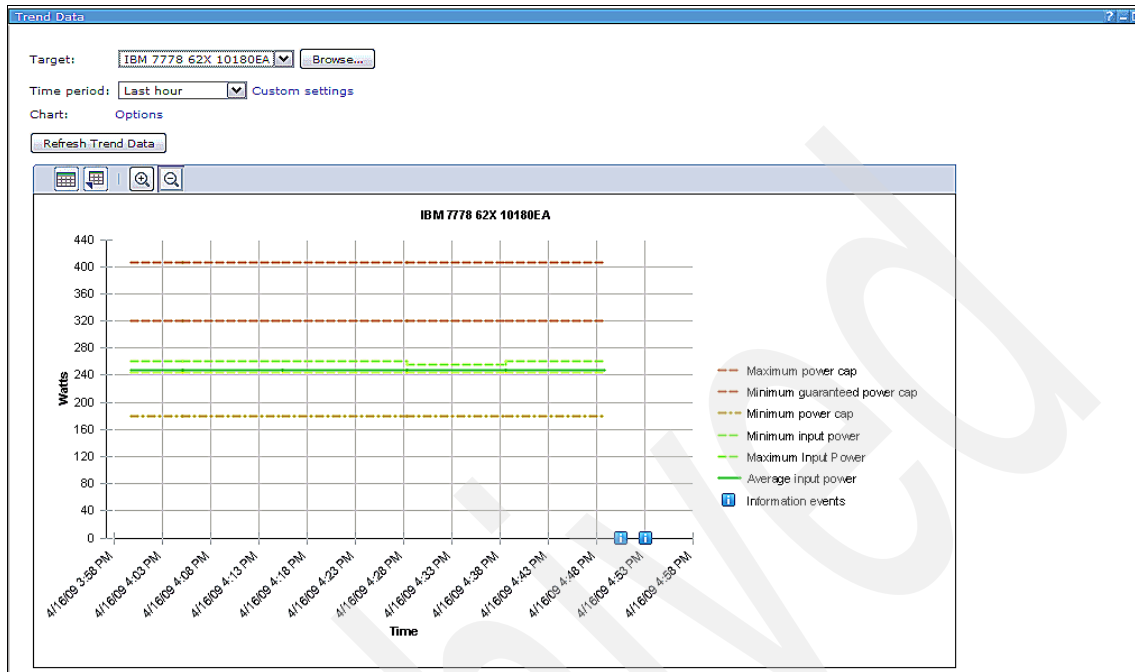


Figure 10-20 Trend Data display

In the trend data panel you can view various power details. You may select a time period from the pull-down menu, or click the **Custom settings** link to change the values. Click **Refresh Trend Data** to see your changes.

Scroll down trend data display to view environment data, such as temperature.

You may also modify chart data by clicking the **Options** link. Use this option to tailor the details on the trend data output based on the data you want to track. Figure 10-21 on page 412 shows an example of the options available.

Chart Options

Select the values to display on the trend data charts.

Power Chart

<input checked="" type="checkbox"/> Input power	<input checked="" type="checkbox"/> Power cap
<input type="checkbox"/> Input power range	<input type="checkbox"/> Power cap range
<input type="checkbox"/> Input power (externally metered)	<input checked="" type="checkbox"/> Critical events
<input checked="" type="checkbox"/> Output power	<input checked="" type="checkbox"/> Warning events
<input type="checkbox"/> Output power range	<input checked="" type="checkbox"/> Information events

Environmental and CPU Chart

<input checked="" type="checkbox"/> Ambient temperature	<input type="checkbox"/> Dew point
<input type="checkbox"/> Ambient temperature range	<input type="checkbox"/> Dew point range
<input checked="" type="checkbox"/> Ambient temperature (externally metered)	<input type="checkbox"/> Dew point (externally metered)
<input checked="" type="checkbox"/> Exhaust temperature	<input checked="" type="checkbox"/> Effective CPU speed
<input type="checkbox"/> Exhaust temperature range	<input checked="" type="checkbox"/> Critical events
<input checked="" type="checkbox"/> Humidity	<input checked="" type="checkbox"/> Warning events
<input type="checkbox"/> Humidity range	<input checked="" type="checkbox"/> Information events
<input type="checkbox"/> Humidity (externally metered)	

OK Cancel

Figure 10-21 Trend data chart options

Information events as noted by the icon will display details about the event if you move the mouse pointer over the icon. In Figure 10-22 you can see that a mode change was made on a resource.




Figure 10-22 Information event details

Trend data may also be exported to your IBM Systems Director server file system. Use the export option and save the file in your preferred location. An

example of this option is shown in Figure 10-23. The file is then viewable by using a spreadsheet program such as Excel®.

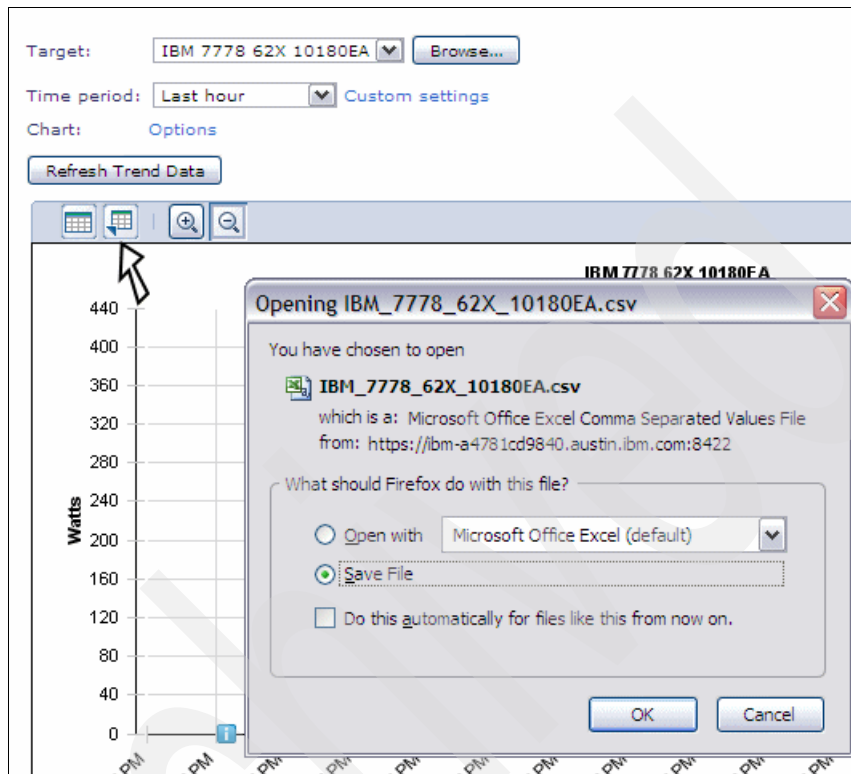


Figure 10-23 Export data

Energy cost calculator

Active Energy Manager has a calculator that can help determine the cost of energy for the monitored resource. Click **Actions** and then, as shown in Figure 10-24 on page 414, select **Energy** → **Energy Cost Calculator**.

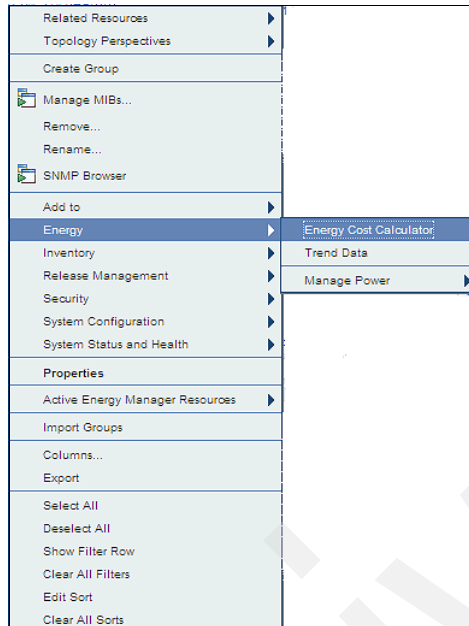


Figure 10-24 Energy calculator option

Set the values for the cost of energy by clicking the cost properties link. Set the values for energy cost, currency type and other values, as shown in Figure 10-25. Click **OK** to save the properties.

 A screenshot of a dialog box titled 'Edit Properties'. The dialog box contains the following fields and values:

- Active Energy Manager support level: Full
- Last time metered: Friday, April 17, 2009 2:46:27 PM CDT
- Nameplate power: 8,800 watts
- Average input power: 1,942 watts
- Average output power: 1,730 watts
- Ambient temperature: 23 Celsius
- Exhaust temperature: 35 Celsius
- Energy price: 0.12
- Currency type: USD (\$) (dropdown menu)
- Cooling rate multiplier: 1.5
- Metering active: True (dropdown menu)
- Metering interval: 5 minutes

 At the bottom of the dialog box are two buttons: 'OK' and 'Cancel'.

Figure 10-25 Energy cost properties

Click the **Calculate Energy Cost** button to display the data, as shown in Figure 10-26.

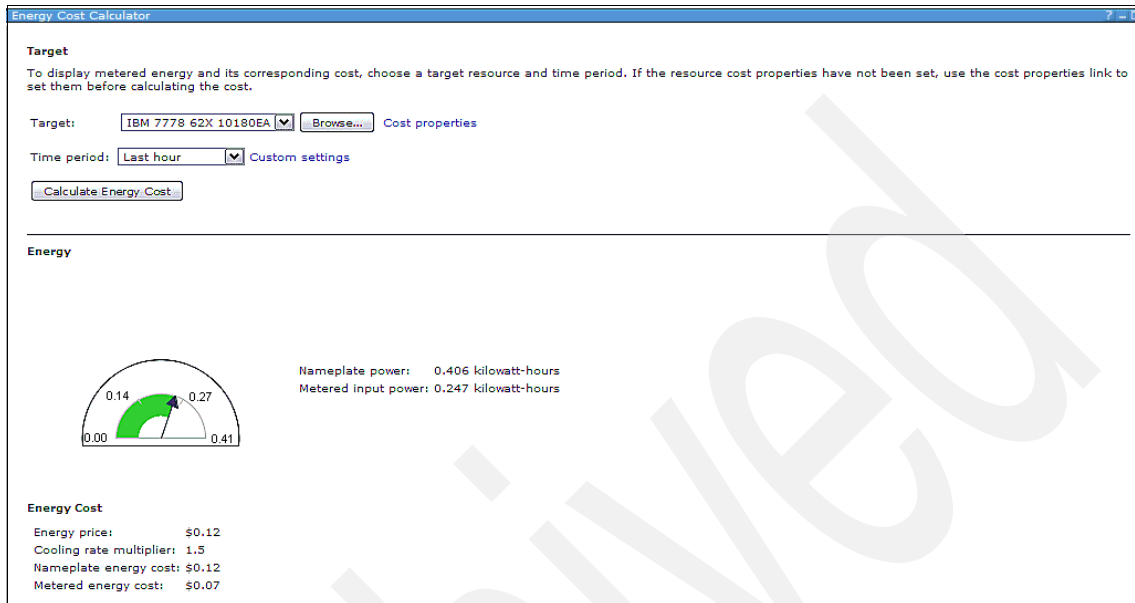


Figure 10-26 Calculated energy cost

10.2.4 Creating power policies

AEM supports the creation and application of power policies to manage energy across a group of systems. This feature allows you to create an energy policy and deploy it across a group or individual supported systems with minimal effort.

While IBM Systems Director is running, the power policies are enforced. This feature is a change from previous versions of Active Energy Manager. Policies can be applied to multiple systems and groups and can be applied immediately or scheduled. Within the policy management options, you can view what policies are in effect and to which systems or groups the policies are applied. Policies can also be removed using the same features as used to create a power policy.

In Active Energy Manager window, scroll down to the Manage section, which has an option to **Work with power policies**, shown in Figure 10-27 on page 416.




Figure 10-27 Work with power policies

Select **Work with power policies**.

The window shown in Figure 10-28 opens. From this window, you may view policies, launch a wizard to create policies, and edit and delete policies. You use this same interface to apply and remove policies after they have been created.

Before you can work with a target or group of targets, you must create a power policy. Click **Browse** to begin the target selection.

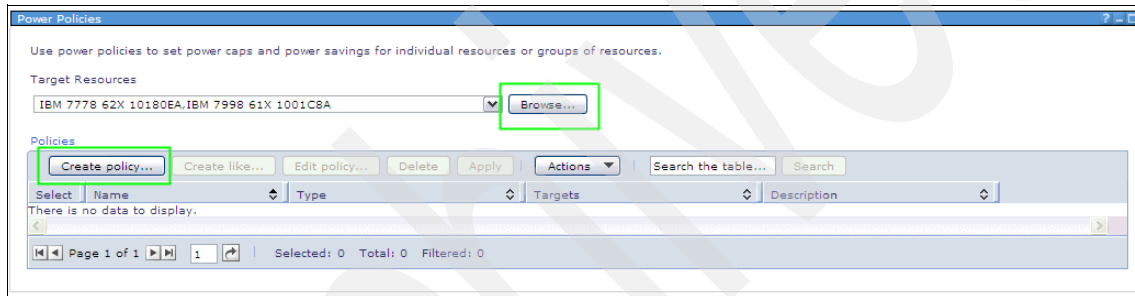


Figure 10-28 Select targets - browse

As shown in Figure 10-29 on page 417, click on the check box to select the intended target or targets. Then, click **Add** to add your selections.

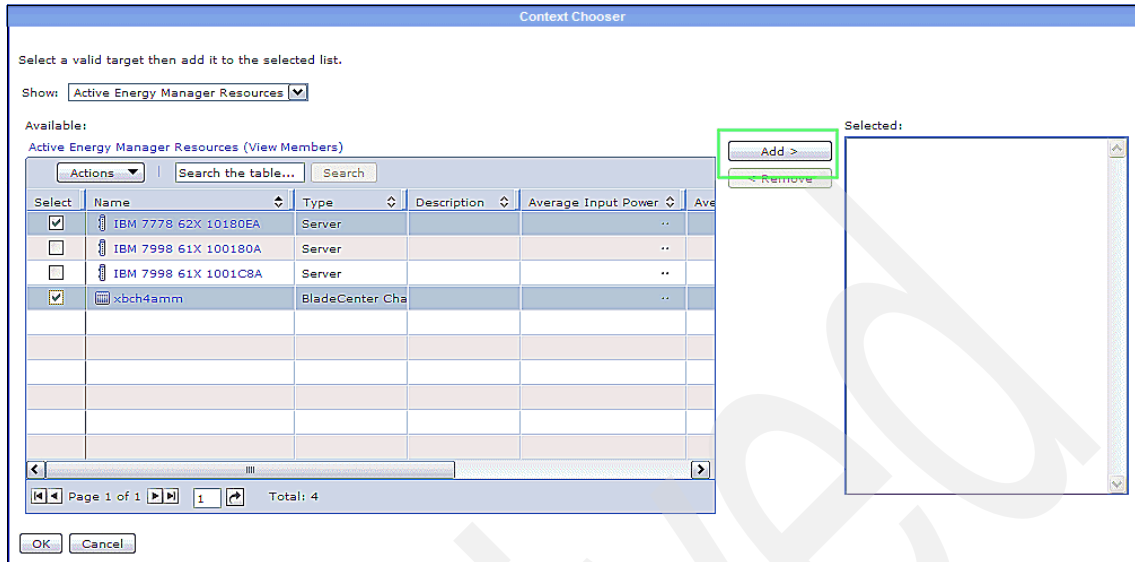


Figure 10-29 Select targets

If you want to create a group policy, first define that policy to a group by using the browse option. Group policies can only be applied to a group and not individual systems.

Note: It is assumed you have created a group prior to using AEM. If a group has not been specified, you can use IBM Systems Director and create a group by navigating resources and using the Create Group wizard. After the group has been created you can add members to the group.

An example of a group selection is provided in Figure 10-30 on page 418. To access groups, use the **Show** pull-down menu and select **Groups**.

Otherwise, to show system resources, select **Active Energy Manager Resources**, as shown in Figure 10-29.

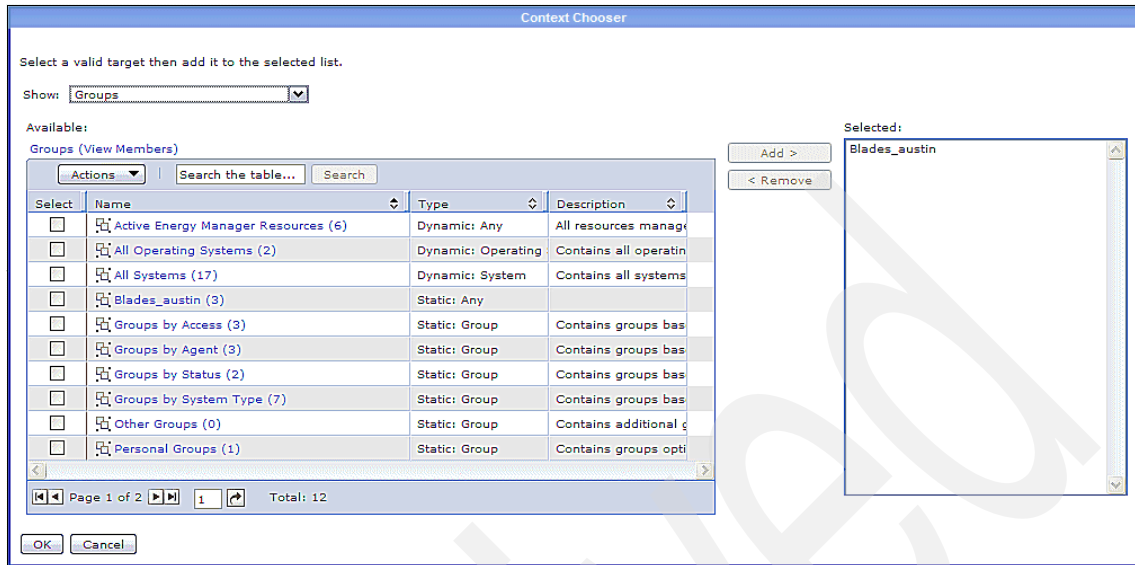


Figure 10-30 Group Select

After your targets are added to the Selected box, as Figure 10-31 indicates, click **OK** to complete your target selection.




Figure 10-31 Selected targets added

After the targets have been defined, you may begin to create a power policy by clicking **Create Policy** as shown in Figure 10-28 on page 416.

A wizard opens and helps you select the options for your policy. Figure 10-32 shows an example of the wizard welcome window and a description of each policy type.

Three policy types can be created: Group Power Capping, System Power Capping, and System Power Savings. Within the policy, you may select to turn on or turn off the feature.

For example, you may create a policy that turns on System Power Savings. Use that policy to turn on power savings for a blade server for use over a weekend or off-shift hours.

Then, create another policy that turns off power savings for the same blade server. Use this new policy to return the blade server to full power mode for week-time or on-shift usage.

To progress through the wizard, select your options, then click **Next**.

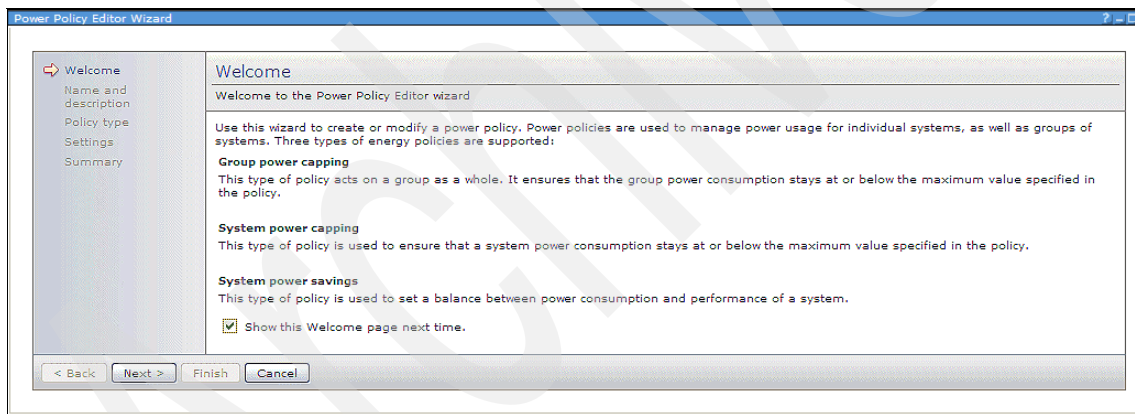


Figure 10-32 Power policy wizard welcome

In the next window, provide a name and description for the policy you are creating. Figure 10-33 on page 420 provides an example of this window. The **Name** field is required, the description field is not required however, it is a good idea to describe what the policy is used for in the description field. Click **Next** to continue.

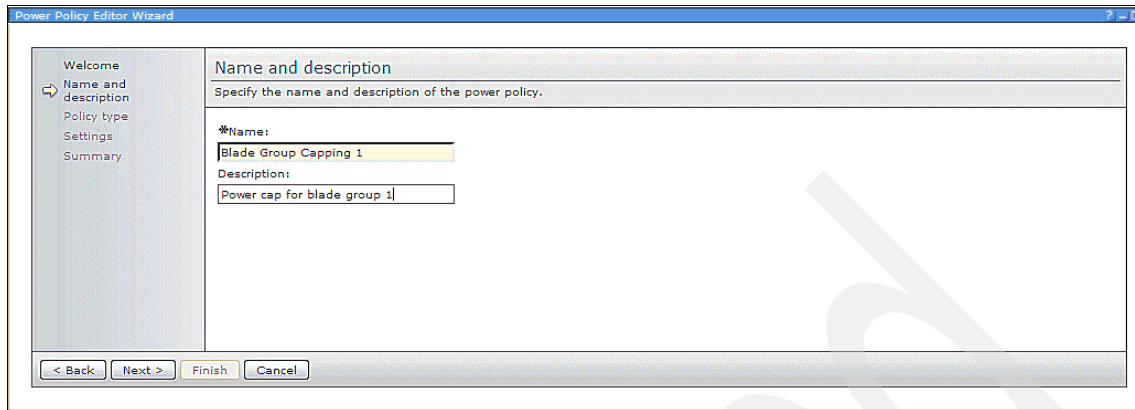


Figure 10-33 Policy name and description

In this next window, you may set the type of power policy by selecting one of the options available. The policy type window shows the three types of policies available and also provides a short description of what the policy can do and what targets it can be used on. Figure 10-34 shows an example of the policy type options.

For our example, we are creating a power policy for a group that will provide *group power capping*.

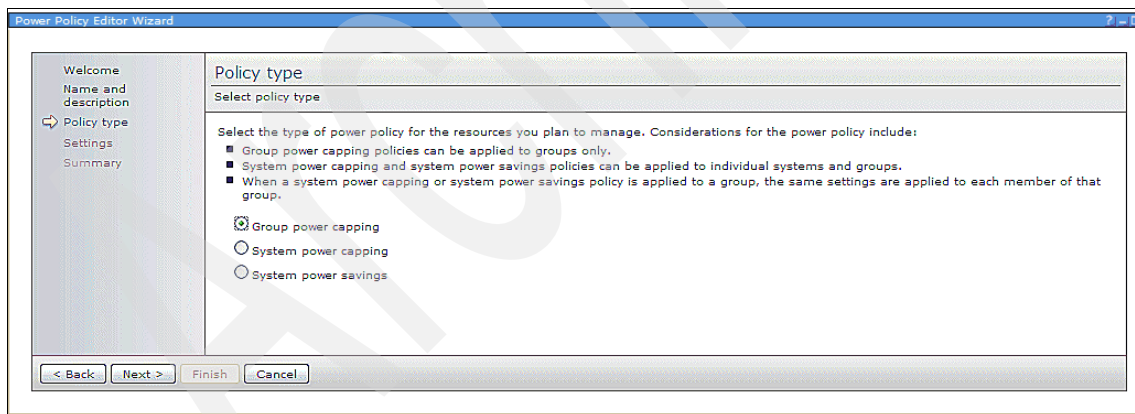


Figure 10-34 Power policy type

As shown in Figure 10-35 on page 421, select the group power capping settings by selecting either the value in watts or use the pull-down menu to change the value to a percentage. In the Power cap value text box, set the value at which you want to cap the group. Click **Next** to continue.

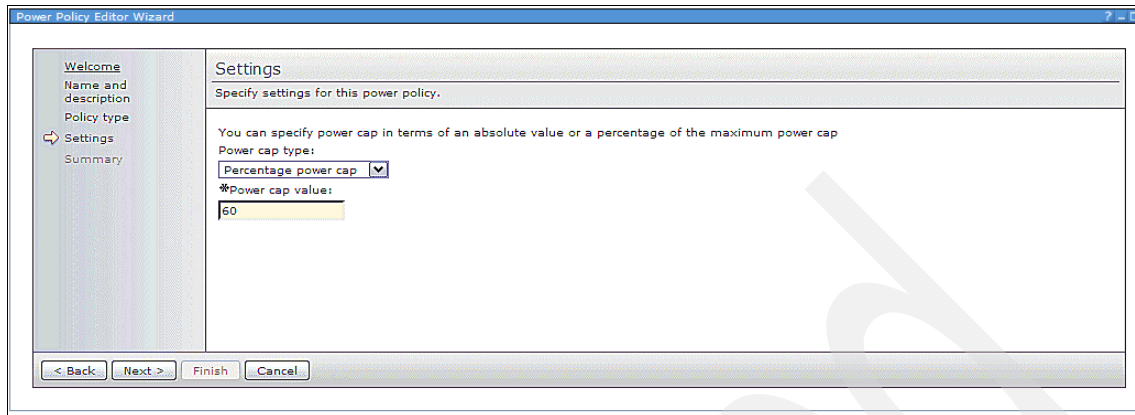


Figure 10-35 Power policy settings

The final window of the wizard provides a summary of your selections. Verify the selections are correct and click **Finish** to complete the policy creation. Figure 10-36 shows an example of the summary panel.

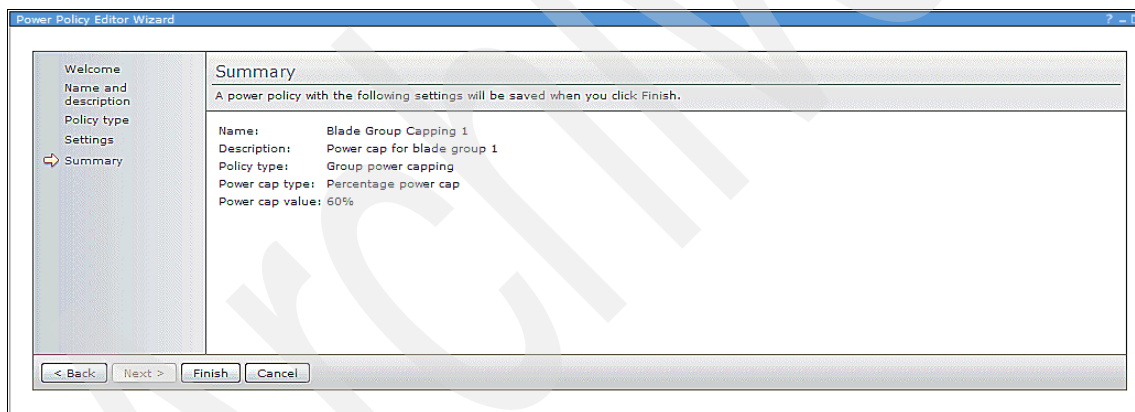


Figure 10-36 Power policy summary

Now that the policy has been created, it can be selected for action. Figure 10-37 on page 422 shows the policy we created with the wizard in the last few screens as well as several other policies we created available for actions.

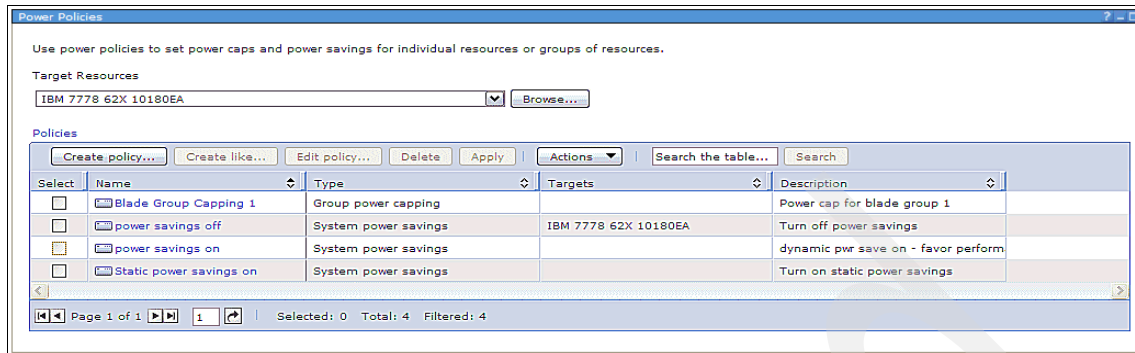


Figure 10-37 Available power policies

Now that power policies are created, we can select a target system or group of systems to apply the power policy to. Earlier, we selected a group of targets using the browse and add features. To apply a power policy to our selected targets, ensure the targets are listed in the **Target resources** selection.

Then, select the power policy and click **Apply**. Figure 10-38 shows an example of the selected targets and the power policy to apply.

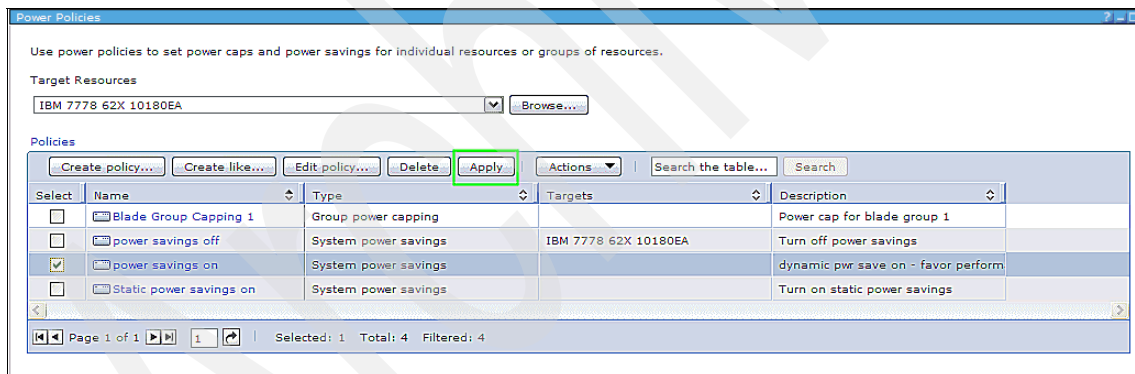


Figure 10-38 Apply power policy

In the next panel, you select when to apply the policy. Figure 10-39 on page 423 shows that the Apply Policy power savings on should be run now.

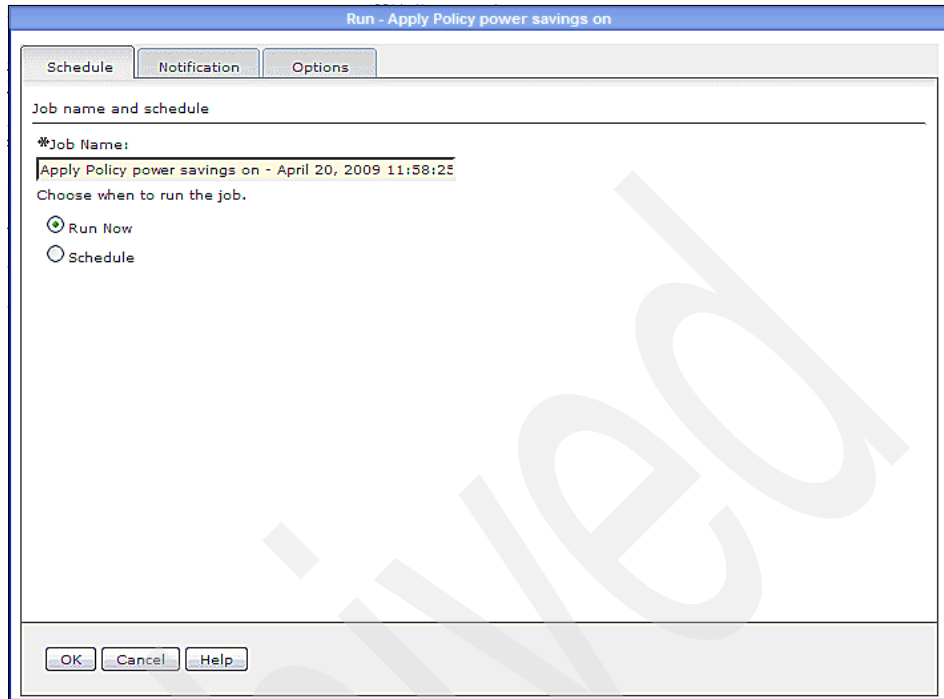


Figure 10-39 Run now - policy apply option

You also have the option of scheduling when to run the power policy. This feature is used to apply a power policy unattended. This would be useful for setting up automatic application of a policy to turn on or turn off power savings for example. Figure 10-40 on page 424 provides an example of the settings to schedule a policy.

Run - Apply Policy power savings on

Schedule Notification Options

Job name and schedule

*Job Name:
Apply Policy power savings on - April 20, 2009 11:58:25

Choose when to run the job.

☐ Run Now
☒ Schedule

*Time:
11:58 AM

*Date:
Apr 20, 2009

Repeat Options

Frequency:

- Once
- Hourly
- Daily
- Weekly
- Monthly
- Yearly
- Custom

OK Cancel Help

Figure 10-40 Policy schedule options

You may also set the system to send you an e-mail when the policy is applied. Select the **Notification** tab and modify the settings for your correct contact information. Figure 10-41 on page 425 shows an example of the Notification tab.

Run - Apply Policy power savings off

Schedule Notification Options

Receive an e-mail notification with the progress of this job.

☒ Notify when this job begins.

☒ Notify when this job is completed successfully.

☐ Notify when this job fails:

☒ Any Error --

☐ Percentage targets with errors: 0

☐ Number of targets with errors: 0

*E-mail address: myemail.email.com

*E-mail server name: mysmtp.server

*E-mail server port number: 80

OK Cancel Help

Figure 10-41 Notification tab

The Options tab enables you to set a time base to use: either management server or local system time. You also have the option to allow the policy action to fail if the system is not available or run when the system becomes available. Figure 10-42 on page 426 shows an example of these settings.

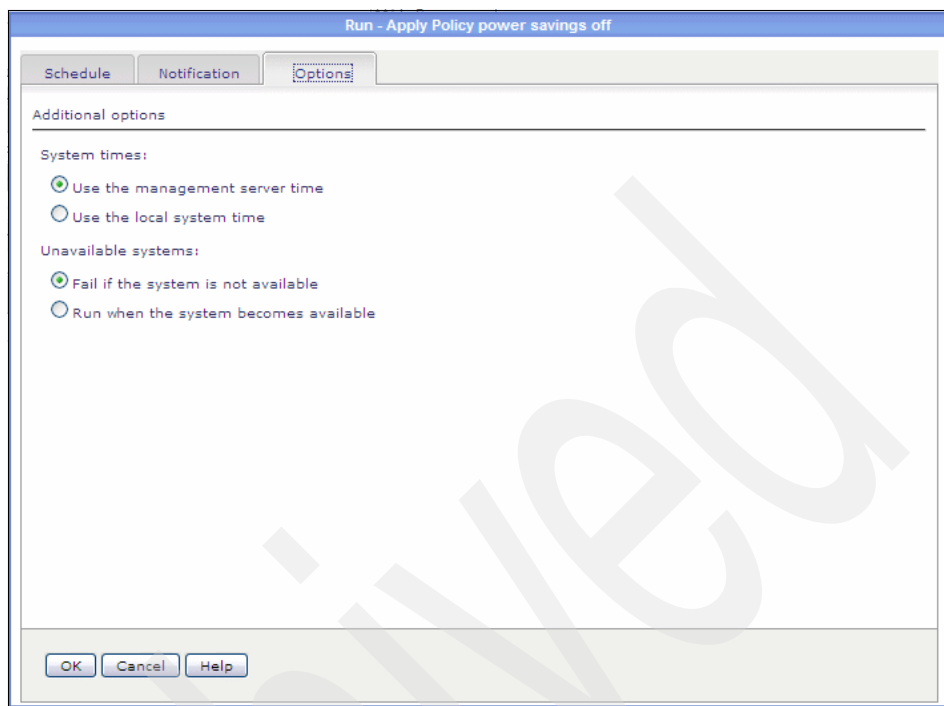


Figure 10-42 Policy options tab

10.2.5 Conclusion

Active Energy Manager can also be controlled through the command-line interface (CLI). Many of the CLIs are useful to the IBM BladeCenter management.

Information about the systems management CLI (smcli) is located at:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/director.cli_6.1/fqm0_r_cli_smcli.html

Information about IBM Systems Director CLI can be found at:

http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/director.cli_6.1/fqm0_r_cli_smcli.html

For more information about IBM Systems Director or Active Energy Manager, be sure to reference the links at the beginning of this chapter.

Performing Live Partition Mobility

This chapter discusses the requirements and configuration procedures to perform Live Partition Mobility between a IBM BladeCenter JS23 and JS43 blades.

This chapter contains the following topics:

- ▶ “Requirements” on page 428
- ▶ “Preparation” on page 432
- ▶ “Migrating the LPAR” on page 442

Additional information about Live Partition Mobility architecture, mechanisms and advanced topics can be found in *PowerVM Live Partition Mobility*, SG24-7460.

For the most current information, documentation, known problems, workarounds and more, go to the Live Partition Mobility support Web site:

<http://www14.software.ibm.com/webapp/set2/sas/f/pm/home.html>

11.1 Requirements

Partition mobility places certain demands on hardware, software, network, and storage configurations. These considerations have to be reviewed early in the setup of an IBM BladeCenter JS23 or JS43 to avoid reconfiguration and rework.

11.1.1 Hardware

The IBM BladeCenter JS23 or JS43 requires a Fibre Channel HBA expansion card for SAN connectivity. All storage assigned to a logical partition must have external backing devices that are visible to both the local and remote Virtual I/O Server (VIOS) systems.

11.1.2 Firmware

In a new blade environment, a best practice is to have the most current system firmware installed. For the hardware configuration used in the writing of this book, the source and target blades had two different levels, the current and previous levels of firmware. This configuration was chosen to represent a customer environment where Live Partition Mobility might be used for system maintenance reasons. IBM BladeCenter JS23 and JS43 firmware updates can be downloaded from:

<http://www.ibm.com>

Select the **Support & downloads**, select **More**, and then choose the product type of **BladeCenter** from the drop-down box. You are directed to the Support for IBM BladeCenter page. Select **BladeCenterJS23** or **BladeCenterJS43** from the Product family drop-down box, then click **Go** to display a list of related blade and BladeCenter updates. Blade system firmware will have a prerequisite AMM firmware that will be identified in the firmware readme file.

11.1.3 VIOS version

Similarly to system firmware, the VIOS version and fix-pack level should be the most recent. To display the current code level from the GUI, click **Updates** from the navigation area. The Management Partition Updates panel opens and the code level is listed, as shown in Figure 11-1 on page 429. If the workstation that you are using has Internet access, the link displayed on the Management Partition Updates panel takes you to a download site for newer updates and fixes, if available. The link to the Virtual I/O Server is also available at:

<http://techsupport.services.ibm.com/server/vios/download>

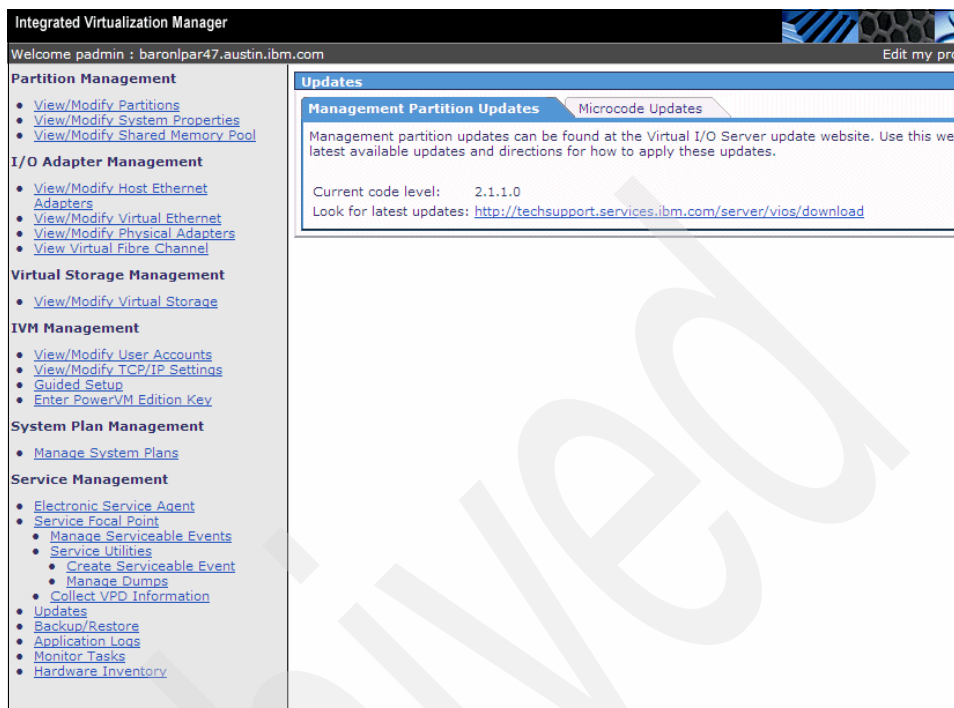


Figure 11-1 Management Partition Updates view

From the CLI use the **ioslevel** command to display the VIOS version and fix-pack level, as shown in Example 11-1. In this example, the VIOS version is 2.1.1.0 and has not had any fix packs installed.

Example 11-1 **ioslevel** command

```
$ ioslevel
2.1.1.0
```

An example of a previous release with a fix pack installed is shown in Example 11-2.

Example 11-2 **ioslevel** command showing fix pack installed

```
$ ioslevel
2.1.0.10-FP-20.1
```

11.1.4 PowerVM Enterprise

PowerVM Enterprise Edition is an optional feature on an IBM BladeCenter JS23 or JS43 and is required to enable Partition Mobility. To determine if this capability is available, use the **lssyscfg** command. Example 11-3 shows the **lssyscfg** returning a value of 1 to indicate active or Live Partition Mobility capability.

Example 11-3 lssyscfg command

```
$ lssyscfg -r sys -F active_lpar_mobility_capable  
1
```

If Live Partition Mobility is not enabled and the feature was purchased with the blade, the activation key should be available on the IBM Capacity on Demand (CoD) Web site:

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

Enter the system type and serial number on the CoD site and submit. A list of available activation codes or keys with a type and description will be displayed. If PowerVM Enterprise Edition was not purchased with the IBM BladeCenter JS23 or JS43, it can be upgraded through the miscellaneous equipment specification (MES) process.

Entering an enablement key through the IVM GUI

The PowerVM key is entered from the GUI by clicking the **Enter PowerVM Edition Key** link in the navigation area. In the window that opens, shown in Figure 11-2 on page 431, enter the PowerVM Edition key number and click **Apply**. When PowerVM Enterprise is enabled, a Mobility section is added to the More Tasks drop-down box on the View/Modify Partitions view.

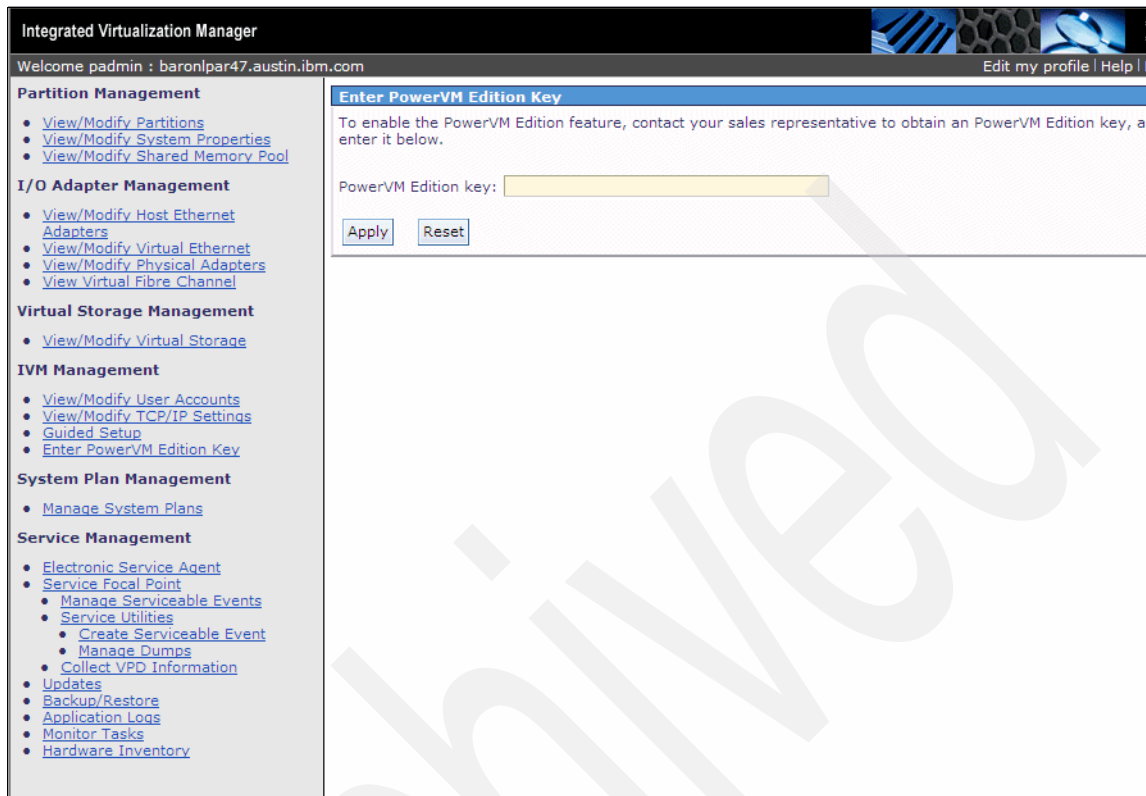


Figure 11-2 PowerVM Enterprise key entry

11.1.5 LPAR OS versions

The running operating system in the mobile partition must be AIX or Linux. The supported operating systems for Live Partition Mobility are:

- ▶ AIX 5L V5.3 with 5300-07 Technology Level or later
- ▶ AIX V6.1 or later
- ▶ Red Hat Enterprise Linux Version 5.1 or later
- ▶ SUSE Linux Enterprise Services 10 (SLES 10) Service Pack 1 or later

11.2 Preparation

This section describes the settings and configurations that must be verified and possibly changed to prepare the local and remote VIOS systems and partitions for partition mobility.

11.2.1 VIOS source and target requirements

We start with VIOS source and target considerations. We look at the memory region size, and storage and hdisk reserve policy.

Memory region size

The memory region size is the smallest block of memory that can be assigned to or changed in an LPAR. To review the current setting, select **View/Modify System Properties** in the navigation area and then select the **Memory** tab, as shown in Figure 11-3 on page 433. The default is dependent on the amount of system memory installed. The value is set to 32 MB (automatic). The memory region size must be the same for the source and target VIOS systems. If the value is changed, a VIOS power down and restart is required (not just a basic shutdown reboot).

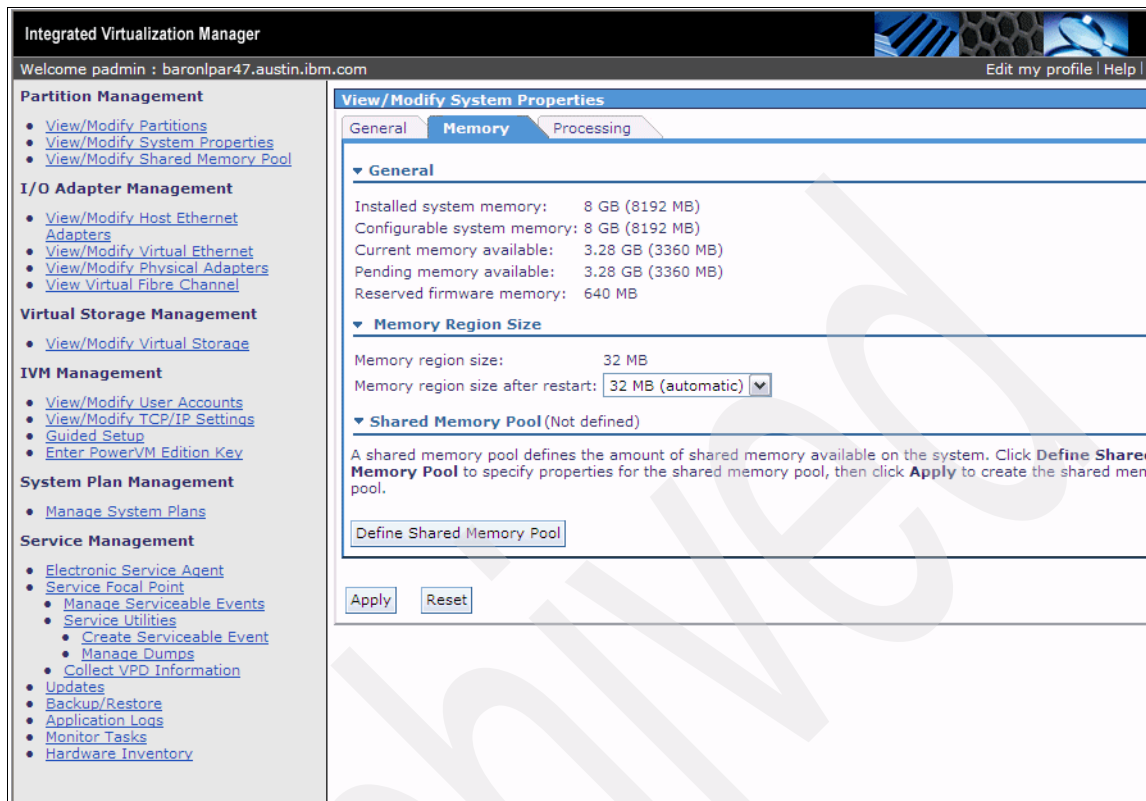


Figure 11-3 Memory region size

Storage and hdisk reserve policy

Only physical volumes (LUNs) visible to the VIOS as an hdisk assigned to an LPAR can be used in mobile partitions. The same physical volumes must also be visible to both the local and remote VIOS systems. The reserve policy of the hdisk must be changed from the default single_path to no_reserve. The reserve policy is changed on an hdisk from both VIOS systems. To review the current policy setting use the following command:

```
lsdev -dev hdisk# -attr | grep reserve
```

The **chdev** command is used to change the attribute of the hdisk. Example 11-4 on page 434 shows hdisk1 with the default reserve policy of single_path, the **chdev** command being issued, and the new reserve policy of no_reserve.

Example 11-4 Changing a hdisk reserve policy

```
$ lsdev -dev hdisk1 -attr | grep reserve
reserve_policy single_path          Reserve Policy
True
```

```
$ chdev -dev hdisk1 -attr reserve_policy=no_reserve
hdisk1 changed
```

```
$ lsdev -dev hdisk1 -attr | grep reserve
reserve_policy no_reserve          Reserve Policy
True
```

Note: The reserve policy cannot be changed on the source VIOS when the disks are assigned to an LPAR. The command fails with the following message:

Some error messages may contain invalid information for the Virtual I/O Server environment.

```
Method error (/etc/methods/chgfcpararray):
    0514-062 Cannot perform the requested function because the
        specified device is busy.
```

If the reserve policy has not been changed, it can quickly be identified by looking on the target VIOS. From the target system IVM GUI navigation area, click the **View/Modify Virtual Storage** link. In the View/Modify Virtual Storage view click the **Physical Volumes** tab and review the size column. If any of the sizes are indicated as Unknown, the reserve policy likely has not been changed and the validation process will fail. Figure 11-4 on page 435 shows hdisk0 and hdisk7 in the Unknown size condition.

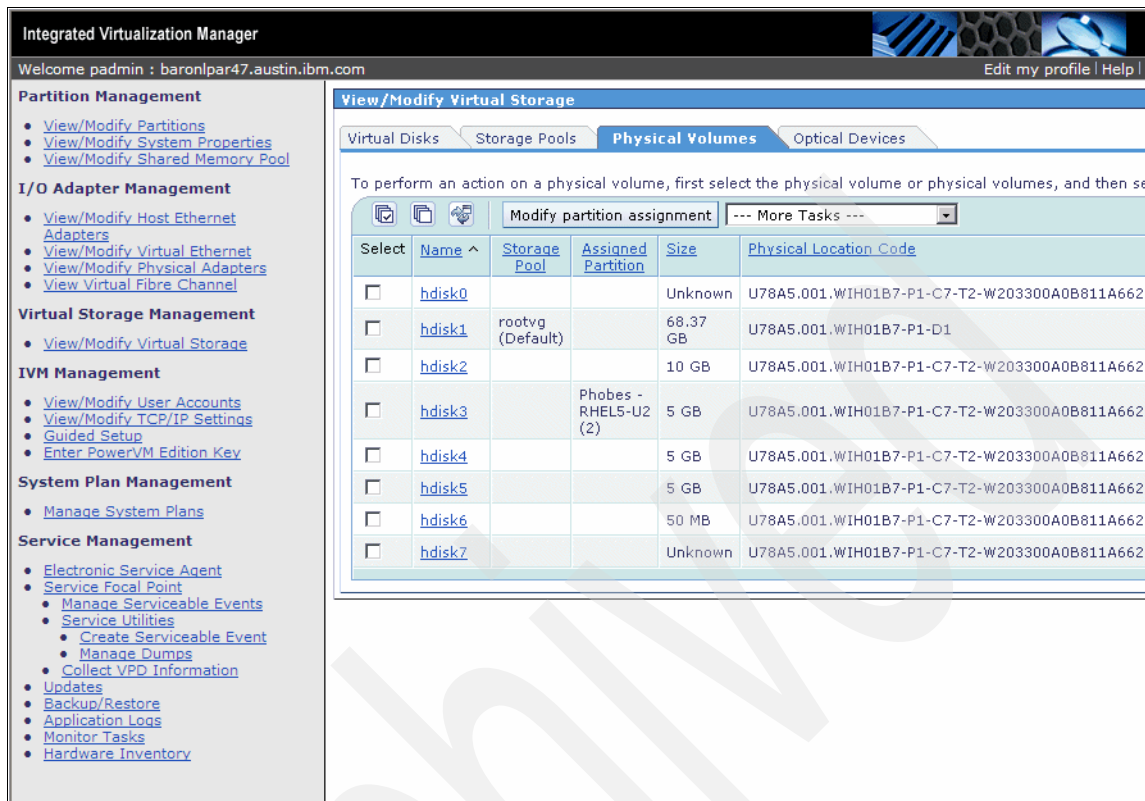


Figure 11-4 The hidsk reserve policy not set correctly

When the validation process is run, an error message similar to the message in Figure 11-5 on page 436 is displayed.

This problem can be resolved by performing the following steps:

1. Shut down the mobile LPAR on the local VIOS if running.
2. Modify the mobile LPAR hdisk assignments on the local VIOS to *none*.
3. Use the `chdev` command to change the hdisks reserve policy to *no_reserve*.
4. Modify the mobile LPAR hdisk assignments to the original assignments.
5. Refresh the View/Modify Virtual Storage view on the remote VIOS.

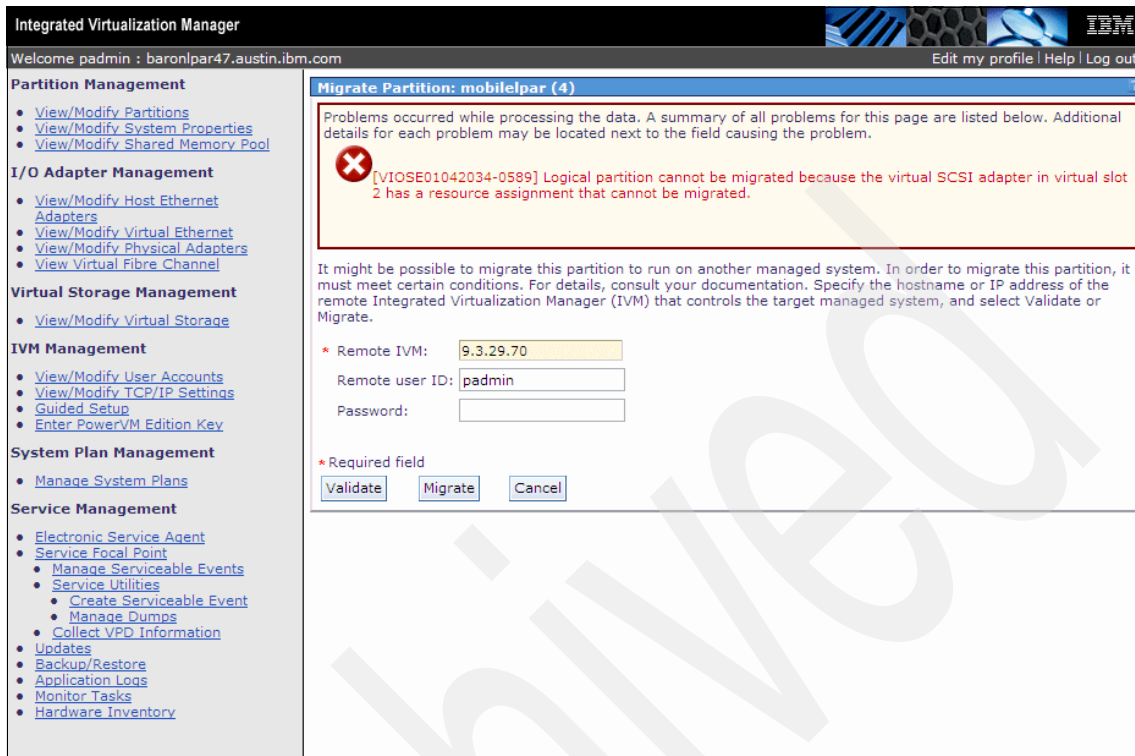


Figure 11-5 Partition Migration validation error message for target storage

11.2.2 Networking

The mobile LPAR external network communication must be through a Shared Ethernet Adapter (SEA). The use of logical ports on a Host Ethernet Adapter (HEA) or physical adapters assigned to the LPAR cannot be used and must be removed if assigned. SEA adapter creation is covered in 4.5.2, “Virtual Ethernet Adapters and SEA” on page 97.

The Resource Monitoring and Control (RMC) daemon must be active on the mobile partition or the validate and migration process will fail. The IVM GUI can be used to verify the status as described in 4.7.11, “Partition properties changes and DLPAR operations” on page 151.

The CLI **lssyscfg** command can also be used to determine the RMC status as shown in Example 11-5.

Example 11-5 lssyscfg command to determine the RMC status

```
$ lssyscfg -r lpar -F name,rmc_state
VIOS-Neptune,active
Phobos - RHEL5-U2,inactive
Mars - AIX 6.1,active
```

Note: Linux partitions must have the Dynamic Reconfiguration Tools package for HMC- or IVM-managed servers installed from the Service and Productivity tools Web site at:

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

Service and Productivity tools are discussed in Appendix D, “Service and productivity tools for Linux” on page 539.

11.2.3 Partition requirements

In this section, we discuss the partition requirements for performing Live Partition Mobility.

Dedicated or shared memory

For a partition that is using dedicated memory, the target VIO Server must have adequate available memory to contain the moving partition.

Note: In cases where the available memory on the target VIO Server and the dedicated memory configured in the mobile partition match, the validation process will fail. The creation of an additional LPAR on the target VIO Server causes firmware to reserve additional memory.

A partition that is using shared memory on the source VIO Server must have a shared memory pool defined on the target VIO Server with adequate available resources for the mobile partition.

Processor compatibility mode

The processor compatibility mode has two settings: the current value and the preferred value for the logical partition. They can be displayed from the Partition Properties Processor tab.

The current value indicates the negotiated compatibility mode for the logical partition. This is the value that the logical partition is currently using.

The following values are possible for the current setting of the processor compatibility mode:

► **POWER6**

This mode is possible for both POWER6 and POWER6 + processor-based servers. This mode indicates that the operating environment for the partition is using all the standard capabilities of the POWER6 processor.

► **POWER6+™**

This mode is possible for POWER6 + processor based servers. This mode indicates that the operating environment for the partition is using all the standard capabilities of the POWER6 + processor.

► **POWER6 Enhanced**

This mode is possible for POWER6 processor based servers. This mode indicates that the operating environment for the partition is using all the standard features of the POWER6 processor and also supports using additional floating-point instructions in the applications that use the POWER6 processor.

► **POWER6+ Enhanced**

This mode is possible for POWER6 + processor based servers. This mode indicates that the operating environment for the partition is using all the standard features of the POWER6 processor and also supports using additional floating-point instructions in the applications that use the POWER6+ processor.

After you select the preferred processor compatibility mode, you must shut down and restart the logical partition so that the hypervisor can check the preferred mode and the operating environment to change the value for the current processor compatibility mode.

For active logical partition migration, the destination server must support both preferred and current processor compatibility modes of the mobile partition. For inactive migrations, the destination server must support only the preferred processor compatibility mode of the mobile partition.

The CPU type on JS23/JS43 is Power6+ and has different compatibility modes that can be selected as shown in Figure 11-6 on page 439.

Partition Properties: mobilelpar (4)

General Memory **Processing** Ethernet Storage Optical/Tape Devices Physical Adapters

Modify the settings by changing the pending values. Before making any changes, use the Retrieve Capability on the General tab to determine whether the partition provides support for making these changes while the active.

Processing Units			Virtual Processors		
Property	Current	Pending	Property	Current	Pending
Minimum	0.1	<input type="text" value="0.1"/>	Minimum	1	<input type="text" value="1"/>
Assigned	0.8	<input type="text" value="0.8"/>	Assigned	8	<input type="text" value="8"/>
Maximum	8	<input type="text" value="8.0"/>	Maximum	8	<input type="text" value="8"/>

General

Property	Current	Pending
Uncapped weight	Medium - 128	<input type="text" value="Medium - 128"/> ▼

Processor compatibility mode:

Current value: POWER6+

Preferred value: ▼

- Default
- POWER6
- POWER6+
- POWER6+ Enhanced

Figure 11-6 Processor compatibility mode on JS23/JS43

JS12 and JS22 blades used POWER6 technology and can be configured for the processor compatibility modes as shown in Figure 11-7 on page 440.

Partition Properties: NIM (2)

General Memory **Processing** Ethernet Storage Optical/Tape Devices Physical Address

Modify the settings by changing the pending values. Before making any changes, use the Retrieve Capability on the General tab to determine whether the partition provides support for making these changes while the active.

Processing Units			Virtual Processors		
Property	Current	Pending	Property	Current	Pending
Minimum	0.1	<input type="text" value="0.1"/>	Minimum	1	<input type="text" value="1"/>
Assigned	0.5	<input type="text" value="0.5"/>	Assigned	1	<input type="text" value="1"/>
Maximum	4	<input type="text" value="4.0"/>	Maximum	6	<input type="text" value="6"/>

General

Property	Current	Pending
Uncapped weight	Low - 64	<input type="text" value="Low - 64"/>

Processor compatibility mode:

Current value: POWER6

Preferred value:

Default
POWER6 Enhanced
POWER6

Figure 11-7 Processor compatibility mode on JS12/JS22

The requirement is that the source and target blades have the ability to match processor compatibility modes. Currently for POWER6 blades the only common processor compatibility mode is POWER6. An LPAR running in POWER6 mode on a JS12 could migrate to a JS23 or JS43. If the JS12 LPAR was running in POWER6 Enhanced migration to a JS23 or JS43 would not be possible without a mode change first on the JS12 to POWER6 mode.

To perform Live Partition Mobility from IBM BladeCenter JS23/JS43 blades to IBM BladeCenter JS12/JS22 the processor compatibility mode on JS23/JS43 must be changed to POWER6 as shown in Figure 11-8 on page 441.

Partition Properties: mobilelpar (4)

General Memory **Processing** Ethernet Storage Optical/Tape Devices Physical Adapter

Modify the settings by changing the pending values. Before making any changes, use the Retrieve Capabilities button on the General tab to determine whether the partition provides support for making these changes while the partition is active.

Processing Units			Virtual Processors		
Property	Current	Pending	Property	Current	Pending
Minimum	0.1	<input type="text" value="0.1"/>	Minimum	1	<input type="text" value="1"/>
Assigned	0.8	<input type="text" value="0.8"/>	Assigned	8	<input type="text" value="8"/>
Maximum	8	<input type="text" value="8.0"/>	Maximum	8	<input type="text" value="8"/>

General

Property	Current	Pending
Uncapped weight	Medium - 128	<input type="text" value="Medium - 128"/>

Processor compatibility mode:

Current value: POWER6

Preferred value:

Figure 11-8 Change the processor compatibility mode on JS23/JS43

Virtual optical devices

All virtual optical devices must be removed from the mobile partition before a successful validation and migration can occur. The example shown in Figure 11-9 on page 442 indicates that the virtual device vtop0 is still assigned to the mobile partition. The device can be removed by unchecking the box and clicking **OK**.

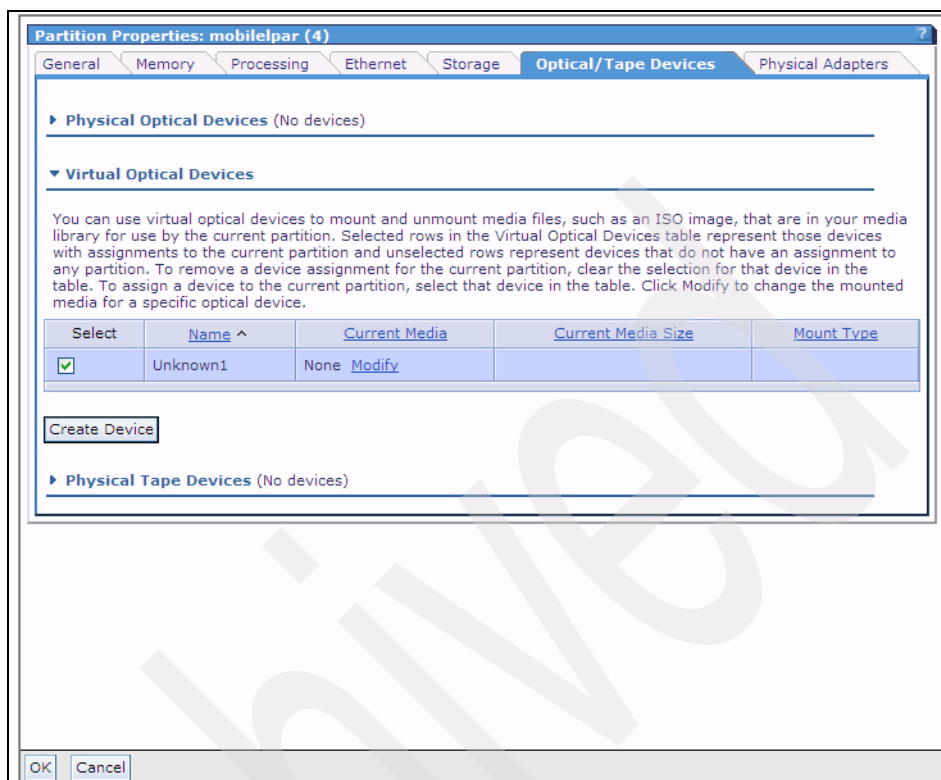


Figure 11-9 Virtual optical device to be removed

11.3 Migrating the LPAR

This section describes how to use the IVM GUI and CLI to validate, migrate, and check status on mobile LPAR.

11.3.1 Using the IVM GUI

Let us first see how we can perform an LPAR migration with IVM.

Validate

The migration process is started by first selecting **View/Modify Partitions** from the navigation area. With this view open, use the check box to select a mobile partition, and select **Migrate** from the More Tasks drop-down box as shown in Figure 11-10 on page 443.

Integrated Virtualization Manager

Welcome padmin : baronlpar47.austin.ibm.com

Edit my profile

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](#)

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:

8 GB

Total processing units:

8

Memory available:

3.28 GB

Processing units available:

8

Reserved firmware memory:

640 MB

Processor pool utilization:

0.02

System attention LED:

Inactive

Partition Details

Create Partition...

Activate

Shutdown

--- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Proc	Pro
<input type="checkbox"/>	1	js43-vios	Running	17.74 Hours	1.5 GB	8	
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1	
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1	
<input checked="" type="checkbox"/>	4	mobilelpar	Running	10.7 Minutes	608 MB	8	

--- More Tasks ---

Open terminal window

Delete

Create based on Operator panel service functions

Reference Codes

--- Mobility ---

Migrate

Status

Properties

Figure 11-10 Partition Migrate option

The Migrate Partition view opens with the mobile partition name appended to the window name. Enter the remote or target IVM-controlled system IP address, remote user ID and password as shown in Figure 11-11 on page 444. Click **Validate** to start the validation process.

Note: The Partition Migration view requests the Remote IVM or HMC IP address. At the time of this publication, IVM to HMC migrations are not supported.

Integrated Virtualization Manager
Welcome padmin : baronlpar47.austin.ibm.com
Edit my profile

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](#)

Migrate Partition: mobilelpar (4)

It might be possible to migrate this partition to run on another managed system. In order to migrate, the target managed system 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 the target managed system.

* Remote IVM:

Remote user ID:

Password:

* Required field

Figure 11-11 Partition Mobility validation

At the end of the successful validation process, the Migrate Partition window is updated, similar to Figure 11-12.

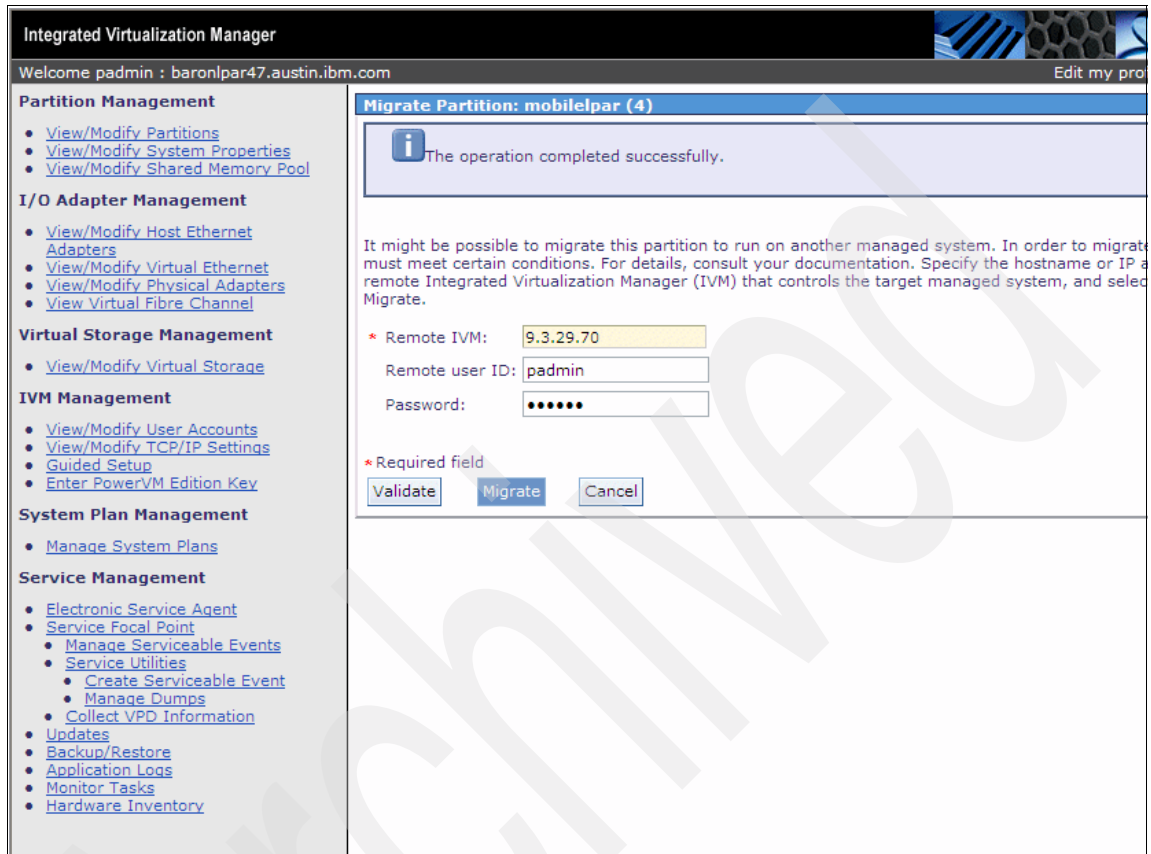


Figure 11-12 Partition Migration validation success

Figure 11-13 shows the results of the validation process, which discovered a problem that prevents a migration. This error message was generated because of a virtual SCSI assignment that could not be migrated. In this example, the problem was the result of a virtual optical device that had an assignment to the mobile partition.

Another example is shown in Figure 11-4 on page 435, where the validation process could not find the required storage on the remote system. The validation process must complete without errors before a partition migration can be attempted.

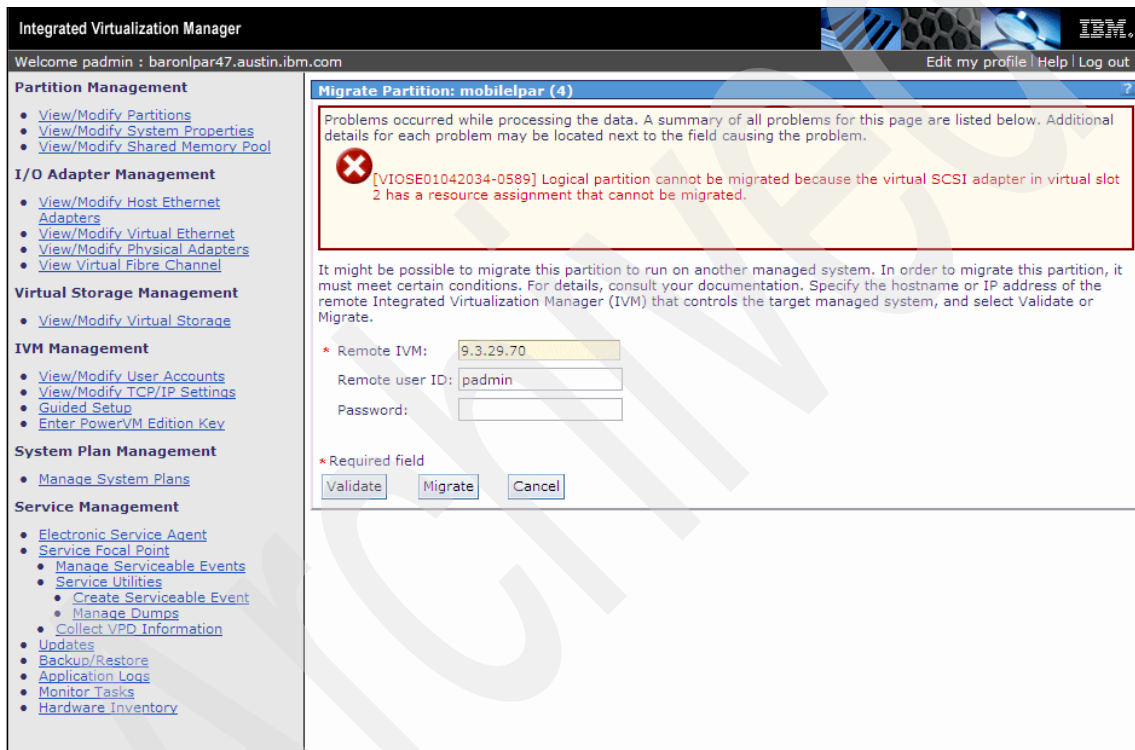


Figure 11-13 Failed validation because of improper virtual SCSI device assignment

Migrate

With a successful completion of the validation process the migrate step can be started. Click **Migrate** to begin the migration process. As part of the migration process, a validation is run again and, at the end of this step, a Migrate Status view is displayed, as shown in Figure 11-14.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. The title bar reads "Integrated Virtualization Manager" and "Welcome padmin : baronlpar47.austin.ibm.com". The sidebar on the left contains the following sections:

- 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](#)

The main panel is titled "Migrate Status" and contains the following text: "The following partitions are currently migrating. You may stop this operation or continue to migrate. If the migration status shows an error, you should select Recover, which will attempt to complete the migration if appropriate." Below this text is a table titled "Partitions Migrating From This System".

Select	Partition ^	Migration Status	Percent Complete	Remote Platform Manager
<input type="checkbox"/>	mobilelpar (4)	Migration Starting	45%	9.3.29.70

At the bottom of the table are "OK" and "Cancel" buttons. Above the table are icons for a folder, a document, and a refresh, along with "Stop" and "Recover" buttons.

Figure 11-14 Migrate Status view

The Migrate Status view can be accessed directly from the View/Modify Partitions window. Check the mobile partition box, then select **Status** under the Mobility section of the More Tasks drop-down box as shown in Figure 11-15. Also note in this same figure that the state of the mobile partition has changed from Running to Migrating - Running.

Integrated Virtualization Manager

Welcome padmin : baronlpar47.austin.ibm.com

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

View/Modify Partitions

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

System Overview

Total system memory:	8 GB	Total processing units:	
Memory available:	3.28 GB	Processing units available:	
Reserved firmware memory:	640 MB	Processor pool utilization:	
System attention LED:	Inactive		

Partition Details

Select	ID ^	Name	State	Uptime	Memory	Pro	More Tasks
<input type="checkbox"/>	1	js43-vios	Running	18.07 Hours	1.5 GB	8	--- More Tasks --- Open terminal window Delete Create based on Operator panel service Reference Codes --- Mobility --- Migrate Status ----- Properties
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1	
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1	
<input type="checkbox"/>	6	mobilelpar	Migrating - Running	30.1 Minutes	608 MB	8	0.8

Figure 11-15 Starting the status view for a mobile partition

Figure 11-16 shows the View/Modify Partitions view on the remote IVM, indicating migration has started.

Note: The mobile partition retains the same LPAR ID number if available on the remote system, otherwise it is assigned the first available ID number.

Integrated Virtualization Manager

Welcome padmin : baronlpar16.austin.ibm.com

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

View/Modify Partitions

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

System Overview

Total system memory:	24 GB	Total processing units:
Memory available:	11.47 GB	Processing units available:
Reserved firmware memory:	704 MB	Processor pool utilization:
System attention LED:	Inactive	

Partition Details

Create Partition... Activate Shutdown --- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units
<input type="checkbox"/>	1	js23-vios	Running	3.11 Hours	2 GB	4	0.4
<input type="checkbox"/>	2	IBMI 2	Not Activated		256 MB	1	0.5
<input type="checkbox"/>	3	JS23AMSlpar3	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	4	JS23DMLpar4	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	5	JS23DPLpar5	Not Activated		1 GB	1	1.0
<input type="checkbox"/>	6	mobilelpar	Migrating - Running		608 MB	8	0.8

Figure 11-16 Remote IVM indicating migration in progress

At the end of the migration process the State of the mobile partition changes from Migrating - Running to Running as shown in Figure 11-17 on the formerly remote system. On the original local system, the mobile partition is removed from the View/Modify Partition view.

Integrated Virtualization Manager

Welcome padmin : baronlpar16.austin.ibm.com

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](#)

View/Modify Partitions

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

System Overview

Total system memory:	24 GB	Total processing units:
Memory available:	11.47 GB	Processing units available:
Reserved firmware memory:	704 MB	Processor pool utilization:
System attention LED:	Inactive	

Partition Details

☐ ☐ ☐ * Create Partition... Activate Shutdown --- More Tasks ---

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units
<input type="checkbox"/>	1	js23-vios	Running	3.14 Hours	2 GB	4	0.4
<input type="checkbox"/>	2	IBMI 2	Not Activated		256 MB	1	0.5
<input type="checkbox"/>	3	JS23AMSlpar3	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	4	JS23DMlpar4	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	5	JS23DPlpar5	Not Activated		1 GB	1	1.0
<input type="checkbox"/>	6	mobilelpar	Running	32.1 Minutes	608 MB	8	0.8

Figure 11-17 Partition migration complete to remote system

11.3.2 From the command line

The IVM **migr1par** command is used to validate and migrate the mobile partition from one IVM-managed system to another. Example 11-6 on page 451 shows the **validate** command and the result when a virtual optical device is still assigned to the mobile partition. A successful validation process will return to the **\$** prompt with no other messages.

Example 11-6 *migr1par* command validate option

```
$ migr1par -o v -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5  
[VIOSE01042034-0418] The partition cannot be migrated because the  
virtual SCSI server adapter has a resource assignment that cannot be  
migrated.
```

Note the following information about flags:

- The -o flag or operation has the following options:

s	Stop
m	Validate and migrate
r	Recover
v	Validate

- The -t flag in Example 11-6 specifies the remote managed system. The -t flag requires a system name and IP address.

Note: The system name is not the same as the host name. The system name can be obtained from the **lssyscfg** command as follows:

```
$ lssyscfg -r sys -F name  
Server-7998-61X-SN7157008
```

The system name can also be viewed or easily changed from the IVM GUI by clicking **View/Modify System Properties** in the navigation area.

- The --id flag specifies the mobile partition to be validated or migrated.

A successful validation will have a zero return code.

Example 11-7 shows the **migr1par** command with the operation flag with the m option for validate and migrate.

Example 11-7 *migr1par* command validate and migrate option

```
$ migr1par -o m -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5  
$
```

The status of the mobile partition can be monitored by using the **lssyscfg** command as shown in Example 11-8 on page 452 or the **ls1parmigr** command used in Example 11-9 on page 452 for the LPAR named Mars. Repeating the **ls1parmigr** command will show a change in the bytes transmitted and the bytes remaining as the migration progresses.

Example 11-8 lssyscfg command used to check migrating partition status

```
$ migr_lpar -o m -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5 &
[1]      24076366

$ lssyscfg -r lpar -F name,state
VIOS-Neptune,Running
Phobes - RHEL5-U2,Running
Mars - AIX 6.1,Migrating - Running
```

Example 11-9 lsparmigr command used to check migrating partition status

```
$ migr_lpar -o m -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5 &
[1]      24228082

$ lsparmigr -r lpar
name=VIOS-Neptune,lpar_id=1,migration_state=Not Migrating
name=Phobes - RHEL5-U2,lpar_id=2,migration_state=Not Migrating
name=Mars - AIX 6.1,lpar_id=5,migration_state=Migration In
Progress,migration_type=active,dest_sys_name=Server-7998-61X-SN7157008,
dest_lpar_id=5,source_msp_name=VIOS-Neptune,source_msp_id=1,dest_msp_na
me=Jupiter,dest_msp_id=1,bytes_transmitted=1117792957,bytes_remaining=1
7162240,remote_manager=172.16.1.100,remote_user=padmin
```

After the migration is complete, the instance of the mobile LPAR no longer appears on the original source VIOS.



System maintenance and diagnostics

This chapter discuss methods and best practices related to important IBM BladeCenter JS23 and JS43 Express maintenance topics, such as:

- ▶ “Firmware updates” on page 454.
- ▶ “System diagnostics” on page 465

12.1 Firmware updates

IBM periodically makes firmware updates available for you to install on the IBM BladeCenter JS23 and JS43 Express, the management module, or expansion cards in the blade server.

IBM BladeCenter JS23 and JS43 Express have a large firmware image, making it impossible to perform firmware updates through the Advanced Management Module (AMM). For these blade servers, one of the following methods should be chosen to update system's firmware:

- ▶ In-band operating system capabilities, such as the `update_flash` command for Linux and AIX, or the command for Virtual I/O Server
- ▶ Firmware update function of AIX diagnostics
- ▶ Firmware update function of stand-alone *Diagnostics* CD
- ▶ Firmware update function of *IBM Installation Toolkit for Linux* (IBMIT) DVD

This section details all steps necessary to update IBM BladeCenter JS23 and JS43 Express firmware image.

12.1.1 Committing new firmware to the TEMP side

IBM BladeCenter JS23 and JS43 Express hold two firmware images in their flash, one located in the temporary (TEMP) side, and the other located in the permanent (PERM) side. Firmware updates *always* take place on the TEMP side.

To commit a newer firmware image to your IBM BladeCenter JS23 and JS43 Express:

1. Make sure you are running the firmware located in the TEMP side. See 12.1.2, “Starting the firmware image from the TEMP side” on page 456 for more details.

Important: Before the installation of the new firmware to the TEMP side begins, the contents of the TEMP side should be committed to the PERM side.

The process provides for a *last known good copy of firmware* that can be used if the new image on the TEMP becomes corrupted.

2. Obtain the new firmware image from the IBM Web site. See 12.1.5, “Getting the latest firmware image from IBM” on page 464 for more details.

3. Copy the new firmware image file to your system, in either `/tmp/fwupdate` or `/home/padmin/fw` for a VIO Server directory. If the directory does not exist, create it by using either of the following command for a VIO Server:

```
mkdir /tmp/fwupdate
mkdir fw
```

4. Log in to the AIX or Linux system as root, or log in to the IVM alpha partition as padmin.

Important: Updates from within an LPAR are not supported. You must be logged in to the VIOS instead.

5. Confirm that the new firmware image file is available in your system by issuing the `ls /tmp/fwupdate` command. This command lists all files in the `/tmp/fwupdate` directory, including any firmware images you copied there. For this example, we use the name `01EAXXX.YYY.ZZZ.img` for our firmware image file. See 12.1.4, “Interpreting the firmware file name” on page 463 to interpret numbers and letters of a real firmware file name.
6. Select one of the following methods to update the firmware image:
 - Install the firmware with the in-band diagnostics of AIX.
 - Update the firmware with the `update_flash` command on AIX (Example 12-1).

Example 12-1 Using the update_flash command inside AIX

```
cd /tmp/fwupdate
/usr/lpp/diagnostics/bin/update_flash -f 01EAXXX.YYY.ZZZ.img
```

- Update the firmware with the `update_flash` command on Linux (Example 12-2).

Example 12-2 Using the update_flash command inside Linux

```
cd /tmp/fwupdate
/usr/sbin/update_flash -f 01EAXXX.YYY.ZZZ.img
```

- Update the firmware with the `ldfware` command on VIOS (Example 12-3).

Example 12-3 Using the ldfware command inside VIOS

```
cd /tmp/fwupdate
ldfware -file /home/padmin/fw/01EAXXX.YYY.ZZZ.img
```

7. Verify that the update was successful by checking the firmware level, as shown in “Verifying current firmware level” on page 462.

12.1.2 Starting the firmware image from the TEMP side

Before running firmware updates, make sure the BladeCenter server is using the firmware located in the TEMP side.

Note: Usually IBM BladeCenter JS23 and JS43 Express are configured to use the TEMP side, leaving the firmware image in the PERM side as a backup.

You may verify which side is being used, and change between firmware sides, from within the SMS menu, and the *Advanced Management Module* (AMM).

Configure to use the TEMP side through AMM

To configure your BladeCenter server to use the firmware image on TEMP side:

1. Access the Advanced Management Module.
2. Select **Blade Tasks** → **Configuration** → **Boot Mode**. Figure 12-1 shows the AMM window for selecting blade's boot mode.

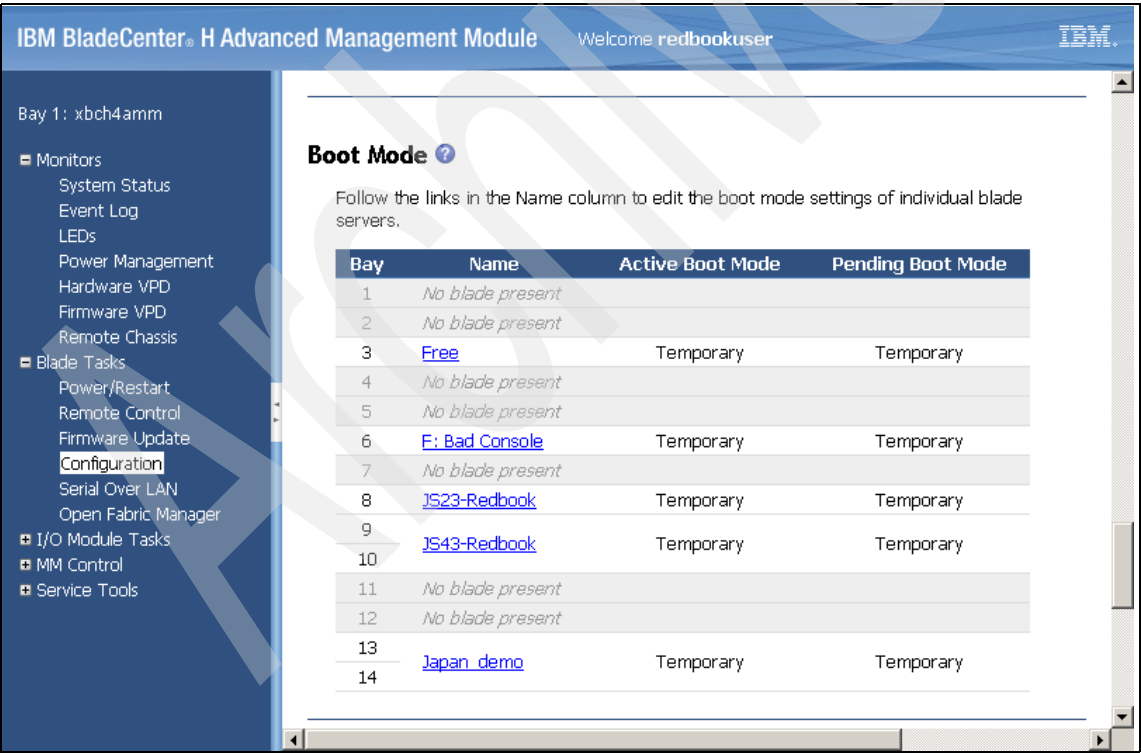


Figure 12-1 Select BladeCenter boot mode main page

3. Select the desired JS23 or JS43 blade server.
4. Select **Temporary**, as shown in Figure 12-2, to force the system to use the firmware image from the TEMP side, and then click **Save**.

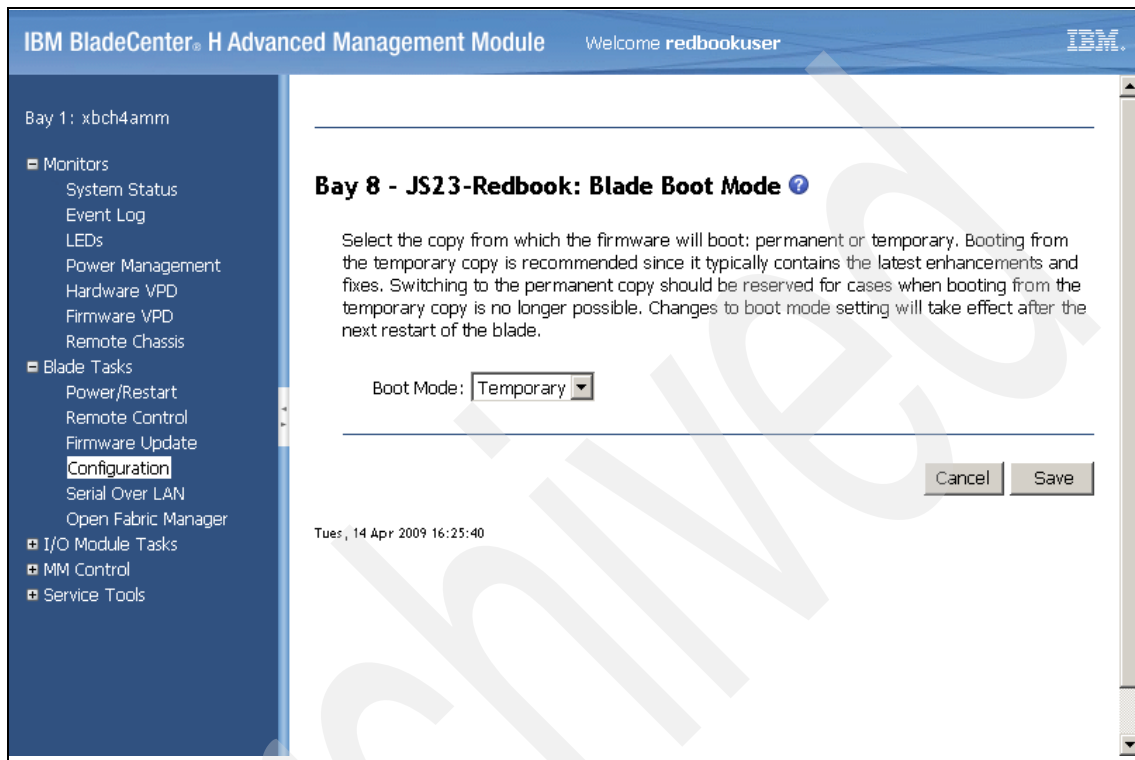


Figure 12-2 Firmware selection page

5. Restart the blade server by selecting **Blade Tasks** → **Power/Restart**. Select a BladeCenter server from the list, then select **Restart Blade** from the Available actions drop-down menu. Finally, click **Perform Action**. Figure 12-3 on page 458 shows the Blade Power/Restart panel.

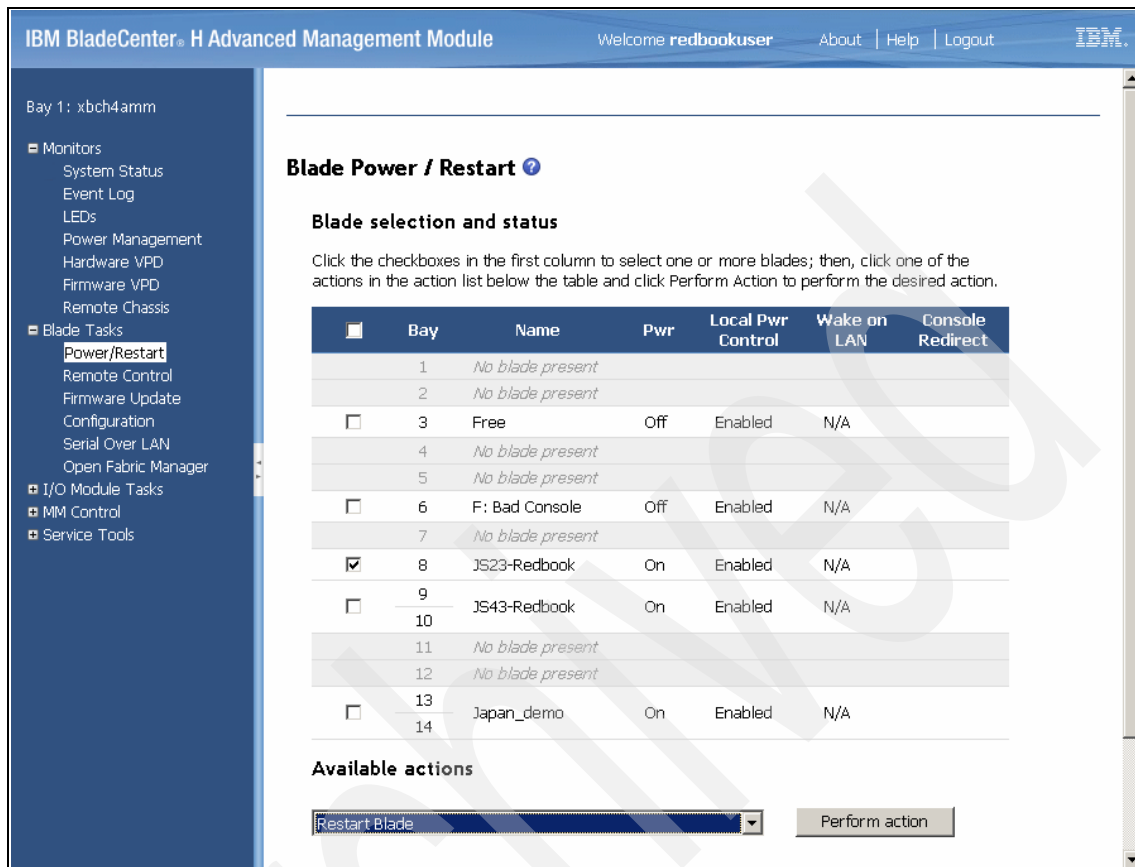


Figure 12-3 Blade Power / Restart

- Verify that the system starts using the firmware image from the TEMP side. It can be done by repeating steps 1 on page 456 and 2 on page 456 (see Figure 12-1 on page 456).

Configure to use the TEMP side through the SMS menu

To configure:

- Boot your blade server and, in the welcome screen, shown in Figure 12-4 on page 459, type the number 1 on the keyboard to select the SMS menu.

Important: Pay attention to the welcome screen shown in Figure 12-4 on page 459. It has a short timeout, and if you miss it, you have to reboot the machine.


```
Version EA340_067
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Main Menu
1.  Select Language
2.  Setup Remote IPL (Initial Program Load)
3.  Change SCSI Settings
4.  Select Console
5.  Select Boot Options
6.  Firmware Boot Side Options

-----

Navigation Keys:

X = eXit System Management Services
-----
Type menu item number and press Enter or select Navigation key:6
```

Figure 12-5 SMS main menu

Important: If your SMS menu does not provide option number 6, you might be in an LPAR. You cannot run firmware updates in IBM BladeCenter JS23 and JS43 Express blade servers from within an LPAR.

3. Figure 12-6 on page 461 shows the SMS Boot Side Option Menu. In the upper left corner you can find the level of firmware being used, and just above options 1 and 2 you can find the firmware side being used. Type 2 and press Enter to select Temporary (the TEMP side) for the next boot.

```
Version EA340_067
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Firmware Boot Side Options Menu

Firmware Boot Side for next boot: Temporary
1. Permanent
2. Temporary

-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management Services
-----
Type menu item number and press Enter or select Navigation key:2
```

Figure 12-6 SMS firmware boot side options

4. Press X and then 1 to restart the system, as shown in Figure 12-7.

```
Version EA340_067
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Are you sure you want to exit System Management Services?
1. Yes
2. No

-----
Navigation Keys:
X = eXit System Management Services
-----
Type menu item number and press Enter or select Navigation key:1
```

Figure 12-7 SMS exit page

12.1.3 Verifying current firmware level

Before updating firmware, you must know which firmware level is running in your IBM BladeCenter JS23 or JS43 Express. Several ways are available to obtain this information:

- ▶ Get firmware level by using AMM.
- ▶ Get firmware level by using SMS menu.
- ▶ Get firmware level by using `lsmcode` command for Linux and AIX, or `lsfware` for Virtual I/O Server.

Get firmware level using the AMM

From within the AMM, select **Monitors** → **Firmware VPD**. Figure 12-8 shows the list with all blade servers and firmware.

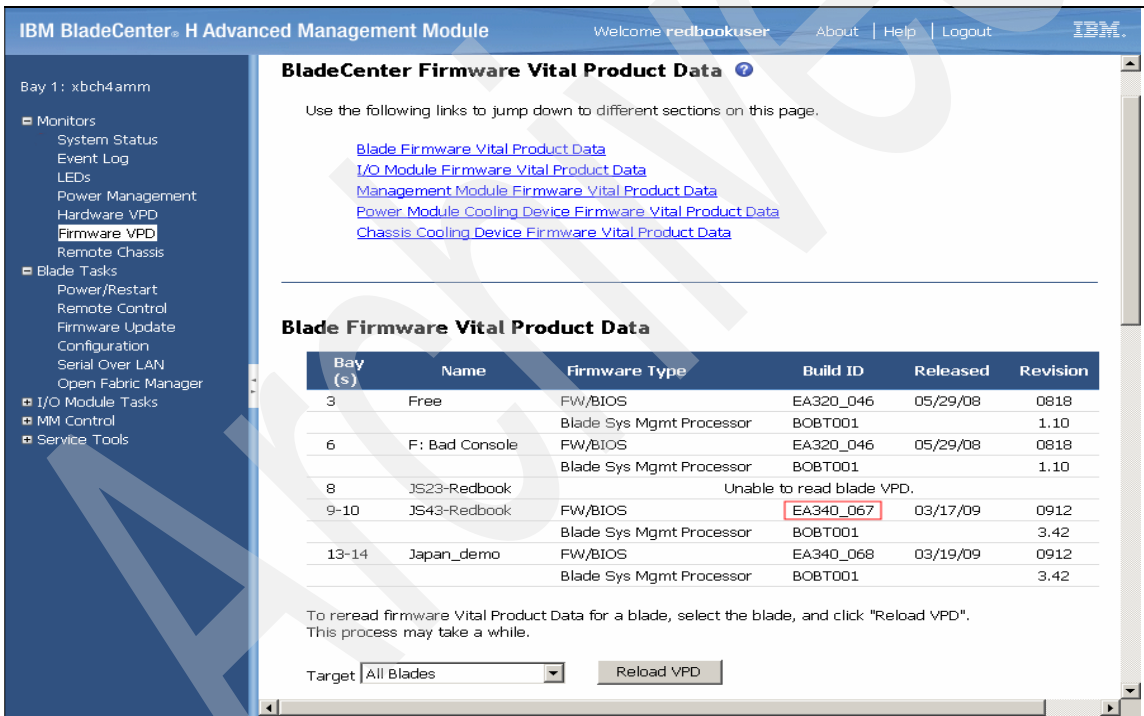


Figure 12-8 Firmware Vital Product Data (VPD) on AMM

Get firmware level using the SMS menu

Boot your blade server and, in the welcome screen, shown in Figure 12-4 on page 459, type the number 1 on the keyboard to select the SMS menu.

Note: Pay attention to the welcome screen shown in Figure 12-4 on page 459. It has a short timeout, and if you miss it you must reboot the machine.

Figure 12-9 shows the SMS main menu. In the left upper corner you can find the current firmware level.

```
Version EA340_067
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Main Menu
1.  Select Language
2.  Setup Remote IPL (Initial Program Load)
3.  Change SCSI Settings
4.  Select Console
5.  Select Boot Options
6.  Firmware Boot Side Options

-----

Navigation Keys:

X = eXit System Management Services

-----
Type menu item number and press Enter or select Navigation key:
```

Figure 12-9 Firmware level inside the SMS main menu

12.1.4 Interpreting the firmware file name

Firmware image files for IBM BladeCenter JS23 and JS43 Express are the same of the previous JS12 and JS22 BladeCenter generation. They have the following versioning format:

EAXXX.YYY.ZZZ

The format indicates:

- | | |
|---------|--|
| EA | This is an identifier. All IBM BladeCenter JS12, JS22, JS23 and JS43 Express firmware files have the EA prefix. |
| XXX | This is the <i>release level</i> . Changes in the release level means major updates in the firmware code. |
| YYY.ZZZ | These are the <i>service pack level</i> and <i>last disruptive service pack level</i> . Values for the service pack and last disruptive service pack are only unique within a release level. |

A firmware installation is always disruptive if:

- ▶ New firmware release level is different from current firmware release level.
- ▶ New firmware service pack level and last disruptive service pack level have the same value.
- ▶ Current firmware's service pack level is lower than the last disruptive service pack level from the new firmware.

IBM BladeCenter server firmware levels have a specific file name notation, which differs from other POWER systems. For a detailed description of the file name format, see:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_naming_convention.pdf

12.1.5 Getting the latest firmware image from IBM

To update the firmware in your IBM BladeCenter JS23 and JS43 Express, obtain the firmware image file from the IBM BladeCenter Support Web site:

<http://www.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000020>

Important: If you have problems accessing the site, use the search function in the IBM Web site home page to find the correct Web address. Search for *Fix Central* or *Support and Downloads*.

When you reach the BladeCenter support site, select the appropriate Product family, Type, Model, and Operating System, to match your system's configuration. Click **Go** when you are ready. Only the Product Family value is required.

You are provided a list with all downloads available for your system. At this time we are looking for the BIOS files (which are actually the available firmware images for your system).

Select the firmware version you want to download. You are redirected to a new page with many available file formats. For this example, we get the *image file* format (with the .img extension).

12.2 System diagnostics

POWER6 processor-based systems contains specialized hardware detection circuits for detecting erroneous hardware operations, and includes extensive hardware and firmware recovery logic. IBM hardware error checkers have the following distinct attributes:

- ▶ Continuous monitoring of system operations to detect potential calculation errors
- ▶ Attempted isolation of physical faults based on runtime detection of each unique failure
- ▶ initiation of a wide variety of recovery mechanisms designed to correct a problem

Machine checks are handled by firmware. When a machine check occurs, the firmware analyzes the error to identify the failing device and creates an error log entry.

In partitioned mode, any error that occurs during partition activity is surfaced to the operating system running in the partition. If an error occurs during POWER hypervisor (PHYP) activities, the system is rebooted by PHYP.

If the system degrades to the point where the service processor cannot reach standby state, then the ability to analyze the error does not exist.

Important: This section is not intended to be a replacement for the information provided in the *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339. Refer to the guide for detailed steps of how to perform diagnostics tasks, determine the root cause of an error, and seek proper support assistance.

12.2.1 Diagnostic tools

This section has a list of tools that can be used to help with diagnosing hardware problems on IBM BladeCenter JS23 and JS43 Express.

Checkpoints and error codes

During system power-on process, the *power-on self-test* (POST) checks the hardware, including certain system components and interconnections, and generates 8-digit checkpoint codes to mark the power-on progress.

If the Service Processor detects a problem during POST, an error code is logged in the AMM event log. Error codes are also logged in Linux syslog or AIX diagnostics log, if possible.

See “Checkpoint code (progress code)” on page 472 for more details.

Light Path and front panel diagnostics

IBM BladeCenter JS23 and JS43 Express comes with the Light Path technology, which helps to determine *customer-replaceable units* (CRU) with problems. Both the Base and MPE planar boards have LEDs for signaling faulty parts. The front panel is also part of the Light Path diagnostics solution, with LEDs to help verify the current state of the system.

Control panel

IBM BladeCenter JS23 and JS43 Express have the same control panel, as shown in Figure 12-10.




Figure 12-10 IBM BladeCenter JS23 and JS43 control panel

The front panel has buttons and LEDs that help on controlling and getting status from your Blade server. Table 12-1 on page 467 has a short description for each button and LED located in the front panel. See *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339 for more information.

Table 12-1 Description of front panel buttons and LEDs

Callout	Description
1	Keyboard/video selection button
2	Media tray selection button
3	Information LED
4	Error LED
5	Power Control button
6	Nonmaskable Interrupt (NMI) reset button
7	Sleep LED. Not used in the IBM BladeCenter JS23 and JS43 Express
8	Power-on LED
9	Activity LED. When lit (green), it indicates activity exists on the hard disk drive or network
10	Location LED

The management module and IBM Director console can help to control certain of front panel LEDs. For instance, a system administrator can turn on or off the Location LED of some blade servers.

To access these functions in the Advanced Management Module Web interface:

1. Access the AMM Web page.
2. Select **Monitors** → **LEDs**. Figure 12-11 on page 468 shows the LEDs monitoring page in AMM.




Figure 12-11 AMM BladeCenter LEDs control and status page

Light Path

Light Path diagnostics is a system of LEDs on the control panel and on your system board (IBM BladeCenter JS43 Express has Light Path LEDs on both boards). When a hardware error occurs, LEDs are lit throughout the blade server.

LEDs are available for many components, such as:

- ▶ Battery.
- ▶ SAS HDD (or SSD) disks, on both Base and MPE planars.
- ▶ Management card on Base planar only.
- ▶ Memory modules on both Base and MPE planars.
- ▶ Expansion cards (1Xe and HSDC daughter cards) on both Base and MPE planars.

The MPE planar has a special LED to help in determining the correct planar where the error is happening. This LED is named *Check card below*, and is used to indicate that the problem is not occurring in the MPE planar, but in the Base planar. This LED is also named the *Not me* LED.

Note: Refer to the *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339. It has more information about how to perform diagnostics using the Light Path technology, and also how to act when known types of problems arise.

Figure 12-12, Figure 12-13 on page 470, and Table 12-2 on page 470 show all Light Path LEDs available on your IBM BladeCenter JS23 and JS43 Express boards.




Figure 12-12 LEDs on the IBM BladeCenter JS23 Express board. This is the same LED configuration of the IBM BladeCenter JS43 Express Base planar




Figure 12-13 LEDs on the IBM BladeCenter JS43 Express MPE planar

Table 12-2 Lightpath LED description.

Callout	Base planar LEDs	Expansion planar LEDs
1	Lightpath power LED	
2	System board LED	
3	SAS hard disk drive LED	
4	DIMM 1-4 error LEDs	None
5	1Xe expansion card connector LED	
6	DIMM 5-8 error LEDs	None
7	PCIe high-speed expansion card LED	
8	Management Card LED	None
9	Battery LED	None
10	None	<i>Check card below</i> LED
11	None	DIMM 9-12 error LEDs
12	None	DIMM 13-16 error LEDs

Diagnostic utilities for the AIX operating system

AIX provides many diagnostic and maintenance functions, such as:

- ▶ Automatic error log analysis.
- ▶ Firmware updates, format disk, and RAID Manager.

For more information about performing diagnostics in your IBM BladeCenter JS23 and JS43 Express using AIX, see:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphau/working01.htm>

Diagnostic utilities for the Linux operating system

Linux on POWER service and productivity tools (RAS Tools) include hardware diagnostic aids and productivity tools for Linux on POWER systems.

To perform error log analysis in Linux, you must have the Automatic Error Log Analysis (diagela) package installed.

Note: SUSE Linux Enterprise Server 11 (SLES11) ships a package named *ppc64-diag* for running log analysis on POWER machines. You should install this package instead of *diagela* on SLES11.

See Appendix D, “Service and productivity tools for Linux” on page 539 for details about RAS Tools for Linux.

Diagnostics utilities for the IBM i operating system

IBM i runs only in logical partitions on IBM BladeCenter JS23 and JS43 Express, therefore it sees virtual devices only. In this situation, the errors associated with hardware failures are not seen inside the operating system, in the normal Product Activity Logs (PAL) or System Activity Logs (SAL). See 12.2.4, “IBM i partition diagnostics and errors” on page 476 for a detailed description of how to gather diagnostic information from an IBM i partition.

Stand-alone diagnostics

When you do not have an operating system installed in your blade server, or you are not able to boot the installed operating systems, you may run diagnostic utilities by using one of the following stand-alone diagnostic solutions:

- ▶ AIX stand-alone *Diagnostics* CD; includes all diagnostic utilities found in AIX.
- ▶ *IBM Installation Toolkit for Linux* DVD, which can be used as a Linux Live DVD, and includes all RAS Tools for Linux. See Appendix D, “Service and productivity tools for Linux” on page 539 for more details about the IBM Installation Toolkit for Linux.

12.2.2 Reference codes

Reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem. IBM BladeCenter JS23 and JS43 Express produces many types of reference codes, in the following types:

- ▶ Progress codes

These 8-digit status codes are generated by the power-on self-test (POST). They are used to show progress when powering on the blade server.

- ▶ Error codes

These are 9-word (8-digit per word) error codes. Error codes are either *system reference codes (SRCs)* or *service request numbers (SRN)*. It can also have a *location code*, which helps on isolating the piece of hardware with problems.

This section describes the types of codes generated by the IBM BladeCenter JS23 and JS43 Express. For a more detailed description of each code type, and a list of possible codes, see the *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339.

For details about how to locate or view the reference codes for a specific blade server, see: 12.2.3, “Diagnostics tasks through the Advanced Management Module” on page 473.

Location code

Location codes identify components of the blade server. They can appear in error codes to help identify which server's component is causing the error.

Location codes exist for components in both Base and MPE planars.

Checkpoint code (progress code)

Checkpoint codes are generated by the power-on self-test (POST), which performs a series of tests to check the operation of the blade server components.

Progress codes do not indicate an error, although in certain cases, the blade server can pause indefinitely (hang). Progress codes for blade servers are 9-word, 8-digit hexadecimal numbers that start with C and D.

A checkpoint might have an associated location code as part of the message. The location code provides information that identifies the failing component during a hanging condition.

System reference code

System reference codes (SRC) are used to identify both hardware and software problems in IBM BladeCenter JS23 and JS43 Express. Those errors can be originated in hardware, in firmware, or in the operating system.

The SRC identifies the component that generated the error code and describes the error. Each SRC has nine words, each word being an eight-digit code.

Usually the first character indicates the type of error, as shown in Table 12-3

Table 12-3 Types of SRCs

First character	Type of error
A	User Attention, which means that user interventions might be required to progress.
B	Error, which means the firmware has identified an error.
C	Checkpoint, which is used to give status on normal boot. Hang indicates unexpected error.
D	Dump, which indicates a dump is in progress. Hang indicates unexpected error.
1	This is a <i>System Power Control Network</i> (SPCN) error,
6	This is a virtual optical device error.

Service request number

Service request numbers (SRN) are error codes generated by the operating system, and they have three digits, a hyphen, then three or four digits after the hyphen. This type of error code can be visualized by the AIX *diagnostics* or the Linux service aid *diagela* (*ppc64-diag* on SLES11) package if it is installed.

12.2.3 Diagnostics tasks through the Advanced Management Module

This section describes how to use some AMM features to diagnostic issues in your IBM BladeCenter JS23 and JS43 Express. For detailed information about how to install, configure, and use the AMM, see:

- *Advanced Management Module User's Guide*

<ftp://ftp.software.ibm.com/systems/support/intellistation/44r5375.pdf>

- *Advanced Management Module Installation Guide*

ftp://ftp.software.ibm.com/systems/support/system_x/44r5269.pdf

Using the AMM to view SRCs

To view the progress codes and error codes for a specific blade server:

Note: IBM BladeCenter JS23 and JS43 does not display the checkpoints in the Serial Over LAN (SOL) console, nor in the graphical console.

1. Access the AMM Web interface.
2. In the menu on the left, select **Service Tools** → **Blade Service Data**, as shown in Figure 12-14.

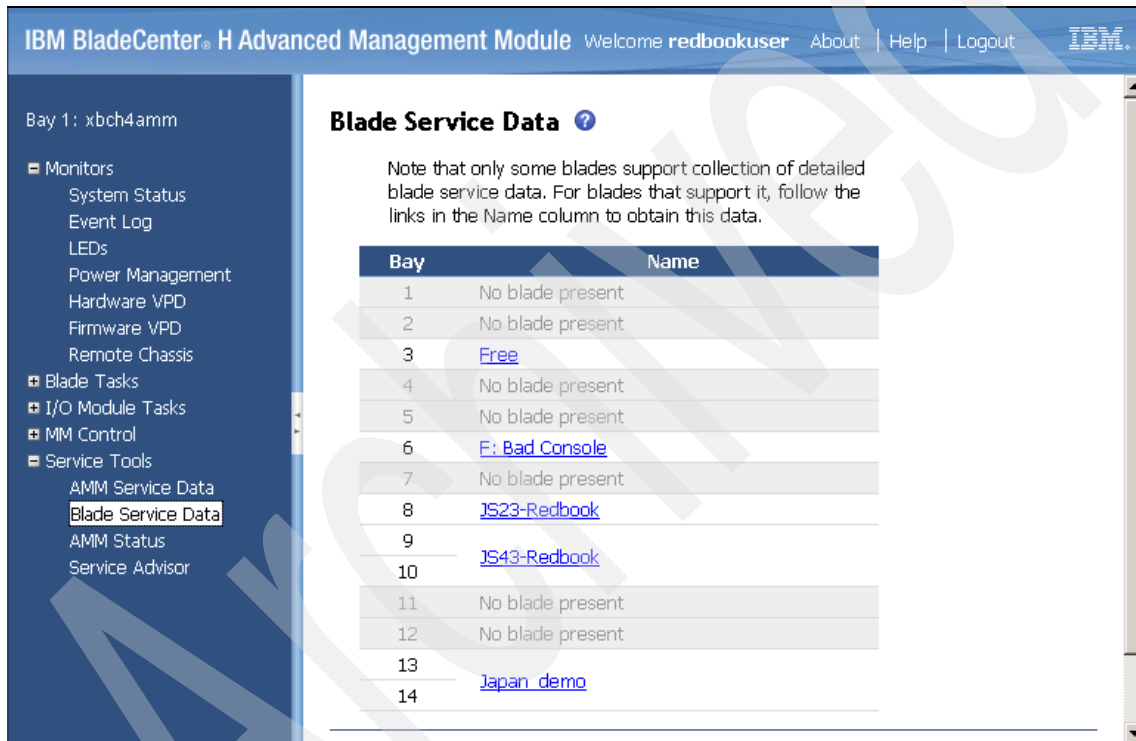


Figure 12-14 AMM Blade Server Data panel

3. Select a blade server. The reference codes are shown for the selected blade server, as shown in Figure 12-15 on page 475.

The Advanced Management Module can display the last 32 reference codes. You may refresh the list to update it.

IBM BladeCenter® H Advanced Management Module Welcome redbookuser <a>About <a>Help <a>Logout IBM			
Bay 1: xbch4amm			
<div> <div>Monitors</div> <div> <div>System Status</div> <div>Event Log</div> <div>LEDs</div> <div>Power Management</div> <div>Hardware VPD</div> <div>Firmware VPD</div> <div>Remote Chassis</div> </div> </div> <div> <div>Blade Tasks</div> <div>I/O Module Tasks</div> <div>MM Control</div> <div>Service Tools</div> <div> <div>AMM Service Data</div> <div>Blade Service Data</div> <div>AMM Status</div> <div>Service Advisor</div> </div> </div>			
System Reference Codes ?			
Follow the links in the System Reference Code column to obtain additional detailed data relating to the particular code.			
Unique ID	System Reference Code	Timestamp	
000000ff	AA00E1A9	2009-04-17 19:11:54	
000000fe	CA00E1A0	2009-04-17 19:11:54	
000000fd	CA00E1B5	2009-04-17 19:11:54	
000000fc	CA00E1F1	2009-04-17 19:11:50	
000000fb	CA00E1F0	2009-04-17 19:11:50	
000000fa	CA00E141	2009-04-17 19:11:50	
000000f9	CA00E1DC	2009-04-17 19:11:42	
000000f8	CA00E1DC	2009-04-17 19:11:42	
000000f7	CA00D008	2009-04-17 19:11:42	
000000f6	CA00E100	2009-04-17 19:11:42	
000000f5	CA00E1FB	2009-04-17 19:11:42	
000000f4	CA00E100	2009-04-17 19:11:42	
000000f3	CA00E1FE	2009-04-17 19:11:42	
000000f2	CA00E1FE	2009-04-17 19:11:42	
000000f1	CA00E1FE	2009-04-17 19:11:42	
000000f0	CA00E1FE	2009-04-17 19:11:42	

Figure 12-15 Power-on checkpoints inside AMM Web interface

Using the AMM to view log messages

You can use the AMM Web interface to view log messages generated by the blade servers within a BladeCenter chassis.

From AMM, select **Monitors** → **Event Log**. Figure 12-16 on page 476 shows the event log interface in AMM.

You may filter the log events by blade server, by severity of event, by date, and also by serviceable flag, which indicates whether the *Call Home* flag was enabled for that event or not. For details about Call Home, see “Service Advisor” on page 476.

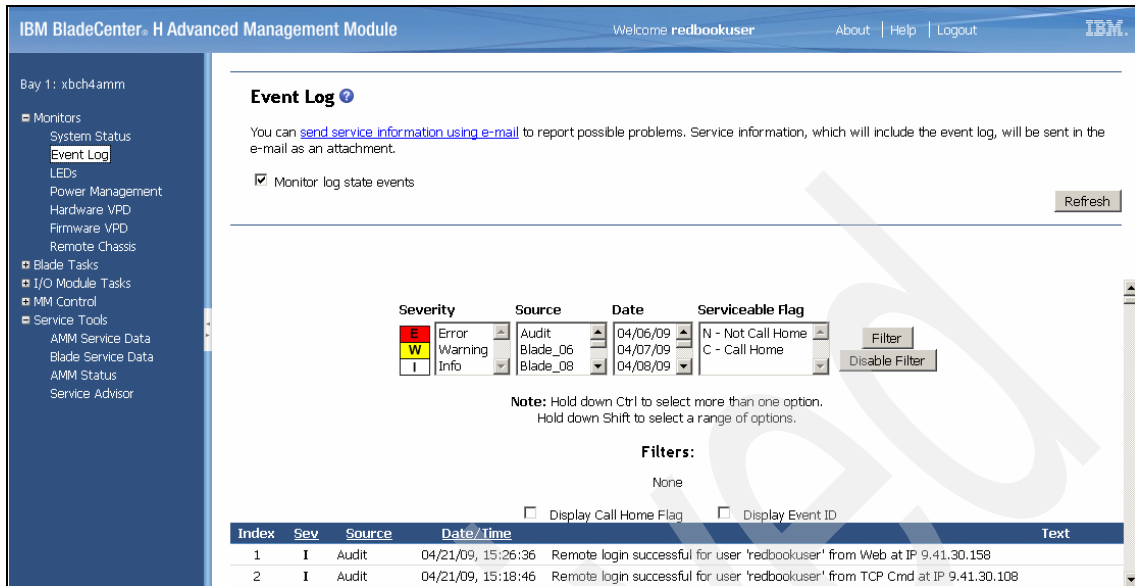


Figure 12-16 AMM event log interface

Service Advisor

The Service Advisor enables the BladeCenter to automatically send hardware and firmware serviceability messages to IBM. Every time a log event with the Call Home flag enabled happens, the AMM's Service Advisor sends a message containing the event log message, BladeCenter unit inventory, and status to IBM Support.

This Call Home feature is disabled by default. You have to accept the Service Advisor Terms and Conditions before enabling it.

For a detailed description of how to enable the Service Advisor in the AMM interface, see:

<ftp://ftp.software.ibm.com/systems/support/intellistation/44r5375.pdf>

Note: No client data from the server or I/O modules is sent within the Call Home contact message.

12.2.4 IBM i partition diagnostics and errors

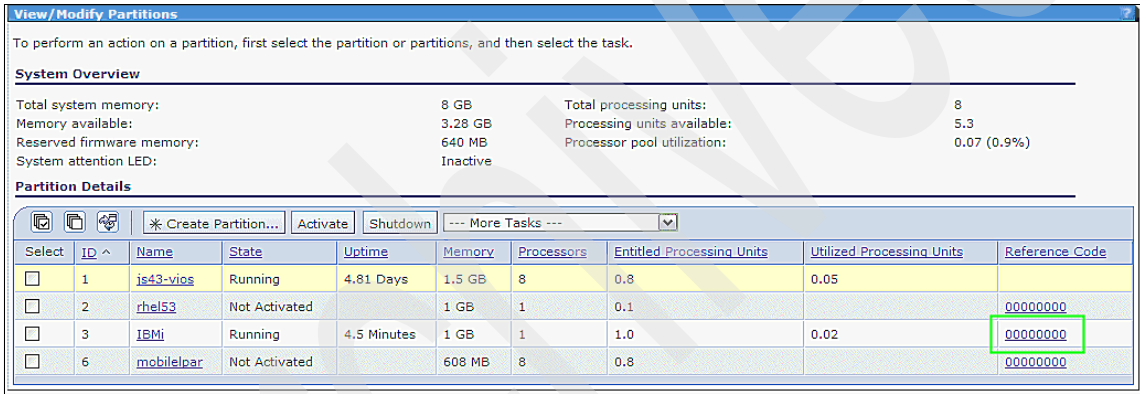
When an IBM i partition is configured on a JS23 and JS43, it will be configured with virtual devices only. Because an IBM i partition on a blade server cannot own any physical resources, the errors associated with hardware failures are not

seen in the normal product activity logs (PAL) or system activity logs (SAL). Most errors that the IBM i partition encounters are related to storage or configuration. Any true hardware errors are reported to the VIOS partition and repaired using VIOS options. In this section, we discuss where to collect error data and configuration information related to an IBM i virtual partition.

Collecting IBM i partition error information

When an IBM i partition encounters a failure, the reference code information is the information that is required to troubleshoot and repair errors. The reference codes can be found using IVM on the View/Modify Partitions panel.

Figure 12-17 shows an example of the partitions view with reference codes. The reference codes are in the column on the right column by default. Clicking the reference code for a partition opens the Reference Code list for the selected partition.



The screenshot shows the 'View/Modify Partitions' window. It includes a 'System Overview' section with memory and processor statistics, and a 'Partition Details' table. The table lists four partitions: 'js43-vios', 'rhel53', 'IBMi', and 'mobilelpar'. The 'Reference Code' column for the 'IBMi' partition is highlighted with a green box, showing the value '00000000'.

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js43-vios	Running	4.81 Days	1.5 GB	8	0.8	0.05	
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1	0.1		00000000
<input type="checkbox"/>	3	IBMi	Running	4.5 Minutes	1 GB	1	1.0	0.02	00000000
<input type="checkbox"/>	6	mobilelpar	Not Activated		608 MB	8	0.8		00000000

Figure 12-17 Partitions view

Another entry point to the reference codes is to use the check box to select the partition, and then select **Reference Codes** from the More Tasks pull-down menu, as shown in Figure 12-18 on page 478.

Figure 12-18 More Tasks - Reference Codes

A new window opens and displays the list of codes for the selected word. In Figure 12-19 on page 479, the codes listed are from the normal with no errors at this time. Selecting any reference code will add additional words to the right of the panel in the Details

A new window opens and displays the list of codes for the selected partition. In Figure 12-19 on page 479, the codes listed are from the last IPL. Everything is normal with no errors at this time. Selecting any reference code displays the additional words to the right of the panel in the Details section.

A new window opens and displays the list of codes for the selected partition. In Figure 12-19 on page 479, the codes listed are from the last IPL. Everything is normal with no errors at this time. Selecting any reference code displays the additional words to the right of the panel in the Details section.

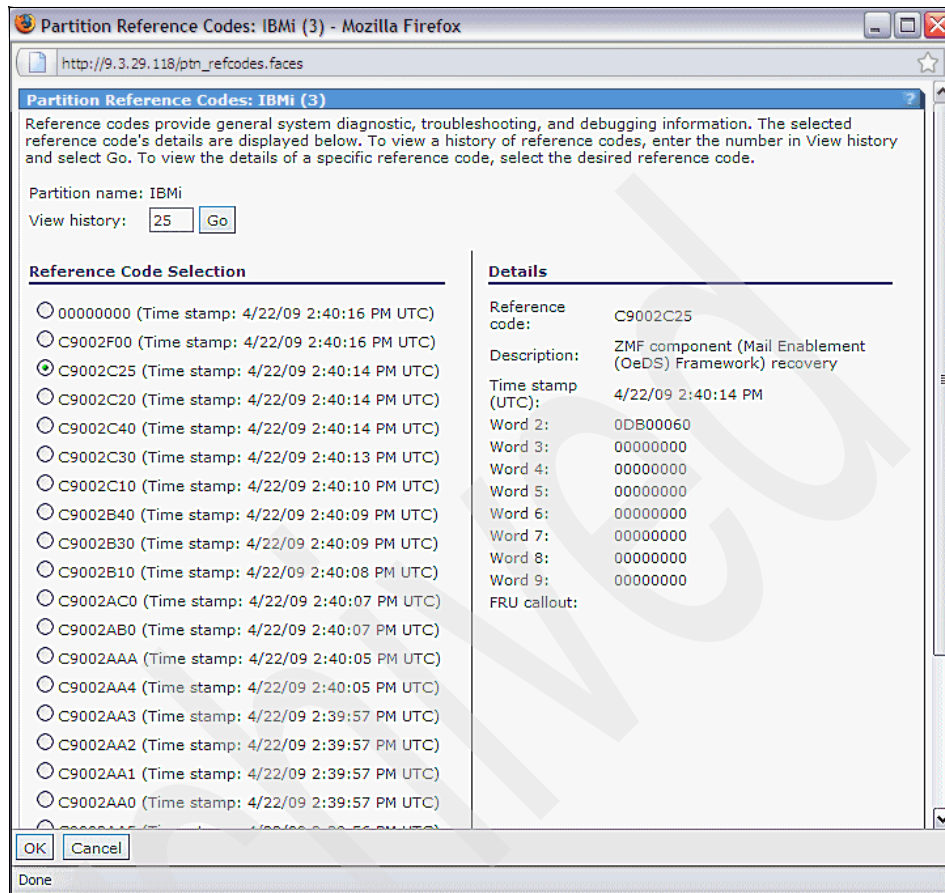


Figure 12-19 Reference Code list - normal IPL

Now, we look at an error condition in the IBM i partition. For this scenario, we assume that the partition was running without problems. Something happened that caused the partition to hang. Users report that the partition is not responding.

There are many ways to troubleshoot and report problems. The intent of this section is *not* to provide procedures beyond collecting data and contacting your next level of support.

Depending on your systems configuration, the chassis can be configured to notify the user of errors. These errors can be included in the notification type so that the user can receive an e-mail indicating that an error on the BladeCenter has occurred.

One place to look for errors is in IVM. Looking at the View/Modify partitions window, we see an error condition on the IBM i partition. In Figure 12-20, notice that the Attention indicator is next to the partition, and that a code is listed in the Reference Code column. Normally we expect to see 00000000 in the Reference Code column if everything is running correctly.

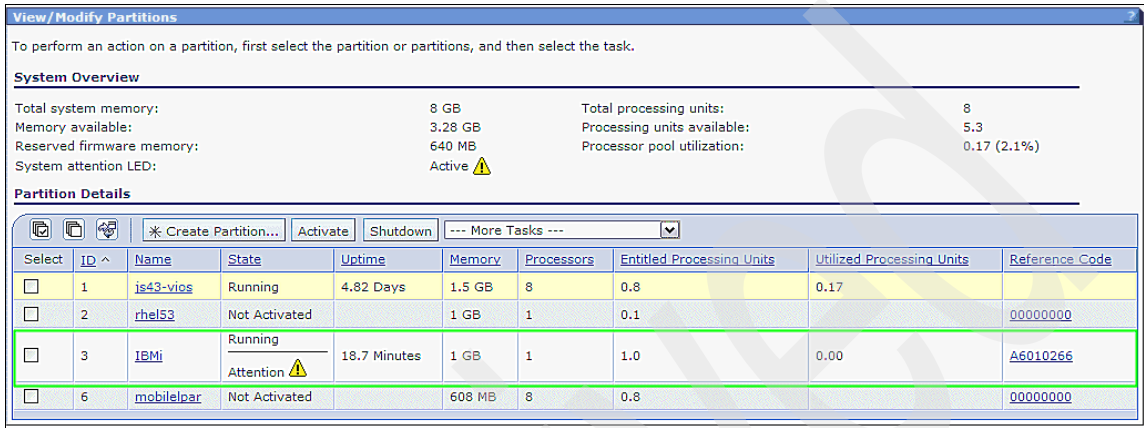


Figure 12-20 IBM i partition error

To determine more about the error, click the reference code or select an option from the More Tasks menu as mentioned previously. Figure 12-21 on page 481 shows an example of the reference codes with the error listed.

Partition Reference Codes: IBMi (3)

Reference codes provide general system diagnostic, troubleshooting, and debugging information. The selected reference code's details are displayed below. To view a history of reference codes, enter the number in View history and select Go. To view the details of a specific reference code, select the desired reference code.

Partition name: IBMi

View history:

Reference Code Selection	Details
<input checked="" type="radio"/> A6050266 (Time stamp: 4/22/09 3:02:09 PM UTC) <input type="radio"/> A6040266 (Time stamp: 4/22/09 3:00:29 PM UTC) <input type="radio"/> A6030266 (Time stamp: 4/22/09 2:58:49 PM UTC) <input type="radio"/> A6020266 (Time stamp: 4/22/09 2:57:09 PM UTC) <input type="radio"/> A6010266 (Time stamp: 4/22/09 2:56:10 PM UTC) <input type="radio"/> 00000000 (Time stamp: 4/22/09 2:40:16 PM UTC) <input type="radio"/> C9002F00 (Time stamp: 4/22/09 2:40:16 PM UTC) <input type="radio"/> C9002C25 (Time stamp: 4/22/09 2:40:14 PM UTC) <input type="radio"/> C9002C20 (Time stamp: 4/22/09 2:40:14 PM UTC) <input type="radio"/> C9002C40 (Time stamp: 4/22/09 2:40:14 PM UTC) <input type="radio"/> C9002C30 (Time stamp: 4/22/09 2:40:13 PM UTC) <input type="radio"/> C9002C10 (Time stamp: 4/22/09 2:40:10 PM UTC) <input type="radio"/> C9002B40 (Time stamp: 4/22/09 2:40:09 PM UTC) <input type="radio"/> C9002B30 (Time stamp: 4/22/09 2:40:09 PM UTC) <input type="radio"/> C9002B10 (Time stamp: 4/22/09 2:40:08 PM UTC) <input type="radio"/> C9002AC0 (Time stamp: 4/22/09 2:40:07 PM UTC) <input type="radio"/> C9002AB0 (Time stamp: 4/22/09 2:40:07 PM UTC) <input type="radio"/> C9002AAA (Time stamp: 4/22/09 2:40:05 PM UTC)	<p>Reference code: A6050266</p> <p>Description: Contact was lost with the device indicated</p> <p>Time stamp (UTC): 4/22/09 3:02:09 PM</p> <p>Word 2: 0DB00062</p> <p>Word 3: 00010004</p> <p>Word 4: 0000002B</p> <p>Word 5: 00000000</p> <p>Word 6: 000100FF</p> <p>Word 7: 00FF0280</p> <p>Word 8: 6B220050</p> <p>Word 9: B6005120</p> <p>FRU callout:</p>

Figure 12-21 Reference Code list - error condition

Using the start of call procedures, this reference code information can be used to complete the Problem Summary Form. This information is used by service and support to troubleshoot the error and provide assistance in resolving the problem. Depending on your skill level, you might be able to navigate through the various Information Center pages to troubleshoot this error further.

Another source for error information is from the AMM. As mentioned previously, the AMM can be configured to send alerts for certain error conditions. Even if the AMM is not configured in this manner, you can still use the event logs to view and collect error information.

Figure 12-22 on page 482 shows an example of the event log on the AMM.

IBM BladeCenter- H Advanced Management Module

Welcome redbookuser

About | Help | Logout

IBM

Bay 1: xbc4amm

Monitors

System Status

Event Log

LEDs

Power Management

Hardware VPD

Firmware VPD

Remote Chassis

Blade Tasks

I/O Module Tasks

MM Control

General Settings

Login Profiles

Alerts

Serial Port

Port Assignments

Network Interfaces

Network Protocols

Chassis Int Network

Security

File Management

Firmware Update

Configuration Mgmt

Restart MM

Service Tools

Event Log

You can [send service information using e-mail](#) to report possible problems. Service information, which will include the event log, will be sent in the e-mail as an attachment.

☒ Monitor log state events

Refresh

Severity

Source

Date

Serviceable Flag

Filter

Disable Filter

Note: Hold down Ctrl to select more than one option.
Hold down Shift to select a range of options.

Filters:

None

☐ Display Call Home Flag ☐ Display Event ID

Index	Sev	Source	Date/Time	Text
1	I	Audit	04/22/09, 10:57:00	Remote logoff successful for user 'redbookuser' from SSH at IP 9.41.30.108
2	I	Audit	04/22/09, 10:57:00	Remote login successful for user 'redbookuser' from SSH at IP 9.41.30.108
3	E	Blade_09	04/22/09, 10:56:27	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038448 A6000266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
4	E	Blade_09	04/22/09, 10:54:48	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038446 A60F0266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
5	E	Blade_09	04/22/09, 10:53:07	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038443 A60E0266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
6	E	Blade_09	04/22/09, 10:51:27	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038440 A60D0266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
7	E	Blade_09	04/22/09, 10:49:48	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (5303843E A60C0266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
8	E	Blade_09	04/22/09, 10:48:08	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038438 A6080266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
9	E	Blade_09	04/22/09, 10:46:28	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038436 A60A0266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
10	E	Blade_09	04/22/09, 10:44:48	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038434 A6090266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
11	E	Blade_09	04/22/09, 10:43:08	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038433 A6080266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
12	E	Blade_09	04/22/09, 10:41:28	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (53038430 A6070266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005
13	E	Blade_09	04/22/09, 10:39:48	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (5303842E A6060266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6822005

Figure 12-22 AMM Event Log

The event log can be filtered to view only events specific to the blade server or other components. Figure 12-23 shows an example of the filter options.

Severity

Source

Date

Serviceable Flag

Filter

Disable Filter

Note: Hold down Ctrl to select more than one option.
Hold down Shift to select a range of options.

Filters:

Blade_09,04/22/09,Error

☐ Display Call Home Flag ☐ Display Event ID

Figure 12-23 Event log filter

In the list of events, you see the error log information. Figure 12-24 on page 483 provides an example of the data in the AMM event log. This data is similar to the data shown on the partition reference code panel as viewed from IVM we looked at earlier.

This data can also be saved by scrolling to the bottom of the event log and clicking **Save Log as Text File**. This data can then be supplied to service and support for further assistance in error determination.

x	Sev	Source	Date/Time	Event ID	Text
	E	Blade_09	04/22/09, 10:56:27	0x10000002 (J543-Redbook)	SYS F/W: Error. Replace SVCDOCS (5303B448 A6000266 00000062 00010004 0000002B 00000000 000100FF 00FF0280 68220050 B6005120)
	E	Blade_09	04/22/09, 10:54:48	0x10000002 (J543-Redbook)	SYS F/W: Error. Replace SVCDOCS (5303B446 A60F0266 00000062 00010004 0000002B 00000000 000100FF 00FF0280 68220050 B6005120)

Figure 12-24 Event log data details

As mentioned earlier, the intention of this book is *not* to explain troubleshooting processes for an IBM i partition. For those of you who are curious about this particular error, it was created by removing the partition assignment of one of the hdisk units for the partition.

Typically, errors of this type in a pure virtual environment are caused by configuration changes made by the administrator or others who might not be familiar with the partition configuration. In our example, we simply added the drive back into the configuration using the Modify Partition Assignment options and the partition recovered. Using IVM, you can see the tasks performed by the Monitor Tasks option. This might provide some indication that a configuration change was made that could have caused the error on the partition. Figure 12-25 shows an example of the Monitor Task option.

Monitor Tasks					
The last 40 tasks that have been run from the web interface on the system are listed below. Select the task to view it's properties.					
Monitored Tasks					
	Properties				
Select	Name	Status	Start Time	Task Duration	User ID
<input checked="" type="radio"/>	Modify partition assignment (physical volume)	Successful	4/22/09 11:04:13 AM	00:00:00	padmin
<input type="radio"/>	Modify partition assignment (physical volume)	Successful	4/22/09 9:55:53 AM	00:00:00	padmin
<input type="radio"/>	Activate	Successful	4/22/09 9:36:43 AM	00:00:00	padmin

Figure 12-25 Monitor tasks

Although the task might not specifically indicate what exactly was modified, at least there is an indication something changed possibly around the time of the partition error.

Important: Removing and adding drives to IBM i partitions can be disastrous to the partition and might cause system-reload conditions if care is not exercised to follow the correct processes.



Part 3

Appendixes

In this part of the book, we provide additional technical support information:

- ▶ Appendix A, “Consoles, SMS, and Open Firmware” on page 487
- ▶ Appendix B, “SUSE Linux Enterprise Server AutoYaST” on page 515
- ▶ Appendix C, “Additional Linux installation configuration options” on page 529
- ▶ Appendix D, “Service and productivity tools for Linux” on page 539

Consoles, SMS, and Open Firmware

This appendix briefly covers the methods to gain access to the console, use the System Maintenance Services Menu (SMS) to select the console to use, and use the Open Firmware prompt to choose fiber channel host bus adapter settings.

This appendix contains the following topics:

- ▶ “Consoles of the IBM BladeCenter JS23 and JS43” on page 488
- ▶ “System Management Services menu” on page 495
- ▶ “Open Firmware interface” on page 503

Consoles of the IBM BladeCenter JS23 and JS43

Like the previous JS12 and JS22 BladeCenter servers, the IBM BladeCenter JS23 and JS43 blades have a graphics adapter. This graphics adapter enables using the KVM switch that is built into the Advanced Management Module to gain access to the console of the blade.

An alternative method to gain access to the console is Serial Over LAN (SOL).

You may use either the graphical console or the SOL console during POST. By default, a JS23/JS43 blade starts the first time with the graphical console, if the configuration for startup was not changed prior to startup. (If you do not change the default configuration, the system starts with the graphical console.) After the first power-on, the blade prompts you to select which console should be used as active console. If you miss this point during POST, you have to enter the System Maintenance Services menu from a graphical console and change the console. See “System Management Services menu” on page 495 for information about how to change the active console.

Graphical console

The graphical console is available by either using the connected keyboard, video, and mouse (KVM) on a management module installed in the BladeCenter chassis, or the remote control function of the Advanced Management Module.

Using keyboard, video, mouse to access graphical console

The physical console connected to a management module is not used most of the time. If you are working locally at the BladeCenter, you may use this console to access blades. One analog console is available that is switched through the KVM switch that is built into the management module between the blades in the chassis. To switch the console, you may use the KVM Select button that is placed on the control panel of each blade. Figure A-1 on page 489 shows the location of the control panel and the control elements on the panel of an IBM BladeCenter JS23 and JS43.




Figure A-1 JS23/JS43 front control panel

Pressing the **KVM Select** button switches the console to the blade on which the button was pressed. Only one blade in a chassis has a lighted KVM Select button.

Note: Be sure that you are using the keyboard, video, and mouse connected to the active Advanced Management Module. Only one management module is active at one time. You can recognize this activity by looking at the LEDs of the management modules. Only one management module has the blinking LED switched on. This applies to installations with redundant management modules only.

You may use the `list -1 2` command in a Telnet or SSH session connected to the Advanced Management Module to identify the active management module.

Key combinations

An alternative method to switching the physical console between blades is by using the following key combination on an IBM Space Saver Keyboard that is typically used in an IBM Rack:

Shift+Num Lock+Num Lock+*bay number*+Enter

Use the key combination as follows:

1. Press and hold the Shift key. On standard keyboards the Shift key is not required.
2. Press Num Lock twice.
3. Release the Shift key.
4. Press the bay number, which is digit in the range of 1 - 14, depending on the chassis you are using.
5. Press Enter.

Using remote control to access graphical console

Remote control is a feature of the management module installed in a BladeCenter chassis. It can connect over an IP connection to the management module and open a browser window that has the graphical console redirected. Only one analog console redirection can be in progress at one time. This remote control shows the same content that is shown on the physical console.

Note: Remote control supports at maximum the standard VESA modes with 1024x768 pixels at 60,70 and 75Hz.

Concurrent KVM (cKVM) allows multiple digital KVM connections on some blades. This feature is not supported on IBM BladeCenter JS23 and JS43.

To open a remote control session:

1. Log on to the management module with your Web browser. The default IP address of the management module is 192.168.70.125/24. Enter the user name and password and click **Log In** as shown in Figure A-2 on page 491. The default account with all administrative rights is USERID with password PASSW0RD (note the number zero).

Note: A good practice is for each administrator to have an account. Authentication against an LDAP server can be used when more accounts are required than are possible to configure on the management module.

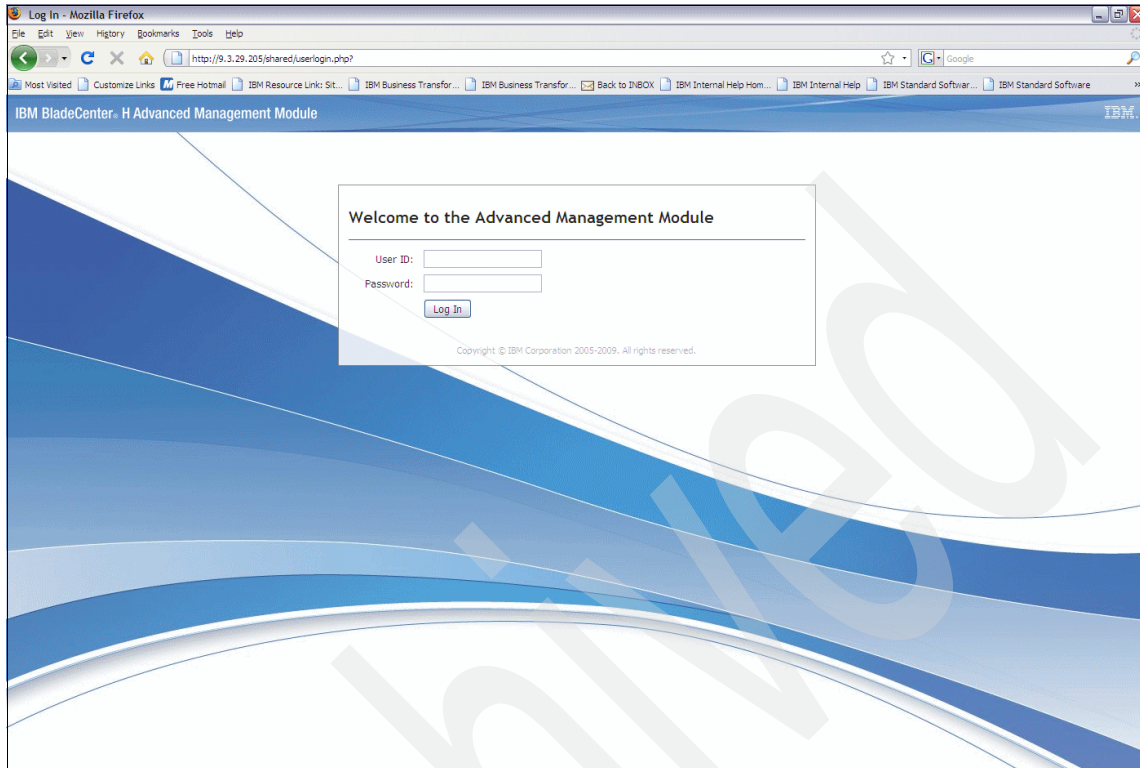


Figure A-2 AMM login panel

2. If prompted, select a timeout value, as shown in Figure A-3 on page 492, to indicate the amount of idle time that the session will be closed. Click **Continue**. Our example has been modified to show no timeout.

The screenshot shows a web interface with a blue header and footer. The main content area is white. At the top, it says "Welcome USERID". Below that, it says "Opening web session to SN#YK148077L1K5". Then, it shows the last login: "Wednesday April 22 2009 9:43 from 9.41.206.112 (Web)". A paragraph explains that the session will expire with inactivity and that data can be refreshed automatically. There are two labels: "Inactive session timeout value:" and "Use automatic refresh:". The "Inactive session timeout value:" has a dropdown menu open, showing options: "no timeout", "1 minute", "5 minutes", "10 minutes", "15 minutes", "20 minutes", and "no timeout" (repeated). Below the dropdown, there is a note: "Note: The inactive session timeout is a global value which will apply to all user logins. To access this global setting, navigate to the Account Security Management section of the Login Profiles page on the main UI." At the bottom, another note says: "Note: To ensure security and avoid login conflicts, always end your sessions using the 'Logout' option."

Welcome USERID

Opening web session to SN#YK148077L1K5

Last login: Wednesday April 22 2009 9:43 from 9.41.206.112 (Web)

Your session will expire with inactivity, and once timed out, the web page will be closed. Select the desired timeout parameter below. In addition, some pages can have their data automatically refreshed if the "Use automatic refresh" box is checked. Click "**Continue**" to start your session.

Inactive session timeout value:

Use automatic refresh: ☐

Note: The inactive session timeout is a global value which will apply to all user logins. To access this global setting, navigate to the **Account Security Management** section of the **Login Profiles** page on the main UI.

Note: To ensure security and avoid login conflicts, always end your sessions using the "**Logout**" option.

Figure A-3 Select timeout parameter

3. After successful login, you see the status page of the AMM. This page gives a short overview of the health of the chassis and the blades. Click **Remote Control** in the menu under Blade Tasks, as shown in Figure A-4 on page 493. Verify that no remote control session is in progress by observing the remote control status. The Refresh button refreshes the status.

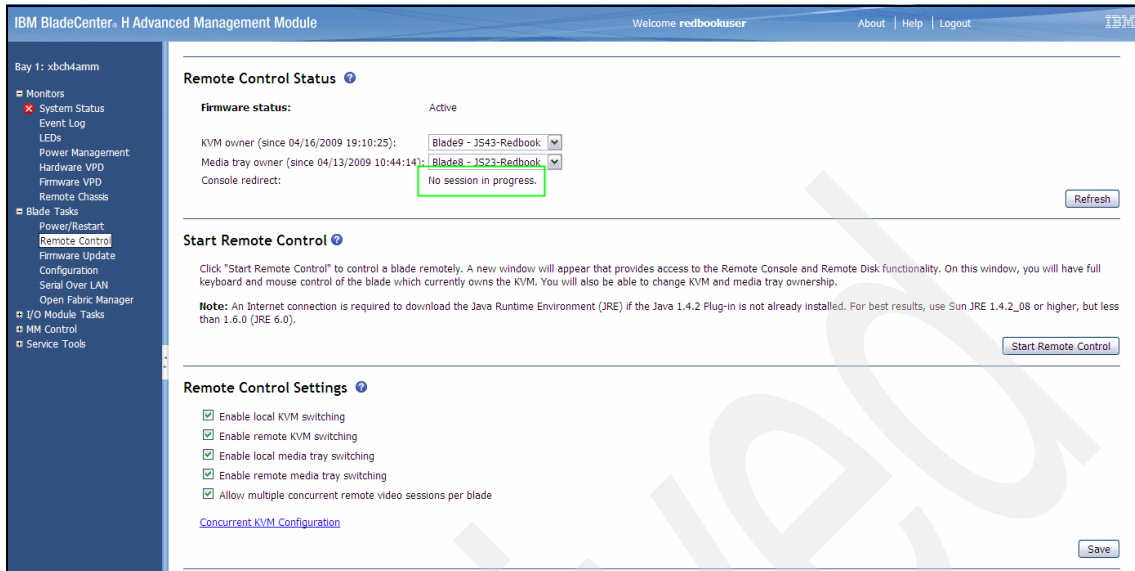


Figure A-4 Blade Remote Control options

4. Scroll down to the Start Remote Control section and click **Start Remote Control** as shown in Figure A-5.

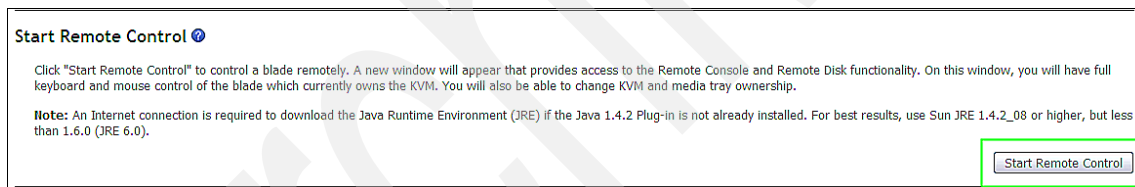


Figure A-5 Start remote control

The remote control Java applet will start in a new window. Be sure that no pop-up blockers are running; or configure the settings to allow the pop-up windows. It can take some time for the window to open and load the applet.

Figure A-6 on page 494 shows remote control with remote media and remote console. Use the KVM drop-down list to switch the console between the available blades. The red A shows that this is an analog video session. As mentioned earlier, only one analog session can be active. The physical console shows the same view as the remote console. Concurrent KVM (cKVM) is shown only with remote control and is represented by a red D. JS23/JS43 does not support cKVM.




Figure A-6 Remote control - remote console and remote disk

Serial Over LAN

Serial Over LAN (SOL) provides a means to manage servers remotely by using a command-line interface (CLI) over a Telnet or secure shell (SSH) connection. SOL is required to manage servers that do not have KVM support. SOL provides console redirection for both BIOS and the blade server operating system. The SOL feature redirects server serial-connection data over a LAN without requiring special cabling. The SOL connection enables blade servers to be managed from any remote location with network access. The advantages of SOL include:

- ▶ Remote administration without keyboard, video, or mouse (headless servers)
- ▶ Reduced cabling and without requiring a serial concentrator
- ▶ Standard Telnet interface, eliminating the requirement for special client software

The IBM BladeCenter management module CLIs provide access to the text-console command prompt on each blade server through an SOL connection, enabling the blade servers to be managed from a remote location.

In the BladeCenter environment, the integrated system management processor (ISMP) and network interface controller (NIC) on each blade server route the serial data from the blade server serial communications port to the network infrastructure of the BladeCenter unit, including an Ethernet-compatible I/O module that supports SOL communication. BladeCenter components are configured for SOL operation through the BladeCenter management module. The management module also acts as a proxy in the network infrastructure to couple a client running a Telnet or SSH session with the management module to an SOL session running on a blade server, enabling the Telnet or SSH client to interact with the serial port of the blade server over the network.

Because all SOL traffic is controlled by and routed through the management module, administrators can segregate the management traffic for the BladeCenter unit from the data traffic of the blade servers. To start an SOL connection with a blade server, first start a Telnet CLI session with the management module. When this Telnet or SSH CLI session is running, you can start a remote-console SOL session with any blade server in the BladeCenter unit that is set up and enabled for SOL operation.

You can establish up to 20 separate Web-interface, Telnet, or SSH sessions with a BladeCenter management module. For a BladeCenter unit, this step enables you to have 14 simultaneous SOL sessions active (one for each of up to 14 blade servers) with six additional CLI sessions available for BladeCenter unit management.

With a BladeCenter S unit you have six simultaneous SOL sessions active (one for each of up to six blade servers) with 14 additional CLI sessions available for BladeCenter unit management. If security is a concern, you can use Secure Shell (SSH) sessions, or connections made through the serial management port that is available on the Advanced Management Module, to establish secure Telnet CLI sessions with the BladeCenter management module before starting an SOL console-redirect session with a blade server.

SOL has the following requirements:

- ▶ An SOL-capable blade server such as the JS23/JS43
- ▶ An Ethernet switch module or Intelligent Pass-Thru Module is installed in bay 1 of a BladeCenter
- ▶ SOL is enabled for those blades that you want to connect to with SOL.
- ▶ The Ethernet switch module must be set up correctly.

For details about how to set up SOL, see the *BladeCenter Serial Over LAN Setup Guide*, which can be found at:

<http://www.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-54666&brandind=5000020>

This guide contains an example of how to establish a Telnet or SSH connection to the management module and then an SOL console.

Also, refer to “Powering on JS23 or JS43 using Telnet or SSH into AMM” on page 497.

System Management Services menu

The System Management Services Menu (SMS menu) is considered the main configuration interface of the IBM BladeCenter JS23/JS43. It is described in *Installation and User's Guide for IBM BladeCenter JS23/JS43*, (Type 7778-23X). This guide is delivered on CD with each JS23/JS43.

One of the first settings you might want to change is the console that you want to use. When the blade starts the first time, it uses the graphical console as the active console if you do not select the SOL console as the active console. The

SOL session cannot be used at this time to access the SMS menu to perform configuration tasks. To switch from the physical console to an SOL console you have to enter the SMS menu over the physical console or Remote Control. See “Graphical console” on page 488 for available consoles and how to use them.

To enter the SMS menu, the blade has to go through the POST. You have to power on the blade or make a restart to be able to enter the SMS menu. As mentioned before, the SMS menu is available only on the active console. The nonactive console can be used only to access the operating system.

Powering on JS23 or JS43 using AMM Web interface

To power on or restart the blade you may use the Advanced Management Module (AMM) Web interface or a Telnet or SSH session to the AMM, as follows:

1. Log on to your AMM Web interface. The default IP address is 192.168.70.125. The default account is USERID with password PASSWORD.

Note: Remember that the 0 in PASSWORD is a zero.

2. In the left menu, click **Power/Restart** under Blade Tasks as shown in Figure A-7.




Figure A-7 Power/Restart in the left menu of the AMM Web interface

The power status of the blade is visible.

3. Click the check box next to the blade. Then, in the Available actions section, select either **Power On Blade** to start the blade or **Restart Blade** to restart the blade. See Figure A-8 on page 497.

IBM BladeCenter. S Advanced Management Module

Welcome USERID

Bay 1: SN#YK148077L1K5

Monitors

System Status

Event Log

LEDs

Power Management

Hardware VPD

Firmware VPD

Remote Chassis

Blade Tasks

Power/Restart

Remote Control

Firmware Update

Configuration

Serial Over LAN

Open Fabric Manager

I/O Module Tasks

Storage Tasks

MM Control

Service Tools

AMM Service Data

Blade Service Data

AMM Status

Service Advisor

Blade Power / Restart

Blade selection and status

Click the checkboxes in the first column to select one or more blades; then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

	Bay	Name	Pwr	Local Pwr Control	Wake on LAN	Console Redirect
<input type="checkbox"/>	1	SN#YL30W7310006	Off	Enabled	N/A	
<input checked="" type="checkbox"/>	2	JS23-Redbook	Off	Enabled	N/A	
<input type="checkbox"/>	3	1GB-2GB Mixed	On	Enabled	N/A	
<input type="checkbox"/>	4	SN#ZK12HJ65L13F	Off	Enabled	On	
<input type="checkbox"/>	5	Free	Off	Enabled	N/A	
	6	No blade present				

Available actions

Power On Blade

Standard actions

Power On Blade

Power Off Blade

Shut Down OS and Power Off Blade

Restart Blade

Restart Blade with NMI

Enable Local Power Control

Disable Local Power Control

Enable Wake on LAN

Disable Wake on LAN

Restart Blade System Mgmt Processor

POWER specific actions

Restart Blade and clear NVRAM

Restart Blade with Diagnostic Boot

Restart Blade with Diagnostic Boot and Default Bootlist

Perform action

Figure A-8 Power/Restart blade options

Note: The **Restart Blade** option performs a power off and a power on of your selected blade and the operating system does not shut down properly. Use this option only when no operating system is running or the blade is in POST, SMS, or Open Firmware prompt.

The blade performs the requested action.

- Refresh this Web page to see a status change.

You may now use the console of your choice to work with the blade. Consoles are described in “Consoles of the IBM BladeCenter JS23 and JS43” on page 488.

Powering on JS23 or JS43 using Telnet or SSH into AMM

To start or power cycle the blade through a Telnet or SSH session, log on with a Telnet or SSH client on the AMM. The default IP address is 192.168.70.125. The

Appendix A. Consoles, SMS, and Open Firmware 497

default account is USERID with password PASSWORD. See Example A-1 on page 499.

Note: Remember that the 0 in PASSWORD is a zero.

Help is available through the command **help** or **help command-name**. Every command can be executed with one of these options to show the online help for the command, for example:

```
env -h
env -help
env ?
```

This example uses the command **env** to show available options to get help. Refer to the *Management Module Command-line Interface Reference Guide* that is located at:

<http://www.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-54667&brandind=5000020>

After logon, you may change the timeout of the Telnet or SSH session with the **telnetcfg** command. The command is issued to the current primary (active) management module.

To get a list of available targets in a BladeCenter chassis, use the command **list -l {number of levels}**. The output in Example A-1 on page 499 shows, at the beginning of the **list -l 2** command, that the first management module is the active one. The **telnetcfg** command uses this active AMM as target to extend the timeout of a Telnet session or switch it off. Use 0 to switch off the timeout or use any value between 1 and 4,294,967,295 seconds. During installation, a good idea is to switch off the timeout so that the console does not disconnect.

Every command that is executed has a target. This target is specified by the **-T** option. To shorten the commands that you work, especially when you have to work for a long time only with a single blade, the environment command **env** is helpful. Specify the target that will be used. In Example A-1 on page 499 we used the blade3 as target. The resulting command is **env -T blade[3]**.

The power state can be managed with the **power** command. To query the power status, use the **-state** parameter. Power the blade on or off with **-on** or **-off** and power cycle the blade with the **-cycle** parameter. The parameter **-c** establishes an SOL session as soon as the blade is SOL ready. As an alternative to the **-c** parameter, you may issue the **console** command to establish an SOL session.

Example: A-1 Use of the power command

```
login as: USERID
Using keyboard-interactive authentication.
password:

Hostname:          moon.ibm.com
Static IP address:  172.16.0.225
Burned-in MAC address: 00:14:5E:DF:AB:28
DHCP:              Disabled - Use static IP configuration.
Last login: Friday June 20 2008 17:37 from 9.3.4.254 (SSH)

system> list -l 2
system
      mm[1]      primary
      power[1]
      power[2]
      power[3]
      power[4]
      blower[1]
      blower[2]
      blower[3]
      blower[4]
      switch[1]
      switch[3]
      switch[4]
      blade[1]   SN#YL30W7310006
      blade[2]   JS23-Redbook
      blade[3]   1GB-2GB Mixed
      blade[4]   SN#ZK12HJ65L13F
mt[1]
      storage[1]
      storage[2]

system>
system> telnetcfg -t 0 -T mm[1]
OK
system> env -T blade[2]
OK
system:blade[2]> power -state
Off
system:blade[2]> power -on -c
Establishing an sol connection.  This may take a few minutes.
```

You may exit from the SOL session and return to the Advanced Management Module CLI by using the key combination ESC+[(left square bracket). This key combination can be defined in the AMM We Interface.

Choosing the active console after first power on

When the blade is powered on for the first time (as described in “Powering on JS23 or JS43 using AMM Web interface” on page 496 or in “Powering on JS23 or JS43 using Telnet or SSH into AMM” on page 497) you see, depending on the console you have chosen, a message to select this console as the active console.

Note: Ensure you select the active console session using the 0 (zero) in a relatively quick-time frame. This session times out and forces you to reboot the blade to re-establish the console session.

Figure A-9 shows an SOL console and Figure A-10 on page 501 shows a picture from the graphics console using Remote Control.

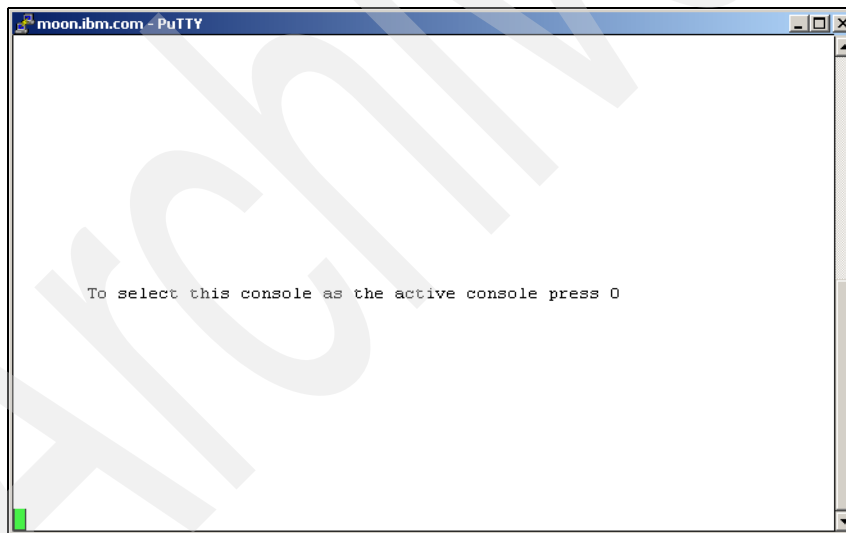


Figure A-9 Serial over a LAN console - select active console

Depending on the console that is open, you have to enter either:

- ▶ 0 (zero) to activate the SOL console
- ▶ 1 to activate the physical console

If you do not select either, the selection defaults to SOL. You have to enter the SMS menu over the physical console to change the active console, in this case as described in the next section.

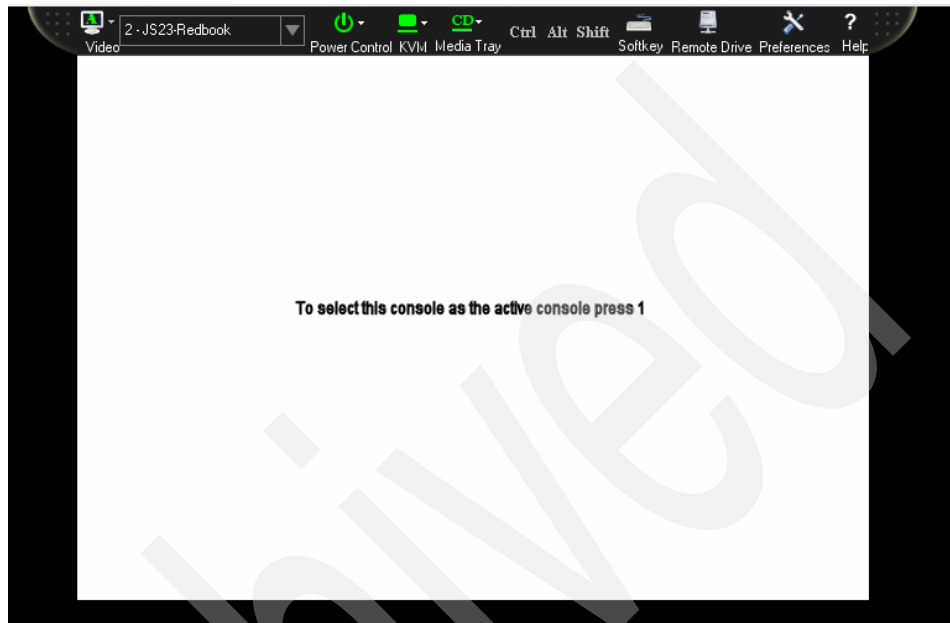


Figure A-10 Physical console shown with remote control - select active console

After a console is selected as the active console, either by the user or automatically, the system shows the power-on self-test (POST).

IBM BladeCenter JS23 and JS43 power-on self-test

As with previous JS2x blades, no system reference codes (SRC) are shown on a console during power-on self-test (POST). The SRCs can be found in the Advanced Management Module under Blade Service Data in the Service Tools section. Click the blade in the table from which you want to see the SRCs. Click **System Reference Codes**. The table with the SRCs can be refreshed by clicking **Refresh**.

The POST prints the words Memory, Keyboard, Network, SCSI, and Speaker as SMS loads. Press 1 to enter the SMS menu, shown in Figure A-11 on page 502. This is the only response you see during power on or restart of the blade on an active console.

During this printout, you are able to choose one of the following options from the menu:

- ▶ (1) Enter the System Maintenance Services Menu
- ▶ (5) Use Default Boot List
- ▶ (6) Use Stored Boot list
- ▶ (8) Enter Open Firmware Prompt

The stored boot list that is used to load the operating system is the default.

Press the number 1 to enter the SMS menu.

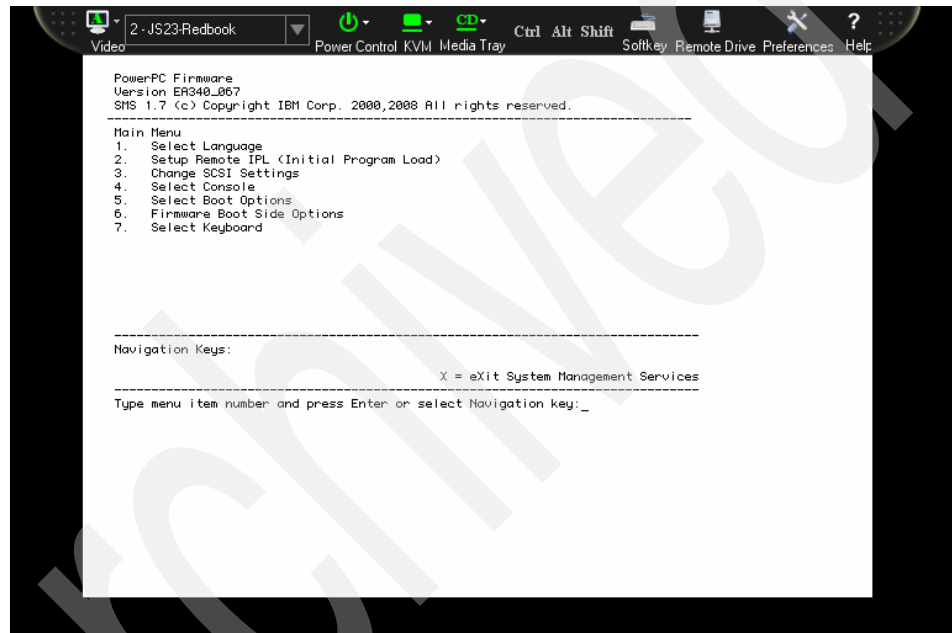


Figure A-11 JS23/JS43 SMS

Selecting active console using the SMS menu

When the blade is going through the POST, you may enter the System Maintenance Services (SMS) menu.

To change the current active console in the SMS menu, press 4 to select the console. See Example A-2 on page 503.

Example: A-2 SMS menu - main menu

Version EA330_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.

Main Menu

1. Select Language
2. Setup Remote IPL (Initial Program Load)
3. Change SCSI Settings
4. Select Console
5. Select Boot Options
6. Firmware Boot Side Options

Navigation Keys:

X = eXit System Management Services

Type menu item number and press Enter or select Navigation key:

After you select a console, you see (on the physical and on the SOL console) a message indicating that you have to enter 0 or 1 depending on the console you use. Figure A-10 on page 501 shows the message on an SOL console and Figure A-9 on page 500 shows the message on the physical console. Switch to the console you want to activate and enter the number as shown on this console. Use either 0 or 1. This console becomes the next active console. If you do not make a select, the system reverts to the original setting.

All other options are similar to other blades with Power CPU or System p servers.

Open Firmware interface

The Open Firmware prompt is a powerful interface into the system firmware of the JS23/JS43. Open Firmware is described in the IEEE 1275 standard. More general information about Open Firmware can be found at:

<http://www.firmworks.com>

The implementation of Open Firmware that is used by IBM on the JS23/JS43 to build the system firmware might have its own extensions and might also not implement each feature or function. The information at this Web site might not apply fully to the IBM BladeCenter JS23/JS43.

In certain situations, using the Open Firmware prompt to change system settings or troubleshoot problems is helpful. Also, the SMS menu is considered the main

configuration interface. We explain several command that help boot SAN with QLogic and Emulex host bus adapters configured. The commands can be used to perform the following tasks:

1. Identify FCode and firmware level.
2. Identify either the World Wide Port Name, World Wide Node Name, or both.
3. Set the connection type.
4. Set the transfer rate.
5. Query available targets.

This appendix contains a section about the QLogic host bus adapter and a section about the Emulex host bus adapter. We start with a description of how to get access to the Open Firmware prompt.

Getting access to the firmware prompt

Use a console of JS23/JS43 and power on or restart the blade. See “Graphical console” on page 488 or “Serial Over LAN” on page 494 about available consoles. “System Management Services menu” on page 495 describes how to power on a blade.

When the blade shows the POST, enter the number 8 to get access to the Open Firmware prompt. During POST the words Memory, Keyboard, Network, SCSI, and Speaker are written to the console, as shown in Example A-3. When this happens, enter 8 before the word Speaker appears. The JS23/JS43 does not show the service request codes on the console, as JS20 or JS21 do. These SRCs are shown in the Web interface of the Advanced Management Module. The content of Example A-3 is the only information shown during POST.

Example: A-3 IBM BladeCenter JS23/JS43 POST

```
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
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

After entering the Open Firmware prompt, you see the command prompt shown in Example A-4.

 $0 >$

Boot settings are stored in the NVRAM of the system. The Open Firmware allows you to verify them with the `printenv` command. The command accepts an optional variable name to display only the content of the given variable. To display your boot device you may use the command `printenv boot-device` as shown in Example A-5. You see an iSCSI boot device entry. The example shows the following information:

ciaddr	iSCSI initiator IP address
giaddr	Gateway IP address
subnet-mask	Subnet mask of the initiator IP address
itname	IQN from the initiator
iport	iSCSI port
ilun	LUN to boot from on the iSCSI target
iname	IQN of the iSCSI target device
siaddr	iSCSI target portal IP address

```
0 > printenv boot-device
----- Partition: common ----- Signature: 0x70 -----
boot-device
/pci@800000020000203/iscsi-toe@1,1:iscsi,ciaddr=172.16.1.44,giaddr=172.16.1.1,subn
et-mask=255.255.255.0,itname=iqn.2008-06.com.ibm.blade5port1,iport=3260,ilun=0,ina
me=iqn.2008-06.com.ibm:storage1.disk1,siaddr=172.16.1.195,2
    ok
0 >
```

The **show-devs** command is used to display all devices in a device path known by the system firmware. A part of the device path might be given as parameter to show only child devices from this device path. Without a given device path, the command returns a listing of all known devices. To figure out the device path of the fibre channel host bus adapter, the full device tree is required. Enter the command **show-devs** on the Open Firmware prompt as shown in Example A-6. The output of the command is shortened to show only the important part of information for the explanation in this section.

Example: A-6 show-devs example output

```
0 > show-devs
00000208dda0: /ibm,serial
00000208eb98: /chosen
.
.
.
00000222fe48: /pci@800000020000202
0000022acb78: /display@1
00000223a0f8: /pci@800000020000204
0000022affe0: /fibre-channel@0
0000022c3da0: /disk
0000022c4790: /fibre-channel@0,1
0000022d8550: /disk
000002244440: /pci@800000020000205
0000022d8f40: /pci@0
0000022daea0: /ethernet@4
0000022eb198: /ethernet@4,1
ok
0 >
```

Look in the device path for the fiber-channel adapters. In a JS23/JS43, there are always dual port host bus adapters used, represented by the following two entries:

Port 1 /pci@800000020000204/fibre-channel@0
Port 2 /pci@800000020000204/fibre-channel@0,1

This information is required in the next steps to change the topology and the link speed of the host bus adapter.

QLogic host bus adapter

This section describes how to

1. Retrieve the World Wide Node Name.
2. Identify FCode and the firmware level.
3. Set the transfer rate.
4. Set the connection mode.
5. Query available targets and LUNs.

The examples in this section were created using a CFFh combo card with the firmware 4.00.24 and FCode 1.25.

Identify your fiber channel host bus adapter as described in Example A-6 on page 506. The device tree in your system might differ from the example shown here. With this information you can build the command to select the device. Enter the following command to select the first host adapter port:

```
" /pci@800000020000204/fibre-channel@0" select-dev
```

The second HBA port is selected with the following command:

```
" /pci@800000020000204/fibre-channel@0,1" select-dev
```

Note that a space exists between the first quotation mark (") and forward slash (/). Example A-7 shows the output of this command. The link of the adapter port comes up and the adapter logs in to the switch. You can now verify the name server of a connected SAN Switch about the World Wide Node and Port Name.

Example: A-7 Select the fiber channel host bus adapter port

```
0 > " /pci@800000020000204/fibre-channel@0" select-dev QLogic QMI3472
Host Adapter Driver(IBM): 1.25 12/11/06
Wait for link up - |/-\|/-\|
Firmware version 4.00.24
ok
0 >
```

To identify the World Wide Node Name you may use the Advance Management module. Alternatively, you can use the command **my-wwn** to retrieve the World Wide Port Name of the host adapter port. To display the retrieved World Wide Port Name enter a dot followed by Enter. Example A-8 on page 508 shows the output of this command.

Example: A-8 Display World Wide Port Name of a QLogic host bus adapter port

```
0 > my-wwn ok
1 > . 2100001b32005216 ok
0 >
```

Firmware version and FCode level of the HBA can be shown with the **version** command. Example A-9 shows the output of this command. The FCode version is on the current level at the time of this writing, but the firmware version can be upgraded to 4.00.27.

Example: A-9 Firmware and FCode versions

```
0 > version QLogic QMI3472 Host Adapter Driver(IBM): 1.25 12/11/06
Firmware version 4.00.24
ok
0 >
```

If an Optical Pass-Thru Module is used, you must change the transfer rate that is set, per default, to auto-negotiation on the 4 GB host bus adapter to a fixed value of 2 GB. The Optical Pass-Thru Module can handle transfer rates only up to 2 GB. auto-negotiation does not work with 4 GB host bus adapters. To change the transfer rate, verify the current settings of the HBA first. Use the command **show-settings**, as shown in Example A-10.

Example: A-10 Settings of the QLogic HBA

```
0 > show-settings QLogic QMI3472 Host Adapter Driver(IBM): 1.25 12/11/06
Firmware version 4.00.24
Serial#
node-wwn 2001001b 32205216
port-wwn 2101001b 32205216
Current HBA Connection Mode: 2 - Loop preferred, otherwise point-to-point
Current HBA Data Rate: Auto-negotiated
Current Fcode Mode: qlc
ok
0 >
```

The **show-settings** command also shows the firmware and FCode versions, node and port names, and data rate. The adapter is currently configured for auto-negotiation. To change the data rate to a fixed data rate of 2 Gb, use the command **set-data-rate** as shown in Example A-11 on page 509.

Example: A-11 Change the data rate of a QLogic HBA

```
0 > set-data-rate Current HBA Data Rate: Auto-negotiated
Do you want to change it? (y/n)
Choose HBA Data Rate:
0 - 1 Gigabit
1 - 2 Gigabit
2 - Auto-negotiated
3 - 4 Gigabit
enter: 1Current HBA Data Rate: 2 Gigabit
ok
0 >
```

Example A-11 shows that the data rate is changed from auto-negotiation to a fixed value of 2 Gb. Enter the number in front of the value that you want defined here. The example uses 1. The changed value will be returned.

Depending on your fiber channel targets and the connectivity that you use to connect to them, you might want to change the connection type to loop or to point-to-point. Use the command **set-connection-mode** to do the change, as shown in Example A-12. The command returns the current setting and lets you change to a new one. The possible options are shown. Select the corresponding number and press Enter.

Example: A-12 Change connection mode of an QLogic HBA

```
0 > set-connection-mode Current HBA Connection Mode: 2 - Loop preferred, otherwise
point-to-point
Do you want to change it? (y/n)
Choose HBA Connection Mode:
0 - Loop Only
1 - Point-to-point only
2 - Loop preferred, otherwise point-to-point
enter: 1Current HBA Connection Mode: 1 - Point-to-point only
ok
0 >
```

The last command that we describe for the QLogic host bus adapter is the **show-children** command, which can be used to show available targets and LUNs to the HBA. Example A-13 on page 510 shows output from one of our Virtual I/O servers. A DS4800 with 7 LUNs is shown. The DS4800 is connected to port 0 of the Switch and JS23/JS43 to port 3.

Example: A-13 List all available devices on the selected QLogic HBA

```
0 > show-children Adapter portID - 610300
***** Fabric Attached Devices *****
  Dev# 0 (0 )   PortID 10000   Port WWN 203200a0b811a662
    LUN 2      DISK  IBM      1815      FASTT 0914
    LUN 3      DISK  IBM      1815      FASTT 0914
    LUN 4      DISK  IBM      1815      FASTT 0914
    LUN 5      DISK  IBM      1815      FASTT 0914
    LUN 6      DISK  IBM      1815      FASTT 0914
    LUN 7      DISK  IBM      1815      FASTT 0914
    LUN 8      DISK  IBM      1815      FASTT 0914

ok
0 >
```

Remember that the described commands require that you have an HBA port selected and that they have effect only on the selected HBA port. You have to perform the necessary actions on both HBA ports.

To leave the Open Firmware prompt and restart the blade, use the command **reset-all**.

When no changes are made, the boot process can be started by leaving the Open Firmware prompt with the commands as shown in Example A-14.

Example: A-14 Leave Open Firmware prompt

```
1 > dev /packages/gui
1 > obe
```

Emulex host bus adapter

This section describes how to

1. Retrieve the World Wide Node Name.
2. Identify the FCode level.
3. Set the link speed.
4. Set the connection mode.

The examples in this section were created using an Emulex CFFv with the FCode 3.10.a0.

Identify your fiber channel host bus adapter as described in Example A-6 on page 506. The device tree in your system can differ from the example shown

here. With this information you can build the command to select the device. Enter the following command to select the first host adapter port:

```
" /pci@800000020000203/fibre-channel@0" select-dev
```

The second HBA port is selected with the following command:

```
" /pci@800000020000203/fibre-channel@0,1" select-dev
```

Note that a space exists between the first quotation mark (") and forward slash (/). Example A-15 shows the output of this command. The link of the adapter port comes up and the adapter logs in to the switch. Now, you are able to verify the name server of a connected SAN Switch about the World Wide Node and Port Name.

Example: A-15 Select fiber channel port

```
0 > " /pci@800000020000203/fibre-channel@1" select-dev ok
0 >
```

World Wide Node and Port Name of the HBA port are shown by the command **host-wwpn/wwnn**. The same information can be retrieved through the Advanced Management Module Web interface under Hardware VPD. Example A-16 shows the WWPN and WWNN of the first port of an Emulex CFFv HBA.

Example: A-16 Display the World Wide Node and Port Name of an Emulex CFFv HBA

```
0 > host-wwpn/wwnn Host_WWPN 10000000 c9660936
Host_WWNN 20000000 c9660936
ok
0 >
```

The installed FCode level on the HBA can be shown with the command **check-vpd** or **.fcode**, as shown in Example A-17.

Example: A-17 Display FCode version of an Emulex CFFv HBA

```
0 > check-vpd
!!! LP1105-BCv Fcode, Copyright (c) 2000-2008 Emulex !!! Version 3.10a0

ok
0 >
0> .fcode Fcode driver version 3.10a0
ok
0>
```

To display the current link speed, use the command **link-speed** as shown in Example A-18 on page 512. The command does not return the current link speed setting.

Example: A-18 Display actual link speed of an Emulex CFFv HBA

```
0 > link-speed
    Current ....

    Link Speed -- 2 Gb/s ok
0 >
```

Link speed or data rate of the Emulex HBA can be set with the command **set-link-speed**, as shown in Example A-19. The command shows the current link speed and the current setting.

Example: A-19 Set link speed of an Emulex CFFv HBA

```
0 > set-link-speed
    Current ....

    Link Speed -- 2 Gb/s
    Link Speed Selected -- Auto Select
```

- 0. Auto Select Link Speed (Default)
- 1. 1 Gb/s Link Speed -- Only
- 2. 2 Gb/s Link Speed -- Only
- 4. 4 Gb/s Link Speed -- Only

Enter <x> to QUIT

Enter a Selection:

Enter the number of your choice and press Enter, as shown in Example A-20 on page 513. The NVRAM of the HBA will be updated.

Example: A-20 Changed link speed in NVRAM of the Emulex CFFv HBA

Enter a Selection: 2

Flash data structure updated.

Signature	4e45504f
Valid_flag	0000004a
Host_did	00000000
Enable_flag	00000005
SFS_Support	00000000
Topology_flag	00000000
Link_Speed_flag	00000002
Diag_Switch	00000000
POST-Linkup	00000000
Boot_id	00000000
Lnk_timer	0000000f
Plogi-timer	00000000
LUN (1 byte)	00000000
DID	00000000
WWPN	0000.0000.0000.0000
LUN (8 bytes)	0000.0000.0000.0000

*** Type reset-all to update. ***

ok

0 >

The connection type or topology setting can be shown with the **.topology** command. This command reads the NVRAM of the Emulex HBA and displays the value as human-readable text. To change the topology setting, use one of the available commands, such as **set-auto-fcal**, **set-auto-ptp**, **set-fc-al**, or **set-ptp**. The command **set-auto-fcal** sets loop preferred, otherwise point-to-point. The command **set-auto-ptp** sets point-to-point, otherwise loop. The two other commands set the connection type to point-to-point or loop. The default setting can be defined with **set-default-mode**. A restart of the blade is required. Example A-21 shows that the topology is set to Point to Point. The set commands return nothing.

Example: A-21 Display connection topology of an Emulex CFFv HBA

1 > .topology Point to Point - Current Mode

Manual Topology

ok

1 >

Remember, the described commands require that you have an HBA port selected and that the commands affect only on the selected HBA port. You have to perform the necessary actions on both HBA ports.

To leave the Open Firmware prompt and restart the blade, use the command **reset-all**. When no changes are made, the boot process can be started by leaving the Open Firmware prompt with the commands shown in Example A-22.

Example: A-22 Leave Open Firmware prompt

```
1 > dev /packages/gui  
1 > obe
```

SUSE Linux Enterprise Server AutoYaST

This appendix describes the SUSE AutoYaST tool to perform automated installations of SUSE Linux Enterprise Server 11.

This appendix contains the following topics:

- ▶ “AutoYaST introduction” on page 516
- ▶ “AutoYaST profile creation methods” on page 516
- ▶ “Creating an AutoYaST profile with YaST Control Center” on page 516

AutoYaST introduction

The AutoYaST configuration tool allows a system administrator to install SUSE Linux Enterprise Server (SLES) on a large number of systems in parallel by using an automated process. The AutoYaST profile is a file written in the Extensible Markup Language (XML). It contains responses to all system configuration questions typically asked during a manual installation. This file is configurable to accommodate the installation of systems with homogeneous and heterogeneous hardware.

This appendix covers AutoYaST profile creation for SUSE Linux Enterprise Server 11.

Note: The procedures covered here are not applicable for previous SLES releases because the XML layouts are different.

AutoYaST profile creation methods

The three methods to create an AutoYaST profile for systems using *identical* hardware are:

- ▶ Clone the installation configuration information from a reference machine that was installed manually. This is discussed in 9.3, “Linux network installation (detailed)” on page 366.
- ▶ Use the YaST Control Center to create and modify the AutoYaST profile.
- ▶ Use an XML editor to create an AutoYaST profile.

We cover the steps for the second method in this appendix.

Creating an AutoYaST profile with YaST Control Center

In this section, we use the AutoYaST configuration tool available in the YaST Control Center with a graphical interface to demonstrate how to create a basic XML file.

Note: This YaST tool can run in graphical or text mode. A mouse can navigate through the graphical version of the tool; the text mode version requires Tab, Enter, Up/Down Arrow, and Spacebar keys to navigate. Otherwise, no difference exists between the two modes and the same configuration options; both result in the same XML file.

Many optional settings are available, but several are mandatory settings or dependencies. We cannot cover every possible configuration option, so we try to provide a general overview to help you become familiar enough with the tool to navigate on your own.

Starting the YaST graphical interface

To start the interface:

1. With a root user ID, run SSH with X11 forwarding enabled into a system running SLES 11. For example:

```
ssh -X root@9.3.20.18
```

Consideration: You must issue this command on a system with a running X Server to use the graphical interface.

2. On the command line type:

```
yast2
```

The YaST Control Center window opens, as shown in Figure B-1 on page 518.

If you are using text mode, see Figure B-2 on page 519



Figure B-1 YaST Control Center in graphics mode

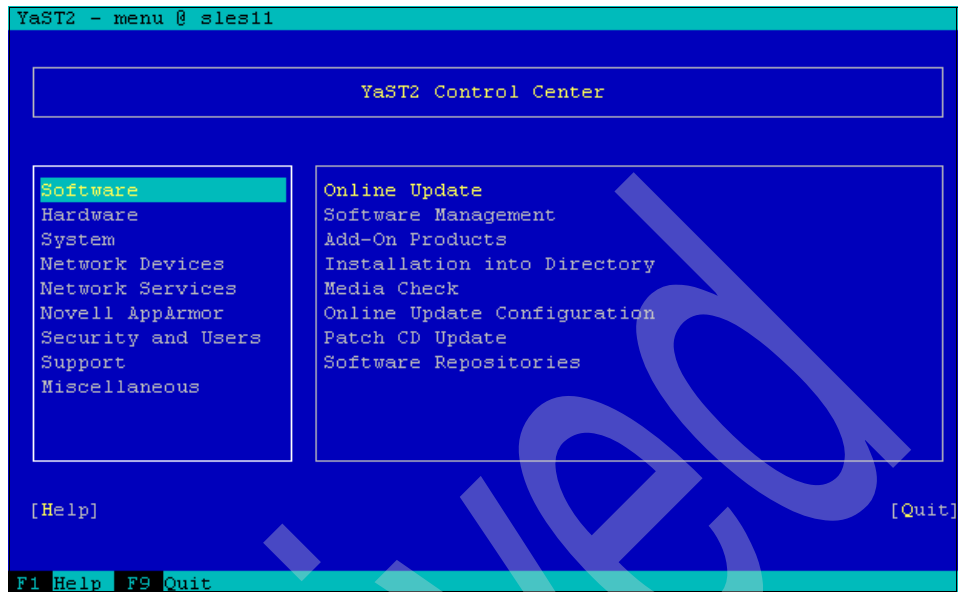


Figure B-2 YaST Control Center in text mode

Navigating the YaST graphical interface

To navigate:

1. Start the YaST application, which opens a window as shown in Figure B-3. In the Miscellaneous section, click the **Autoinstallation** applet.



Figure B-3 Selecting the Autoinstallation option

The main AutoYaST configuration window opens as shown in Figure B-4.

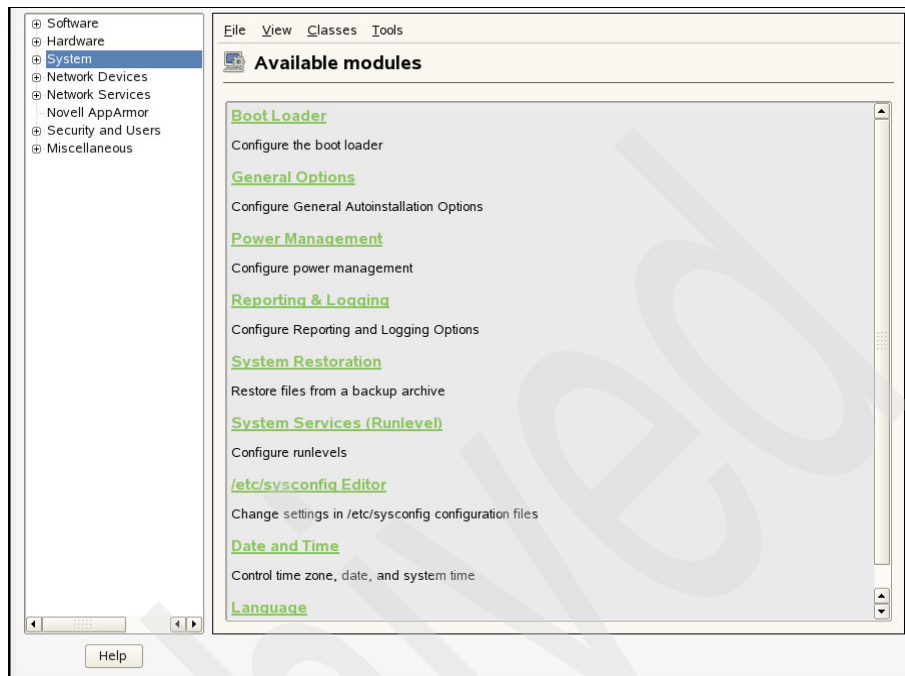


Figure B-4 Main AutoYaST menu (SLES 11)

2. Clone the configuration of the installation server by selecting **Tools** → **Create Reference Profile**, as shown in Figure B-5.




Figure B-5 Create a reference profile

3. A second window opens, as shown in Figure B-6. In addition to the default resources such as boot loader, partitioning, and software selection, you may add other aspects of your system to the profile by checking the specific items in the Select Additional Resources section. When you are ready, click **Create** so YaST can collect the system information and create the AutoYaST profile.

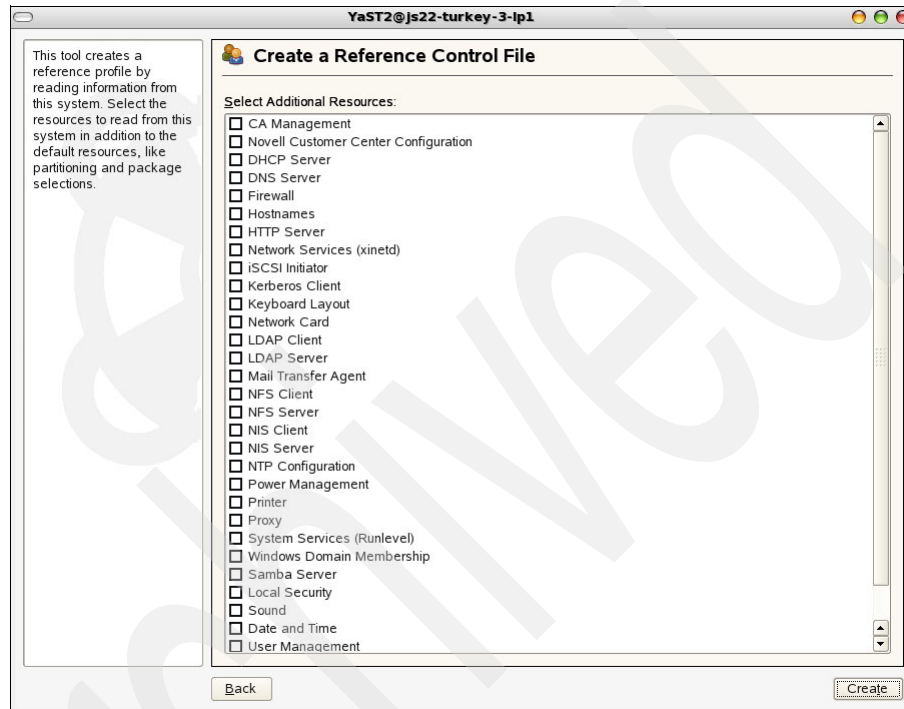


Figure B-6 Selecting additional resources

4. If the profile is complete and meets your requirements, select **File** → **Save** and enter a file name such as:

sles11_autoinst.xml

5. You may also adjust certain options that are provided on the left side of the AutoYaST main menu. Options include:

- Software

Use these options to select and configure the Online Update and Package Selection sections. Figure B-7 on page 524 shows that we selected the **Software** → **Package Selection**. We chose the minimum software configuration to save time during installation.

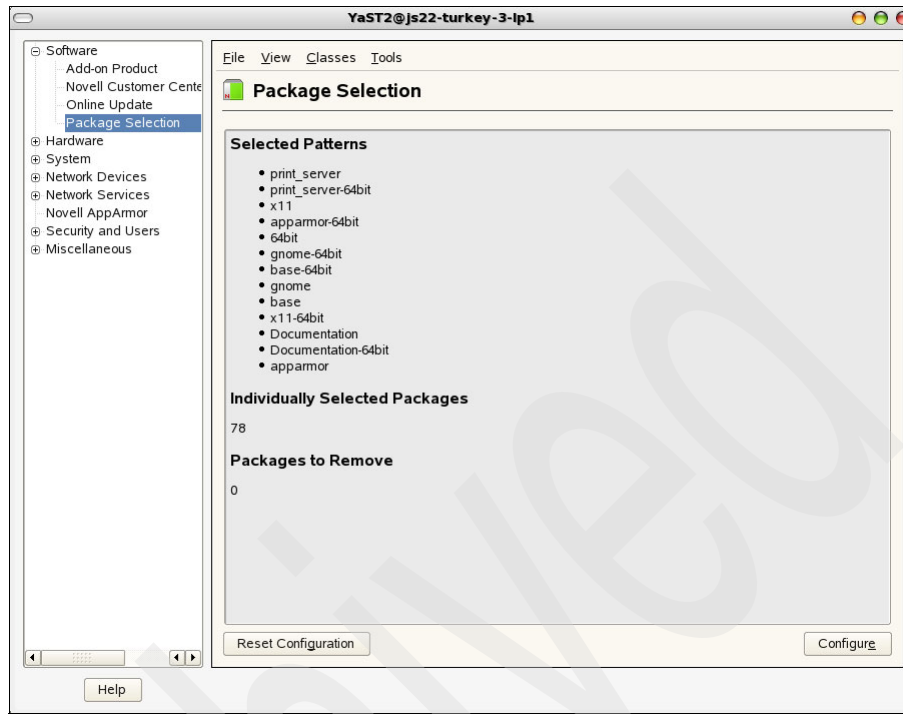


Figure B-7 AutoYaST software selection

– Hardware

This option configures partitioning, sound, printer, and graphics card and monitor, if necessary. The Partitioning settings are critical for this configuration to work, so verify that they match your hard disk environment and that each partition meets the minimum SUSE partition size requirements.

– System

This option sets the general system information such as language configuration, time zone, other locale-related settings, logging, and run-level information in this option. The most important configuration is Boot Loader, which is cloned from the `/etc/lilo.conf` file of the system running the YaST tool so it might require editing. Edit it by selecting **Configure** → **Edit**. Click **Finish** when you are done.

– Network Devices

This option sets the network adapter information. You can set network module information and IP details here, as follows:

- v. Click **Configure** to open the menu.
- vi. The Network cards configuration main menu opens. Select **Configure** to add interfaces.
- vii. Select **Traditional Method with ifup** → **Next**.
- viii. Remove any static IP configurations on the next panel and click **Add**. Certain selections are already configured, such as Device Type: Ethernet. Specify `ehea`, for example, as the module name for the adapter and click **Next**.
- ix. In the host name and name server section, choose **DHCP** for the hostname and domain name (global) and also choose **DHCP** for name servers and the domain search list.
- x. Click **OK** → **Next**. Interface `eth0` is ready now.

To create interface `eth1`, repeat the steps. However, the DHCP settings are automatically taken from the `eth0` configuration and the interface name automatically changes to `eth1`.

– Network Services

This option configures network clients and daemons using network services. You may choose from more than 15 daemons; all are optional.

– Security and Users

This option creates users and configures security policies.

To allow SSH through the enabled firewall:

- i. Click **Firewall**.
- ii. Select **Start-Up** and change the firewall setting to **When Booting**.
- iii. Select **Interfaces** and change `eth0` to **External Zone**.
- iv. Change `eth1` to **Internal Zone**.
- v. In the Allowed Services section, select **Secure Shell** to enable Port 22 connections through the firewall. Disable **Protect from Internal Zone**.

To set the password requirements:

- i. Select **Local Security** from the Security and Users menu.
- ii. Click **Configure**.
- iii. Adapt the Security settings. The default is seven characters for the maximum password length.

A mandatory step is to define the root user password to log in a root through SSH. To set the password:

- i. Select **User Management** from the left menu.
- ii. Click **Configure**.
- iii. Select **Users**, select the **root** row, and change the filter by clicking **Set Filter** → **System Users**; see in Figure B-8.




Figure B-8 Configure the root user

- iv. Select the **root** row again and click **Edit**.
- v. Add the root user password. This password is saved as encrypted in the XML file. Click **Accept**.
- vi. Click **Finish** to return to the AutoYaST main menu.

— Misc

This option enables you to add complete configuration files, or to add special scripts to run before and after the installation.

6. Remember to save the edits with **File** → **Save**. See the created file in Example B-1 on page 527.

Example: B-1 Part of newly created XML file

```
<?xml version="1.0"?>
<!DOCTYPE profile>
<profile xmlns="http://www.suse.com/1.0/yast2ns"
xmlns:config="http://www.suse.com/1.0/configns">
  <add-on>
    <add_on_products config:type="list"/>
  </add-on>
  <bootloader>
    <global>
      <activate>true</activate>
      <boot_chrp_custom>/dev/sda1</boot_chrp_custom>
      <default>Linux</default>
      <lines_cache_id>2</lines_cache_id>
      <timeout config:type="integer">80</timeout>
    </global>
  <uid>104</uid>
  .
  .
  .
  .
  .
    <user_password>*</user_password>
    <username>pulse</username>
  </user>
</users>
<x11>
  <color_depth config:type="integer">4</color_depth>
  <display_manager>gdm</display_manager>
  <enable_3d config:type="boolean">>false</enable_3d>
  <monitor>
    <display>
      <max_hsync config:type="integer">42</max_hsync>
      <max_vsync config:type="integer">72</max_vsync>
      <min_hsync config:type="integer">30</min_hsync>
      <min_vsync config:type="integer">50</min_vsync>
    </display>
    <monitor_device>Unknown</monitor_device>
    <monitor_vendor>Unknown</monitor_vendor>
  </monitor>
  <resolution>640x480 (VGA)</resolution>
  <window_manager>gnome</window_manager>
</x11>
</profile>
```

Additional Linux installation configuration options

This appendix describes several other options to install Linux natively or on an LPAR.

This appendix contains the following topics:

- ▶ “Basic preparations for a Linux network installation” on page 530
- ▶ “Virtual optical device setup and installation” on page 538

Basic preparations for a Linux network installation

This section provides all the basic information to set up services for a Linux network installation. In principle, this is not bound to a specific operating system or distribution that runs on the infrastructure server to provide the necessary services. Nevertheless, all descriptions in this section are based on general Linux services, commands, and parameters. We presume that the required files for all the services are already installed and that all the commands are issued with superuser rights.

Services required for installing Linux using the network

Performing a network installation requires the following services:

- ▶ A running Bootstrap Protocol (BOOTP) service or a Dynamic Host Configuration Protocol (DHCP) service that includes BOOTP support to configure the network interface of a BladeCenter JS23.
- ▶ A running Trivial File Transfer Protocol (TFTP) service to serve the boot image to a BladeCenter JS23.
- ▶ The setting up of *one* of the following services to provide the installation packages for a network installation after the boot image is loaded:
 - File Transfer Protocol (FTP)
 - Hypertext Transfer Protocol (HTTP)
 - Network File System (NFS)

The next sections describe how to configure BOOTP or DHCP, TFTP, and NFS services.

Note: If a firewall is running on the installation server, update the settings to allow traffic for the installation protocol.

Configuring a BOOTP or DHCP service

DHCP is an extension to the original BOOTP specification. As a result, you can use DHCP to provide the BOOTP information for booting by using the network. The standard DHCP daemon is named **dhcpcd**, but other DHCP daemons exist.

Note: The directory you use for the configuration files depends on the distribution. The following directories are possible examples:

- ▶ `/etc/`
- ▶ `/etc/sysconfig/`
- ▶ `/etc/default/`
- ▶ `/etc/xinet.d/` (eXtended InterNET daemon configuration files)

The examples in this appendix use the most common directories. In general, the name of a configuration or script file is related to the name of the installed package. For example, if a DHCP daemon is named `dhcpcd3-server`, you can find:

- ▶ The configuration is in `/etc/dhcpcd3-server.conf` and `/etc/sysconfig/dhcpcd3-server`
- ▶ The start/stop script is in `/etc/init.d/dhcp3-server`.

The standard DHCP daemon is configured through two files:

- ▶ `/etc/sysconfig/dhcpd`
Stores the basic configuration.
- ▶ `/etc/dhcpd.conf`
Contains configuration information for each registered client.

For a running service, the configuration actually used in most cases is copied in a subdirectory of `/var/`. See Example C-1 on page 532 for a simple client configuration stored in `dhcpd.conf`.

Note: Keep in mind that Example C-1 contains environment-specific Internet Protocol (IP) and Media Access Control (MAC) address information. One way to learn the MAC address of a JS23 BladeCenter is to use the BladeCenter management module. Select **Monitors** → **Hardware VPD** from the left-side options, select your blade bay, and then select the **Ports** tab on the right side. For more information about how to customize `dhcpd.conf`, see the man pages of the `dhcpd` service and use the `man dhcpd` at the command prompt.

Example: C-1 Sample dhcpd.conf

```
ddns-update-style none;
allow booting;
allow bootp;
always-reply-rfc1048 true;

shared-network TEST {
    option routers 172.16.1.1;
    subnet 172.16.1.0 netmask 255.255.255.0 {
        option broadcast-address 172.16.1.255;
        range dynamic-bootp 172.16.1.68 172.16.1.80;
        default-lease-time 444;
        next-server 172.16.1.197;
    }

    host JS23 {
        hardware ethernet 00:1a:64:44:21:53;
        fixed-address 172.16.1.79;
        filename "install";
    }
}
```

You can find the start and stop scripts of Linux services in the `/etc/init.d/` directory. To start the standard DHCP daemon, use the `/etc/init.d/dhcpd start` command. To restart the DHCP daemon, use the `/etc/init.d/dhcpd restart` command.

The filename `"install"`; lines in the `dhcpd.conf` file point to a directory in the `/tftpboot` directory. We show how to configure it in the next section.

Tip for Linux beginners: The following tasks help you double-check or troubleshoot a configuration in general:

1. To trace messages of running services, type the following command to get the last 10 messages and auto-update if new messages exist:
`tail -f -n 10 /var/log/messages`
2. Connect to a running service with a local client, remote client, or both these clients and try to receive the data that you want.
3. Make sure a changed configuration is activated by restarting a service *directly* after editing, for example:
 - a. `vi /etc/dhcpd.conf`
 - b. `/etc/init.d/dhcpd restart`

Configuring a Trivial File Transfer Protocol service

You can use the TFTP to provide a bootable image during a network installation. Several implementations of TFTP daemons are available. The standard TFTP daemon is named **tftpd**. In general, the **xinetd** or **inetd** super daemons are used to create a TFTP daemon. You can also run a TFTP daemon without one of the super daemons.

Install the **tftp-server** rpm if it was not installed by default, create a **/tftpboot** directory if does not exist, and set the **tftp** service to **disable=no**. For example, on SLES11:

1. Type **yast** on the command line to start the YaST tool.
2. Scroll down to **Network Services**.
3. Select **TFTP Server** on the right side and press Enter.
4. In the new window, select **Enable** and create **/tftpboot** in the Boot Image Directory section as shown in Figure C-1.

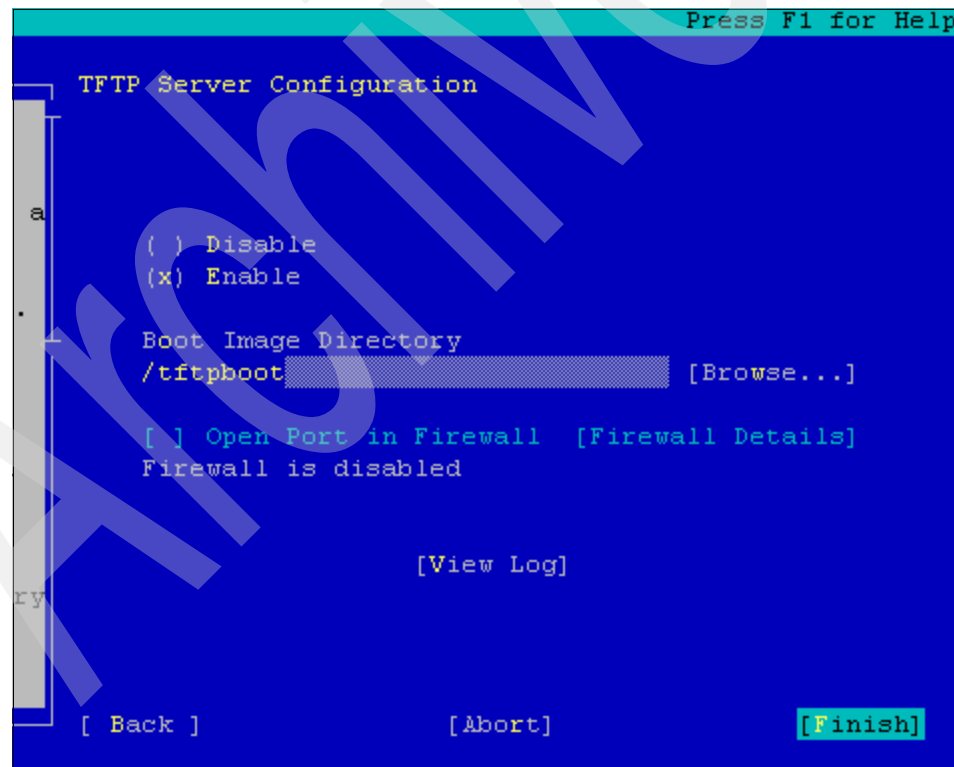


Figure C-1 Configure a TFTP server in SLES11

5. Finally, scroll down to **[Finish]** and press the Enter key.

Example C-2 shows a TFTP daemon configuration for xinetd that is stored in `/etc/xinet.d/tftpd`.

Example: C-2 Configuring a TFTP daemon in the `/etc/xinet.d/tftp` file on SLES11

```
# default: off
# description: tftp service is provided primarily for booting or when a
\
#      router need an upgrade. Most sites run this only on machines
acting as
#      "boot servers".
service tftp
{
    socket_type      = dgram
    protocol         = udp
    wait             = yes
    user             = root
    server           = /usr/sbin/in.tftpd
    server_args      = -s /tftpboot
    disable          = no
}
```

Tip: You can edit the `etc/xinet.d/tftp` file by using a text editor also.

Red Hat Enterprise Linux and SLES create TFTP from the xinetd daemon. Restart the xinetd service after the tftp configuration is complete, as follows:

```
venus:/ # service xinetd restart
Shutting down xinetd: done
Starting INET services. (xinetd) done
```

The next step is to copy the bootable installation kernel into the `/tftpboot` directory. The bootable kernel must match the Linux distribution that you install. It must also match the *file name* listed in the `dhcpd.conf`, as shown in Example C-1 on page 532. The process is slightly different for Red Hat Enterprise Linux and SLES, so we document both separately here.

Copying the SLES11 install kernel

To copy the SLES11 install kernel:

1. Mount the SLES11 DVD1 on the system running the tftp server. For example, on a system running SLES, type:

```
mount /dev/sr0 /mnt
```

2. Enter the following command:

```
cp /mnt/suseboot/inst64 /tftpboot/install
```

Copying the Red Hat Enterprise Linux 5 install kernel

To copy the Red Hat Enterprise Linux 5 install kernel:

1. Mount the Red Hat Enterprise Linux 5.2 DVD1 on the system running the tftp server. For example, on a system running Red Hat Enterprise Linux 5, type:

```
mount /dev/cdrom /mnt
```

2. Enter the following command:

```
cp /mnt/images/netboot/ppc64.img /tftpboot/install
```

The next step is the preparation of the installation source directory and the corresponding service.

Configuring a Network File System Protocol service

With DHCP and TFTP protocols configured, you have to configure the installation source.

On SLES

This section shows how to set up a Network File System (NFS) server using the Installation Server utility provided in the YaST tool. Keep in mind that this is just one way to set up an NFS server. For this setup, our installation server is running SLES 11.

To configure on SLES:

1. Type **yast** on the command line.
2. In the YaST window, scroll down to Miscellaneous and select **Installation Server** on the right side.
3. Select **Configure as NFS Source** and enter the desired source location directory. This example uses `/install` as shown in Figure C-2 on page 536.




Figure C-2 Initial setup of SLES NFS installation server

4. Select **[Next]**.
5. Leave the defaults for **Host Wild Card** and **Options**.
6. Select **[Next]**. With this, an NFS server serving /install is set up automatically.
7. Select **Add** to configure an installation source.
8. As Source Name, enter the desired name for this installation source, for example, sles11. This creates a subdirectory sles11 under /install.
9. At the Source Configuration window, if you have a DVD of the operating system, click the check box. Otherwise, if you have ISO images of the DVD contents, select the second check box and browse with Select Directory to the directory that contains all ISO images of all CDs (see Figure C-3 on page 537).
10. Select **[Next]** when finished.




Figure C-3 Source configuration window

11. If you chose the Read CD or DVD Medium option given in Figure C-3, you are prompted to insert the first DVD.
12. Insert SLES11 DVD1 into the BladeCenter media tray and press **[Continue]**. The data from DVD1 is copied to the `/install/sles11/CD1` directory.

Note: If you used the CD option instead of a DVD, you are prompted for the other CDs at this step.

13. Select **[Finish]** after all the data is copied. The installation server is now ready.

On Red Hat Enterprise Linux

To prepare the NFS server on a system running Red Hat Enterprise Linux 5.2:

1. Move the DVD ISO image to an export directory:

```
mount -o loop /location/to/disk/RHEL5.2.iso /mnt/  
cp -a /mnt/* /install/RHEL5.2/  
umount /mnt/
```

2. Make sure the export directory is exported through NFS entry in /etc/exports. For example:

```
/install/RHEL5.2 *(ro, async, no_root squash)
```

3. Restart the NFS daemon with:

```
/sbin/service nfs start  
/sbin/service nfs reload
```

Virtual optical device setup and installation

This installation option uses the virtual optical device on the Integrated Virtual Manager (IVM) to perform a CD/DVD installation of a Linux operating system image. The Linux image is stored in the IVM's virtual media library. Follow the steps given in 4.6.4, "Optical and tape devices" on page 127 to copy the CD/DVD image onto the VIOS virtual media library. After the CD/DVD image is copied onto the VIOS hard disk, assign the virtual optical device to the LPAR. The installation process is the same as in 8.2, "Linux LPAR installation using DVD" on page 329 for Red Hat Enterprise Linux and 8.3, "Linux network installation" on page 333 for SLES, but remember to select the *virtual optical device* as the boot device in the SMS menu and not the physical media tray.

Note: PowerVM must be installed and properly configured on the JS23 BladeCenter for this to work.

This installation is much faster than a traditional CD/DVD installation. The other benefit is that a read-only image in the VIOS media library is concurrently accessible to all the LPARs on the same IVM so you can run simultaneous installations of an operating system.

Service and productivity tools for Linux

This appendix describes how to install IBM service diagnostic aids and productivity tools for the Linux operating system running on BladeCenter or IVM-managed servers for the JS23 and JS43 BladeCenter.

This appendix contains the following topics:

- ▶ “Overview” on page 540
- ▶ “Install tools for Red Hat Enterprise Linux 5 or SLES 11” on page 542

Overview

The IBM service diagnostic and productivity packages for Linux on POWER architecture provide the latest system diagnostic information such as reliability, availability, and serviceability (RAS) functions, and the ability to modify logical partition (LPAR) profiles with hotplug, dynamic LPAR (DLPAR), and Live Partition Mobility capabilities.

The service and productivity packages available for your configuration are dependent on the system environment and the Linux operating system that is installed. The decision tree in Figure D-1 on page 541 shows how to determine the appropriate packages for your environment.

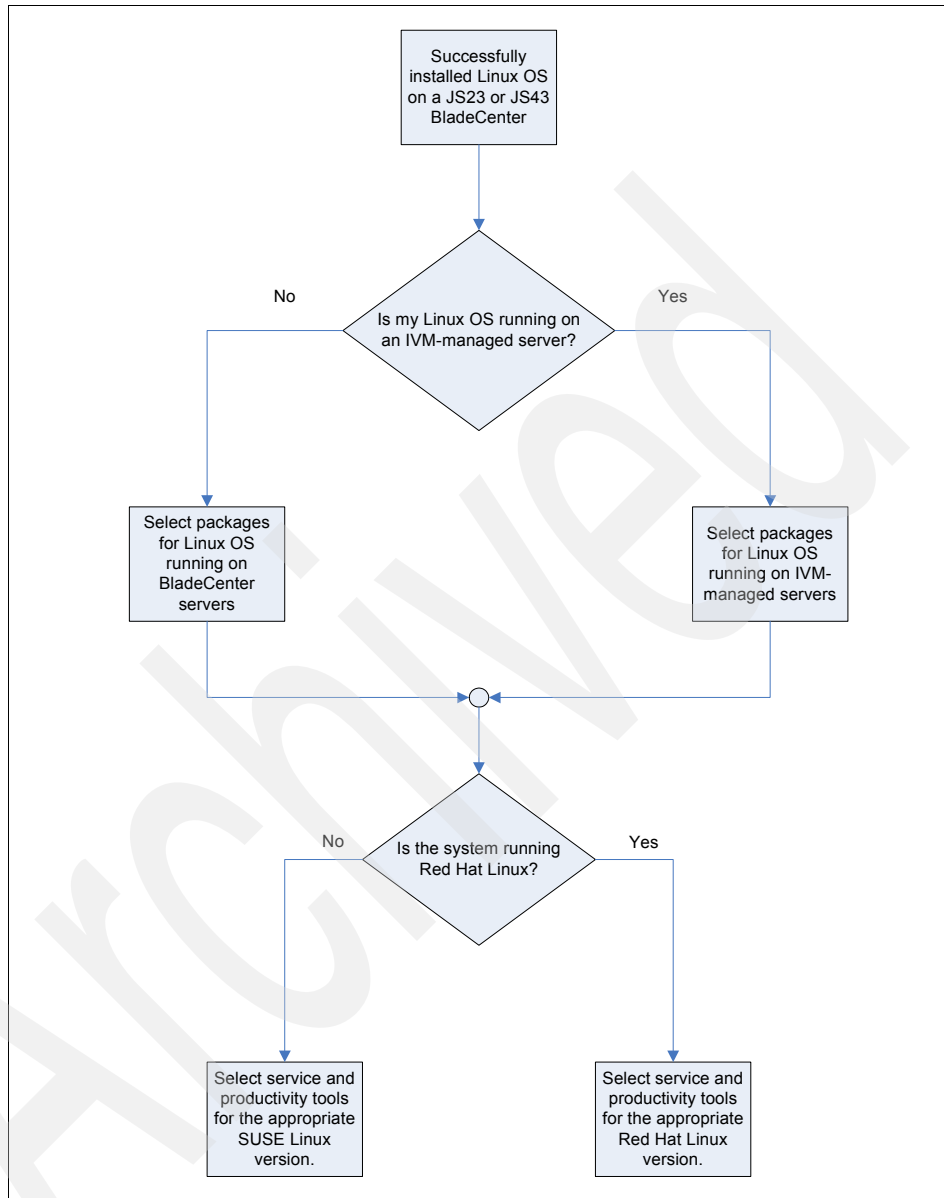


Figure D-1 Service diagnostic and productivity packages decision tree

Install tools for Red Hat Enterprise Linux 5 or SLES 11

This section describes how to install tools for Red Hat Enterprise Linux 5 or SLES 11 that is running on:

- ▶ BladeCenter servers
- ▶ IVM-managed servers

Installing tools on BladeCenter servers

This section describes the steps to configure a JS23 BladeCenter running on a BladeCenter server with the service diagnostic and productivity tools. These steps are applicable for systems running a native Red Hat Enterprise Linux 5/SLES 11 (or later) installation environment:

1. Use a Web browser to connect to:
<https://www14.software.ibm.com/webapp/set2/sas/f/1opdiags/home.html>
2. On this Web site, click the **on BladeCenter servers** link under the Red Hat or SUSE Linux distribution headings, as shown in Figure D-2.

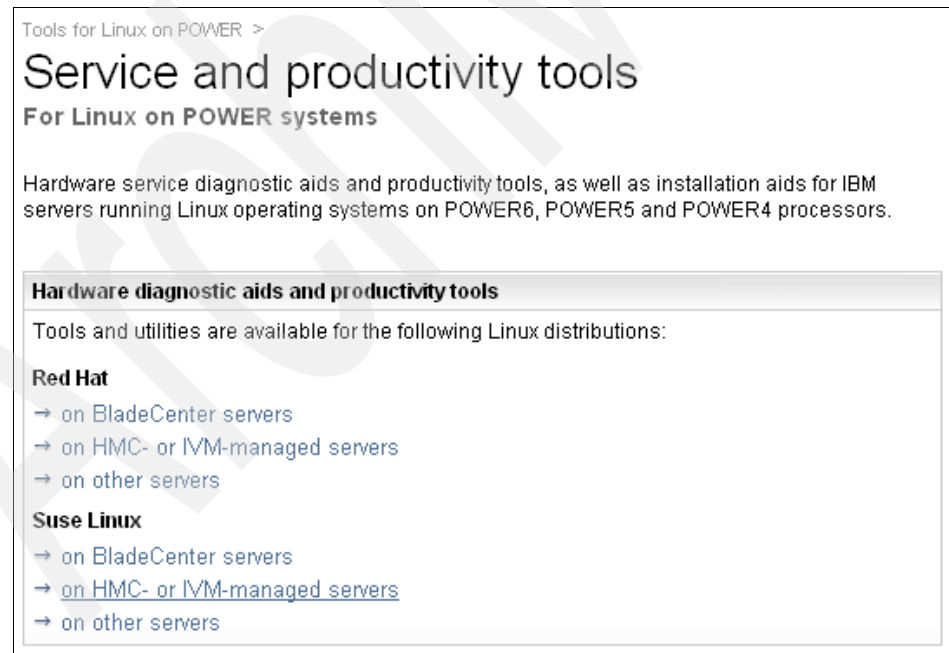


Figure D-2 Select "on BladeCenter servers"

3. Click the tab that matches your Linux operating system (OS) level. In our example in Figure D-3, we select Red Hat Enterprise Linux 5 (RHEL 5). The packages under each tab are unique to that Linux OS level.

Tools for Linux on POWER >

Service and productivity tools

For Red Hat Linux on Blade servers

RHEL 5

RHEL 4

RHEL 3

The following tools are available for BladeCenter servers running Red Hat Linux RHEL 5. The Tool name link provides a brief description of the tool.

Tool packages must be installed in the order listed in the table.

Figure D-3 OS level selection tabs example

4. Under the Package downloads column, click and save each of the packages. At the time of this publication, available packages are listed in Figure D-3/

RHEL 5 tools		
Tool name	Package downloads	Last update
Platform Enablement Library	librtas-1.3.4-0.ppc64.rpm	2009.03.05
Hardware Inventory	lsvdp-1.6.5-1.ppc64.rpm libvdp-2.1.0-1.ppc64.rpm (prerequisite for lsvdp RPM)	2009.03.05
Service log	servicelog-1.0.1-1.ppc64.rpm libservicelog-1.0.1-1.ppc64.rpm (prereq for servicelog RPM) libservicelog-devel-1.0.1-1.ppc64.rpm (prereq for libservicelog RPM)	2009.03.10
Error Log Analysis	diagela-2.2.3-0.ppc64.rpm	2009.02.24

Figure D-4 Available packages for Red Hat on BladeCenter servers

Tip: Click the links under the Tool name column for the latest detailed description of each tool. We explain these packages in Table D-1.

Table D-1 Description of blade server packages

Tool	Description
Platform Enablement Library	Enables applications to access certain functionality provided by platform firmware.
Hardware Inventory	Provides vital product data (VPD) about hardware components to higher-level serviceability tools.
Service Log	Creates a database to store system-generated events that might require service.
Error Log Analysis	Provides automatic analysis and notification of errors reported by the platform firmware.

5. Use a transfer protocol such as FTP or SCP to send each .rpm package to the target system or save these .rpm packages to a CD or DVD and mount the device (see “Tip1: If you place the .rpm files on a CD and DVD, follow these steps to access the files:” on page 545).
6. Install each .rpm package with:

```
rpm -Uvh <packagename>.rpm
```

Important: These packages *must* be installed in the order listed in the table. Otherwise, dependency failures can occur.

7. Perform a system shutdown and restart after installing all the packages.

The service aids and productivity tools section is complete.

Tip1: If you place the .rpm files on a CD and DVD, follow these steps to access the files:

1. Assign the JS23 media tray to the appropriate blade bay.
2. Mount the media tray to a directory on the system such as /mnt:

On Red Hat:

```
mount /dev/cdrom /mnt/
```

On SUSE Linux:

```
mount /dev/sr0 /mnt/
```

3. Move to the mounted directory with:

```
cd /mnt
```

4. Install each rpm with:

```
rpm -Uvh <packagename>.rpm
```

Tip2: Place these .rpm files in a yum repository to quickly update or install these tools on a large number of machines.

Installing tools on IVM-managed servers

Use these steps to configure a JS23 BladeCenter LPAR running on an IVM-managed server with the service aids and productivity tools:

1. Use a Web browser to connect to:
<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>
2. On this Web site, click on **HMC- or IVM-managed servers** under the Red Hat or SUSE Linux distribution headings, as shown in Figure D-2 on page 542, depending on what OS is running on the LPAR.
3. Click the tab (as shown in Figure D-3 on page 543) that matches your Linux operating system (OS) level. The packages under each tab are unique to that Linux OS level.
4. Click on and save each of the packages under the Download column. At the time of this publication, the packages are as listed in Figure D-5 on page 546.

RHEL 5 tools	
Tool name	Download
Platform Enablement Library	librtas-1.3.3-0.ppc64.rpm
SRC	src-1.3.0.2-07305.ppc.rpm
RSCT utilities	rsct.core.utils-2.5.0.1-07305.ppc.rpm
RSCT core	rsct.core-2.5.0.1-07305.ppc.rpm
CSM core	csm.core-1.7.0.1-57.ppc.rpm
CSM client	csm.client-1.7.0.1-57.ppc.rpm
ServiceRM	devices.chrp.base.ServiceRM-2.2.0.0-6.ppc.rpm
DynamicRM	DynamicRM-1.3.2-0.ppc64.rpm
Hardware Inventory	lsydpd-0.15.1-1.ppc.rpm
Service log	servicelog-0.2.9-0.ppc64.rpm
Error Log Analysis	diagela-2.2.2-0.ppc64.rpm
PCI Hotplug Tools	rpa-pci-hotplug-1.0-30.ppc64.rpm
Dynamic Reconfiguration Tools	rpa-dlpar-1.0-47.ppc64.rpm
Inventory Scout	IBMinvscout-2.2-5.ppc.rpm

Figure D-5 Available packages for Red Hat/SUSE Linux on IVM-managed server

Tip: Click the links under the Tool name column for the latest detailed description of each tool.

These packages are described in Table D-2 on page 547.

Table D-2 Description of IVM-managed server packages

Tool	Description
Platform Enablement Library	Allows application to access certain functionality provided by platform firmware.
SRC	Manages daemons on the systems.
RSCT utilities	RSC packages provide the Resource Monitoring and Control (RMC) functions and infrastructure necessary to monitor and manage one or more Linux systems
RSCT core	See description for RSCT utilities.
CSM core	CSM packages provide for the exchange of host-based authentication security keys.
CSM-client	See description for CSM core.
ServiceRM	Service Resource Manager is a Reliable, Scalable, Cluster Technology (RSCT) resource manager that creates the Serviceable Events from output of Error Log Analysis Tool (diagela).
DynamicRM	Dynamic Resource Manager is a Reliable, Scalable, Cluster Technology (RSCT) resource manager that allows an IVM to dynamically add or remove processors or I/O slots from a running partition and perform certain shutdown operations on a partition.
Hardware Inventory	Provides vital product data (VPD) about hardware components to higher-level serviceability tools.
Service Log	Creates a database to store system-generated events that might require service.
Error Log Analysis	Provides automatic analysis and notification of errors reported by the platform firmware.
PCI Hotplug Tools	Allows PCI devices to be added, removed, or replaced while the system is in operation.
Dynamic Reconfiguration Tool	Allows the addition and removal of processors and I/O slots from a running partition.
Inventory Scout	Surveys one or more systems for hardware and software information.

5. Use a transfer protocol such as FTP or SCP to send each .rpm package to the target system or save them to a CD or DVD and mount the device (see “Tip1: If you place the .rpm files on a CD and DVD, follow these steps to access the files:” on page 545).

6. Install each .rpm package with:

```
rpm -Uvh <packagename>.rpm
```

Important: These packages *must* be installed in the order listed in the table. Otherwise, dependency failures can occur.

7. Perform a complete LPAR shutdown and reactivation after installing all packages.
8. Log on to your IVM and ensure that the General tab of the LPAR properties shows all DLPAR services enabled with Yes, as Figure D-6 shows.

Dynamic Logical Partitioning (DLPAR)	
Partition hostname or IP address:	172.16.1.196
Partition communication state:	Active
Memory DLPAR capable:	Yes
Processing DLPAR capable:	Yes
I/O adapter DLPAR capable:	Yes

Figure D-6 DLPAR and Live Partition mobility services are enabled

See Chapter 4, “System planning and configuration using VIOS with IVM” on page 65 for more information about IVM options and functions.

Installation of the service aids and productivity tools is complete.

Tip: Place these .rpm files in a yum repository to quickly update or install these tools on a large number of machines.

Abbreviations and acronyms

ABR	Automatic BIOS recovery	CCSP	Cisco Certified Security Professional
AC	alternating current	CD-ROM	compact disc read only memory
ACL	access control list	CDP	Cisco Discovery Protocol
AES	Advanced Encryption Standard	CE	Conformité Européene
AMD™	Advanced Micro Devices™	CLI	command-line interface
AMM	Advanced Management Module	CNA	Cisco Network Assistance
API	application programming interface	CNS	Cisco Network Services
APV	Advanced Power Virtualization	COG	configuration and option guide
ARP	Address Resolution Protocol	CPM	Copper Pass-thru Module
AS	Australian Standards	CPU	central processing unit
ASF	Alert Standard Format	CRU	customer replaceable units
ASIC	application-specific integrated circuit	CSM	Cluster Systems Management
ASR	automatic server restart	CTS	clear to send
BASP	Broadcom Advanced Server Program	DASD	direct access storage device
BBI	browser-based interface	DC	domain controller
BCM	Broadcom	DDM	Deployment and Management
BE	Broadband Engine	DDR	Double Data Rate
BGP	Border Gateway Protocol	DHCP	Dynamic Host Configuration Protocol
BIOS	basic input output system	DIMM	dual inline memory module
BMC	Baseboard Management Controller	DIP	destination IP
BNT™	BLADE Network Technologies™, Inc	DMAC	destination MAC address
BOFM	BladeCenter Open Fabric Manager	DNS	Domain Name System
BPDU	Bridge protocol data unit	DP	dual processor
BSE	BladeCenter Storage Expansion	DPOD	Dynamic Ports on Demand
BSMP	blade system management processor	DSA	Dynamic System Analysis
BTU	British Thermal Unit	DSCP	Differentiated Services Code Point
CCDA	Cisco Certified Design Associate	DSM	disk storage module
CCNP	Cisco Certified Network Professional	DSUB	D-subminiature
		DTP	Dynamic Trunking Protocol
		DVI	Digital Video Interface

DVMRP	Distance Vector Multicast Routing Protocol	HSDC	high speed daughter card
DVS	Digital Video Surveillance	HSESM	high speed Ethernet switch module
ECC	error correction code	HSFF	high-speed form factor
EDA	Electronic Design Automation	HSIBPM	high-speed InfiniBand pass-thru module
EIGRP	Enhanced Interior Gateway Routing Protocol	HSIBSM	high speed InfiniBand switch module
EMC	electromagnetic compatibility	HSRP	Hot Standby Routing Protocol
EMEA	Europe, Middle East, Africa	HT	Hyper-Threading
EOT	Enhanced object tracking	HTTP	Hypertext Transfer Protocol
EPOW	emergency power-off warning	I/O	input/output
ESD	electrostatic discharge	IB	InfiniBand
ESM	Ethernet switch modules	IBBM	InfiniBand bridge module
ETSI	European Telecommunications Standard Industry	IBM	International Business Machines
FAN	Fabric Address Notification	ICMP	Internet control message protocol
FB-DIMM	Fully Buffered DIMMs	ICPM	Intelligent Copper Pass-thru Module
FBDIMM	Fully Buffered DIMM	ID	identifier
FC	Fibre Channel	IDE	integrated drive electronics
FCP	Flow Control Packet	IEC	International Electro-technical Commission
FCSM	Fibre Channel Switch Module	IEEE	Institute of Electrical and Electronics Engineers
FDD	floppy diskette drive	IGESM	Intelligent Gigabit Ethernet Switch Module
FDX	full duplex	IGMP	Internet Group Management Protocol
FSB	front-side bus	IGRP	Interior Gateway Routing Protocol
FTP	File Transfer Protocol	IM	instant messaging
FTSS	Field Technical Sales Support	IME	Integrated Mirroring Enhanced
GB	gigabyte	IOS	Internetwork Operating System
GUI	graphical user interface	IP	Internet Protocol
HA	high availability	IPM	Intelligent Pass-thru Module
HBA	host bus adapter	IPMI	Intelligent Platform Management Interface
HCA	host channel adapter	IPTV	Internet Protocol Television
HD	high definition	IRDP	ICMP Router Discovery Protocol
HDD	hard disk drive	IS	information store
HH	half high		
HPC	high performance computing		
HS	hot swap		

ISL	Inter-Switch Link	MVR	Multicast VLAN registration
ISMP	Integrated System Management Processor	NAT	Network Address Translation
ISP	Internet service provider	NDCLA	Non-Disruptive Code Load Activation
IT	information technology	NEBS	Network Equipment Building System
ITS	IBM Integrated Technology Services	NGN	next-generation network
ITSO	International Technical Support Organization	NIB	network interface backup
IVM	Integrated Virtualization Manager	NIC	network interface card
KB	kilobyte	NMI	non-maskable interrupt
KVM	keyboard video mouse	NOS	network operating system
LACP	Link Aggregation Control Protocol	NPIV	N_Port ID Virtualization
LAN	local area network	NSF	Notes Storage File
LED	light-emitting diode	NTP	Network Time Protocol
LLDP	Link Layer Discovery Protocol	OBFL	On-board failure logging
LPAR	logical partitions	ODPA	On-Demand Port Activation
LPH	low profile handle	OFED	OpenFabrics Enterprise Distribution
LR	long range	OFM	Open Fabric Manager
LTO	Linear Tape-Open	OPM	Optical Pass-thru Module
LUN	logical unit number	OS	operating system
MAC	media access control	OSPF	Open Shortest Path First
MAN	metropolitan area network	PBR	Policy-based routing
MB	megabyte	PC	personal computer
MDS	Multilayer DataCenter Switch	PCI	Peripheral Component Interconnect
MIB	management information base	PDF	Portable Document Format
MIO	Memory and I/O	PDU	power distribution unit
MM	Management Module	PFA	Predictive Failure Analysis
MMF	Multi Mode Fiber	POST	power-on self test
MP	multiprocessor	PPP	Point-to-Point Protocol
MPE	Multi Processor Expansion	PVST	Per-VLAN Spanning Tree
MPI	Message Passing Interface	PXE	Preboot Execution Environment
MSDP	Multicast Source Discovery Protocol	RAC	Real Application Clusters
MSIM	Multi-Switch Interconnect Module	RADIUS	Remote Authentication Dial In User Service
MSTP	Multiple Spanning Tree Protocol		
MTM	machine-type-model		

RAID	redundant array of independent disks	SFF	Small Form Factor
RAM	random access memory	SFP	small form-factor pluggable
RAS	remote access services; row address strobe; reliability, availability, and serviceability	SIMD	single instruction multiple data
RDAC	Redundant Disk Array Controller	SIO	Storage and I/O
RDC	Remote Desktop Connection	SIP	source IP
RDIMM	registered DIMM	SLB	Server Load Balancing
RDM	Remote Deployment Manager	SLES	SUSE Linux Enterprise Server
RDMA	Remote Direct Memory Access	SMAC	source MAC address
RETAIN®	Remote Electronic Technical Assistance Information Network	SMI-S	Storage Management Initiative - Specification
RHEL	Red Hat Enterprise Linux	SMP	symmetric multiprocessing
RIP	Routing Information Protocol	SMS	System Management Services
RMCP	Remote Management Control Protocol	SNMP	Simple Network Management Protocol
RMON	Remote Monitoring	SOL	Serial Over LAN
RP	route processor	SPORE	ServerProven Opportunity Request for Evaluation
RPF	reverse path forwarding	SR	short range
RPM	revolutions per minute	SRP	Storage RDMA Protocol
RPQ	Request Per Qualification	SRR	shaped round robin
RSA	Remote Supervisor Adapter	SSCT	Standalone Solution Configuration Tool
RSCN	Registered State Change Notification	SSD	solid state drive
RSTP	Rapid Spanning Tree Protocol	SSH	Secure Shell
SAN	storage area network	SSL	Secure Sockets Layer
SAS	Serial Attached SCSI	SSP	Serial SCSI Protocol
SASCM	SAS Connectivity Module	STP	Spanning Tree Protocol
SATA	Serial ATA	TACACS	Terminal Access Controller Access Control System
SBB	Sales Building Block	TB	terabyte
SCM	Supply Chain Management	TCO	total cost of ownership
SCSI	Small Computer System Interface	TCP	Transmission Control Protocol
SDD	Subsystem Device Driver	TCP/IP	Transmission Control Protocol/Internet Protocol
SDK	Software Developers' Kit	TFTP	Trivial File Transfer Protocol
SDR	Single Data Rate	TSM	Tivoli Storage Manager
SDRAM	static dynamic RAM	TX	transmit

UDLD	UniDirectional link detection
UDP	user datagram protocol
ULP	upper layer protocols
URL	Uniform Resource Locator
USB	universal serial bus
UTF	Universal Telco Frame
UTP	unshielded twisted pair
VBS	Virtual Blade Switch
VGA	video graphics array
VIOS	Virtual I/O Server
VLAN	virtual LAN
VLP	very low profile
VM	virtual machine
VMPS	VLAN Membership Policy Server
VNC	Virtual Network Computing
VOIC	Virtual I/O Client
VOIP	Voice over Internet Protocol
VOIS	Virtual I/O Server
VPD	vital product data
VPN	virtual private network
VQP	VLAN Query Protocol
VRRP	virtual router redundancy protocol
VSAN	Virtual Storage Area Network
VT	Virtualization Technology
VTP	VLAN Trunking Protocol
WAN	wide area network
WOL	Wake on LAN
WTD	Weighted tail drop
WWN	World Wide Name
XDR	extreme data rate
XM	extended memory

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 560. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *IBM BladeCenter Products and Technology*, SG24-7523
- ▶ *IBM System i and System p System Planning and Deployment: Simplifying Logical Partitioning*, SG24-7487
- ▶ *IBM System Storage DS4000 and Storage Manager V10.30*, SG24-7010
- ▶ *PowerVM Live Partition Mobility*, SG24-7460
- ▶ *Managing OS/400 with Operations Navigator V5R1 Volume 1: Overview and More*, SG24-6226
- ▶ *Managing OS/400 with Operations Navigator V5R1 Volume 5: Performance Management*, SG24-6565
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- ▶ *Implementing IBM Director 5.20*, SG24-6188
- ▶ *Going Green with IBM Systems Director Active Energy Manager*, REDP-4361
- ▶ *Integrated Virtualization Manager on IBM System p5*, REDP-4061
- ▶ *Implementing the IBM BladeCenter S Chassis*, SG24-7682
- ▶ *IBM System i Overview: Models 515, 525, 550, 570, 595, and More*, REDP-5052
- ▶ *IBM BladeCenter JS12 and JS22 Implementation Guide*, SG24-7655

Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM Systems Director Active Energy Manager Version 3.1.1 is an IBM Director extension.
<http://www.ibm.com/systems/management/director/extensions/actengmrg.html>
- ▶ IBM periodically releases maintenance packages for the AIX 5L operating system. These packages are available on CD-ROM, or you can download them from the Web.
<http://www.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>
- ▶ In AIX 5L V5.3, the **suma** command is also available, which helps the administrator to automate the task of checking and downloading operating system downloads.
<http://www.ibm.com/systems/p/os/aix/whitepapers/suma.html>
- ▶ Information is available about features and external devices supported by Linux on the IBM BladeCenter JS23 and IBM BladeCenter JS43 servers.
<http://www.ibm.com/systems/power/software/linux/index.html>
- ▶ SUSE Linux Enterprise Server 11 Web site has more information.
<http://developer.novell.com/yessearch/Search.jsp>
- ▶ Red Hat Enterprise Linux Web site has more information.
<https://hardware.redhat.com/?pagename=hcl&view=certified&vendor=4&class=8>
- ▶ Many features described in this document are operating system dependent and might not be available on Linux. For more information, visit the IBM Linux Web site.
http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html
- ▶ For more information about IBM i V6.1 operating systems running on IBM BladeCenter JS23 and IBM BladeCenter JS43 see IBM Power Blade servers Web site.
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>
- ▶ To download IBM Director, see IBM Systems Director Downloads Web site.
<http://www.ibm.com/systems/management/director/downloads/>
- ▶ Complete VIOS configuration and maintenance information is in *Using the Virtual I/O Server*.
http://publib.boulder.ibm.com/infocenter/iserics/v1r3s/en_US/info/ip hb1/iphb1.pdf

- ▶ The *IBM BladeCenter Interoperability Guide* contains helpful information.
<https://www.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-5073016&brandind=5000020>
- ▶ The Virtual I/O Server data sheet gives an overview of supported storage subsystems and the failover driver that is supported with the subsystem. Also, use the data sheet to verify which components, supported by the blade, are supported by the Virtual IO server.
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>
- ▶ All supported hardware and operating systems are listed on IBM ServerProven.
<http://www.ibm.com/servers/eserver/serverproven/compat/us/>
- ▶ Verify the supported operating systems on the blade by using the link to NOS Support on the Compatibility for BladeCenter products page of ServerProven.
<http://www.ibm.com/servers/eserver/serverproven/compat/us/eserver.html>
- ▶ DS8000 interoperability matrix is on the IBM System Storage DS8000 site.
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>
- ▶ DS6000 interoperability matrix is on the IBM TotalStorage DS6000 site.
<http://www.ibm.com/servers/storage/disk/ds6000/interop.html>
- ▶ DS4000 interoperability matrix is on the DS4000 interoperability matrix site.
<http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html>
- ▶ DS3000 interoperability matrix is available as a PDF.
<http://www.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf>
- ▶ ESS interoperability matrix is on the Enterprise Storage Server family site.
<http://www.ibm.com/servers/storage/disk/ess/interop-matrix.html>
- ▶ The N series interoperability matrix is on the Network attached storage site.
<http://www.ibm.com/systems/storage/nas/interophome.html>
- ▶ SAN Volume Controller support matrix is on the Support for SAN Volume Controller site.
<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskind=3&brandind=5000033&familyind=5329743>
- ▶ The SAN switch interoperability matrix is on the Storage Area Network site.
<http://www.ibm.com/systems/storage/san/index.html>

- ▶ The System Storage Interoperation Center (SSIC) helps to identify a supported storage environment.
<http://www.ibm.com/systems/support/storage/config/ssic>
- ▶ DS3000 and DS4000 support Web pages usually provide update packages for the supported adapters that contain the settings required for the HBA.
<http://www.ibm.com/systems/support/storage/config/hba/index.wss>
- ▶ The Storage Configuration Manager (SCM) can help to create an individual configuration if you are not familiar with using the SAS I/O module command-line interface. Download SCM software from the support site.
<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ndocid=MIGR-5502070&brandind=5000008>
- ▶ System i Access Web site has information about IBM System i Access for Windows V6R1.
<http://www.ibm.com/systems/i/software/access/index.html>
- ▶ System i Access Web site also has software for IBM System i Access for Windows.
<http://www.ibm.com/systems/i/software/access/caorder.html>
- ▶ Service tools user IDs are created through DST or SST and are separate from IBM i V6.1 user profiles. See the Service tools user IDs Web site.
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzamh/rzamhwhatuserids.htm>
- ▶ IBM Workload Estimator is available on the IBM Systems Workload Estimator Web site.
<http://www.ibm.com/systems/support/tools/estimator/index.html>
- ▶ Performance actions related to disk formatting are described in the IBM i Information Center.
<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/topic/rzahg/icmain.htm>
- ▶ Language feature codes are available in the Information Center.
<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/index.jsp?topic=/rzahc/rzahcnlvfeaturecodes.htm>
- ▶ Use the IBM i Recommended fixes Web site to get a list of the latest recommended PTFs.
http://www-912.ibm.com/s_dir/slkbases/recommendedfixes

- ▶ Fix Central is the primary Web site for downloading fixes for all operating systems and applications.
<http://www-912.ibm.com/eserver/support/fixes>
- ▶ The Navigator for i Web site has IBM Systems Director Navigator for i functionality information.
<http://www.ibm.com/systems/i/software/navigator/index.html>
- ▶ The Virtual I/O Server site is good source for processor and memory requirements for PowerVM partitions based on I/O requirements.
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>
- ▶ The Red Hat Recommended Partitioning Scheme is available on the Red Hat documentation site.
http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/ch11s03.html
- ▶ Red Hat installation instructions are available on the Red Hat documentation Web site.
http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/pt-install-info-ppc.html
- ▶ The Novell Web site has additional installation preparation information for SLES10 Sp2.
<http://www.novell.com/documentation/sles10/index.html>
- ▶ SLES installation and administration information can assist with the completion of the SLES installation.
<http://www.novell.com/documentation/sles11/#administration>
- ▶ The Virtual I/O Server site is has download information.
<http://techsupport.services.ibm.com/server/vios/download>
- ▶ Linux partitions must have the Dynamic Reconfiguration Tools package for HMC or IVM managed servers installed from the Service and productivity tools Web site.
<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>
- ▶ *Serial Over LAN (SOL) Setup Guide - BladeCenter* contains details about setting up SOL.
<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-54666&brandind=5000020>

- ▶ *Advanced Management Module / Management Module command-line interface Reference Guide - IBM BladeCenter E, H, T, HT, S* has detailed information.
<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-54667&brandind=5000020>
- ▶ IBM Power Blade servers Web site has a technical overview and full details, as well as latest updates on IBM i on Power blades. See the *Read Me First*.
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

/etc/dhcpd.conf 531
/etc/init.d/dhcpd restart command 532
/etc/init.d/dhcpd start command 532
/etc/xinet.d/tftpd 534
/sbin/service nfs reload command 538
/sbin/service nfs start command 538

Numerics

128-bit vector execution unit 41
64-bit POWER6 processors 10

A

active console 500
Active Energy Manager (AEM) 10, 60
 console 402
Active Memory Sharing (AMS) 173
 configuration using IVM 178
 DLPAR operations 203
 planning considerations 174
 consolidation factors 178
 CPU resource sizing 176
 memory weight 177
 paging devices 176
 prerequisites 175
 workload 175
 Shared Memory Pool 175
add_vlan_ids 110
Advanced Management Module (AMM) 244, 392
Advanced Management Web interface 59
Advanced Peer-to-Peer Networking (APPN) 305
Advanced Program-to-Program Communication (APPC) 305
AIX 57
 install kernel load 229
AIX 6.1
 Installation 205
 LPAR Installation on VIOS 225
Alert Standard Format (ASF) 14
alternate processor retry 38
Altivec SIMD technology 41
Anaconda 333

Anchor Smartchip 45
automated monitoring 61
automatic server restart (ASR) 15
automatic service processor reset 15
AutoYaST 387, 515–516
 profile creation 516
AutoYaST main menu options 523
 Hardware 524
 Misc 526
 Network Devices 525
 Network Services 525
 Security and Users 525
 Software 523
 System 524
Available Virtual Disks 268

B

backup and restore
 IBM i V6.1 307
Backup, Recovery, and Media Services
 See BRMS
baseboard management controller (BMC) 14
basic Kickstart configuration file 355
Basic Zone Permission Table 248
bit steering 21
blade server features 14
 baseboard management controller (BMC) 14
 disk drive support 14
 I/O expansion 15
 IBM Director 14
 integrated network support 14
 large system memory capacity 15
 light path diagnostics 15
 power throttling 15
 reliability, availability and serviceability (RAS) 15
BladeCenter
 benefits 6
 highlights 4
BladeCenter H 23
BladeCenter JS23 16
 Internal disk 18
 Memory DIMMs 18

- Memory features 18
- Processor features 17
- BladeCenter JS43 19
 - Internal disk 21
 - Memory DIMMs 21
 - Memory features 21
 - Processor features 20
- BladeCenter Open Fabric Manager 7
- BladeCenter Power Domain 395
- BladeCenter Storage (BCS) 53
- boot device 77
- boot net command 388
- Bootstrap Protocol (BOOTP) 530
- bridge 218
- BRMS 307, 309

C

- cache availability 38
- Call Home
 - contact message 476
- CD-ROM
 - virtual 222
- centralized BladeCenter environment 6
- CFFh 70
- CFFv 70
- chassis 22
- checkstop analysis 15
- chhwres 110
- Chipkill
 - memory features 21
 - technology 10
- chsyscfg 70
- chsysstate 165, 170
- client partition
 - create using IVM 212
- Cluster 1350™ portfolio 10
- Cluster Systems Management
 - See CSM
- CMM 176
- Collaborative Memory Manager (CMM) loan policy 176
- Combined Form Factor horizontal (CFFh) 70
- Combined Form Factor vertical (CFFv) 70
- commands
 - AMM
 - env 400
 - fuelg 398
 - chsyscfg 70

- IBM i
 - PWRDWNSYS 326
 - SHTDWNSYS 167
- install_assist 233
- IVM
 - chsyscfg 93
 - chsysstate 165
 - LPAR shutdown 170
 - lpcfgop 81
 - lshwres 93
 - lslparmigr 451
 - lsrefcode 165
 - lssyscfg, determine RMC status 437
 - lssyscfg, determine system name 451
 - lssyscfg, LPAR state 166
 - lssyscfg, migrate status 451
 - lsvet 430
 - migrpar 450
- license 81
- lsfware 462
- lsmcode 462
- lssw 324
- migrpar 451
- power 498
- printenv 505
 - printenv boot-device 505
- show-devs 506
- suma 58
- telnetcfg 498
- update_flash 454
- validate 450
- VIOS
 - cfgassist 81–82
 - chdev, change hdisk reserve policy 433
 - entstat 70
 - help 79
 - ioslevel 429
 - license 81
 - lsdev 118
 - lsdev, check hdisk reserve policy 433
 - lspv 118
 - lstcpip 81, 84
 - mktcpip 81–82
 - mkvt 163
 - rmtcpip 87
- commit new firmware 454
- configuration file
 - Kickstart 355
- connect the System i LAN console 273

- Configuration wizard 274
- dedicated service tools (DST) 281
- hosts file 277
- service tool device ID 281
- system serial number 279
- target partition 280
- control panel 466
- controlled end delay time 326
- Create Partition wizard 190
- Create Storage Pools 206
- Create virtual media library for backup 309
- creating IBM i partition 264
 - load source and console identification 270
- creating IBM i V6.1 partition 263
- CSM 61
- CSM Cluster Systems Management (CSM) 59
- current firmware level 462

D

- decimal floating point 39
- default ID 80
- degraded boot support 15
- DHCP 530
 - AutoYaST configuration 525
 - configuration 531
 - restarting 532
- diagela 471
- diagnostics 61
- digital KVM 490
- disk configuration in BladeCenter S 247
 - Basic Zone Permission Table 248
 - individual BladeCenter S disk configuration 252
 - Predefined Config 251
 - Storage Configuration Manager (SCM) 558
 - Zoning 248
- disk considerations in BladeCenter S 245
- Disk Storage Modules (DSM) 246
- disk storage types 74
- DLPAR 151, 203, 358, 389, 540
 - capabilities, retrieve 152
 - ethernet 156
 - memory 153
 - optical devices 159
 - physical adapters 160
 - processing 154
 - storage 158
- Dynamic Host Configuration Protocol
 - See DHCP

- dynamic logical partition
 - See DLPAR
- Dynamic Power Saver Mode 42
- Dynamic Reconfiguration Tools package 437

E

- eConfigurator 10
- EnergyScale 44, 391
 - definition of 43
 - technology 42
- enhanced SMT features 40
- entstat 70, 116
- error correction 21
- Ethernet bridge 218
- Ethernet I/O module 244
- expansion cards 70
 - CFFh, Qlogic combo card 71
 - CFFv, QLogic 4 Gbs fibre channel HBA 71
 - HSFF, Cisco 4x Infiniband 73
- Extended Error Handling (EEH) 16
- external storage 74

F

- fibre channel storage 74
- File Transfer Protocol (FTP)
 - installation server 530
- firmware 487
 - download updates for JS12/22 428
 - file name 463
 - latest image 464
 - level 462
 - prompt, accessing 504
- first-failure data capture (FFDC) 6, 16
- Flexible Support Processor (FSP) 44
- front panel (control panel) 466
- fuelg 398, 400

G

- Gigabit Ethernet 23

H

- hdisk reserve policy 433
- help 79
- High Speed Form Factor (HSFF) 70
- high-performance computing (HPC) applications 5
- HMC 263
- host bus adapters (HBA) 77

- Host Ethernet Adapter (HEA) 50, 70, 95, 436
 - configuring as a SEA 97
- hotplug 540
- Hypertext Markup Language (HTML) 78
- Hypertext Transfer Protocol (HTTP) 530
- hypervisor 438

I

- I/O hotplug 358, 389
- IBM BladeCenter chassis 22
 - BladeCenter H 23
 - BladeCenter HT 30
 - BladeCenter S 26, 74
- IBM Cluster Systems Management
 - See CSM
- IBM Director 59–60
- IBM Director benefits 36
 - increased productivity of IT personnel and users 36
 - reduced downtime 36
 - reduced power consumption 36
 - reduced service and support costs 36
- IBM i 59
 - storage considerations 76
- IBM i V6.1 59
 - backup/restore
 - create virtual media library for backup 309
 - IBM SAS LTO 307
 - install operating system 296
 - shutdown and restart 325
 - shutdown IBM i V6.1 partition 325
 - start an IBM i V6.1 partition 326
 - TCP/IP setup 305
- IBM i V6.1 installation 235
 - disk configuration in BladeCenter S 247
 - disk considerations in BladeCenter S 245
 - hardware environments 238
 - media preparation 272
 - preparation 236
 - PTFs 304
 - IBM ID 305
 - software installation process 236
 - storage consideration BladeCenter H 245
- IBM i V6.1 IPL types 284
 - IPL type A 284
 - IPL type B 284
 - IPL type C 284
 - IPL type D 284
- IBM i V6.1 Navigator for i 305
- IBM ID 305
- IBM ServerProven 75, 557
- IBM service and productivity tools 358, 389
- IBM System Access for Windows V6R1 253
 - install System i Access for Windows 254
 - preparation 254
- IBM Tivoli Storage Manager (TSM) 324
 - client check 324
- IBM Total Storage 74
- IBM Total Storage N Series 75
- IBM Workload Estimator (IBM WLE) 177
- iee_virtual_eth 110
- IEEE 1275 standard 503
- InfiniBand 23
- install
 - IBM AIX 6.1 206
 - IBM i V6.1 operating system 296
 - Licensed Internal Code (LIC) 284
 - Linux
 - configuring a BOOTP or DHCP service 531
 - network installation 530
 - virtual optical device setup and installation 538
 - Red Hat Enterprise Linux 5.3 329
 - automated installation 346
 - Red Hat Enterprise Linux 5.3 over network 333
 - Red Hat Enterprise Linux 5.3 using DVD 329
 - SLES 11 over the network 366
 - SLES 11 using a DVD 361
 - SUSE Linux Enterprise Server 11 360
 - System i Access for Window 254
 - System i Access for Windows 254
- install_assist 233
- installation assistant 233
- installation summary
 - AIX 230
- integrated Ethernet Switch module 244
- integrated network support 14
- Integrated Virtual Ethernet (IVE) 50
- Integrated Virtualization Manager (IVM) 56, 366
 - command-line interface 79
 - Create Partition button 140
 - Guided Setup view 88
 - Partition Properties 90
 - General tab 90
 - Memory tab 90
 - Processing tab 90
 - user interface 78, 189

- navigation area 78
- work area 78
- View/Modify Host Ethernet Adapters 95
 - properties 96
 - Connected Partitions tab 96
 - General tab 96
- View/Modify Partitions link 89
- View/Modify Physical Adapters link 103
- View/Modify Virtual Ethernet, SEA configuration 99
- View/Modify Virtual Storage 117
 - Physical Volumes tab 117
 - Storage Pools tab 119
- VIOS management 77
- Integrated Virtualization Manager V1.4 324
- Intelligent Copper Pass-thru Module 70
- Intelligent Platform Management Interface (IPMI) 14
- Internal SAS 74
- interoperability matrix 76
- interposer card 73
- IPv6 87
- is_trunk 110

J

- JS23 Express and JS43 Express, legend 3

K

- Kickstart
 - configuration file 355
 - installation 357
 - profile creation 347
- Kickstart Configurator tool 348
- KVM 78, 490

L

- last known good firmware 454
- Licensed Internal Code (LIC) 284
- Light Path 468
- Light Path diagnostics 15–16
- light-emitting diodes (LEDs) 16
- Linux 58
- Linux network installation 530
- live partition migration 358, 389, 540
- Live Partition Mobility 57, 75, 427
 - LPAR OS versions 431
 - Migrate option 447

- migrate using CLI 451
- migrate using IVM 447
- migrating the LPAR 442
 - from the command line 450
 - using the IVM UI 442
- PowerVM Enterprise 430
- preparation 432
 - memory region size 432
 - networking 436
 - storage and hdisk reserve policy 433
- requirements 428
 - firmware 428
 - hardware 428
 - VIOS version 428
- status using CLI 451
- status using IVM 448
- validate using CLI 450
- validate using IVM 442
- logical over-commit 175
- logical partition (LPAR)
 - non-dynamic operations 77
 - See also* LPAR or DLPAR 540
- logical partitioning 56
- logical volume base 206
- LPAR activation
 - from the CLI 165
 - from the UI 163
- LPAR creation
 - optical devices 147
 - partition Ethernet 144
 - partition memory 142
 - partition name 141
 - partition processors 143
 - partition storage types and assignments 145
 - partition summary 149
 - physical adapters 148
- LPAR shutdown 166
 - from the CLI 170
 - from the UI 167
- lpcfgop 81
- lpcfgop command 81
- lslparmigr 451
- lsrefcode 165
- lssyscfg 166, 437, 451
- lstcpip 84

M

- machine checks 465

- man dhcpd command 531
- media 78
- Media Access Control (MAC) 531
- media library 317
 - adding new media 132
 - create 130
 - create blank media 134
 - delete 132
 - extend 132
 - modify media assignment 136
 - size 317
- media tray 78
- memory
 - placement rules 46
- memory controller 18, 21
- memory region size 432
- memory scrubbing 21
- memory subsystem 45
- Micro-Partitioning 56
- migrpar 450
- mobile partition 449
- monitoring
 - automated 61
 - resources 61
- multiprocessor expansion unit (MPE) 19
- Multi-Switch Interconnect Module (MSIM) 73

N

- native Red Hat Enterprise Linux 5.2 installation 346
- native SLES10 SP2 installation 387
- Navigator for i 305
- network considerations
 - AMM 244
 - Ethernet I/O module 244
 - IBM i V6.1 installation 244
 - PC for LAN Console 245
 - V6R1 LAN console 245
 - V6R1 production interface 245
 - VIOS/IVM 244
- Network File System (NFS) 333, 366, 530, 535
- network installation 61
 - configuration file 531
 - installation source 530
 - Linux 530
 - Trivial File Transfer Protocol (TFTP) 533
- network interface backup (NIB) 70
- network interface controller (NIC) 494
- NIM server 78

- non over-commit 175
- Normal Mode Boot 335
- Novell SUSE Linux Enterprise Server 57
- Novell Web site 361

O

- Online Update and Package Selection 523
- Open Firmware 487
 - boot parameters 357, 388
 - interface 503
 - prompt 503
- Operating system support
 - IBM i 59
 - Linux 58
- operating system support 57
- optical devices 127, 207
 - physical 127
- Optical Pass-thru Module 72

P

- paddle card 53
- paging devices 176
- paging storage pool 179
- partition
 - creating a partition
 - AIX 215
 - IBM i V6.1 263
 - preparation 212
 - preparing PowerVM client partition 212
 - View/Modify Partitions page 213
- Partition Management 215
- Partition Mobility 70
- PCI Express (PCIe) 48
- PCI-x 48
- physical optical device 209
- physical optical remote media 127
- physical over-commit 175
- physical volumes 117
 - assigning 220
- Point-to-Point Protocol (PPP) 14
- port_vlan_id 110
- POST 16, 501
- Power cap 44
- Power Capping 42
- Power Saver Mode 42
- Power Trending 42
- POWER® Hypervisor (PHYP) 465
- POWER6 438

- POWER6 Enhanced 438
- POWER6 Hypervisor (PHYP) 16
- POWER6 processor 38–39, 454
 - Altivec (SIMD) 41
 - decimal floating point 39
 - EnergyScale technology 42
 - simultaneous multi-threading 40
- POWER6 processor-based blade family 10
- POWER6+ 438
- POWER6+ Enhanced 438
- power-on self-test
 - See POST
- PowerVM 56, 430
 - editions for BladeCenter JS23 and JS43 57
 - key 430
- printenv 505
 - printenv boot-device 505
- Processor Core Nap 43
- processor instruction retry 38
- PWRDWN SYS 326

R

- RAID SAS Switch Modules (RSSM) 54
- RAS 15, 358, 389, 540
- Red Hat Enterprise Linux 57
 - configuration tool 349
 - Kickstart 357
 - Version 5.2 automated installation 346
- Red Hat Enterprise Linux for POWER
 - Version 4.6 328
 - Version 5.1 328
- Redbooks Web site 560
 - contact us xviii
- reliability, availability, and serviceability
 - See RAS
- reliable double data rate 2 (DDR2) 5
- remote command execution 61
- remote control 490
 - Java applet 493
- Renesas Technology 13
- Reserved Firmware Memory value 178
- reset-all 510
- resource balancing 40
- resource monitoring 61
- Resource Monitoring and Control (RMC) 152, 436
- Restart Blade option 497
- rootvg 131
- rpm -Uvh command 544, 548

S

- SAS 10, 51
 - controller 53
 - modules 244
 - Serial Attached SCSI
 - See SAS
- SAS Disk Array Manager for AIX 53
- SAS/SATA 74
- SCM 252, 558
- Secure Shell
 - See SSH
- Serial Over LAN
 - See SOL
- ServerProven 75
- service and productivity tools for Linux 539
- Service Processor 44
- service request number (SRN) 472–473
- service tool device ID 281
- service xinetd restart command 534
- set-connection-mode 509
- Shared Dedicated Capacity 57
- Shared Ethernet Adapter (SEA) 97, 436
- shared memory partitions
 - definition of 174
- shared memory pool 203
- show-devs 506
- shutdown IBM i V6.1 partition 325
- simultaneous multi-threading
 - See SMT
- single processor checkstopping 38
- single rack-mounted chassis 4
- single-threaded (ST) operation 41
- smit sasdam 53
- SMS 487
 - options 495
- SMT 40
 - enhanced features 40
- SOL 14, 78–79, 494
- Solid State Disk Technology 5
- Solid state drive (SSD) 14
- specialized hardware detection circuits 465
- SSH 489
 - connection. 494
- SSIC 77
- Standard Form Factor 70
- Start an IBM i V6.1 partition 326
- Start Remote Control 493
- starting the firmware image
 - from the TEMP side 456

- Static Power Saver Mode 42
- storage area network (SAN) 333, 366
- Storage Configuration Manager
 - See SCM
- storage consideration BladeCenter H 245
- storage considerations 75
- Storage Management 206
- storage pool 119, 179, 206
 - delete 121
 - reduce 121
- suma 58
- superuser rights 530
- support matrixes 75
- SUSE Linux Enterprise Server (SLES) 57, 359
- SynapSense Wireless Sensor Node 61
- system diagnostics 465
 - diagnostic tools 465
 - checkpoints and error codes 465
 - for the AIX operating system 471
 - for the IBM i operating system 471
 - for the Linux operating system 471
 - Light Path and Front Panel diagnostics 466
 - stand-alone diagnostics 471
 - IBM i partition diagnostics and errors 476
 - reference codes 472
 - error codes 472
 - progress codes 472
 - Service Advisor 476
 - using Advanced Management Module 473
- System i Access for Windows 253–254
- System Management Service
 - See SMS
- system reference code (SRC) 472–473
- system specification
 - BladeCenter JS23 and JS43 chassis support 22
 - BladeCenter JS23 disk support 18
 - BladeCenter JS23 memory support 18
 - BladeCenter JS23 processor support 18
 - BladeCenter JS43 disk support 21
 - BladeCenter JS43 memory support 21
 - BladeCenter JS43 processor support 20
 - BladeCenter support 22
 - JS23 standard configuration 16
 - JS43 configuration options 19
- System Storage Interoperation Center 77
- system-config-kickstart command 349
- Systems management
 - Advanced Management Web interface 59

- IBM Director 60
- systems management 59
 - Cluster Systems Management (CSM) 59, 61

T

- Telnet 489
- telnetcfg 498
- tftpboot 533
- Thermal Power Management Device (TPMD) 43
 - POWER6 processor 43
- thread priority 41
- Trivial File Transfer Protocol (TFTP) 530
 - configuration 533
 - network installation 530
- trunk_priority 110
- TSM client check 324

U

- Universal Serial Bus (USB) subsystem 50
- Unmount All option 323
- USB 50

V

- Very Low Profile (VLP) RDIMM packages 18
- VIOC (VIO Client) 114
- VIOS 56, 65
 - data sheet 75, 557
 - default user 87
 - fix pack download location 428
 - planning 67, 174
- VIOS/IVM 244
- virtual CD-ROM 222
- virtual disks 122
 - assigning 220
 - create 122
 - delete 126
 - extend 124
- Virtual Ethernet
 - Integrated Virtualization Manager (IVM) 78
- Virtual Ethernet Adapters 97
- Virtual I/O Server
 - See VIOS
- Virtual I/O Server (VIOS)
 - default user 80
- virtual media library 206
- virtual optical devices 147, 441, 538
- Virtual Processor default settings 93

Virtual Storage Management 116
virtual tape 147
virtual terminal 161
vital product data (VPD) 16
VMLibrary 209

W

Windows 253
WRKOPTVOL 314

Y

yast command 533
YaST graphical interface 517
yast2 command 517

Z

Zone Group ID 248
Zoning 248



Redbooks

IBM BladeCenter JS23 and JS43 Implementation Guide

(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages



IBM BladeCenter JS23 and JS43 Implementation Guide



**Featuring
installation
techniques for the
IBM AIX, IBM i, and
Linux operating
systems**

**Showing Live
Partition Mobility
scenarios**

**Detailing AMS, IVM,
and power
management**

This IBM Redbooks publication provides a detailed technical guide for configuring and using the IBM BladeCenter JS23 and IBM BladeCenter JS43 servers. These IBM Power Blade servers feature the latest IBM POWER6 processor technology. This book teaches you how to set up the latest Power Blade servers to run AIX, i, and Linux operating systems in the IBM BladeCenter architecture.

This book is useful for data center system architects; network, storage and facilities engineers; system administrators; and application architects that want a detailed system start-up guide and integration guide for IBM Power Blade servers hosting virtualized operating system partitions. We include procedures for installation and configuration of Virtual Input/Output Server (VIOS), Integrated Virtualization Manager (IVM), IBM AIX Version 6.1, IBM i V6.1 (i5/OS V6R1), Red Hat Enterprise Linux, and SUSE Enterprise Linux.

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