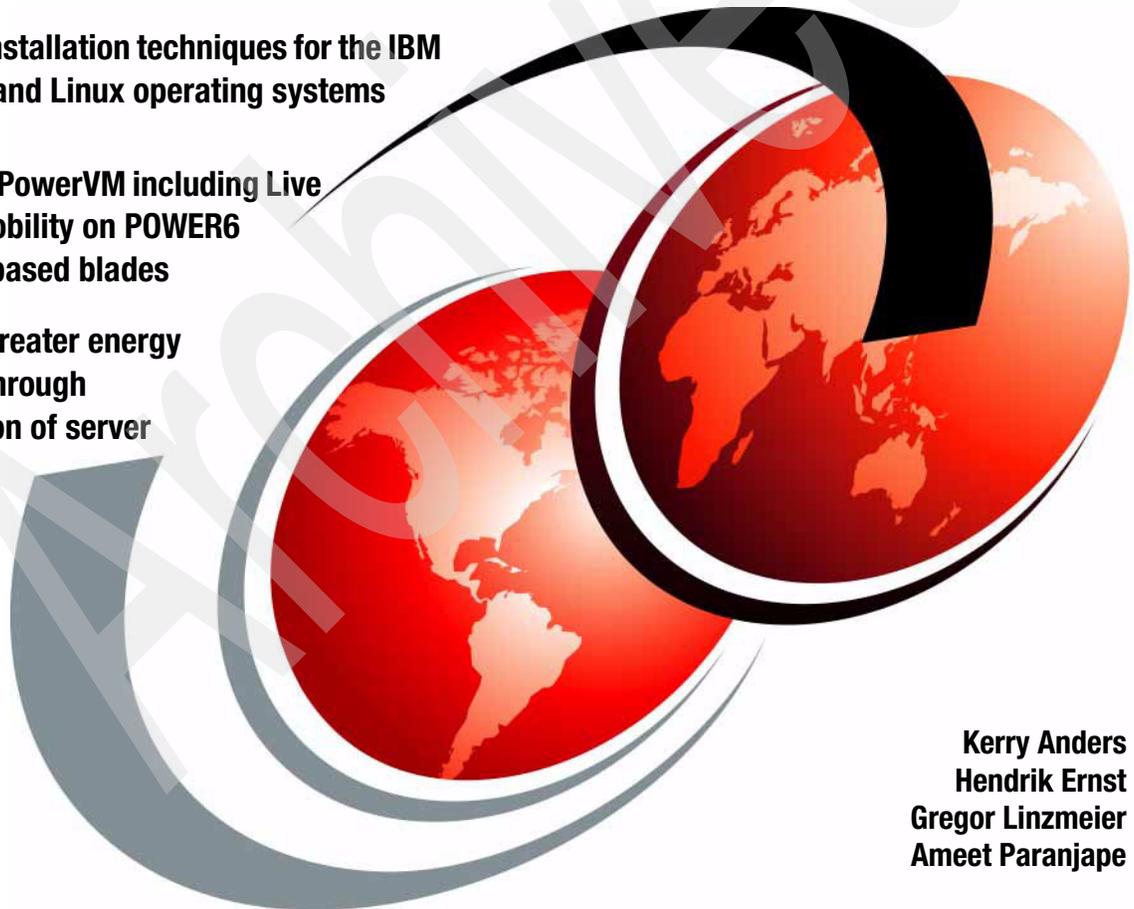


IBM BladeCenter JS12 and JS22 Implementation Guide

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

Supporting PowerVM including Live
Partition Mobility on POWER6
processor-based blades

Achieving greater energy
efficiency through
consolidation of server
workloads



Kerry Anders
Hendrik Ernst
Gregor Linzmeier
Ameet Paranjape



International Technical Support Organization

**IBM BladeCenter JS12 and JS22 Implementation
Guide**

December 2008

Archived

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

First Edition (December 2008)

This edition applies to IBM BladeCenter JS12, IBM BladeCenter JS22, IBM AIX Version 6.1, IBM i 6.1, Red Hat Enterprise Linux for POWER Version 5.2, SUSE Linux Enterprise Server 10 Service Pack 2 (SLES 10 SP2) for POWER.

© Copyright International Business Machines Corporation 2008. 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	ix
Trademarks	x
Preface	xi
The team that wrote this book	xi
Acknowledgements	xii
Become a published author	xiii
Comments welcome	xiii
Part 1. The IBM BladeCenter JS12 and JS22 servers	1
Chapter 1. General description	3
1.1 System specification	4
1.2 Physical specifications of the BladeCenter JS22 Express	5
1.3 Physical specifications of the BladeCenter JS12 Express	5
1.4 Physical packaging	6
1.5 Minimum and optional features of BladeCenter JS22	6
1.5.1 Processor features	7
1.5.2 Memory features	8
1.5.3 Memory DIMMs	8
1.5.4 Internal disk	8
1.6 Minimum and optional features of BladeCenter JS12	9
1.6.1 Processor features	10
1.6.2 Memory features	10
1.6.3 Memory DIMMs	10
1.6.4 Internal disk	11
1.7 IBM BladeCenter chassis	11
1.7.1 BladeCenter H	12
1.7.2 BladeCenter S	16
1.7.3 BladeCenter E	20
1.7.4 BladeCenter T	24
1.7.5 BladeCenter HT	27
Chapter 2. Technical description of the hardware architecture	31
2.1 POWER6 processor	32
2.1.1 Decimal floating point	33
2.2 Simultaneous Multi Threading (SMT)	34
2.3 AltiVec (SIMD)	35
2.4 IBM EnergyScale technology	36

2.4.1 Thermal Power Management Device (TPMD)	37
2.5 Memory subsystem	38
2.5.1 Memory description of IBM BladeCenter JS22 Express	38
2.5.2 Memory placement rules for IBM BladeCenter JS22 Express	39
2.5.3 Memory description of IBM BladeCenter JS12 Express	40
2.5.4 Memory placement rules for IBM BladeCenter JS12 Express	40
2.6 I/O subsystem	41
2.6.1 PCI express.	42
2.6.2 Universal Serial Bus (USB) subsystem.	42
2.6.3 Serial Attached SCSI (SAS) storage subsystem.	43
2.6.4 Integrated Virtual Ethernet (IVE).	43
2.7 Expansion cards	44
2.8 PowerVM.	46
2.9 Operating system support	47
2.9.1 AIX	47
2.9.2 Linux	48
2.9.3 IBM i	49
2.10 Systems management.	49
2.10.1 BladeCenter Advanced Management Web interface	49
2.10.2 IBM Director	50
2.10.3 Cluster Systems Management (CSM).	50
Part 2. System configuration and management.	53
Chapter 3. System configuration using VIOS and IVM	55
3.1 Planning considerations	56
3.1.1 General considerations	56
3.1.2 VIOS installation considerations	57
3.1.3 Internal and external storage considerations	57
3.1.4 External storage preparation.	66
3.2 VIOS system management using IVM	70
3.2.1 IVM user interface.	71
3.2.2 VIOS/IVM command line interface	72
3.3 First VIOS login	73
3.3.1 Password set.	73
3.3.2 License acceptance	73
3.3.3 Initial network setup	74
3.4 First IVM connection	80
3.4.1 Connecting to IVM.	80
3.4.2 Verify and modify VIOS partition memory and processors	82
3.5 VIOS network management and setup	86
3.5.1 Host Ethernet Adapters.	86
3.5.2 Virtual Ethernet Adapters and SEA.	89

3.5.3 Physical adapters	94
3.6 VIOS Storage Configuration and Management	98
3.6.1 Physical volumes	99
3.6.2 Storage pools	101
3.6.3 Virtual disks	103
3.6.4 Optical devices	108
3.7 Partition configuration for Virtual I/O Client (VIOC)	118
3.7.1 Partition Mobility considerations	118
3.7.2 Partition Name and Environment	118
3.7.3 Partition Name	119
3.7.4 Partition Memory	120
3.7.5 Partition Processors	121
3.7.6 Partition Ethernet	122
3.7.7 Partition Storage Types and Assignments	123
3.7.8 Optical devices	125
3.7.9 Physical adapters	126
3.7.10 Partition Summary	127
3.7.11 Partition Properties Changes and DLPAR operations	129
3.8 Console access and activating a partition	138
3.8.1 Opening a virtual terminal	138
3.8.2 Activating a partition	140
3.8.3 Shutdown a partition	144
Chapter 4. IBM AIX Installation	149
4.1 Install IBM AIX 6.1 in a PowerVM client partition	150
4.1.1 Create a virtual media library	150
4.1.2 Prepare the PowerVM client partition	156
4.1.3 Install AIX 6.1 in a logical partition of the Virtual IO Server	169
4.2 Install IBM AIX 6.1 on IBM BladeCenter JS12 or JS22 using a iSCSI target	177
4.2.1 Prepare the iSCSI target	178
4.2.2 Perform the installation of AIX V6.1 from DVD onto a iSCSI target	189
4.3 Boot from iSCSI LUN	200
Chapter 5. IBM i 6.1 installation	209
5.1 Preparing for installation	210
5.1.1 Software installation process	210
5.1.2 Hardware environments	212
5.1.3 BladeCenter hardware preparation	214
5.1.4 VIO Server software environments	216
5.1.5 Network considerations	216
5.1.6 Storage consideration BladeCenter H	217
5.1.7 Disk consideration in BladeCenter S	217

5.1.8	Disk configuration in BladeCenter S	219
5.1.9	Individual BladeCenter S disk configuration	224
5.2	IBM System Access for Windows V6R1	225
5.2.1	Preparation System i Access for Windows	225
5.2.2	Installation System i Access for Windows	226
5.2.3	Connect the System i LAN console.	236
5.3	Creating an IBM i 6.1 partition	246
5.3.1	IBM i 6.1 minimum requirements	246
5.3.2	VIO Server configuration	247
5.3.3	Creating an IBM i partition.	247
5.3.4	IBM i 6.1 install media preparation	249
5.3.5	IBM i 6.1 IPL types	250
5.4	Installing the IBM i 6.1 Operating System	261
5.4.1	IBM i 6.1 installing PTFs	269
5.4.2	IBM i 6.1 TCP/IP setup	270
5.4.3	IBM i 6.1 Navigator for i.	270
5.5	IBM i 6.1 Backup/Restore	271
5.5.1	Creating a virtual media library for backup	272
5.5.2	IBM Tivoli Storage Manager	278
5.5.3	IBM i 6.1 shutdown and restart	279
	Chapter 6. Red Hat Enterprise Linux installation	283
6.1	Supported Red Hat operating systems	284
6.1.1	Considerations and prerequisites	284
6.2	Linux LPAR installation using DVD	285
6.3	Linux network installation (detailed)	289
6.4	Native Red Hat Enterprise Linux 5.2 installation	300
6.5	Red Hat Enterprise Linux 5.2 automated installation	300
6.6	IBM service and productivity tools.	304
	Chapter 7. SUSE Linux Enterprise Server installation	305
7.1	Supported SUSE operating systems.	306
7.2	Considerations and prerequisites	306
7.2.1	PowerVM LPAR considerations	306
7.2.2	SUSE Linux Enterprise Server 10 SP2 installation considerations.	306
7.3	Linux LPAR installation using DVD	307
7.4	Linux network installation (detailed)	311
7.5	Native SLES10 SP2 installation	325
7.6	SLES 10 SP2 automated installation	326
7.7	IBM service and productivity tools.	327
	Chapter 8. Performing Live Partition Mobility	329
8.1	Requirements	330
8.1.1	Hardware.	330

8.1.2	Firmware	330
8.1.3	VIOS Version	330
8.1.4	PowerVM Enterprise	332
8.1.5	LPAR OS versions	333
8.2	Preparation	334
8.2.1	VIOS (source and target)	334
8.2.2	Networking	339
8.3	Migrating the LPAR	340
8.3.1	Using the IVM UI	340
Chapter 9. JS12 and JS22 power management using EnergyScale technology		351
9.1	Power management through the AMM	352
9.1.1	Using the AMM Web UI for blade power management	352
9.1.2	Using the AMM CLI UI for blade power management	361
9.2	Power management through AEM	364
9.2.1	Active Energy Manager Console	364
9.3	Power Saver Mode example	369
9.3.1	Enabling Power Saver Mode from the AMM	370
9.3.2	Power Saver Mode enabled from IBM Director/AEM	373
9.3.3	Scheduling Power Saver Mode with IBM Director and AEM	377
Part 3.	Appendixes	389
Appendix A. Consoles, SMS, and Open Firmware		391
Consoles of the IBM BladeCenter JS12 and JS22		392
	Graphical console	392
	Serial Over LAN	397
System Management Services menu		399
	Power on the IBM BladeCenter JS12 or JS22 using the Advanced Management Module Web interface	400
	Power on the IBM BladeCenter JS12 and JS22 using a Telnet or SSH session into the Advanced Management Module	402
	Choosing the active console after first power on	404
	IBM BladeCenter JS12 and JS22 Power On Self Test (POST)	405
	Select the active console using the System Maintenance Services menu	406
Open Firmware interface		407
	Get access to the firmware prompt	408
	QLogic host bus adapter	410
	Emulex host bus adapter	414
Appendix B. Red Hat Enterprise Linux Kickstart		419
Kickstart introduction		420
Kickstart profile creation methods		420

Create a Kickstart file using Kickstart Configurator	420
Appendix C. SUSE Linux Enterprise Server AutoYaST	429
AutoYaST introduction	430
AutoYaST profile creation methods	430
Create an AutoYaST profile using YaST Control Center	430
Appendix D. Additional Linux installation configuration options	449
Basic preparations for a Linux network installation	450
Installing Linux using the network - General remarks	450
Configuring a BOOTP or DHCP service	451
Configuring a Trivial File Transfer Protocol service	453
Configuring a Network File System Protocol service	455
Virtual optical device setup and installation	458
Appendix E. VIOS/IVM Command Help	459
The help command	459
Using the -h flag	461
Using the man command	462
Appendix F. Service and productivity tools for Linux	465
IBM service diagnostic aids and productivity tools overview	466
Install tools on Red Hat Enterprise Linux 5/SLES 10 running on BladeCenter servers	468
Install tools on Red Hat Enterprise Linux 5/SLES 10 running on IVM-managed servers	471
Related publications	477
IBM Redbooks	477
Online resources	477
How to get Redbooks	484
Help from IBM	484
Index	485

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.

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

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming

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

Trademarks

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

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

1350™	i5/OS®	POWER6™
Advanced Peer-to-Peer Networking®	IBM Systems Director Active Energy Manager™	PowerVM™
AIX 5L™	IBM®	Redbooks®
AIX®	IntelliStation®	Redbooks (logo)  ®
BladeCenter®	iSeries®	RS/6000®
Calibrated Vectors Cooling™	Micro-Partitioning™	ServerProven®
Chipkill™	OS/400®	System i®
DS4000™	POWER™	System p®
DS6000™	POWER Hypervisor™	System Storage™
DS8000™	Power Systems™	System x™
EnergyScale™	POWER5™	Tivoli®

The following terms are trademarks of other companies:

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

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.

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

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

Microsoft, 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.

Preface

This IBM® Redbooks® publication provides a detailed technical guide for configuring and using the IBM BladeCenter® JS12 and IBM BladeCenter JS22 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 will be useful for data center system architects; network, storage and facilities engineers; system administrators; and application architects that need a detailed system bring-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 6.1 (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, Austin Center.

Kerry Anders is a Consultant in System p® Lab Services for the IBM Systems and Technology Group, based in Austin, Texas. He supports clients implementing POWER™-based 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.

Hendrik Ernst has worked for Computer Service GmbH in Erfurt since 1998 performing Intel® server presales support. In 1999, he joined the IBM German Country Postsales Support Team, specializing in the IBM System x™, IBM BladeCenter, Microsoft® Cluster Services, and storage solutions. He previously authored *IBM TotalStorage DS300 and DS400 Best Practices Guide*, SG24-7121 and *IBM System Storage DS3000: Introduction and Implementation Guide*, SG24-7065.

Gregor Linzmeier is an IBM Advisory IT Specialist for IBM Power Systems™ workstation and entry servers as part of the Systems and Technology Group in Mainz, Germany supporting IBM sales, business partners, and clients with pre-sales consultation and implementation of client/server environments. He has worked for more than 15 years as an infrastructure specialist focused on RT, RS/6000®, IBM IntelliStation® POWER, AIX in large CATIA client/server projects, POWER-based BladeCenter architectures, and Partition migration using AIX Thin server technology and HA License Use Management (HALUM).

Ameet Paranjape is a Software Engineer for the Linux Technology Center in the IBM Systems and Technology Group at Austin, Texas. He has experience developing and debugging Red Hat and SUSE Linux distributions on the IBM POWER architecture. Additionally, he was part of a core group of Linux Technology Center members who helped test, debug, and support key Linux features on BladeCenter JS22 and JS12.

Production of this IBM Redbooks publication was managed by:

Chris Almond is an ITSO Project Leader and IT Architect based at the ITSO Center in Austin, Texas. In his current role, he specializes in managing technical content development projects focused on Linux, System p/AIX systems engineering, various IBM Software Group products, and innovation program topics. He has a total of 18 years of IT industry experience, including the last eight with IBM.

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:

Catherine Kostetsky, Amartey Pearson, Barry Barnett, Vess Natchev, Kyle Wurgler, Greg McIntire, Helmut Stein, Sabine Jordan, Norbert Masennek, Nathan Fontenot, Jose Otavio Rizzatti Ferreira, Christian Schroeder, Hermann Depebrock, Steven Royer, Bernhard Zeller, Steven Royer, and Michael Strosaker.

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, and *Techniques to Build a Diskless Boot Linux Cluster of JS21 Blades*, REDP-4238:

Trina Bunting, Wayne Kimble, Stephen Hochstetler, Benjamin Ebrahimi, Tom Junkin, Bernhard Zeller, and Bob Maher.

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, Ghadir Gholami, Robert Amos, John C. Henderson, Kaena Freitas, and Khaja Mohammad

Finally, the team would also like to acknowledge the support for this project provided by **Nicolette McFadden**, Technical Support Management, IBM Systems and Technology Group; **Jim Cook**, ITSO Project Leader; **Scott Vetter**, ITSO System p Team Leader; and our book editor, **Alfred Schwab**, 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

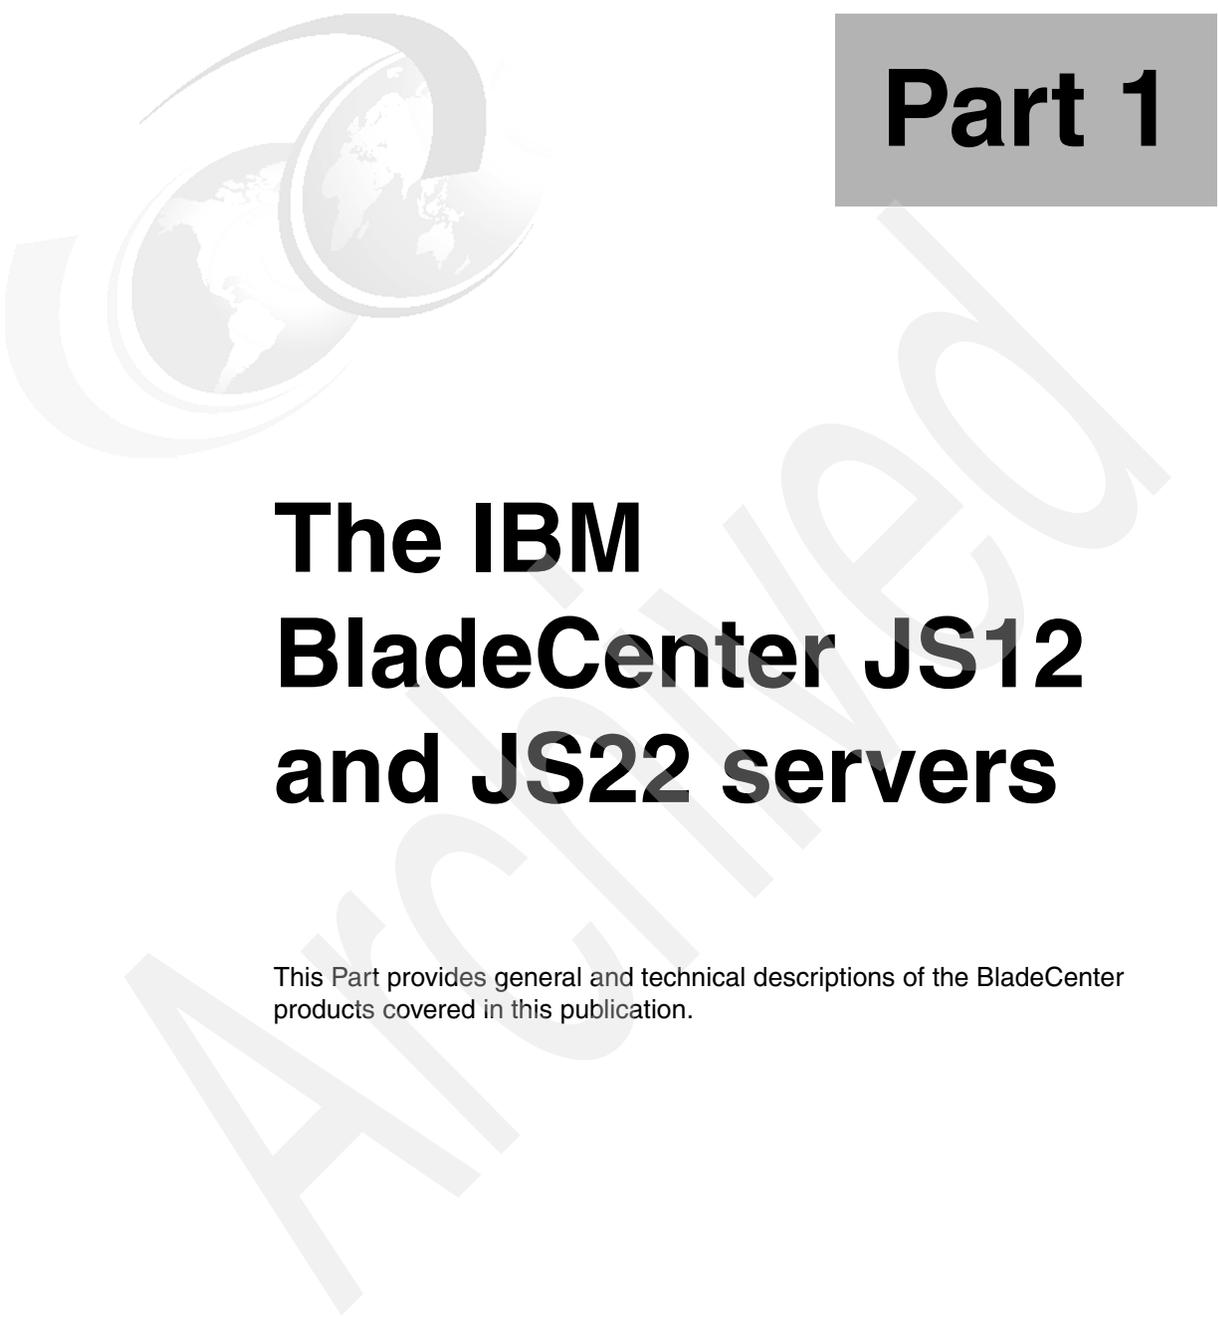
- ▶ Send your comments in an e-mail to:

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

Archived



Part 1

The IBM BladeCenter JS12 and JS22 servers

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

Archived

General description

The IBM BladeCenter POWER6 processor-based blade consists of two models: The IBM BladeCenter JS22 and JS12 Express blade servers deliver many new leading-edge technologies. They support AIX V6, Linux, and IBM i operating systems, IBM EnergyScale™ technology, IBM PowerVM™ Virtualization and Altivec SIMD acceleration functionality in a single, high-performance blade server.

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 BladeCenter JS22 Express blade server refreshes every major subsystem to the latest industry-standard functionality. With the energy-efficient 4.0 GHz high performance dual-core, 64-bit POWER6 processors along with the improved reliability of a SAS disk drive bay supporting a 73 GB or 146 GB disk drive, the JS22 combines leadership performance per watt and enterprise-class reliability features. The BladeCenter JS22 Express supports up to 32 GB of DDR2-based memory in four DIMM slots. In addition to standard support for ECC and IBM Chipkill™ technology designed for improved memory fault protection, the BladeCenter JS22 Express also offers memory running at up to 667 MHz for outstanding performance.

The BladeCenter JS12 Express blade has been preconfigured and tested by IBM, utilizing 3.8 GHz 64-bit IBM POWER6 processor cores; it is available in a

two-core configuration and is designed to deliver great performance. With faster and more reliable double data rate 2 (DDR2) memory options, including expandability up to 64 GB and a SAS disk subsystem supporting up to two on-board hard drives with a total capacity of 292 GB, the BladeCenter JS12 Express is designed for increased uptime and enhanced performance. Along with built-in support for IBM PowerVM virtualization capabilities, the BladeCenter JS12 Express offers an ideal blade server solution for a diverse range of needs.

The BladeCenter JS22 Express (7998-61X), BladeCenter JS12 Express (7998-60X), BladeCenter S Chassis (7779-BCS), and BladeCenter H Chassis (7989-BCH) are part of the Cluster 1350™ portfolio. The BladeCenter H or S Chassis must be ordered with at least one BladeCenter JS22 Express or one BladeCenter JS12 Express.

Built on the IBM BladeCenter family of products, integrated platforms with a high degree of deployment flexibility, scalability, and manageability, the new BladeCenter JS12 Express is the premier Small and Medium Business (SMB) blade, or the BladeCenter JS22 Express is a medium-enterprise data center blade. They represent one of the most cost-efficient solutions for UNIX®, IBM i, and Linux deployment, further enhanced by their ability to be installed in the same chassis with other BladeCenter LSxx, HSxx, QSxx, and JSxx blade servers.

The JS22 and JS12 blades can be configured in the IBM eConfigurator similar to other IBM System p servers.

1.1 System specification

Table 1-1 describes the characteristics of the BladeCenter JS22 Express and BladeCenter JS12 Express server.

Table 1-1 BladeCenter JS22 and JS12 Express system specification

Description	Range
Operating temperature	10 to 35 degrees C (50 to 95 degrees F) ^a 10 to 32 degrees C (50 to 90 degrees F)
Relative humidity	8% to 80%
Maximum altitude	2,133 m (7000 ft)
Power consumption	350 watts maximum
Operating frequency	50 to 60 Hz

a. Depending on the altitude

1.2 Physical specifications of the BladeCenter JS22 Express

Table 1-2 lists the physical specifications of the BladeCenter JS22 Express.

Table 1-2 BladeCenter JS22 physical specifications

Description	7998-61X
Processor	Single-wide, 4-core, 2-socket with two cores per socket, 4.0 GHz POWER6 processor with AltiVec SIMD accelerator.
Number of cores	4
L1 cache	64 KB I-Cache and 32 KB D-Cache per core
L2 cache	4 MB per core
Memory	2 GB minimum (2 x 1 GB) 667 MHz DDR2 32 GB maximum (4 x 8 GB) 533 MHz DDR2
Internal disk	One 73 GB or 146 GB SAS disk

1.3 Physical specifications of the BladeCenter JS12 Express

Table 1-3 lists the physical specifications of the BladeCenter JS12 Express.

Table 1-3 BladeCenter JS12 physical specifications

Description	7998-60X
Processor	Single-wide, 2-core, 1-socket, 3.8 GHz POWER6 processor with AltiVec SIMD accelerator.
Number of cores	2
L1 cache	64 KB I-Cache and 32 KB D-Cache per core
L2 cache	4 MB per core
Memory	2 GB minimum (2 x 1 GB) 667 MHz DDR2 64 GB maximum (4 x 8 GB) 400 MHz DDR2
Internal disk	2 x 73 GB or 146 GB SAS disks

1.4 Physical packaging

The following sections discuss the physical attributes found on a BladeCenter JS22 and JS12 Express server, as shown in Table 1-4.

Table 1-4 BladeCenter JS22 & JS12 Express physical packaging

Dimension	Value
Width	29 mm (1.14 inch)
Depth	445.8 mm (17.55 inch)
Height	245.2 mm (9.6 inch)
Weight	4.35 kg (9.6 lb)

1.5 Minimum and optional features of BladeCenter JS22

The BladeCenter JS22 blade is based on a modular design of two dual-core 64-bit POWER6 processors and integrated L2 cache soldered directly to the system planar board. Table 1-5 shows a standard JS22 configuration.

Table 1-5 JS22 standard configuration

7998-61X	Processor	L2	Memory	Ethernet	Disk
#8400	2-socket, 4-core, 4.0 GHz, POWER6	4 MB per core	2 GB - 32 GB	Dual Gigabit	73 GB or 146 GB

Figure 1-1 shows processor placement, memory slots, SAS disk, and the expansion option connectors.

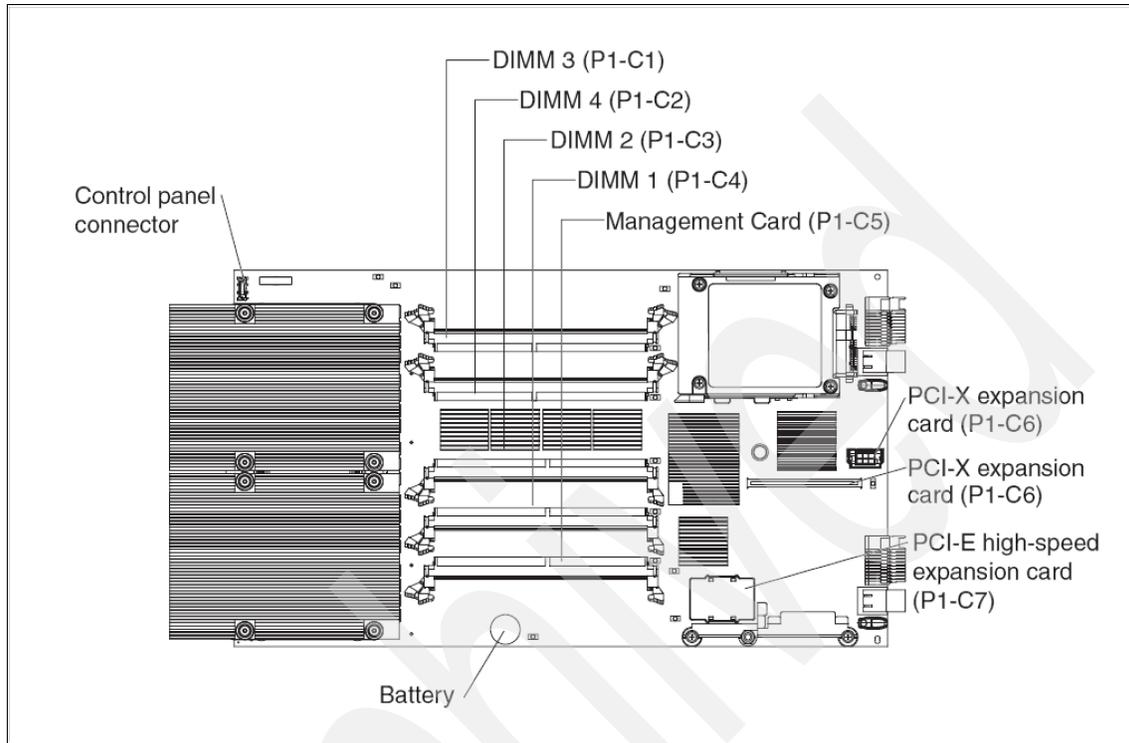


Figure 1-1 Detailed view of the JS22 blade design

1.5.1 Processor features

The key processor features are as follows:

- ▶ The BladeCenter JS22 blade provides the support for a 2-socket, 4-core, POWER6 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, 32-KB Data-Cache, and 4 MB of L2 cache.

Table 1-6 shows the supported processor on a BladeCenter JS22 blade.

Table 1-6 BladeCenter JS22 processor support

Feature	Description
#8400	IBM BladeCenter JS22 4-core 4.0 GHz

1.5.2 Memory features

The integrated memory controller supports four pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 2 GB (2x1 GB) and the maximum is 32 GB (2x16 GB). All the memory features support memory scrubbing, error correction, chipkill, and bit steering.

1.5.3 Memory DIMMs

Table 1-7 provides a list of supported memory on a BladeCenter JS22 blade.

Table 1-7 BladeCenter JS22 memory support

Feature	Description
#8232	2 GB (2x1 GB) DDR2 667 MHz DIMMs
#8233	4 GB (2x2 GB) DDR2 667 MHz DIMMs
#8234	8 GB (2x4 GB) DDR2 667 MHz DIMMs
#8235	16 GB (2x8 GB) DDR2 533 MHz DIMMs

1.5.4 Internal disk

Table 1-8 provides a list of supported disks on a BladeCenter JS22 blade.

Table 1-8 BladeCenter JS22 disk support

Feature	Description
#8237	73 GB SAS 10K SFF hard disk drive
#8236	146 GB SAS 10K SFF hard disk drive

1.6 Minimum and optional features of BladeCenter JS12

The BladeCenter JS12 Express blade is based on a design of one dual-core 64-bit POWER6 processor and integrated L2 cache soldered directly to the system planar board. Table 1-9 shows a standard BladeCenter JS12 Express configuration.

Table 1-9 JS12 standard configuration

7998-60X	Processor	L2	Memory	Ethernet	Disk
#8442	1-socket, 2-core, 3.8 GHz, POWER6	4 MB per core	2 GB - 64 GB	Dual Gigabit	Two 73 GB or 146 GB

Figure 1-2 shows processor placement, memory slots, SAS disk, and the expansion option connectors.

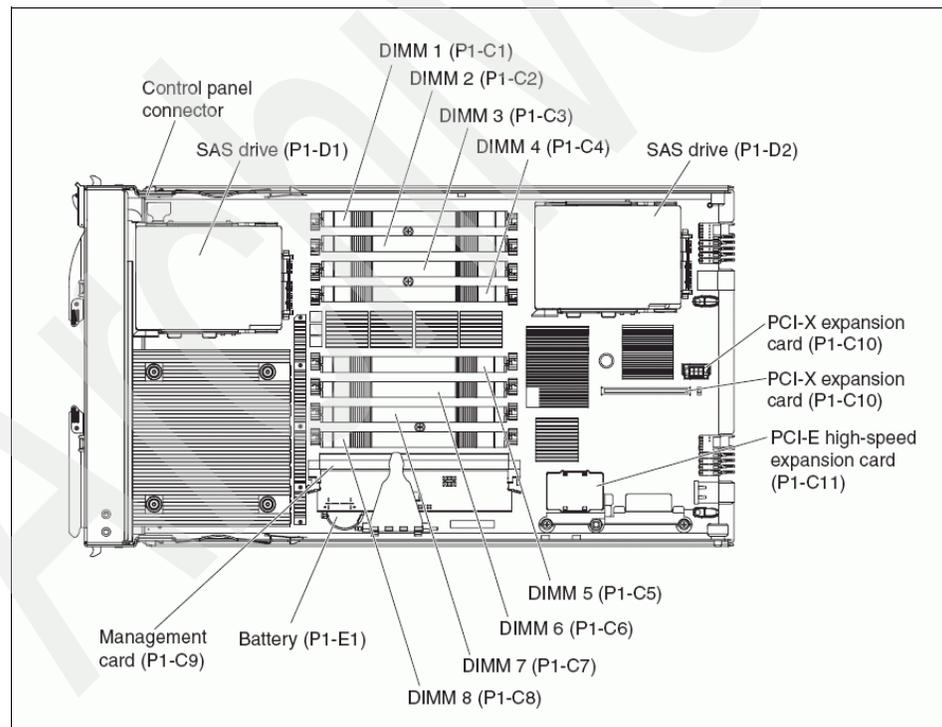


Figure 1-2 Detailed view of the JS12 blade design

1.6.1 Processor features

The major processor features are as follows:

- ▶ The BladeCenter JS12 Express blade provides the support of a 1-socket, 2-core, POWER6 processor implementation.
- ▶ Each processor is directly mounted to the system planar board, providing multi-processing capability.
- ▶ The processor core includes a 64-KB Instruction-Cache, 32-KB Data-Cache, and 4 MB of L2 cache.

Table 1-10 shows the supported processor on a BladeCenter JS12 blade.

Table 1-10 BladeCenter JS12 processor support

Feature	Description
#8442	IBM BladeCenter JS12 2-core 3.8GHz

1.6.2 Memory features

The integrated memory controller supports four pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 2 GB (2x1 GB) and the maximum is 64 GB (4x16 GB). All the memory features support memory scrubbing, error correction, chipkill, and bit steering.

1.6.3 Memory DIMMs

Table 1-11 provides a list of supported memory on a BladeCenter J12 Express blade server.

Table 1-11 BladeCenter JS12 memory support

Feature	Description
#8220	2 GB (2x1 GB) DDR2 667 MHz DIMMs
#8229	4 GB (2x2 GB) DDR2 667 MHz DIMMs
#8239	8 GB (2x4 GB) DDR2 667 MHz DIMMs
#8245	16 GB (2x8 GB) DDR2 400 MHz DIMMs

1.6.4 Internal disk

Table 1-12 provides a list of supported disks on a BladeCenter JS12 blade.

Table 1-12 BladeCenter JS12 disk support

Feature	Description
#8237	73 GB SAS 10K SFF hard disk drive
#8236	146 GB SAS 10K SFF hard disk drive

1.7 IBM BladeCenter chassis

The BladeCenter JS22 and BladeCenter JS12 Express blade are supported in the BladeCenter chassis as shown in Table 1-13. Note that operating system selection (in particular IBM i) should be taken into account for overall system support (see 5.1.2, “Hardware environments” on page 212).

Table 1-13 BladeCenter JS22 and JS12 chassis support

Blade	BC S	BC H	BC HT	BC T	BC E
JS22	YES	YES	YES	NO	NO
JS12	YES	YES	YES	YES	YES

Five chassis available in the BladeCenter chassis family:

- ▶ IBM BladeCenter E provides the greatest density and common fabric support.
- ▶ IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments.
- ▶ IBM BladeCenter T models are designed specifically for telecommunication network infrastructures and other rugged environments.
- ▶ IBM BladeCenter S 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 models are designed for high-performance flexible telecommunication environments by supporting high-speed internetworking technologies such as 10 Gb Ethernet.

Table 1-14 provides a list of supported BladeCenter chassis and the total number of JS22 blades installable into a chassis.

Table 1-14 BladeCenter support

Chassis	Number of blades
BladeCenter E chassis	14
BladeCenter T chassis	8
BladeCenter H chassis	14
BladeCenter HT chassis	12
BladeCenter S chassis	6

1.7.1 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 1-3 and Figure 1-4 on page 14 display the front and rear view of an IBM BladeCenter H.

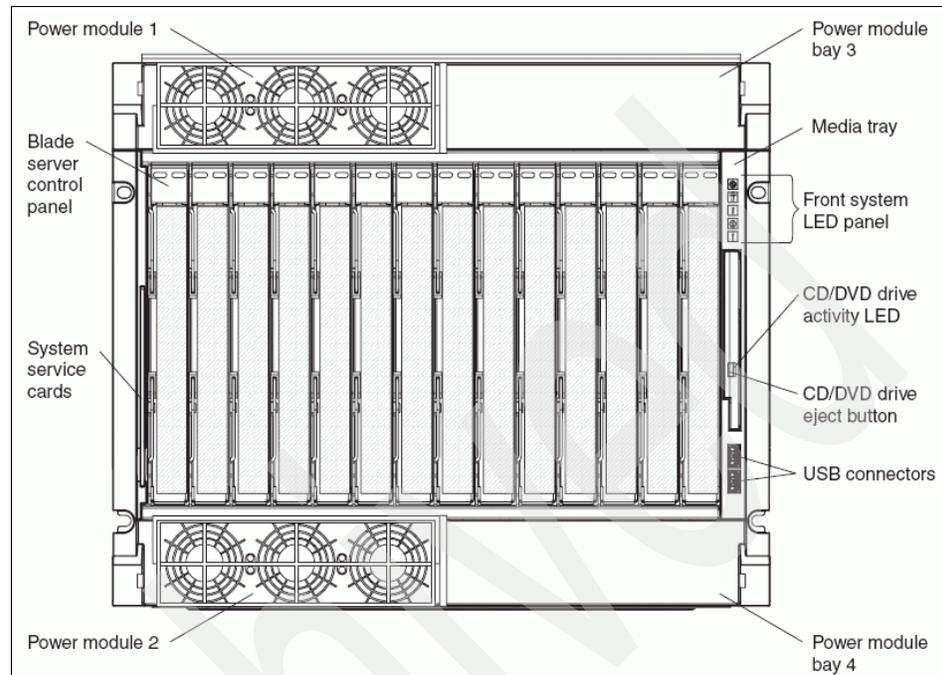


Figure 1-3 Front view of BladeCenter H

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

- ▶ A media tray at the front right, with a DVD drive, two USB v2.0 ports, and a system status LED panel.
- ▶ One pair of 2,900-watt power modules. An additional power module option (containing two 2,900 W power modules) is available.
- ▶ Two hot swap fan modules (two extra hot swap fan modules are included with the additional power module option).
- ▶ 14 hot swap blade server bays supporting different blade server types.

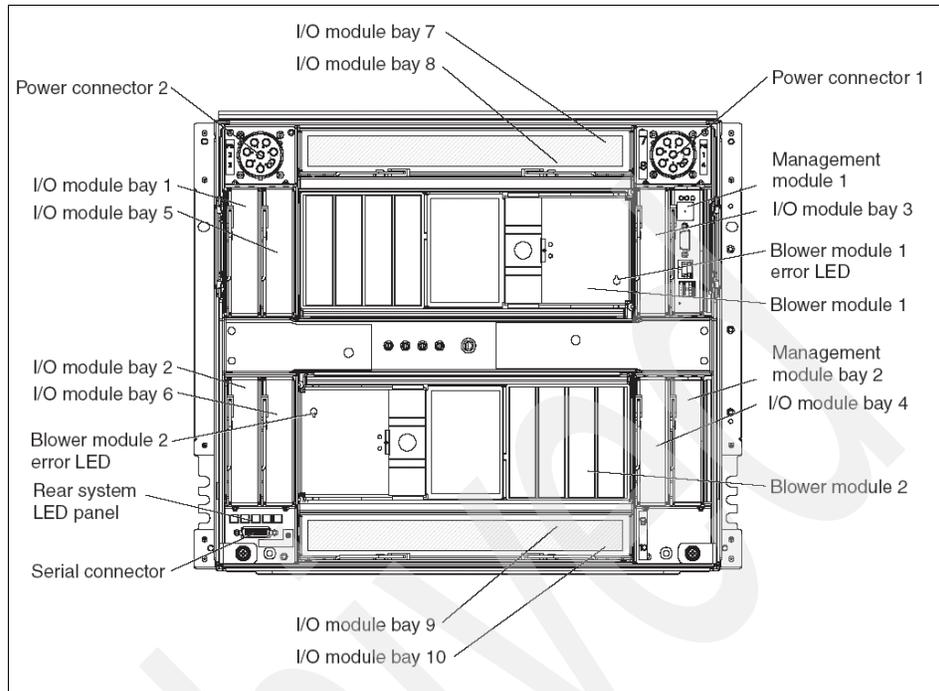


Figure 1-4 Rear view of BladeCenter H

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 different 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 needs. 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 comes standard with no power cord. Power cords need 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 2,900-watt 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 or FRU P/N 03N4962.

Table 1-15 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

Feature	Specification
Number of power supplies (standard/maximum)	2 / 4 ^a
Number of blowers (standard/maximum)	2 / 2
Dimensions	Height: 15.75 inch (400 mm) Width: 17.40 inch (422 mm) Depth: 28.00 inch (711 mm)

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

1.7.2 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 1-5 on page 17 and Figure 1-6 on page 18 display the front and rear view of an IBM BladeCenter S.

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.

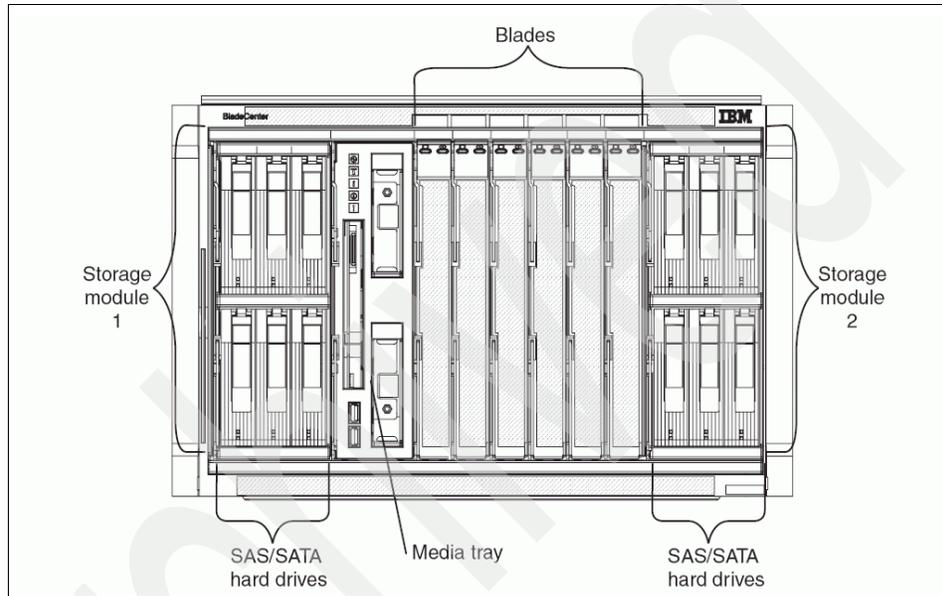


Figure 1-5 BladeCenter S front view

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).

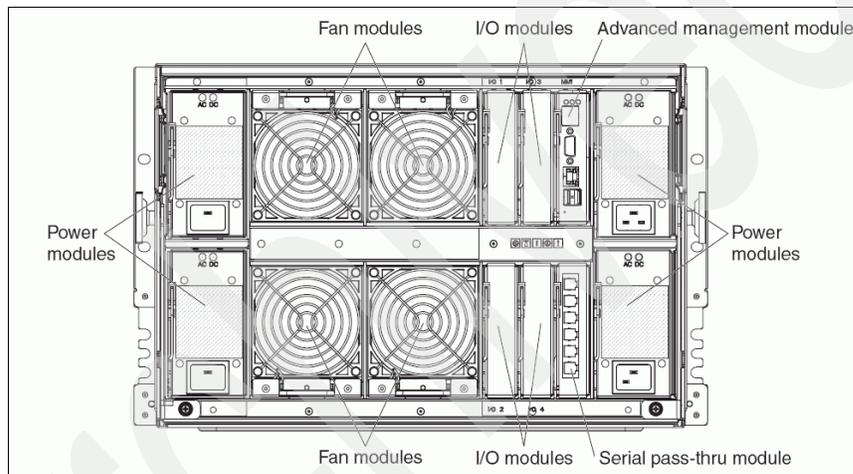


Figure 1-6 BladeCenter S rear view

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 needs. 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 will

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 to 127 v or 200 to 240 v AC power and can be attached to standard office power outlets.

The BladeCenter S chassis ships standard 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 1-16 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 (30mm blade servers)
Number of switch module slots	4 hot-swap (3 x standard, 1 x standard for future use)
Switch modules (std/max)	0 / 4
Storage modules (std/max)	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
Number of blowers (standard/maximum)	4 / 4

Feature	Specification
Dimensions	Height: 12.00 inch (306.3 mm) Width: 17.50 inch (440 mm) Depth: 28.90 inch (733.4 mm)

1.7.3 BladeCenter E

The IBM BladeCenter E (machine type 8677) is designed to be a highly modular chassis to accommodate a range of diverse business requirements. BladeCenter supports not only blade servers, but also a wide range of networking modules, including Gigabit Ethernet, Fibre Channel, and InfiniBand for high-speed connectivity to the client's existing network environment.

BladeCenter E also supports a redundant pair of Management Modules for comprehensive systems management. Providing a wide selection of integrated switching options, BladeCenter systems help you lower the Total Cost of Ownership (TCO) by eliminating the need to purchase additional keyboards, videos, and mice (KVM), Ethernet and Fibre Channel switches, or the cumbersome and expensive cabling the switches require.

BladeCenter E's superior density and feature set are made possible by using an innovative chassis architecture. Because BladeCenter E uses super energy-efficient components and shared infrastructure architecture, clients can realize lower power consumption when compared to their most likely alternative, non-blade server designs.

BladeCenter E's lower power consumption and Calibrated Vectors Cooling™ allow more servers to fit in a tight power or cooling environment.

The following are key features on the front of IBM BladeCenter E chassis:

- ▶ A media tray at the front with a diskette drive, a DVD drive (slim IDE), one USB V1.1 port, and a system-status panel
- ▶ 14 hot swap blade bays supporting different blade server models

Figure 1-7 and Figure 1-8 on page 22 show the front and rear view of the IBM BladeCenter E.

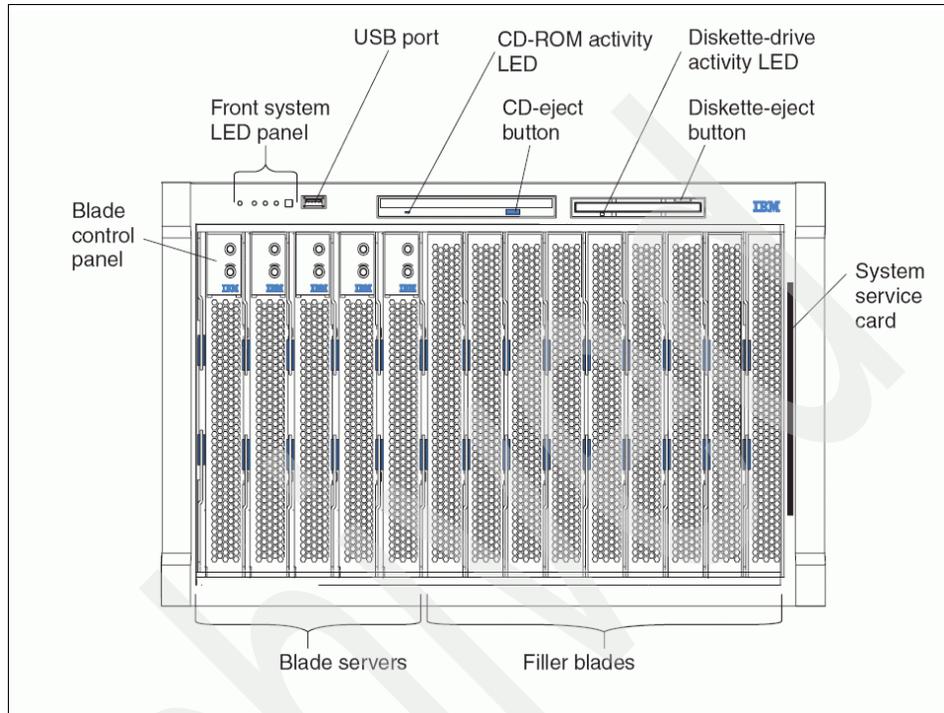


Figure 1-7 IBM BladeCenter E front view

The key features on the rear of IBM BladeCenter E are:

- ▶ Four hot-swap I/O module bays
- ▶ Two hot-swap Management Module bays—with one Management Module as standard
- ▶ Four hot-swap power module bays—with two power modules as standard
- ▶ Two hot-swap blowers

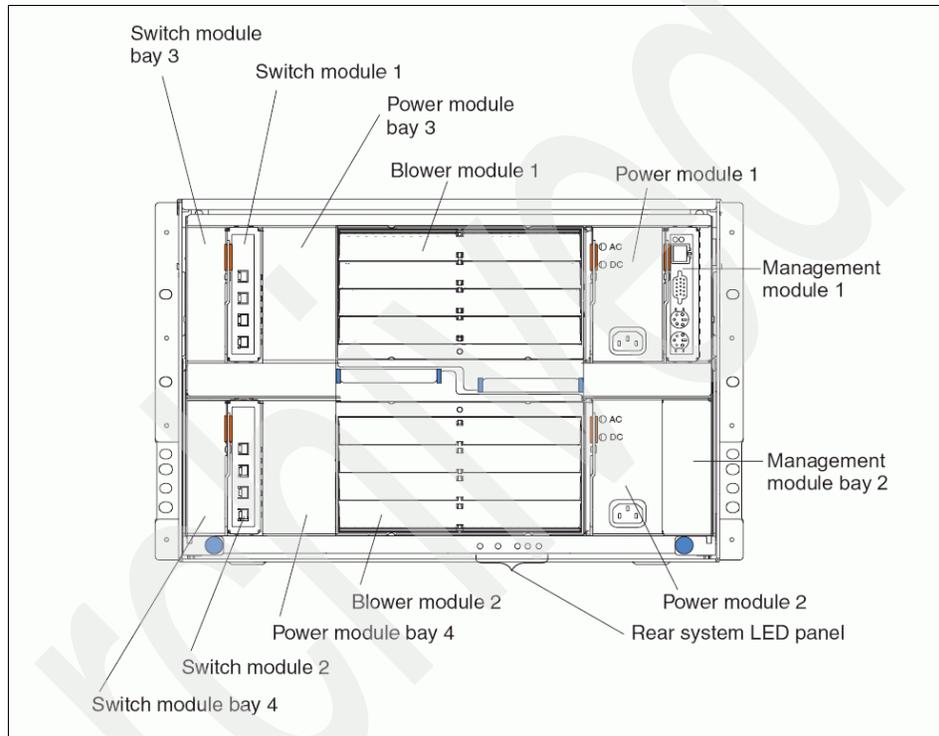


Figure 1-8 IBM BladeCenter E rear view

The BladeCenter E chassis allows for either 14 single-slot blade servers or seven double-slot blade servers. In specific configurations some blades can use three or more slots. However, you can mix different blade server models in one chassis to meet your requirements, subject to power and cooling requirements, and some other configuration rules that we discuss later.

Note: Installing blades that use more than one slot should preferably be installed in one power domain. BladeCenter E has two power domains, which means you should use slot 1 to 6 or slot 7 to 14 to install the blade.

BladeCenter E does not ship standard with any I/O modules. You need to choose these I/O modules depending on your connectivity needs. An Ethernet Switch Module (ESM) 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 BladeCenter E chassis ships standard with:

- ▶ One Advanced Management Module
- ▶ Two hot-swap power supply modules
- ▶ Two hot-swap blower modules
- ▶ One USB v1.1 port
- ▶ One DVD-ROM drive
- ▶ One 1.44 MB diskette drive

Table 1-17 lists the features of IBM BladeCenter E.

Table 1-17 BladeCenter E specifications

Feature	Specification
Machine type	8677-3Ru, 3RX
Rack dimension	7U x 28.9 inches (711 mm)
DVD/CD standard drive	1 x DVD ROM (in Media Tray)
Diskette drive	1 x 1.44 Diskette drive (in Media Tray)
Number of blade slots	14 (30mm blade servers)
Number of switch module slots	4 hot-swap
Switch modules (std/max)	None
Power supply	2000 W AC
Number of power supplies (standard/maximum)	2 / 4 ^a
Number of blowers (standard/maximum)	2 / 2
Dimensions	Height: 12.00 inch (306.3 mm) Width: 17.50 inch (440 mm) Depth: 28.90 inch (733.4 mm)

a. Four power supplies are required to use blade servers in slot 6 to 14

1.7.4 BladeCenter T

The IBM BladeCenter T is a compact eight-server blade chassis designed for high-density server installations, typically for telecommunications use. This 8U high chassis with DC or AC power supplies provides a cost-effective, high-performance, high-availability solution for telecommunication network and other “rugged” non-telco environments. The IBM BladeCenter T chassis is positioned for expansion, capacity, redundancy, and carrier-grade NEBS level/ETSI compliance in DC models.

BladeCenter T provides a solid foundation for next-generation networks, enabling service providers to become on demand providers. Coupled with technological expertise within the enterprise data center, IBM is leveraging the industry know-how of key Business Partners to jointly deliver added value within service provider networks.

The key features on the front of BladeCenter T are:

- ▶ A media tray at the front with a DVD/CD-RW drive (slim IDE), two USB v1.1 ports, and a system-status panel
- ▶ Two hot-swap management module bays
- ▶ Four hot-swap power-module bays
- ▶ Front bezel with changeable filter
- ▶ Eight hot-swap blade bays supporting different blade server types

Figure 1-9 and Figure 1-10 on page 26 show the front and rear view of IBM BladeCenter T.

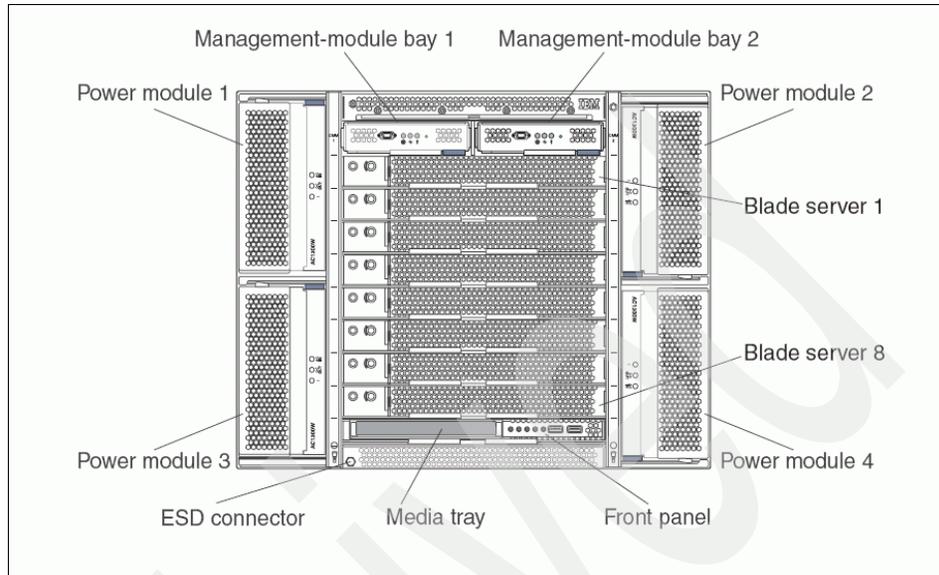


Figure 1-9 IBM BladeCenter T front view

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

- ▶ Four hot-swap I/O module bays
- ▶ One hot-swap KVM (keyboard, video, and mouse) module
- ▶ One hot-swap LAN module
- ▶ Four variable speed hot-swap blowers
- ▶ AC power (machine type 8730) or DC power (machine type 8720) connectors

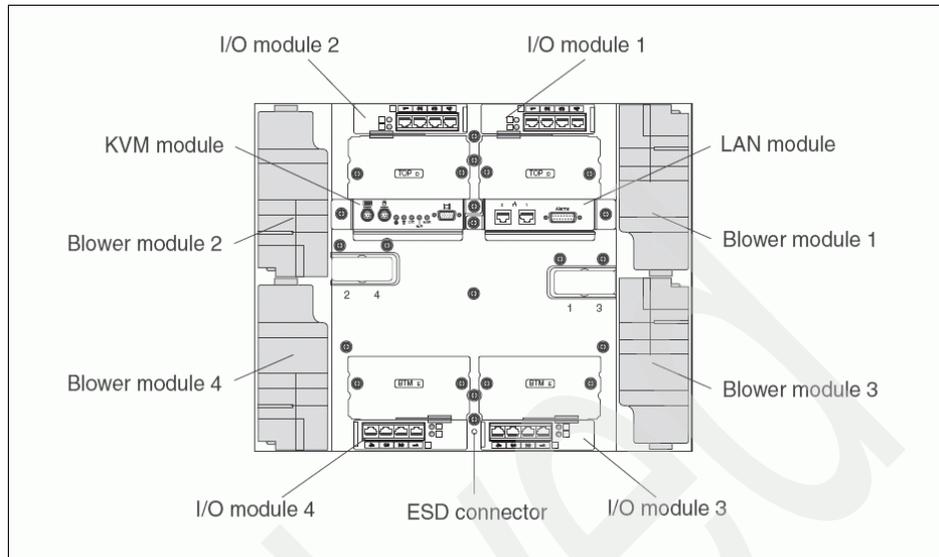


Figure 1-10 IBM BladeCenter T rear view

BladeCenter T delivers rich telecom features and functionality, including integrated servers, storage and networking, fault-tolerant features, optional hot swappable redundant DC or AC power supplies and cooling, and built-in system management resources in a 20 inch (50.8 cm) deep chassis. The result is a Network Equipment Building Systems (NEBS-3) and ETSI-compliant server platform optimized for next-generation networks.

Table 1-18 lists the features of IBM BladeCenter T.

Table 1-18 BladeCenter T specifications

Feature	Specification
Machine type	8720-2RX (DC) 8730-2RX (AC)
Rack dimension	8U x 20 inches (598 mm)
DVD/CD standard drive	1 x DVD ROM (in Media Tray)
Diskette drive	None
Number of blade slots	8 (30mm blade servers)
Number of switch module slots	4
Switch modules (std/max)	None

Feature	Specification
Power supply	2000 W AC
Number of power supplies (standard/maximum)	2 / 4
Number of blowers (standard/maximum)	4 / 4
Dimensions	Height: 13.75 inch (349 mm) Width: 17.50 inch (440 mm) Depth: 20.00 inch (508 mm)

1.7.5 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 GbEthernet 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-telco 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 1-11 and Figure 1-12 on page 29 show the front and rear view of the IBM BladeCenter HT.

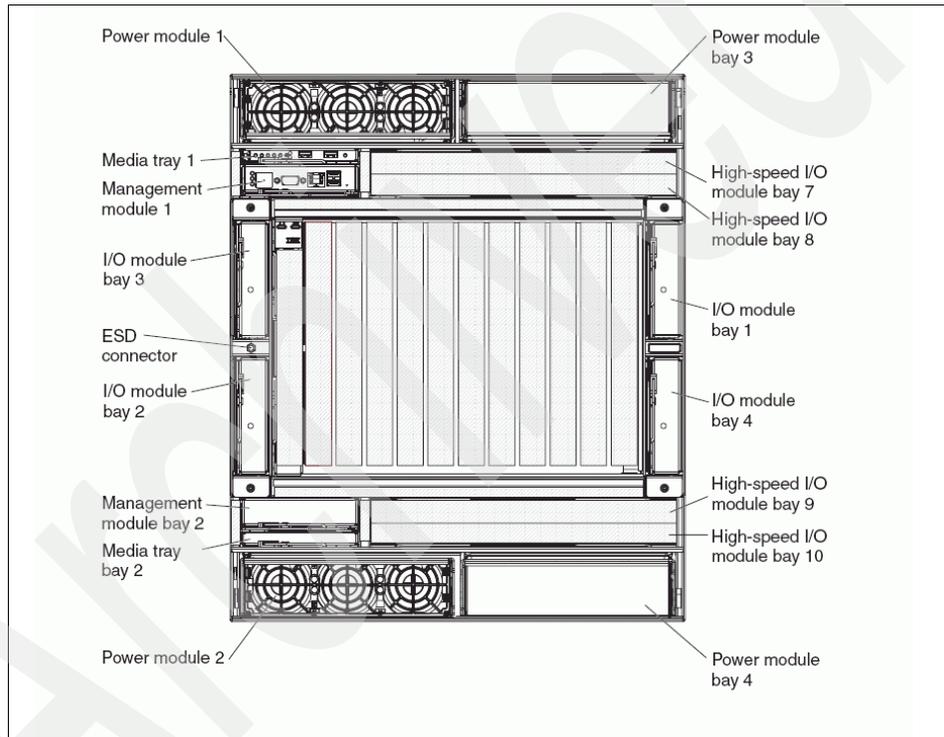


Figure 1-11 IBM BladeCenter HT front view

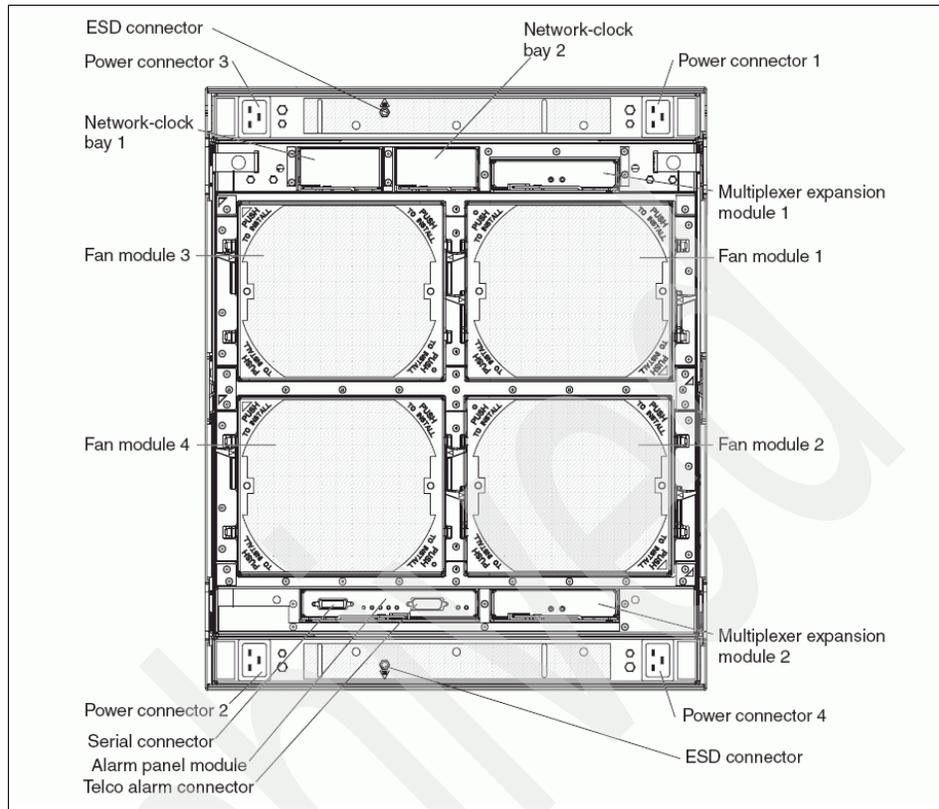


Figure 1-12 IBM BladeCenter HT rear view

Table 1-19 lists the features of the IBM BladeCenter E.

Table 1-19 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 (30mm blade servers)
Number of switch module slots	4

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

- 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 need to choose these I/O modules depending on your connectivity needs. 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.



Technical description of the hardware architecture

This section discusses the hardware architecture, including:

- ▶ POWER6 processor
- ▶ Simultaneous Multi Threading (SMT)
- ▶ AltiVec (SIMD)
- ▶ IBM EnergyScale technology
- ▶ Memory subsystem
- ▶ I/O subsystem
- ▶ Expansion cards
- ▶ PowerVM
- ▶ Operating system support
- ▶ Systems management

2.1 POWER6 processor

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

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

▶ **Processor instruction retry**

Soft failures in the processor core are transient errors. When an error is encountered in the core, the POWER6 processor will first automatically retry the instruction. If the source of the error was truly transient, the instruction will succeed and the system will continue 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 will 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 new feature in the BladeCenter JS22 blade 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 will be 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.

Figure 2-1 shows a high-level view of a POWER6 processor. On the left side is the schematic for IBM BladeCenter JS22; on the right side is the schematic for the IBM BladeCenter JS12 Express server.

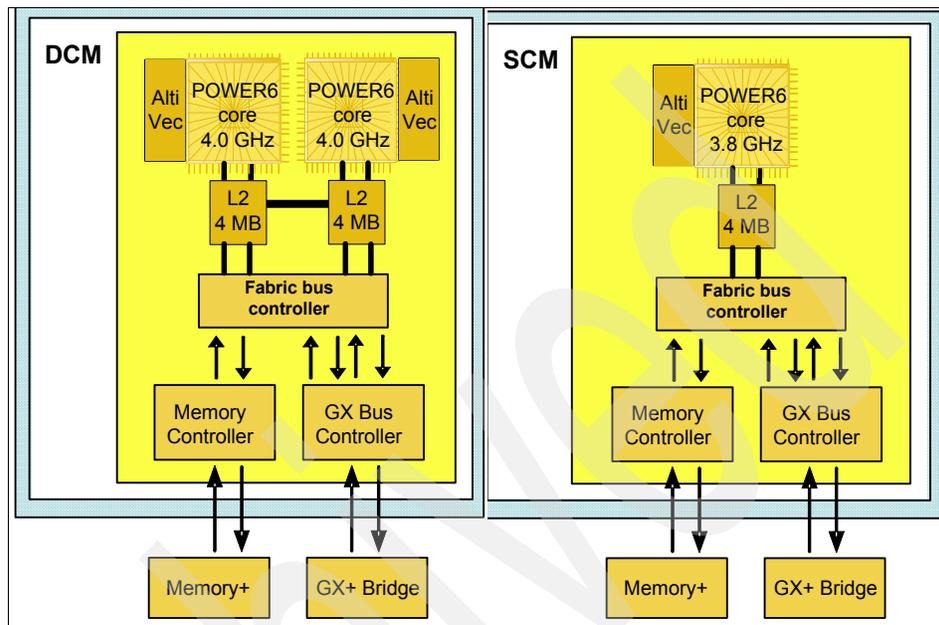


Figure 2-1 POWER6 processor

2.1.1 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 are provided to perform arithmetic, compare, test, quantum-adjustment, conversion, and format operations on operands held in FPRs or FPR pairs.

Arithmetic instructions These instructions perform addition, subtraction, multiplication, and division operations.

Compare instructions These instructions perform a comparison operation on the numerical value of two DFP operands.

Test instructions These instructions test the data class, the data group, the exponent, or the number of significant digits of a DFP operand.

Quantum-adjustment instructions These instructions convert a DFP number to a result in the form that has the designated exponent, which may be explicitly or implicitly specified.

Conversion instructions These instructions perform conversion between different data formats or data types.

Format instructions These instructions facilitate composing or decomposing a DFP operand.

For example, the SAP® NetWeaver 7.10 ABAP™ kernel introduces a new SAP ABAP data type called 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 allows for highly simplified ABAP-coding while increasing numeric accuracy, and with a potential for significant performance improvements.

2.2 Simultaneous Multi Threading (SMT)

The POWER6 processor core has been designed to support both enhanced Simultaneous Multi Threading (SMT) and Single Threaded (ST) operation modes. The IBM BladeCenter JS22 and JS12 server supports the SMT technology in both systems.

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 (ST) 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 ST operation is ideal.

2.3 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.

The AltiVec technology is not comparable to the IBM POWER6 processor implementation, which uses the Simultaneous Multithreading functionality.

2.4 IBM EnergyScale technology

IBM EnergyScale technology is featured on the 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 provides the administrators with the information to predict power consumption across their infrastructure and to react to business and processing needs. For example, an administrator could adjust server consumption to reduce electrical costs. To collect power data for the BladeCenter JS22 blade you do not need any additional hardware because it collects the information internally.

Power Saver Mode Power Saver Mode reduces the voltage and frequency by a fixed percentage. 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 CPU 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.

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 datacenter. 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 whenever they are unlicensed again.

2.4.1 Thermal Power Management Device (TPMD)

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. The BladeCenter JS22 blade implementation uses an integrated circuit called Thermal Power Management Device (TPMD), placed on the management card.

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.

IBM Systems Director Active Energy Manager Version 3.1.1 is an IBM Director extension. For more information about the IBM Active Energy Manager, see:

<http://www.ibm.com/systems/management/director/extensions/actengmrg.html>

2.5 Memory subsystem

For IBM BladeCenter JS22 and IBM BladeCenter JS22 Express server, 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 DDR-II memory DIMMs.

2.5.1 Memory description of IBM BladeCenter JS22 Express

Any POWER6 chip interfaces two memory buffer chips, so the fully configured IBM BladeCenter JS22 Express has two POWER6 chips (for a total of four POWER6 processor cores) with the related memory buffer chips as shown in Figure 2-2 on page 39. Each channel from a POWER6 chip to a memory buffer chip allows 2 bytes of Read data and 1 byte of Write data.

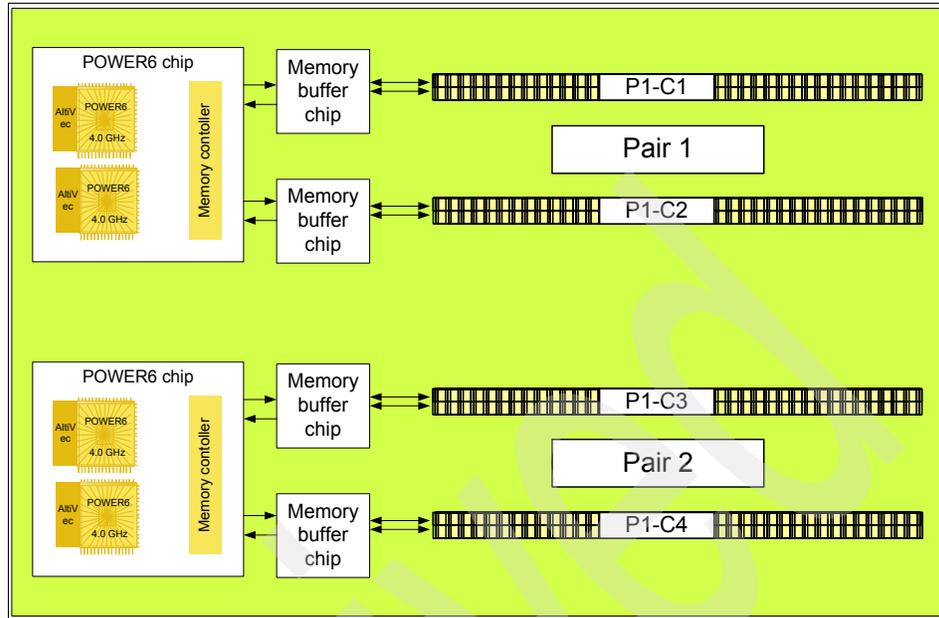


Figure 2-2 IBM BladeCenter JS22 memory subsystem

Table 2-1 provides a list of supported memory DIMMs on an IBM BladeCenter JS22 server.

Table 2-1 IBM BladeCenter JS22 Express memory support

Feature	Description
#8232	2 GB (2x1 GB) DDR2 667 MHz DIMMs
#8233	4 GB (2x2 GB) DDR2 667 MHz DIMMs
#8234	8 GB (2x4 GB) DDR2 667 MHz DIMMs
#8235	16 GB (2x8 GB) DDR2 533 MHz DIMMs

A minimum of two DIMMs of the same size must be plugged into connectors at location codes P1-C1 and P1-C2 or P1-C3 and P1-C4 as shown in memory layout Figure 2-2.

2.5.2 Memory placement rules for IBM BladeCenter JS22 Express

Both memory DIMMs in a pair must have the same size and the same speed. The two different pairs are separated and independent and can have different size and speed specifications.

2.5.3 Memory description of IBM BladeCenter JS12 Express

The POWER6 chip interfaces four memory buffer chips. The BladeCenter JS12 Express has one POWER6 chip (for a total of two POWER6 processor cores) with the related memory buffer chips, as shown in Figure 2-3. Each channel from the POWER6 chip to memory buffer chip allows 2 bytes of Read data and 1 byte of Write data.

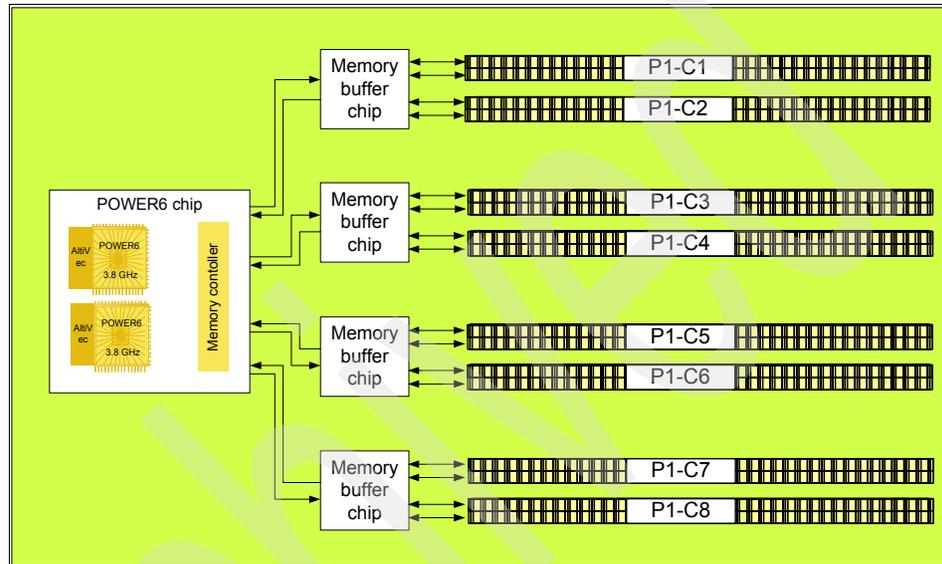


Figure 2-3 IBM BladeCenter JS12 memory subsystem

2.5.4 Memory placement rules for IBM BladeCenter JS12 Express

The minimum memory capacity for an IBM BladeCenter JS12 Express server at initial order is 2 GB. The memory placement rules are provided in Figure 2-4 on page 41. To populate the remaining memory slots with additional memory DIMMs, use the following additional rules:

- ▶ DIMMs must be installed in pairs.
- ▶ DIMMs pairs must have the same size for pair 1 and 3, and for pair 2 and 4.
- ▶ Using 8 GB DIMMs, all DIMMs must be 8 GB.

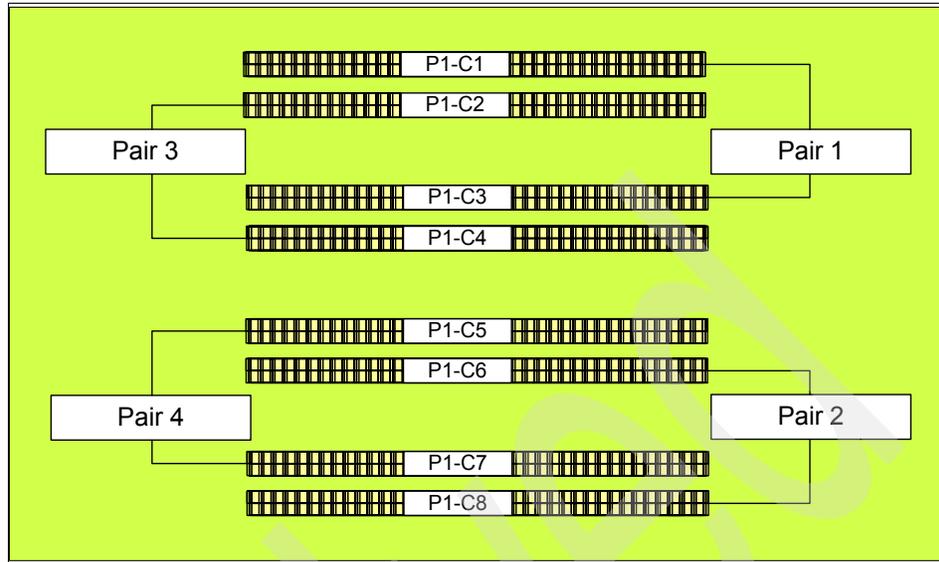


Figure 2-4 IBM BladeCenter JS12 memory placement rules

Table 2-2 provides a list of supported memory DIMMs on a BladeCenter JS12 Express server.

Table 2-2 IBM BladeCenter JS12 memory support

Feature	Description
#8220	2 GB (2x1 GB) DDR2 667 MHz DIMMs
#8229	4 GB (2x2 GB) DDR2 667 MHz DIMMs
#8239	8 GB (2x4 GB) DDR2 667 MHz DIMMs
#8245	16 GB (2x8 GB) DDR2 400 MHz DIMMs

A minimum of two DIMMs of the same size must be plugged into connectors at location codes P1-C1 and P1-C3 as shown in the memory placement rules of Figure 2-4.

2.6 I/O subsystem

A mixture of both PCIe and PCI-x cards are supported in the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server. The IBM BladeCenter JS22

and IBM BladeCenter JS12 Express supports PCI-e x4, PCI-e x8, and PCI-x 1.0 64-bit extension cards.

2.6.1 PCI express

PCI-Express (PCI-e) 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
- ▶ 2.6 Gbits per pin pair in each direction (speed scale in future generations)
- ▶ 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

2.6.2 Universal Serial Bus (USB) subsystem

The USB bus connects USB devices to a USB host. The USB buses used in the BladeServer chassis conform to the Universal Serial Bus 2.0 specification.

The IBM BladeCenter JS22 and IBM BladeCenter JS12 Express uses 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 via the mid-plane.

The USB bus routed to the Management Module is used for the keyboard on the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express servers. This allows the Management Module to connect the keyboard to the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server and control the routing of keyboard inputs to one of fourteen IBM BladeCenter JS22 or IBM BladeCenter JS12 Express server installed in a BladeCenter H chassis. The other USB bus is routed to the Media Tray, which connects to the DVD-ROM drive and diskette drive. Both USB buses are controlled by the Management Module independent of each other. This provides the capability to assign the CD and floppy disk drive (on the BCE) to one IBM BladeCenter JS22 or IBM BladeCenter JS12 Express server while the keyboard is assigned to another processor blade. However, this should not restrict the assignment of shared resources of both USB buses to a single IBM BladeCenter JS22 or IBM BladeCenter JS12 Express server.

2.6.3 Serial Attached SCSI (SAS) storage subsystem

The embedded SAS controller is used on the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server. The SAS controller host interface is connected to a 32-bit PCI-X bus on the host bridge chip running at 67 MHz. There is no hot plug function on the IBM BladeCenter JS22 or IBM BladeCenter JS12 Express server.

2.6.4 Integrated Virtual Ethernet (IVE)

IVE is the name given to the collection of hardware components (including the Host Ethernet Adapter or 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 1Gbps ports
- ▶ External network connectivity for LPARs using dedicated ports without the need of 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 or PC-x bus. This method provides IVE with the high throughput and low latency of a bus imbedded in the I/O controller. IVE also includes special hardware features that provide logical Ethernet adapters. The IVE feature implemented on the JS12 and JS22 provides two 1-Gbps physical ports.

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 need for both the virtual Ethernet and the SEA. It provides most of the functionality of each.

Therefore, this eliminates the need 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.

2.7 Expansion cards

Each IBM BladeCenter JS22 and IBM BladeCenter JS12 Express system contains expansion options connectors that can expand the external connectivity to the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server. Table 2-3 lists the feature numbers and form factors of supported expansion cards for an IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server.

Table 2-3 IBM BladeCenter JS22 and JS12 Expansion Cards form factors

Feature code/ PN	I/O Expansion cards	PCI-E	PCI-X	StFF ^a	CFFv ^b	CFFh ^c	HSFF ^d
8238 32R1923	Qlogic iSCSI Dual Port Expansion Card		X	X			
8248 41Y8527	Qlogic 4 Gb Fibre Channel Dual Port Expansion Card		X		X		
8249 39Y3910	1 Gb Ethernet Dual Port Expansion Card		X		X		
8250 39Y9190	SAS Dual Port Expansion Card		X		X		
8251 43W6859	Emulex 4 Gb Fibre Channel Dual Port Expansion Card		X		X		
8252 39Y9306	4 Gb Fibre Channel / 1Gb Ethernet Expansion Card	X				X	
8256 43W4921	Cisco 4x Infiniband Dual Port Expansion Card	X					X
8258 44W4423	Mellanox 4x Infiniband Dual Port DDR Expansion Card	X					X

a. Standard Form Factor

b. Combined Form Factor vertical

c. Combined Form Factor horizontal

d. High Speed Form Factor

Table 2-4 provides a list of available expansion cards for IBM BladeCenter JS22 and IBM BladeCenter JS12 Express and related operating system support.

Table 2-4 OS and adapter support

Feature code	Expansion card	VIO 1.5	AIX 5.3	AIX 6.1	IBM i	Suse 10	RH 4.6	RH 5.1
8238	Qlogic iSCSI Dual Port Expansion Card	YES	YES	YES	NO	YES	YES	YES
8248	Qlogic 4 GB Fibre Channel Dual Port Expansion Card	YES	YES	YES	NO	YES	YES	YES
8249	1 GB Ethernet Dual Port Expansion Card	YES	YES	YES	NO	YES	YES	YES
8250	SAS Dual Port Expansion Card	YES	YES	YES	YES	YES	YES	YES
8251	4 GB Fibre Channel Dual Port Expansion Card	YES	YES	YES	NO	YES	YES	YES
8252	4 GB Fibre Channel / 1GB Ethernet Expansion Card	YES	YES	YES	YES	YES	YES	YES
8256	4x Infiniband Dual Port Expansion Card	NO	YES	YES	NO	NO	YES	YES
8258	4x Infiniband Dual Port DDR Expansion Card	NO	YES	YES	NO	NO	YES	YES

To facilitate the support and common description of available expansion cards, Table 2-5 provides a list of feature to part number conversions.

Table 2-5 BladeCenter JS22 and JS12 feature to part number conversion table

Expansion Card	Feature	Part number
Qlogic iSCSI Dual Port Expansion Card	8238	32R1923
Qlogic 4 Gb Fibre Channel Dual Port Expansion Card	8248	41Y8527
1 Gb Ethernet Dual Port Expansion Card	8249	39Y9310
SAS Dual Port Expansion Card	8250	39Y9190
4 Gb Fibre Channel Dual Port Expansion Card	8251	43W6859
4 Gb Fibre Channel / 1 Gb Ethernet Expansion Card	8252	39Y9306
4x Infiniband Dual Port Expansion Card	8256	43W4921
4x Infiniband Dual Port DDR Expansion Card	8258	44W4423

Note: The Qlogic iSCSI Expansion Card for IBM BladeCenter JS22 Express is a standard, or large, form factor expansion card. The installation of this card precludes support for the SAS internal disk drive.

2.8 PowerVM

The PowerVM platform is the family of technologies, capabilities, and offerings that provide virtualization capabilities on the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express systems. PowerVM is the new umbrella branding term for Power Systems Virtualization (Logical Partitioning, 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 2-6 on page 47 shows the licensed features of each of the two different editions of PowerVM for IBM BladeCenter JS22 and IBM BladeCenter JS12 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 JS22 and IBM BladeCenter JS12 Express server.

This section provides information about the virtualization capabilities of the PowerVM Express Edition, Standard Edition, and Enterprise Edition, which are available on this system. Upgrade from one PowerVM Standard Edition to the Enterprise Edition is possible and completely undistruptive. The upgrade does not even require the installation of additional software. Enter a key code in the hypervisor in order to unlock the next level of function.

Table 2-6 PowerVM editions for BladeCenter JS22 and JS12

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

2.9 Operating system support

The IBM BladeCenter JS22 and IBM BladeCenter JS12 Express supports the IBM AIX, Novell® SuSE Linux Enterprise Server, Red Hat Enterprise Linux, and IBM i operating systems. Using the PowerVM Standard Edition on an IBM BladeCenter JS22 and IBM BladeCenter JS12 Express running a VIOS, a combination of IBM i, IBM AIX, and Linux partitions can be deployed.

The following sections describe in detail the recommended operating support versions and levels.

2.9.1 AIX

The following versions of AIX 5L™ are supported on the BladeCenter JS22:

- ▶ AIX 5L V5.3 with 5300-06 Technology Level and Service Pack 7
- ▶ AIX 5L V5.3 with 5300-07 Technology Level and Service Pack 4

- ▶ AIX 5L V5.3 with 5300-08 Technology Level, or later
- ▶ AIX V6.1 with 6100-00 Technology level and Service Pack 5
- ▶ AIX V6.1 with 6100-01 Technology level, 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, which 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>

2.9.2 Linux

The following versions of Linux are supported on BladeCenter JS12:

- ▶ Novell SUSE Linux Enterprise Server 10 for IBM POWER Service Pack 2
- ▶ Red Hat Enterprise Linux 4 for IBM POWER Series Update 6
- ▶ Red Hat Enterprise Linux 5 for IBM POWER Series Update 1

For information about the features and external devices supported by Linux on the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server, visit:

<http://www.ibm.com/systems/power/software/linux/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

2.9.3 IBM i

IBM i 6.1 is offered with per processor and per user entitlement. When IBM i 6.1 is selected on an IBM BladeCenter JS22 and IBM BladeCenter JS12 Express, at least one IBM i 6.1 processor entitlement must be ordered and at least 10 user entitlements. Additional processor entitlements can be ordered to a maximum of four on the IBM BladeCenter JS22 and two on the IBM BladeCenter JS12 Express. Also, additional user entitlements in blocks of 10 users can be ordered.

IBM i 6.1 on IBM BladeCenter JS22 and IBM BladeCenter JS12 Express uses IBM PowerVM Standard Edition, which includes the 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 JS22 and JS12 blades.

Use PowerVM product number 5761-SS1, Version 6.1 plus the latest PTF Package and the following individual PTFs:

- ▶ SI30649
- ▶ MF43772
- ▶ SI30583
- ▶ MF43935

For more information about IBM i operating systems running on IBM BladeCenter JS22 and IBM BladeCenter JS12 Express, see:

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

2.10 Systems management

Several tools are available to manage IBM BladeCenter JS22 and IBM BladeCenter JS12 Express servers. The BladeCenter Web interface can effectively manage the blades. In addition, IBM Director and Cluster Systems Management (CSM) are management tools that can manage your environment.

2.10.1 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.

The BladeCenter Web interface allows the following:

- ▶ A System Administrator can easily and effectively manage up to 14 blade servers from an integrated interface.
- ▶ Power the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express on or off.
- ▶ Perform firmware management.
- ▶ 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 JS22 and IBM BladeCenter JS12 Express server.

2.10.2 IBM Director

IBM Director is a graphical user interface (GUI)-based management tool that provides management functions for the BladeCenter. IBM Director enables you to remotely manage many IBM and non-IBM servers, including the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server. The IBM Director console allows System Administrators to manage multiple BladeCenter chassis in a heterogeneous environment or environments where a Director infrastructure exists.

IBM Director V5.2 supports the following functions on the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express:

- ▶ Events
- ▶ Resource monitoring
- ▶ Inventory (limited)
- ▶ Remote session
- ▶ Software distribution
- ▶ File transfer
- ▶ Process management
- ▶ MPA

Visit the following Web site to download IBM Director 5.2:

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

2.10.3 Cluster Systems Management (CSM)

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 V1.5 supports the following functions on the IBM BladeCenter JS22 and IBM BladeCenter JS12 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/cluster/software/csm/>

Archived



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 JS12 and JS22 up and running the supported operating systems, and several other management-oriented topics.

Archived

System configuration using VIOS and IVM

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

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 this link:

<http://publib.boulder.ibm.com/infocenter/systems/topic/iphb1/iphb1pdf.pdf>

The *PowerVM Editions Operations Guide* also contains useful information:

<http://publib.boulder.ibm.com/infocenter/systems/topic/iphdx/sa76-0100.pdf>

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

3.1 Planning considerations

Planning considerations extend across the IBM BladeCenter JS22 or JS12, IBM BladeCenters and storage together.

3.1.1 General considerations

Individual logical partitions must be designed so that their requirements (and desired management operations, such as Partition Mobility) fit the CPU and memory available on the blade. IVM, unlike HMC-managed systems, does not allow the overcommitment of resources even if an LPAR is not active.

The IBM BladeCenter JS22 has a single internal SAS drive. Disk mirroring is highly recommended for the operating system. The mirror can be between the internal drive and a SAN LUN, or between two SAN LUNs.

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 requires the deletion of the LPAR and the creation of a new one, or the use of the `chsyscfg` command. The `chsyscfg` command can be used to switch when the partition is in the Not activated state.

Note that EtherChannel configuration (if used), as well as BladeCenter switch module type and bay location, require pre-planning. The HEA port cannot be mixed with other Ethernet devices to create the EtherChannel device. The use of the HEA ports to create an EtherChannel will require the use of the Intelligent Copper Pass-thru Module for IBM BladeCenter in switch module bay one to retain Serial over LAN (SOL) capability.

Note: The HEA adapter implementation on JS12/22 blades will always show a link status of Up, when configured 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 network interface backup (NIB) configuration, the link state cannot be used as the failover determination. NIB should be set up using the ping option.

3.1.2 VIOS installation considerations

The Virtual I/O Server installation is performed like a native install of AIX. The basic requirements are:

- ▶ Console access, provided by Serial over LAN (SoL) or KVM
- ▶ 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

The AIX process can be reviewed in the white paper *Installing AIX6 on the POWER6 BladeCenter JS22 using Keyboard Video Mouse or Serial Over LAN*, which is available at:

ftp://w3.ibm.com/support/americas/pseries/js22AIX61_v2_122607.pdf

3.1.3 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 JS12 and JS22 blades. For information about supported I/O expansion cards, along with their feature codes and option part numbers, refer to 2.7, “Expansion cards” on page 44.

There are currently four types of I/O expansion cards supported in a JS12 or JS22. They are:

- ▶ Standard Form Factor (StFF) I/O expansion card
- ▶ Combined Form Factor vertical (CFFv) I/O expansion cards
- ▶ High Speed Form Factor (HSFF) I/O expansion cards
- ▶ Combined Form Factor horizontal (CFFh) I/O expansion cards

StFF and CFFv adapter cards are always connecting to bay three and four of a BladeCenter chassis when installed in a JS12 or JS22. Refer to BladeCenter documentation for details of the bay numbering.

Figure 3-1 on page 58 shows the QLogic® 4 Gbs Fibre Channel HBA as CFFv I/O expansion card.

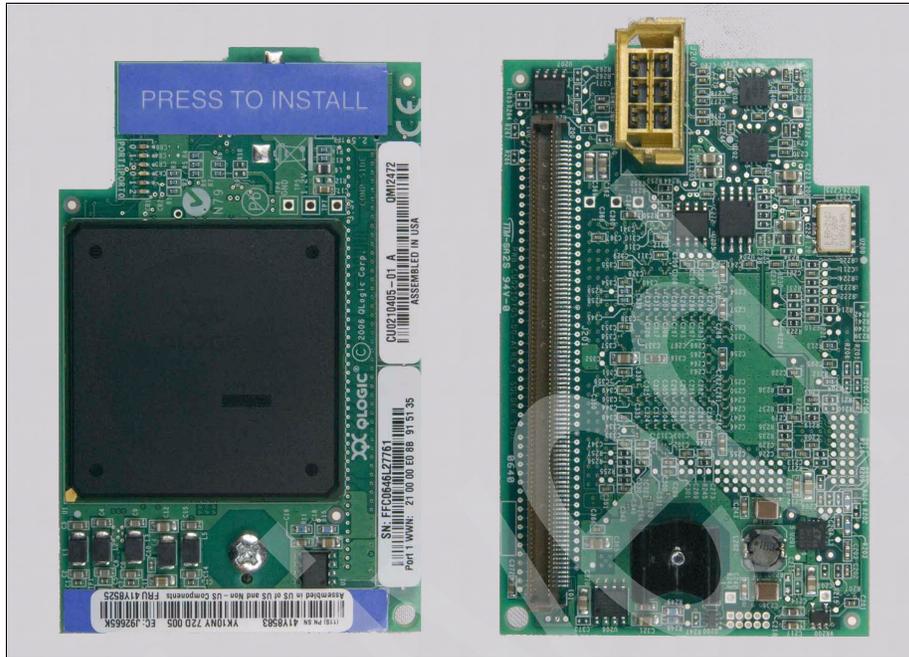


Figure 3-1 Qlogic 4 Gbs FC HBA as CFFv I/O expansion card

In a BladeCenter H chassis, for example, 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 2¹.
- ▶ You must install a module that supports the I/O connectivity of the installed CFFv or StFF I/O expansion cards in Bay 3 and Bay 4.
- ▶ You must use the Optical Pass -thru Module to connect to Ethernet or Fibre Channel using CFFv or the StFF cards in the blades.

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.

CFFh and HSFF cards can only be used when the blade is installed in a BladeCenter H or HT. Only these BladeCenter models have the high speed bays

¹ BladeCenter S does not support a module in Bay 2 at this time. 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.

7, 8, 9, and 10. These bays have a horizontal orientation. (The standard module bays have a vertical orientation.)

HSFF cards are typically Infiniband host channel adapters that are currently only supported by Linux. These I/O expansion cards connect to Bay 7 and Bay 9 of a BladeCenter H or HT.

Some blade types (currently *not* the JS12 or JS22) support the 10 Gb Ethernet adapters available as HSFF cards. These cards connect to Bay 8 and Bay 10 of a BladeCenter H or HT.

Figure 3-2 shows an example of an HSFF I/O expansion card.



Figure 3-2 HSFF I/O Expansion card - Cisco 4x Infiniband HCA

The CFFh card is a special card which utilizes the high speed bays, but uses standard modules for connectivity. This connectivity is achieved by using the Multi-Switch Interconnect Module (MSIM). One MSIM² gets installed in Bays 7/8 and in Bays 9/10. Each MSIM can be populated with a standard Fibre Channel switch module and an Ethernet switch module for the BladeCenter.

Figure 3-3 on page 60 shows a diagram of the MSIM and the type of switches that can be installed in the module bays. Refer to *BladeCenter Interoperability Guide* for details about which switches are supported in the MSIM. The guide can be found at:

<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5073016&brandind=5000020>

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.

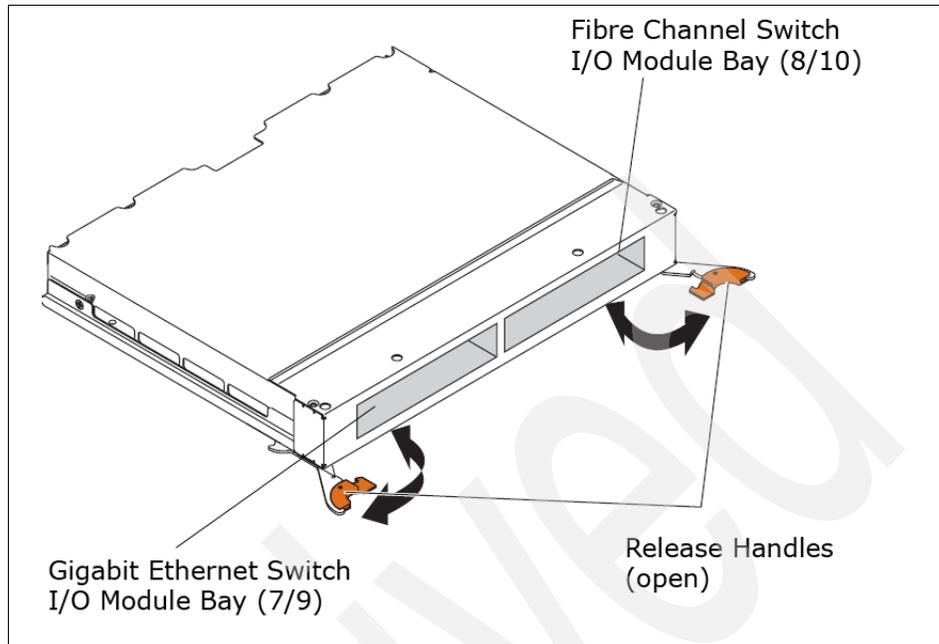


Figure 3-3 Multi-Switch Interconnect Module

The CFFh 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 PCI-e bus of the blade. Figure 3-4 on page 61 shows a CFFh card.

Together with an installed CFFh I/O expansion card, it is also possible to install the CFFv I/O expansion card. This gives, in addition to the two onboard network interfaces, six more I/O ports. These six ports can be four FC ports and two gigabit Ethernet ports or two Fibre Channel ports and four gigabit Ethernet ports depending on the type of CFFv card installed.

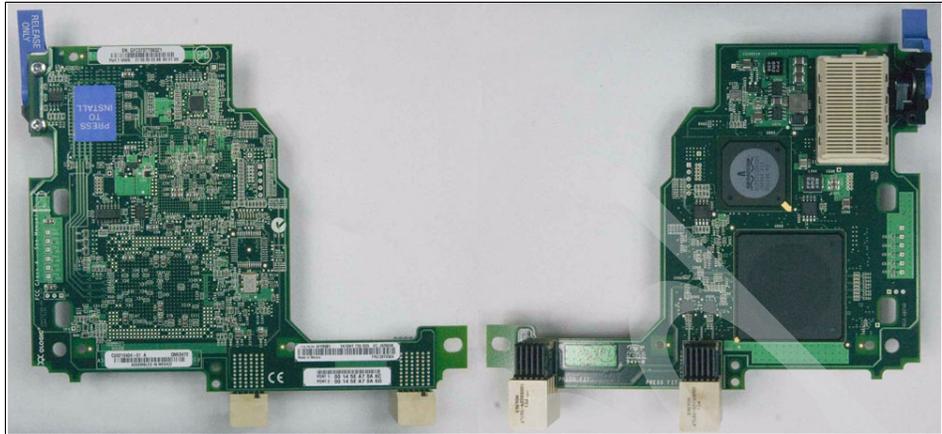


Figure 3-4 Qlogic 4 Gbs/Broadcom 1 Gbs CFFh I/O Adapter card

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

JS12/JS22 storage

There are currently four different types of storage available:

- ▶ Internal SAS disk storage
- ▶ External SAS/SATA disk storage
- ▶ External iSCSI storage
- ▶ External Fibre Channel storage

Internal SAS storage

For internal storage, each blade has one or two internal hard drives available. The IBM BladeCenter JS12 and JS22 uses SAS hard drives. In the JS22, one 73 GB or 146 GB SAS hard drive can be installed. The JS12 supports two drives.

The SAS controller is embedded on the planar of each blade and has no RAID functionality. For a small installation, this might be enough storage. However, disk redundancy cannot be achieved in a JS22 without utilizing external storage. In such a configuration, external storage might be connected using SAS, iSCSI, or Fibre Channel.

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 a SAS Expansion Card for IBM BladeCenter in the JS12 or JS22 and, depending on the number of storage modules installed in the chassis, one or two SAS Connectivity modules.

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 recommended due to limited network port availability. You have only the two onboard network interfaces of the blade available.

IBM has made a statement of direction to support the DS3200 storage subsystem.

Also, the BladeCenter Boot Disk subsystem is not an alternative because there is no support for the data. This means that there is a requirement to install an additional adapter to get access to data storage that cannot be connected in a BladeCenter E, T, or S. Bay 3 and Bay 4 are the only available bays for storage I/O connectivity, and they are populated with SAS connectivity modules.

External iSCSI storage

External iSCSI storage is currently supported with some N series storage subsystems. iSCSI storage can be accessed only with the iSCSI host bus adapter.

The IBM BladeCenter JS12 and JS22 have support for the iSCSI Expansion Card for IBM BladeCenter, as shown in Figure 3-5 on page 63. This adapter card is a standard form factor card. It requires the removal of the internal SAS drive and the tray on the JS22 and the internal drive and tray in the "D2" position on a JS12 (see 1.6, "Minimum and optional features of BladeCenter JS12" on page 9) that is installed next to the backplane connectors.

Consequently, you will have only one disk left in the JS12 and no disk in the JS22. The JS22 can only boot from an iSCSI storage subsystem in this configuration.

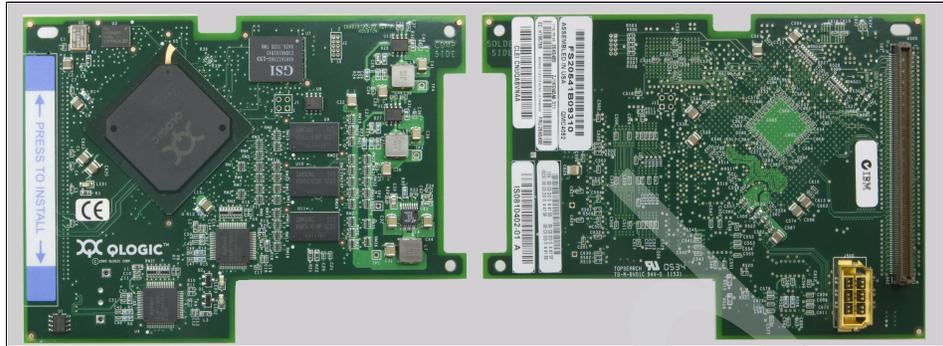


Figure 3-5 QLogic iSCSI HBA as StFF

Note: iSCSI boot and the network interfaces or host bus adapters used depend on the operating system you plan to boot. For example:

- ▶ Virtual I/O Server supports iSCSI boot only with the QLogic iSCSI HBA.
- ▶ AIX V6.1 can be used with regular network interfaces or with the QLogic iSCSI HBA.
- ▶ POWER6 processor-based blades do not support the firmware iSCSI initiator or the configuration tool if the iSCSI firmware initiator contains HS or LS blades.

For iSCSI storage area networks, it is strongly recommended that you use a separate Ethernet network infrastructure. This requires the installation of a Pass-thru Module or Ethernet Switch in Bay 3 and Bay 4 of a BladeCenter chassis.

The use of Ethernet Switch modules reduces the cabling effort but requires some configuration. Pass-thru Modules entail greater cabling effort, but allow you to use different I/O adapter cards in the different blades connected to Bay 3 and Bay 4 of the BladeCenter chassis.

This is especially valid for the Optical Pass-thru Module, which allows you to connect Ethernet and Fibre Channel to external switches using optical cables. The maximum bandwidth per port is 2 Gbs.

Note: iSCSI software initiator is currently not supported with Virtual I/O Server. You may set up a software initiator inside your logical partition as long as the operating system in the LPAR supports this.

External Fibre Channel storage

Various IBM and non-IBM Fibre Channel storage subsystems can be used to connect to a JS12 or JS22 installed in a BladeCenter, including models from the following series:

- ▶ IBM Total Storage DS8000™ series
- ▶ IBM Total Storage DS6000™ series
- ▶ IBM Total Storage DS4000™ series
- ▶ IBM Total Storage DS3000 series
- ▶ IBM Total Storage N series

They are supported by the Virtual IO Server. 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, there will be zoning limitations if different switch vendors are used.

General storage considerations and support matrixes

You need to check a few support matrixes to plan your live partition mobility installation. This section points to support matrixes you can use to build a fully supported solution.

1. Start with the blade itself. All supported hardware and operating systems are listed on IBM ServerProven®. ServerProven can be found at:

<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/>

Click the middle of the page on the BladeCenter picture for blade- and BladeCenter-related information. You can find consolidated information from different sources in *BladeCenter Interoperability Guide* at:

<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ndocid=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-03.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 matrices of IBM storage products.

The DS8000 interoperability matrix can be found at:

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

The DS6000 interoperability matrix can be found at:

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

The DS4000 interoperability matrix can be found at:

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

The DS3000 interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf>

The ESS interoperability matrix can be found at:

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

The N series interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

The SAN volume Controller support matrix can be found at:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskind=3&brandind=5000033&familyind=5329743>

The SAN switch interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/san/index.html>

6. If you plan to implement IBM i on a JS12 or JS22, then in addition to the previous steps, verify the support matrix of IBM i on blades. Some restrictions apply when IBM i is used.

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

http://www-03.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 System Storage™ Interoperation Center (SSIC) helps to identify supported storage environments. You find this Web-based tool at:

<http://www-03.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 may not have internal disks available. Redundancy may be not available at boot time for a boot disk. Manual interaction may be required and an administrator may need to be trained for situations where interaction with the system is required.

After completing these steps, you should be familiar with the supported environments. Start with setup of the storage. 3.1.4, “External storage preparation” on page 66 describes, as an example, the setup of a DS4800 as shared storage used for live partition mobility.

3.1.4 External storage preparation

Preparing storage depends on the type of storage subsystem that you use. This section describes, at a high level, the storage partitioning of an IBM Total Storage DS4800.

For more detailed information about storage preparation for DS4000, refer to the product documentation and to the IBM Redbooks publication *IBM System Storage DS4000 and Storage Manager V10.10*, SG24-7010, which is available from the following site:

<http://www.redbooks.ibm.com/abstracts/sg247010.html?Open>

To set up your storage, apply the supported cabling and attach the storage to your BladeCenter. Our setup used a IBM BladeCenter H. In the BladeCenter H is a Multi-Switch Interconnect Module (MSIM) installed in Bay 7 and 8 and a second MSIM in Bays 9 and 10. The MSIM uses two high-speed slots to provide the ability to install two standard I/O modules, one Ethernet and one Fibre Channel switch module.

To use the switch models installed in the MSIMs, you need to install a CFFh I/O expansion card in the blade. This I/O expansion card contains a dual port QLogic 4 Gbs Fibre Channel host bus adapter and a 1 Gbs Broadcom gigabit Ethernet adapter. Each blade has, with the CFFh I/O expansion card, four Ethernet ports and two Fibre Channel ports available.

There are two blades that use a CFFh I/O expansion card, and one blade that uses a CFFv card. The blade with the CFFv FC host bus adapter installed uses Bays 3 and 4 to connect to the storage subsystem.

Figure 3-6 on page 68 shows the basic SAN cabling and an example of each blade configuration. Unused internal connections between the blade and the BladeCenter are shown with dashed grey lines.

It is important to note that shared storage between Virtual I/O Servers and the use of the same IP subnet between source and target are required for partition mobility. This means that all blades use the same gigabit Ethernet networks connected via Bays 1 and 2 of the BladeCenter chassis. Only two of the blades can utilize the networks provided via Bays 7 and 9.

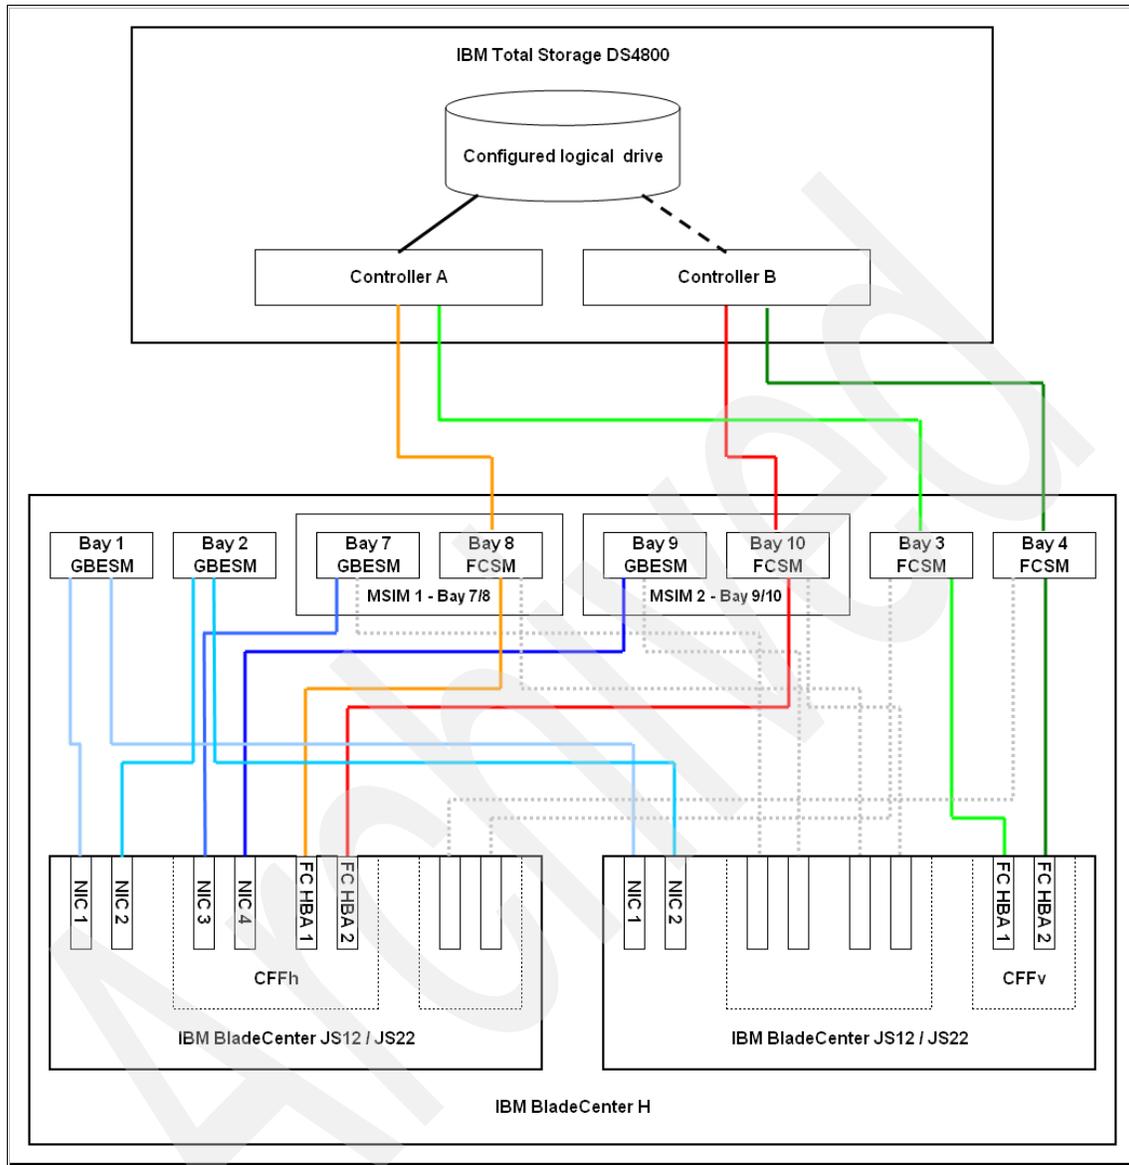


Figure 3-6 Basic SAN setup

After all required logical drives are created, the storage partitions need to be configured. The host group JS22-Partition-Mobility was created with three hosts: JS22-1-Jupiter-VIOS, JS22-2-Saturn-VIOS, and JS22-3-Neptune-VIOS. Each host has two host ports defined with the worldwide port names of the host bus adapter ports. The host ports are defined as type AIX.

Logical drives can be assigned to the host group JS22-Partition-Mobility or to a host. Assigned logical drives to a host group are accessible by any host in the host group. Logical drives that are assigned to a host in the host group are only accessible by the host itself. The logical drives that are used for the Virtual I/O Server installation and the media library are not shared between different Virtual I/O Server hosts.

When the host group, host, and host ports are defined, map only the logical drives to the hosts that are required to install the Virtual I/O Servers. All shared logical drives can be mapped later.

Three logical drives are mapped to each of the hosts in the host group. They are mapped as LUN 0, 1, and 2. The first two are used for the Virtual I/O Server itself. A mirrored rootvg can be created to have a second copy of the Virtual I/O Server installation with these first two disks.

The logical drives used for the Virtual I/O Server installation need to be large enough to hold the Virtual I/O Server installation and future updates. It is recommended that you create a logical drive with at least 30 GB capacity. An expansion of the logical drive (as this is supported by the Total Storage DS4800) is not supported for the disk containing the rootvg in Virtual I/O Server or AIX. However, logical drives that contain other volume groups or storage pools can be extended. An alternate install might be used in case the disk space is not sufficient on the logical drive containing the rootvg.

The third logical drive is used for the media library. The media library cannot be shared between different Virtual I/O Servers. For that reason, the logical drive that will be used for the media library is also mapped to the host only. The size of this logical drive depends on the number of CD or DVD images you plan to store in the media library. Additional space is required when you plan to use the media library as backup, as you need to do this with logical partitions running IBM i.

After the Virtual I/O Server is installed and configured on each host that will be used for partition mobility, map the logical drives that will be used as virtualized disks for the logical partitions.

Figure 3-7 on page 70 shows the finalized storage partitioning for three JS22s that are used for partition mobility. In the Accessible By column, notice that the logical drives mapped as LUN 0, 1, and 2 are mapped to the host itself. These are the non-shared logical drives or hdisk in the Virtual I/O Server.

The Virtual I/O Server will boot over the storage area network from LUN 0. LUNs 3 - 8 are shared among all three hosts (JS22-1 Jupiter-VIOS, JS22-2-Saturn-VIOS, and JS22-3-Neptune-VIOS) because they are mapped to the host group JS22-Partition-Mobility.

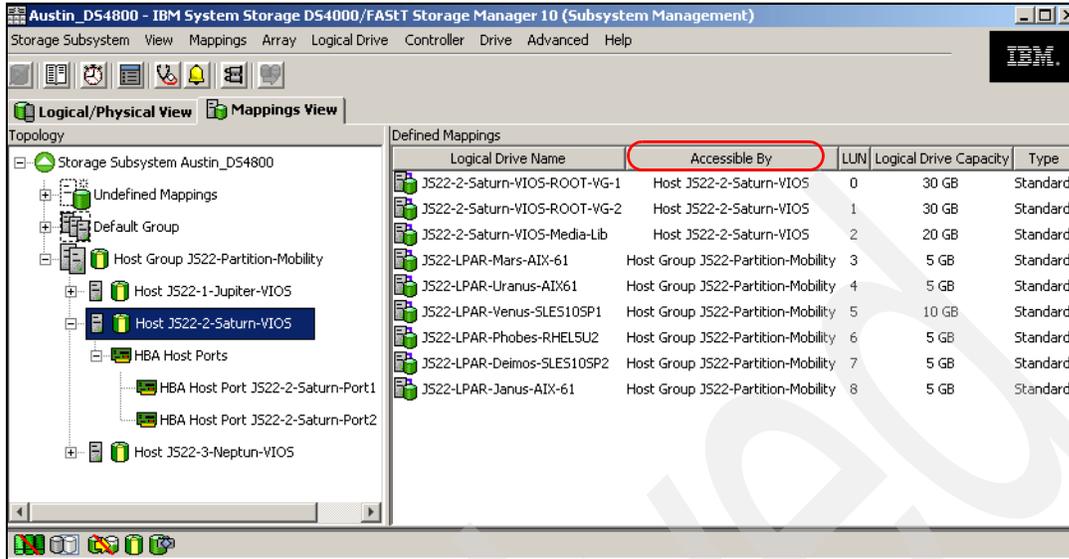


Figure 3-7 DS4000 Storage Manager with storage partitioning for partition mobility

Proceed with the configuration of Virtual I/O Server, as described in the next section.

3.2 VIOS system management using IVM

Using IVM is useful for when two or more partitions are needed, or 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 user interface (UI) designed to be as simple and intuitive as possible, incorporating partition management, storage management, serviceability, and monitoring capabilities
- ▶ A command line interface (CLI), to enable 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.

3.2.1 IVM user interface

The IVM user interface is a Hypertext Markup Language (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 UI consists of several elements. The following two elements will be used most frequently:

- | | |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigation area | The navigation area on the left side of the window displays the tasks and links that you can access in the work area. |
| Work area | The work 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 3-8 on page 72 shows the navigation and works of the IVM UI. UI help can be obtained from any window by clicking the ? link in the upper right corner of any window or view.

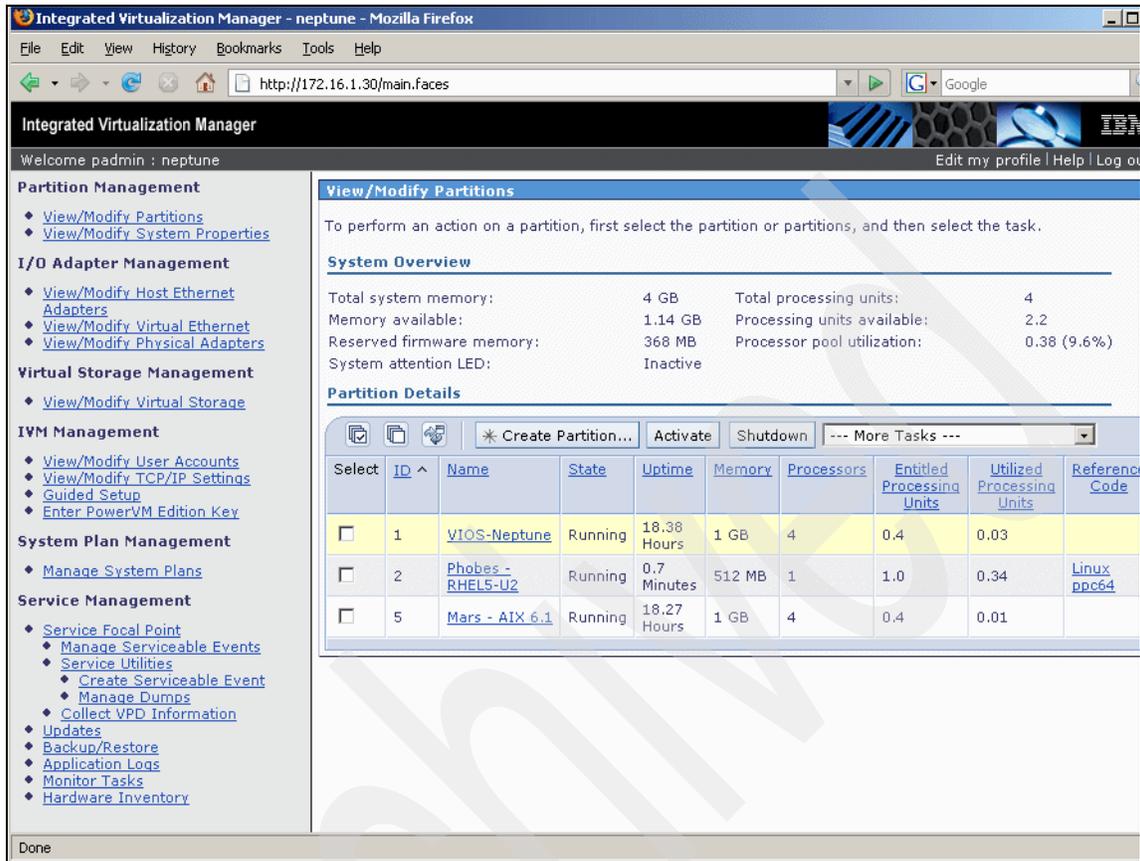


Figure 3-8 IVM navigation and work areas

The login to the UI is described in 3.4.1, “Connecting to IVM” on page 80

3.2.2 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 \$ prompt. A list of available commands can be displayed by using the **help** command. Individual command help can be viewed by using the **-h** flag. Detailed command help can be shown using the **man** command.

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

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

3.3 First VIOS login

During the first login to the VIOS running on a IBM BladeCenter JS22 or JS12, you will be 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 a Serial Over Lan (SoL) or KVM session.

3.3.1 Password set

From the login prompt, use the default user ID of `padmi n`. You will be required to enter a password and then re-enter the same password for confirmation, as shown in Example 3-1.

3.3.2 License acceptance

The licensing agreement must be accepted prior to starting any configuration tasks. After setting the password, you will be presented three license options as shown in Example 3-1. Enter `v` to view the license, `d` to decline, or `a` to accept.

Example 3-1 Initial login screen showing password set and licensing options

```
IBM Virtual I/O Server
```

```
login: padmi n
```

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

```
padmi n'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 enter a, enter the **license -accept** command as shown in Example 3-2.

Example 3-2 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 3-3.

Example 3-3 The license status

```
$ license
The license has been accepted
en_US Jun 9 2008, 18:33:08 8(padmin)
```

3.3.3 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 different methods: by using the **mktcpip** command, or 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 3-4 on page 75.

Example 3-4 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]

        [-nsrvdomain Domain] [-plen prefixLength]
```

The `-cabletype` and `-plen` flags are the only two optional parameters. Example 3-5 shows how to configure a VIOS network interface.

Example 3-5 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 9.3.5.139
-nsrvdomain ibm.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 3-6 on page 76 shows the initial `cfgassist` menu. You start the tool by entering the `cfgassist` command.

Example 3-6 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
```

```
Esc+1=Help      Esc+2=Refresh    Esc+3=Cancel    F8=Image
F9=Shell        F10=Exit        Enter=Do
```

By selecting **VIOS TCP/IP Configuration**, you will be presented a list of available network interfaces as shown in Example 3-7.

Example 3-7 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 3-8, you enter the TCP/IP configuration by pressing the Enter key. This completes the initial TCP/IP configuration of the VIOS.

Example 3-8 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) [9.3.5.139]
    DOMAIN Name [ibm.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 3-9 on page 78 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.

Example 3-9 *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 Groups	Gateway	Flags	Refs	Use	If	Exp
-----------------------	---------	-------	------	-----	----	-----

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.ibm.com	UHSb	0	0	en6	-
- =>						
172.16.1/24	saturn.ibm.com	U	2	5628596	en6	-
-						
saturn.ibm.com	localhost	UGHS	0	39074	lo0	-
-						
172.16.1.255	saturn.ibm.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    9.3.5.139
    domain    ibm.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 reset the system to a newly installed condition, use the **rmtcpip -f -all** command.

3.4 First IVM connection

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

3.4.1 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 3.3.3, “Initial network setup” on page 74).

A Welcome window that contains the login and password prompts opens, as shown in Figure 3-9. The default user ID is `padmin`, and the password is the one you defined during the VIOS installation.

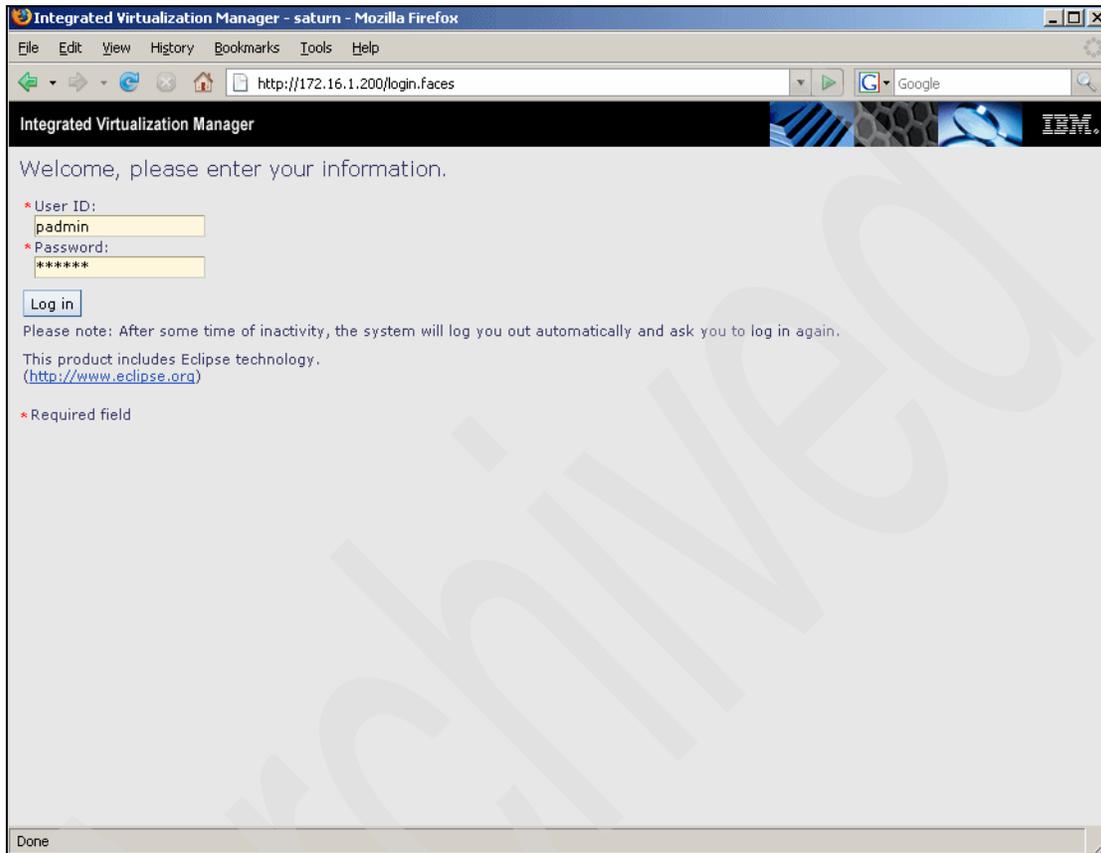


Figure 3-9 The Welcome window

The first connection to the IVM UI will display the guided setup window as shown in Figure 3-10 on page 82. 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.

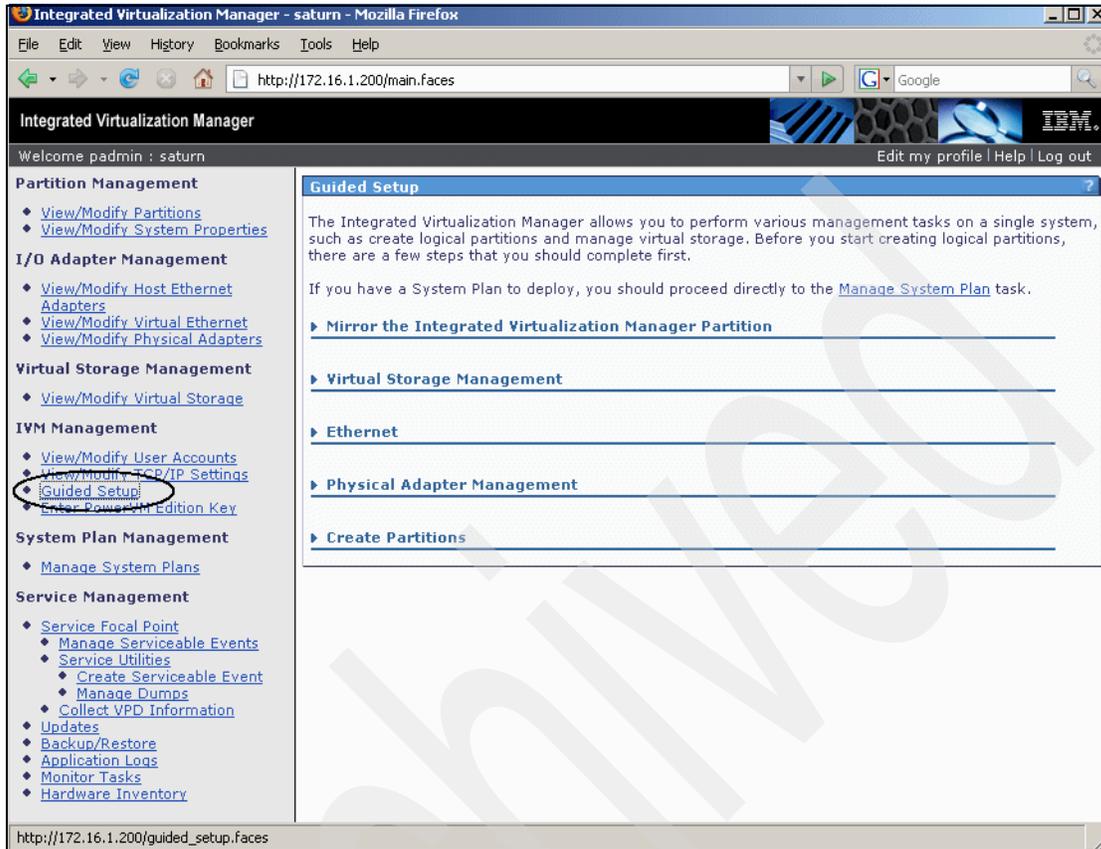


Figure 3-10 Guided Setup window

3.4.2 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 the maximum value between 512 MB and one-eighth of the installed system memory.
- ▶ The number of (virtual) processors is equal to the number of physical processors, 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 window as shown in Figure 3-11. This window can be accessed by the **View/Modify Partitions** link in the navigation area.

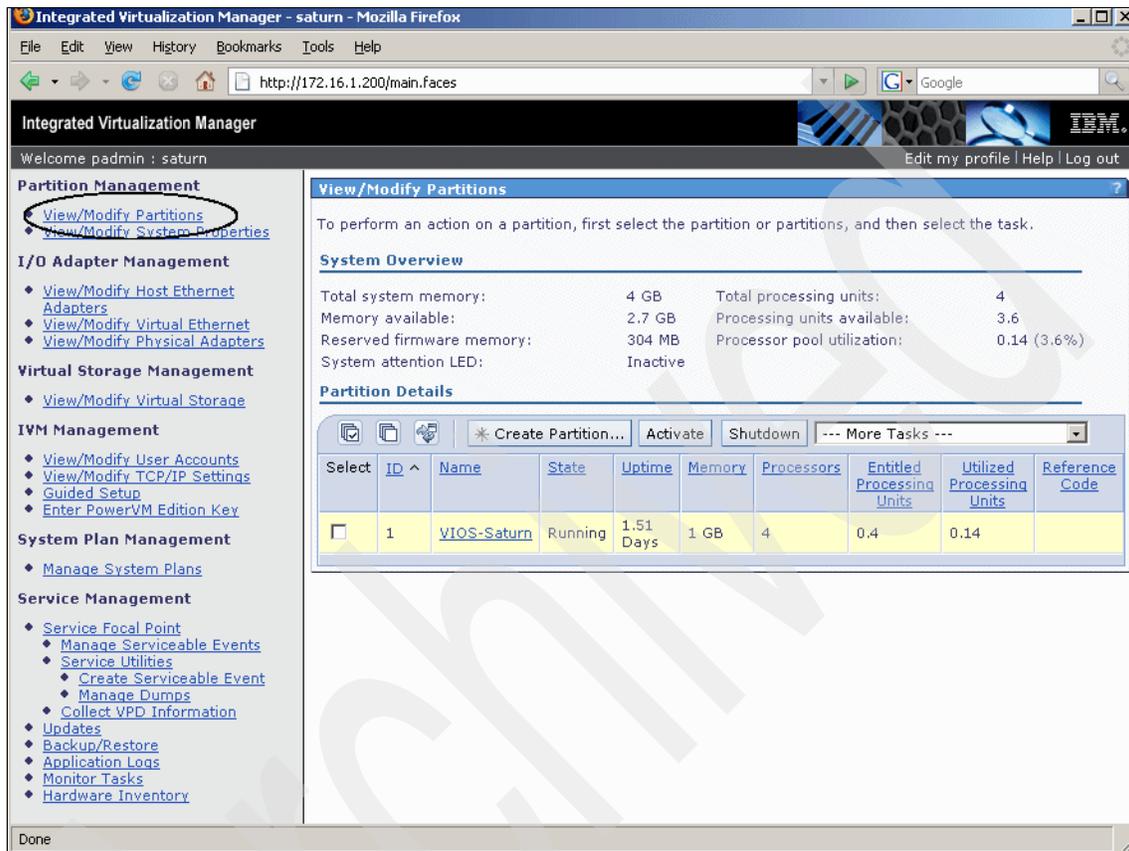


Figure 3-11 View/Modify Partitions window

Administrators can change properties of the VIOS LPAR, including memory or processing units allocation, by using the IVM UI. 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 3-12 on page 84. The name of the VIOS can be changed from the General tab, if desired. The Memory and Processors tabs are used to view or change the allocations. Figure 3-13 on page 84 shows the General tab.

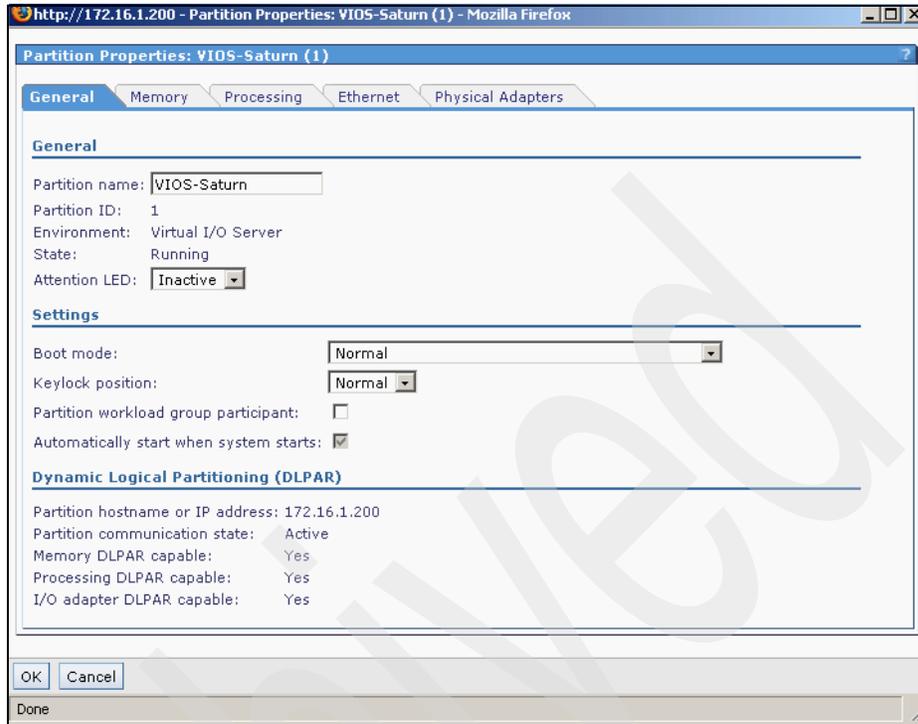


Figure 3-12 Partition Properties, General tab

Figure 3-13 shows the Memory tab.

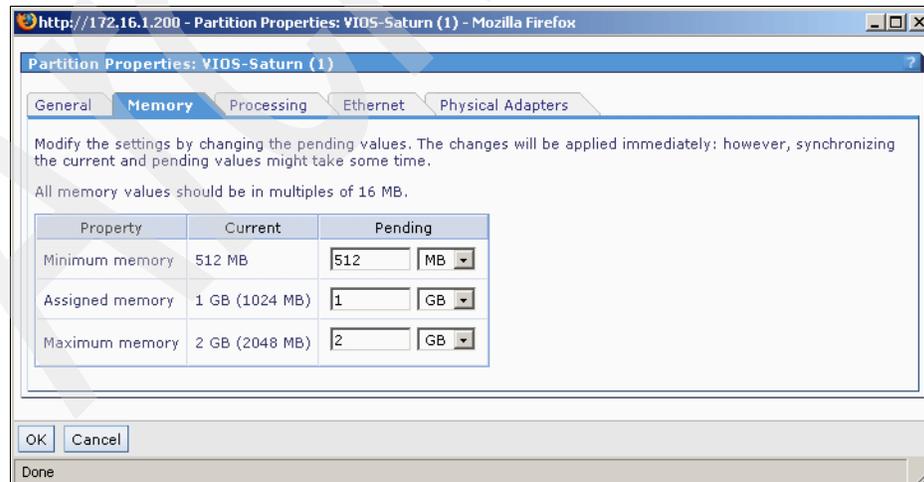


Figure 3-13 Partition Properties, Memory tab

The default memory configuration for the LPAR may need to be modified if your VIOS defaults to 512MB of memory and 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. The Assigned memory value should have an initial starting value of 1GB with the Maximum memory value set to 2 GB.

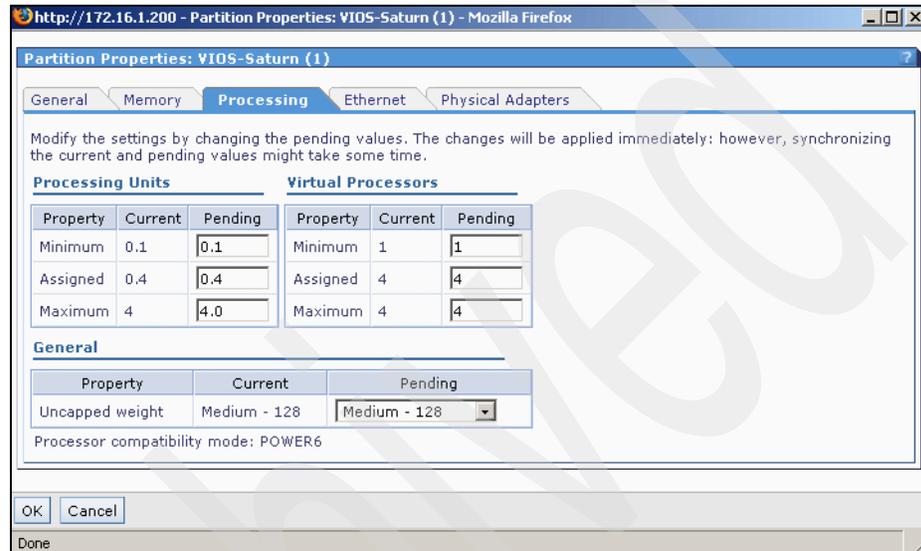


Figure 3-14 Partition Properties, Processing tab

Processing unit allocations for the VIOS are recommended to remain at the install 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 3-10.

Example 3-10 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 3.5.2, “Virtual Ethernet Adapters and SEA” on page 89. Physical Adapters tabs are discussed in 3.5.3, “Physical adapters” on page 94.

3.5 VIOS network management and setup

When you are preparing the IBM BladeCenter JS22 or JS12 and VIOS for additional LPARs, and planning for the ability for those LPARs to reach a physical network, you need to understand the networking options that are available. The three choices are:

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

3.5.1 Host Ethernet Adapters

Using the Host Ethernet Adapters (HEA) is the recommended method for providing a network interface to LPARs. The individual logical ports on the same physical port can be assigned to an LPAR without requiring a bridge. However, if Partition Mobility is a consideration, then this option *cannot* be used. The assignment of a logical port from the HEA can be done during the LPAR creation process. Refer to 2.6.4, “Integrated Virtual Ethernet (IVE)” on page 43 for additional technical details about the HEA.

You configure the HEA port mode by selecting **View/Modify Host Ethernet Adapters** from the navigation area. This displays the UI window, as shown in Figure 3-15 on page 87.

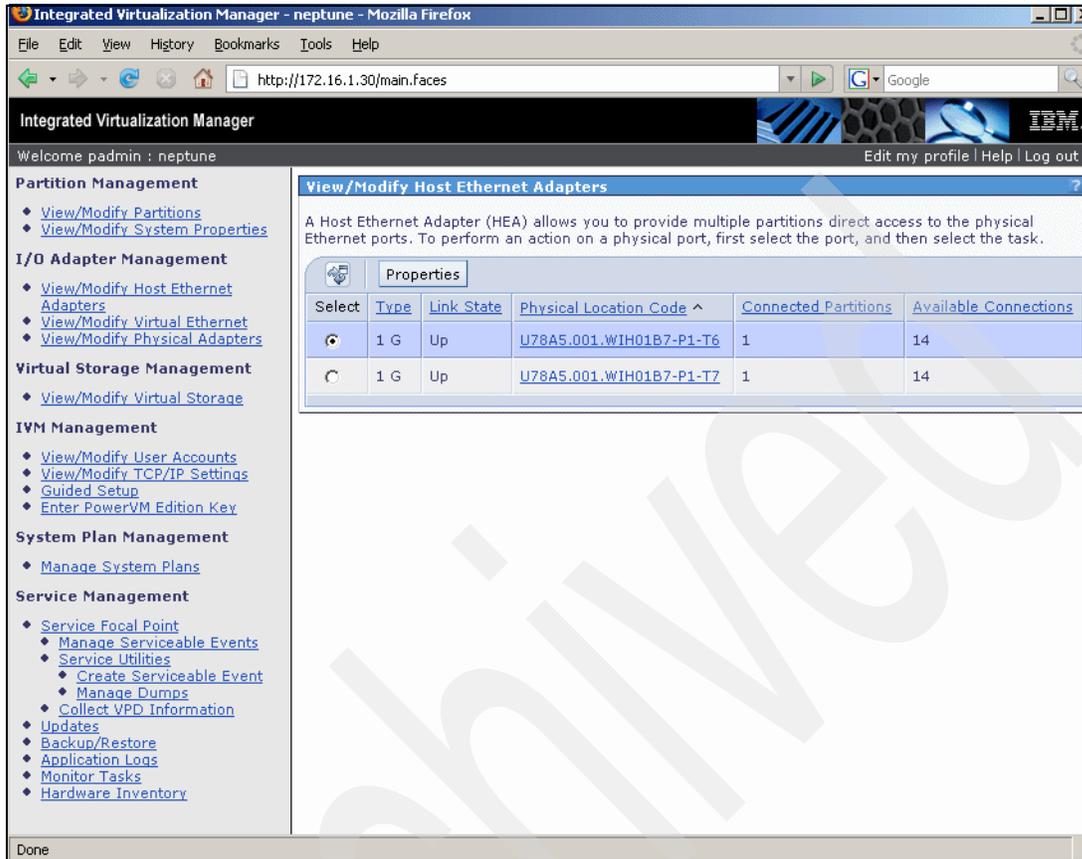


Figure 3-15 View/Modify Host Ethernet Adapters window

Both HEA ports are shown. The default configuration is port sharing with 14 logical connections available per port; this can readily be determined by the Available Connections column.

To view and modify the port properties, select one of the ports and click **Properties**. The HEA Physical Port Properties window will be displayed in a new window; Figure 3-16 on page 88 shows the General tab already selected.

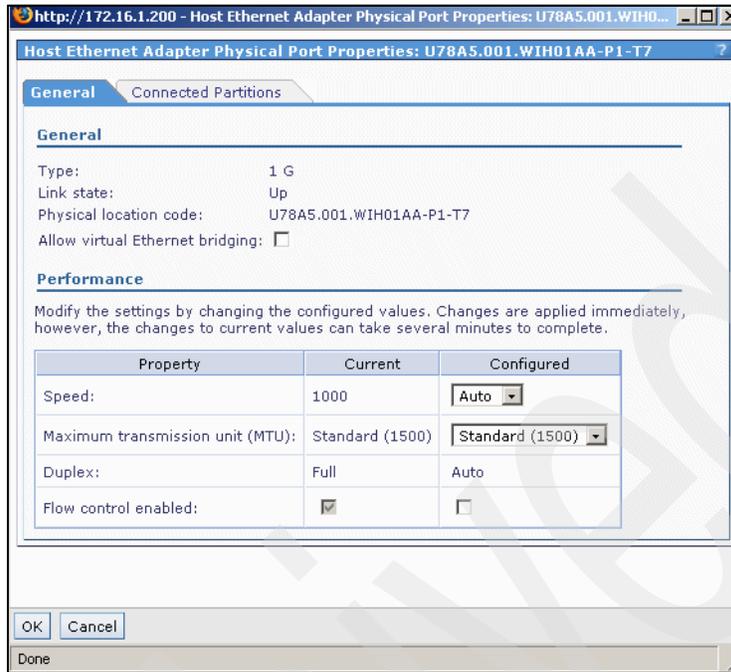


Figure 3-16 HEA Port Properties

You can display a list of connected partitions (if any) and MAC addresses by selecting the **Connected Partitions** tab, as shown in Figure 3-17 on page 89.

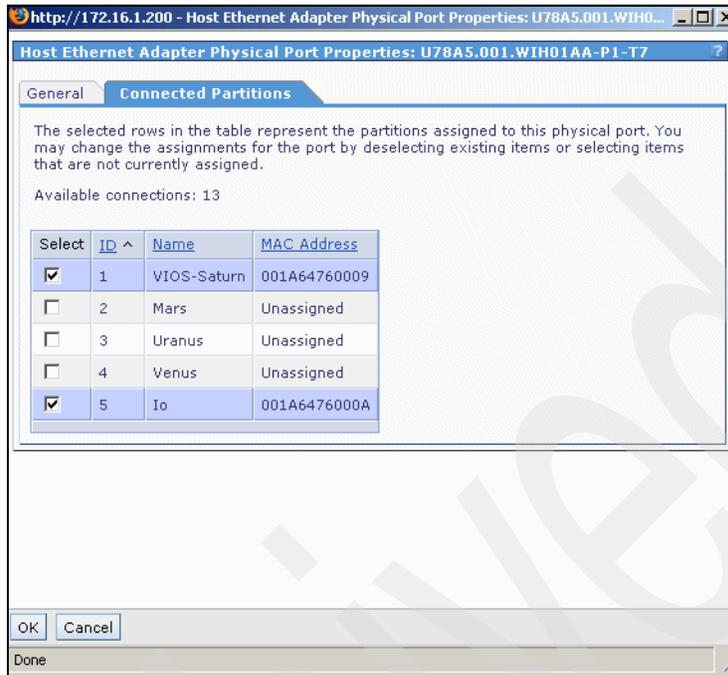


Figure 3-17 HEA Port Properties, Connected Partitions

3.5.2 Virtual Ethernet Adapters and SEA

Virtual adapters exist in the hypervisor that allows LPARs to communicate with each other without the need for 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 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 a 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 3-18 on page 90. 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.

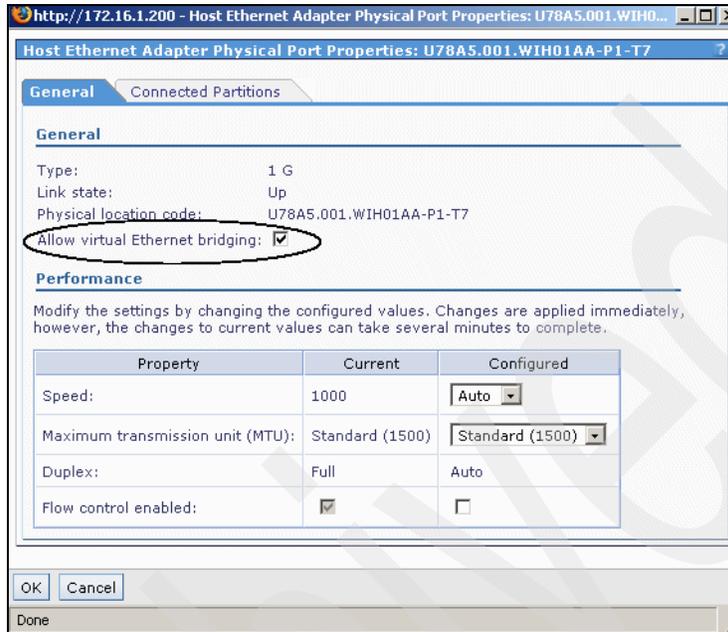


Figure 3-18 HEA port setting for Ethernet bridging

Physical Ethernet ports on an expansion card 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. Figure 3-19 on page 91 shows the four virtual Ethernet adapters that are created by default on the VIOS.

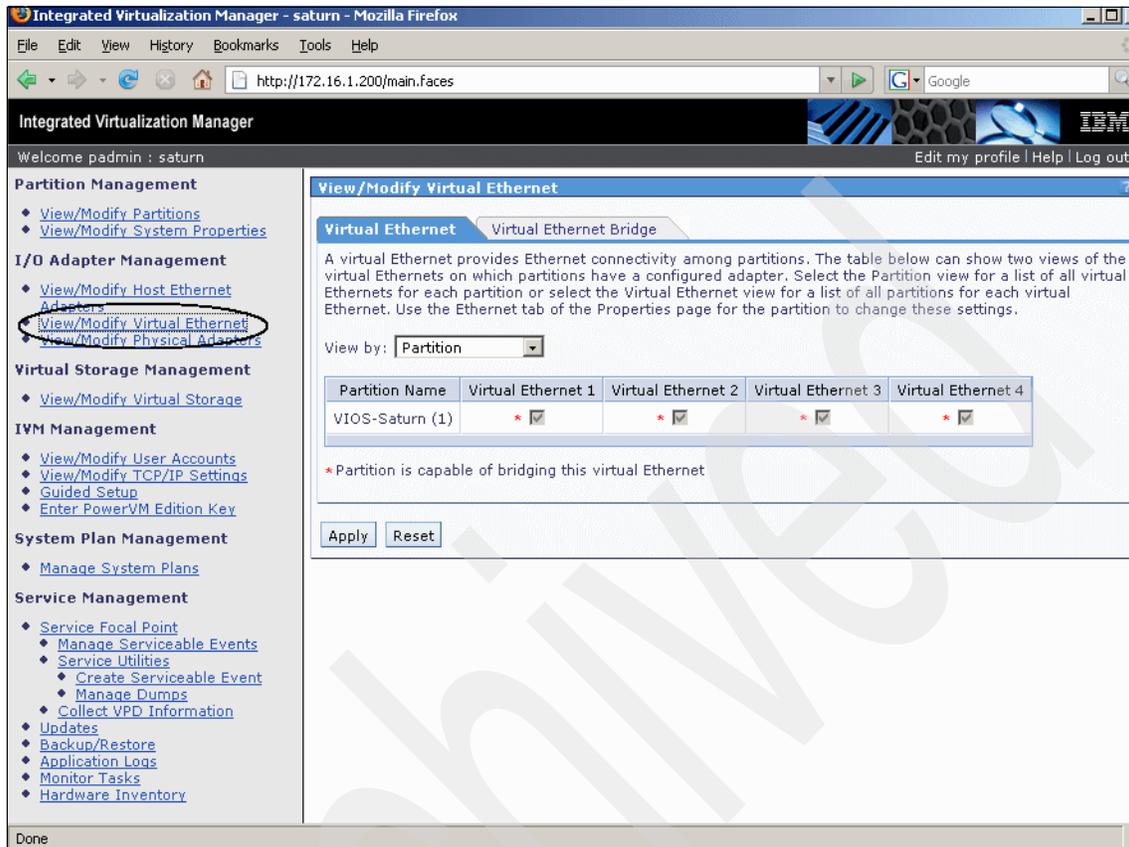


Figure 3-19 View/Modify Virtual Ethernet window

Use the **Virtual Ethernet Bridge** tab to display the virtual to physical options for creating an SEA, as shown in Figure 3-20 on page 92. 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. The drop-down box in the UI will allow the selection of the same adapter for another virtual Ethernet ID, but an error message will be generated when the **Apply** button is clicked.
- ▶ EtherChannel adapters, if created, will also be listed in the drop-down box.

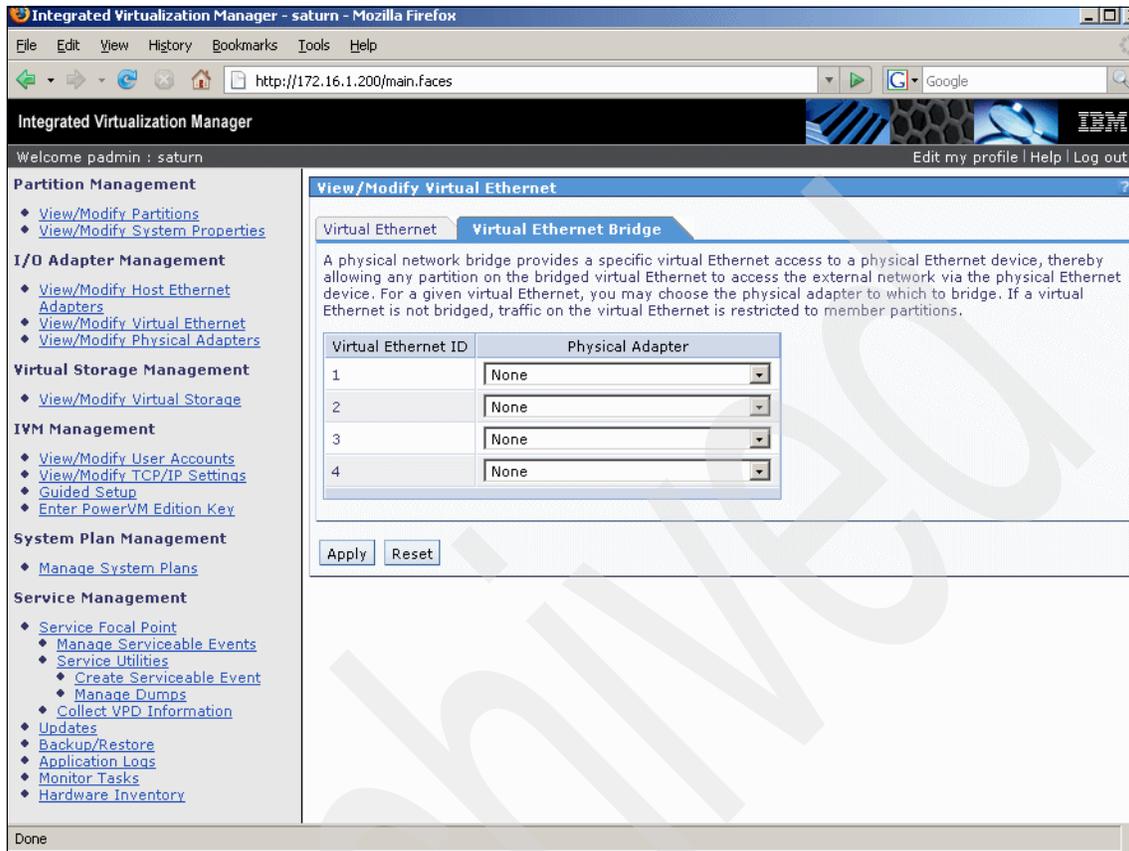


Figure 3-20 View/Modify Virtual Ethernet Bridge tab

Figure 3-21 on page 93 shows a physical adapter selection.

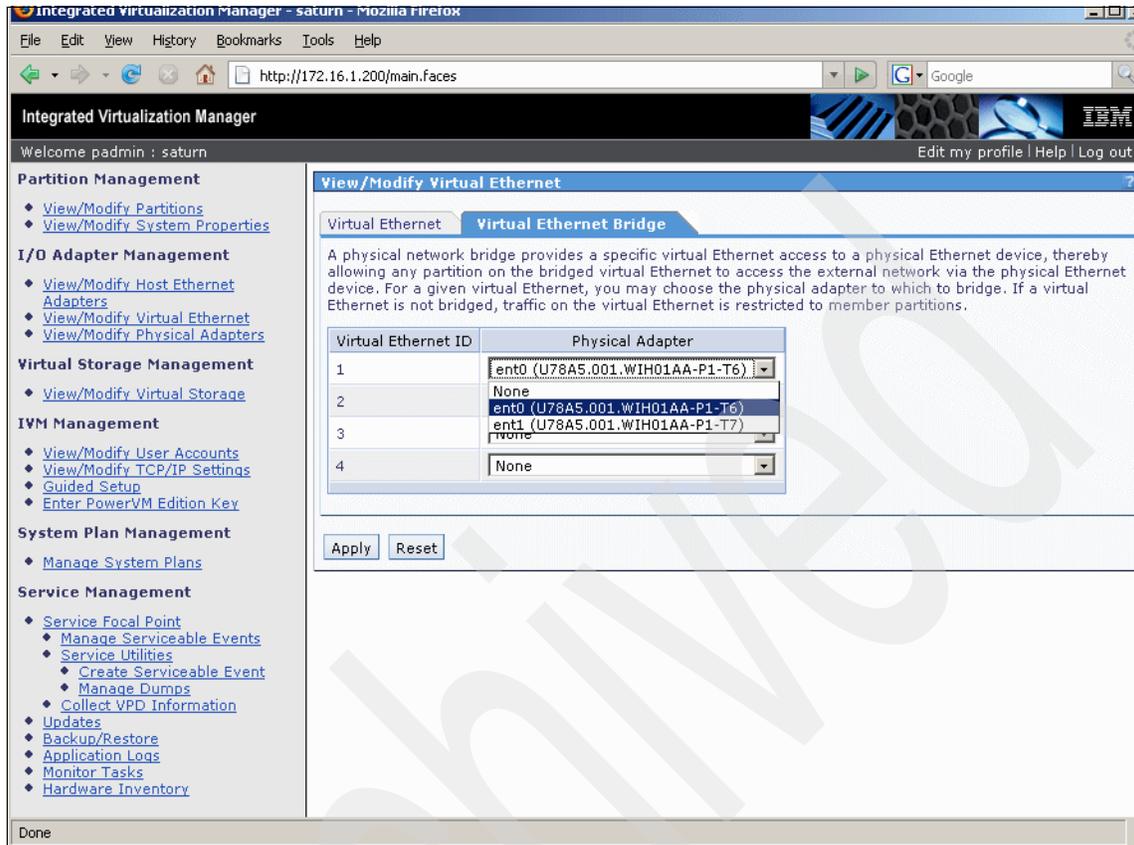


Figure 3-21 Physical adapter selection for SEA creation

Figure 3-22 on page 94 indicates the successful creation of the SEA.

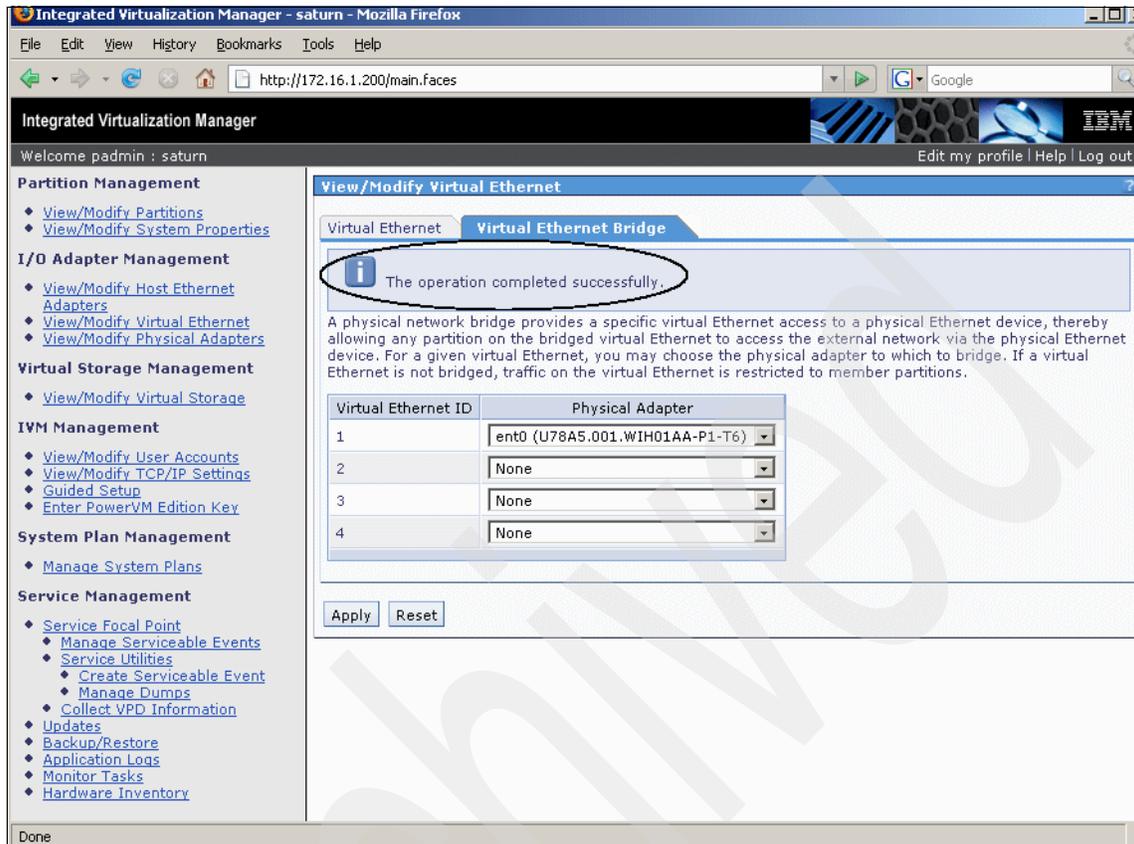


Figure 3-22 Successful SEA creation result

3.5.3 Physical adapters

With the IBM BladeCenter JS22 and JS12, 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 partitions, the resources must be purely virtual.

To assign a physical adapter, select the **View/Modify Physical Adapters** link from the navigation area to display the window shown in Figure 3-23 on page 95.

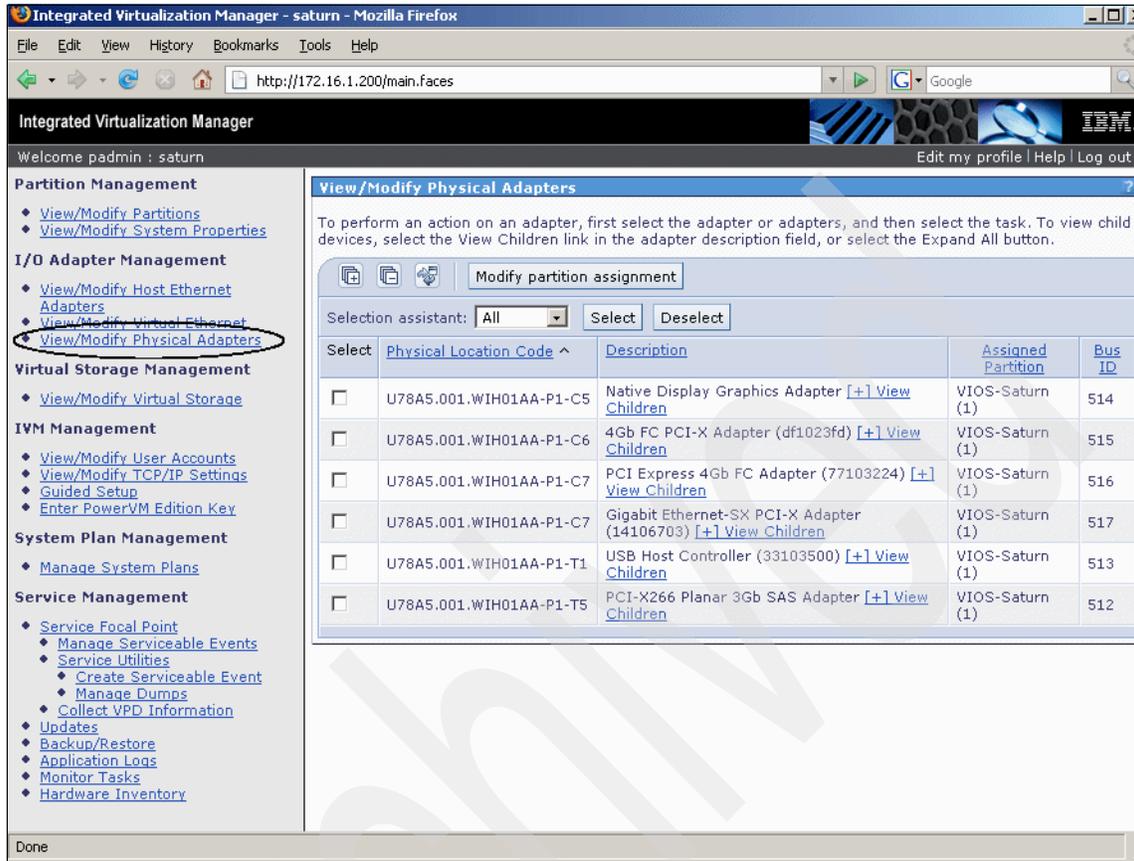


Figure 3-23 View/Modify Physical Adapters window

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 3-24 on page 96, the Gigabit Ethernet expansion card ports are being reassigned to partition 2.

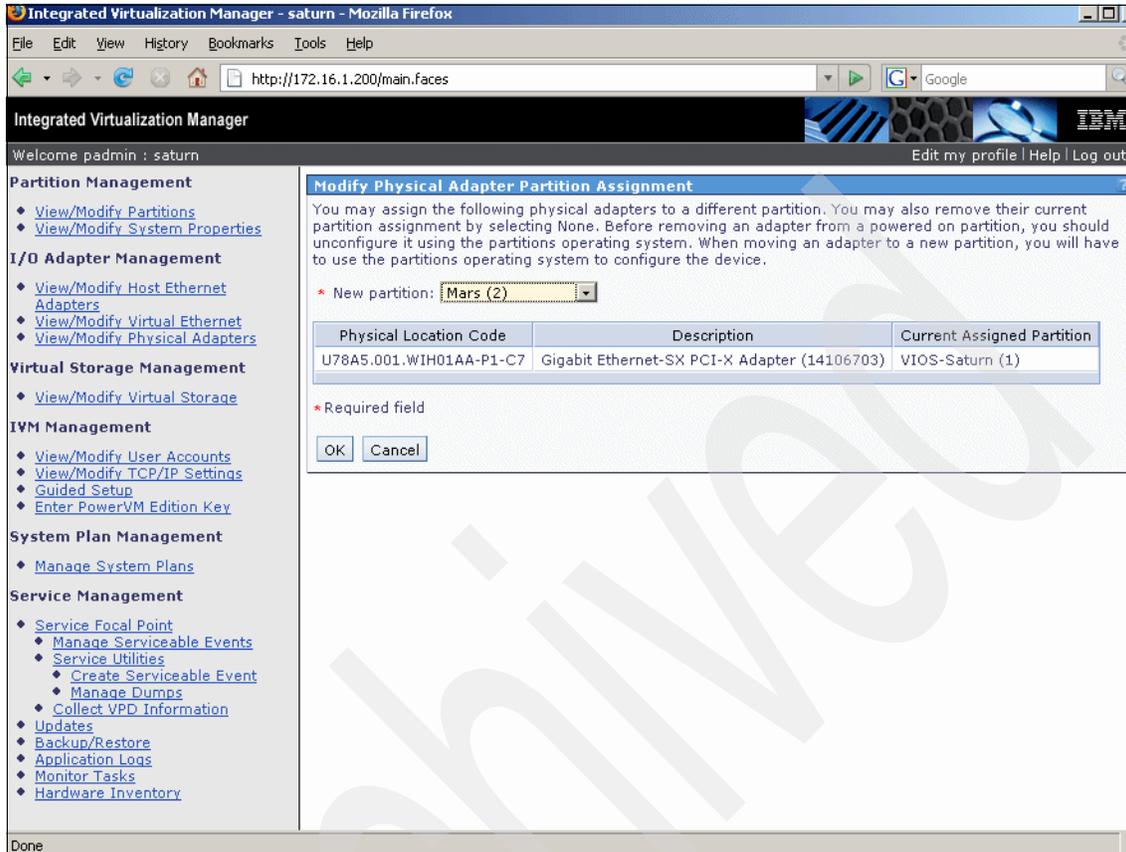


Figure 3-24 Physical Adapter assignment to new partition

Figure 3-25 on page 97 shows the change in partition ownership.

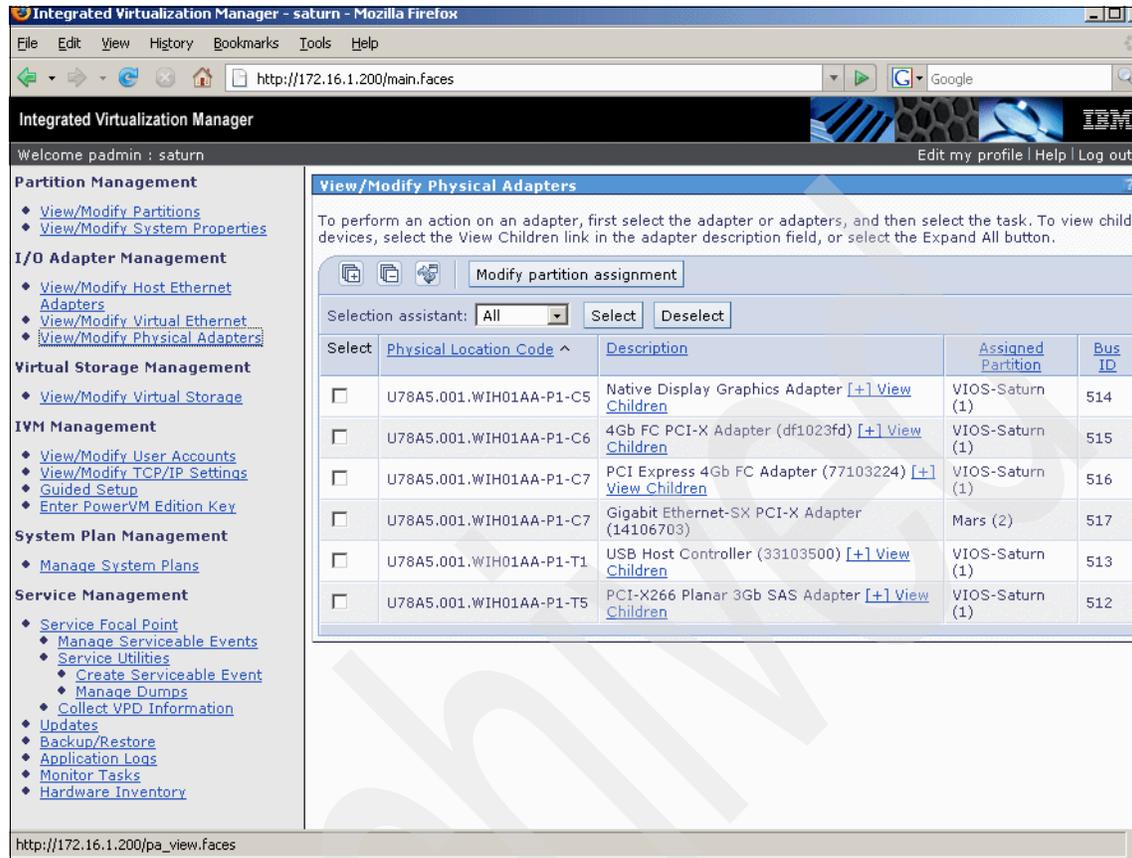


Figure 3-25 View/Modify Physical Adapter window showing change of ownership of Gigabit Ethernet Adapter

Example 3-11 shows the changes in adapter availability in a logical partition, starting with the original virtual Ethernet adapter through the addition of the two physical ports from the IBM BladeCenter JS22 or JS12 expansion card.

Example 3-11 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
#
```

3.6 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 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 3-26.

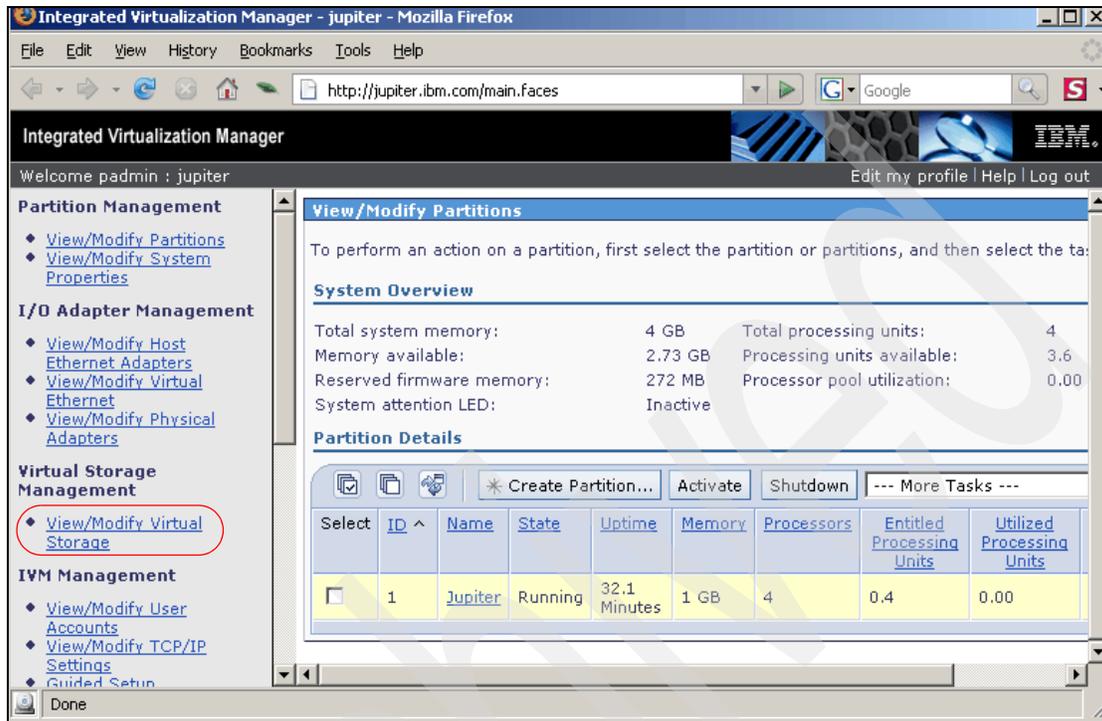


Figure 3-26 View and modify virtual storage

3.6.1 Physical volumes

Physical volumes are the hard drives that are available to the VIOS. They can be installed locally in the IBM BladeCenter JS22 or JS12 blades, SAS drives available from IBM BladeCenter S chassis, or LUNs available from an iSCSI or Fibre Channel storage area network subsystem.

A physical volume is shown as hdisk0, hdisk1 and so on the Virtual I/O Server. Logical partitions can be assigned complete physical volumes. However, they will appear as a virtual SCSI disk drive on the LPAR. This direct assignment is a requirement if you are planning for Live Partition Mobility or using IBM i. Physical volumes can also be used to build storage pools.

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 3-27 on page 100. This displays the list of the physical volumes available to the VIOS.

View/Modify Virtual Storage					
Virtual Disks		Storage Pools	Physical Volumes	Optical Devices	
To perform an action on a physical volume, first select the physical volume or physical volumes, and then select the task.					
<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>		Modify partition assignment		--- More Tasks ---	
Select	Name ^	Storage Pool	Assigned Partition	Size	Physical Location Code
<input type="checkbox"/>	hdisk0	rootvg (Default)		30 GB	U78A5.001.PARM125-P1-C7-T1-W203200A0B811A662-L0
<input type="checkbox"/>	hdisk1			50 MB	U78A5.001.PARM125-P1-C7-T1-W203200A0B811A662-L1000000000000
<input type="checkbox"/>	hdisk2			20 GB	U78A5.001.PARM125-P1-C7-T1-W203200A0B811A662-L2000000000000
<input type="checkbox"/>	hdisk3			5 GB	U78A5.001.PARM125-P1-C7-T1-W203200A0B811A662-L3000000000000
<input type="checkbox"/>	hdisk4			5 GB	U78A5.001.PARM125-P1-C7-T1-W203200A0B811A662-L4000000000000
<input type="checkbox"/>	hdisk5			5 GB	U78A5.001.PARM125-P1-C7-T1-W203200A0B811A662-L5000000000000

Figure 3-27 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 3-12 shows the output of the **lsdev -type disk** command.

Example 3-12 Physical volumes found with **lsdev**

```

$ lsdev -type disk
name          status      description
hdisk0        Available  1815    DS4800 Disk Array Device
hdisk1        Available  1815    DS4800 Disk Array Device
hdisk2        Available  1815    DS4800 Disk Array Device
hdisk3        Available  1815    DS4800 Disk Array Device
hdisk4        Available  1815    DS4800 Disk Array Device
hdisk5        Available  1815    DS4800 Disk Array Device
hdisk6        Available  1815    DS4800 Disk Array Device
hdisk7        Available  1815    DS4800 Disk Array Device
hdisk8        Available  1815    DS4800 Disk Array Device

```

Example 3-13 shows the output of the **lspv -size** command.

Example 3-13 Physical volumes found with **lspv -size**

```

$ lspv -size
NAME          PVID          SIZE(megabytes)
hdisk2        0000059a7372d3a2  20480
hdisk0        0000059a4b9d273c  30720
hdisk1        0000059a573d00cf  30720
hdisk3        0000059a9d6e8947  5120

```

hdisk4	0000059a9d56323e	5120
hdisk5	none	10240
hdisk6	none	5120
hdisk7	none	5120
hdisk8	000570081e960af8	5120

3.6.2 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 the 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, click the **Storage Pools** tab from the View/Modify Virtual Storage window. Figure 3-28 shows a list of all available storage pools.



Figure 3-28 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) that will be used 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).

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**.

Note: Do not use lengthy names for your volume groups, because they will not be shown in their full length within the IVM Web interface. Use 12 or fewer characters.

Figure 3-29 shows that, in this case, *hdisk2* was chosen.

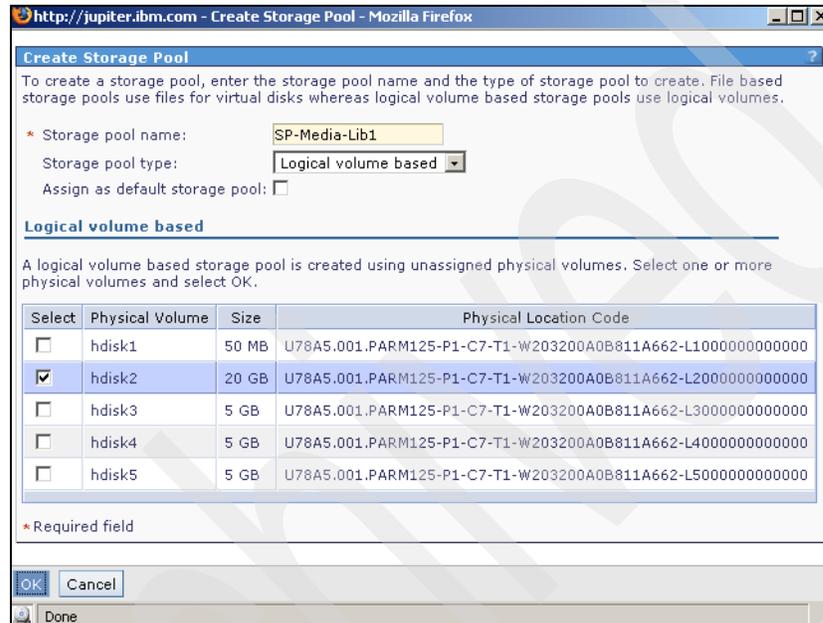


Figure 3-29 Create new storage pool

Figure 3-29 shows the new storage pool.

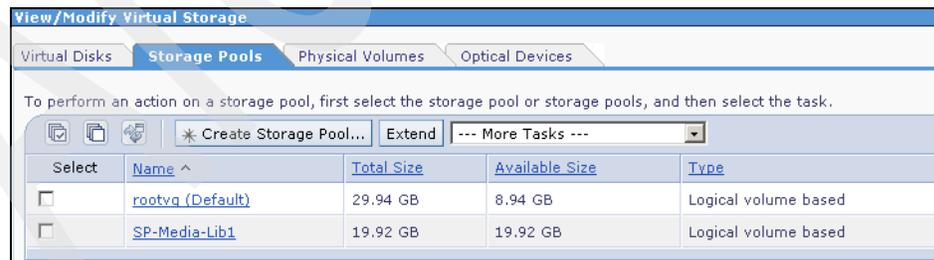


Figure 3-30 Newly created storage pool shown in IVM

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 3-31. A dialog opens that guides you through the modification of the storage pool.

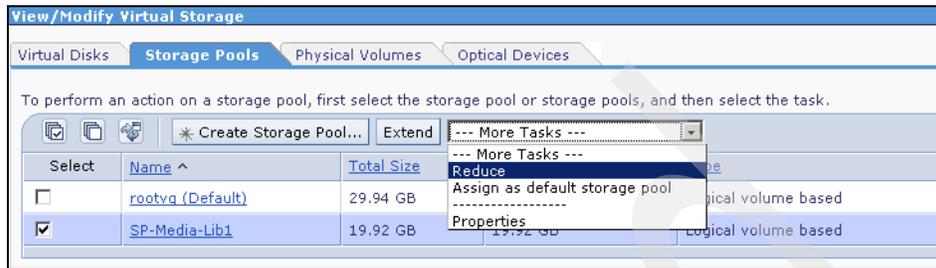


Figure 3-31 Reduce or delete a storage pool

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 3-32.



Figure 3-32 Delete storage pool.

3.6.3 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 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 3.7.2, “Partition Name and Environment” on page 118, 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 3.6.2, “Storage pools” on page 101, 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 the **Create Virtual Disk** button as shown in Figure 3-33; this will create a new virtual disk.

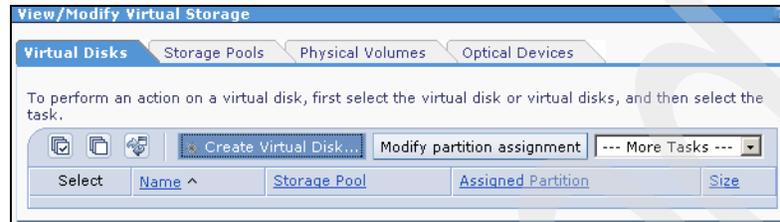


Figure 3-33 Create virtual disk

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 3-34 on page 105. 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}`.

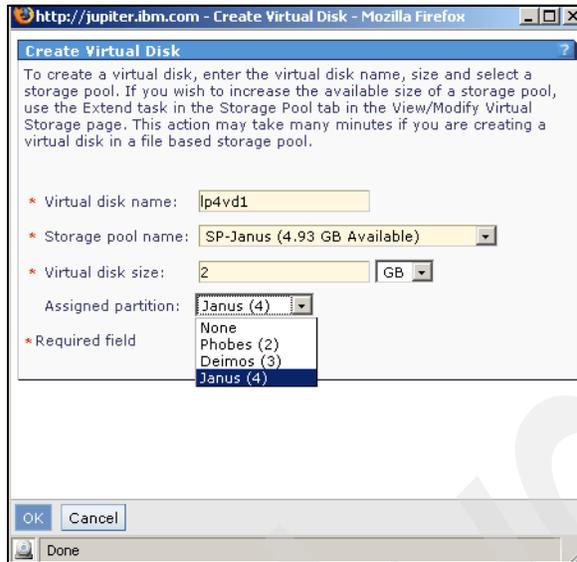


Figure 3-34 Virtual disk settings

The newly created virtual disk appears in the list, as shown in Figure 3-35.

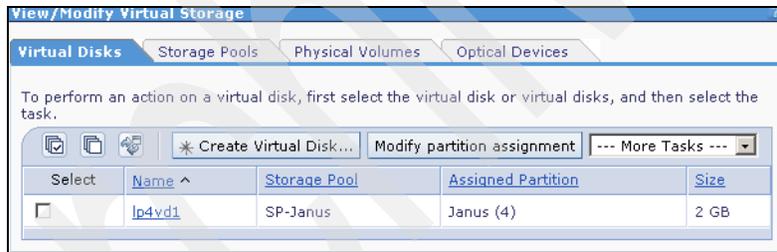


Figure 3-35 The newly-created virtual disk

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 as long as 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. Select the **More Tasks...** drop-down box, and then select **Extend**, as shown in Figure 3-36 on page 106.

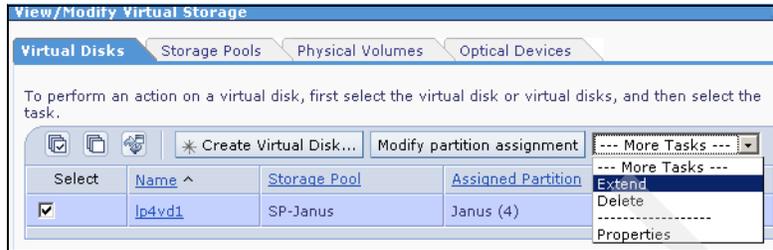


Figure 3-36 Extend virtual disk

Specify the amount of space that the virtual disk will be extended, then click **OK** as shown in Figure 3-37. 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 will be generated, alerting the administrator. To continue, select the **Force extend on running partition** check box and click the **OK** button again.

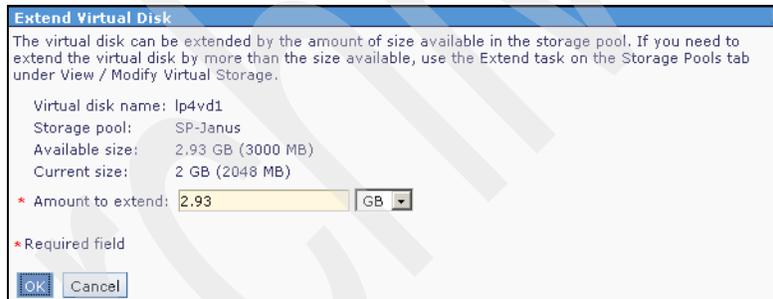


Figure 3-37 Virtual disk extension settings

The new size is shown in the list of available virtual disks when the extension is complete, as shown in Figure 3-38 on page 107.

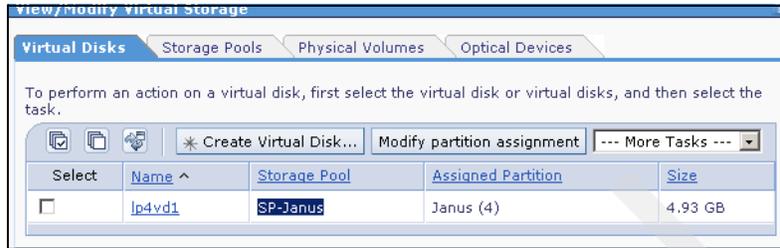


Figure 3-38 Extended virtual disk

The next section explains how to delete an 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 attempt to delete a virtual disk on a running partition, a warning message will be generated, alerting the administrator. To continue, select the **Force device removal from a running partition** check box and click the **OK** button 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 select **Delete** from the **More Tasks...** drop down box as shown in Figure 3-39.

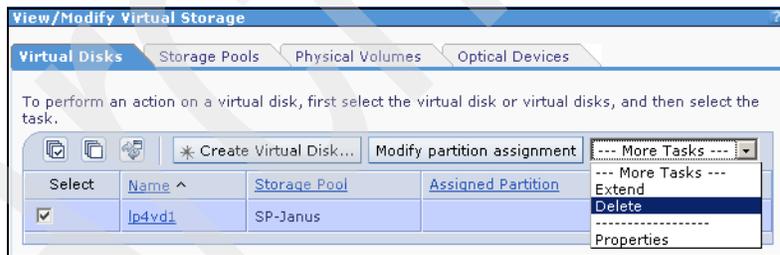


Figure 3-39 Delete virtual disk

Confirm the deletion of the virtual disk by clicking **OK**, as shown in Figure 3-40 on page 108.



Figure 3-40 Confirm deletion of the virtual disk

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

3.6.4 Optical devices

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

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

Physical optical devices

Physical optical devices are the CD or DVD drives installed in the media tray of a 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 an CD or DVD in your laptop or desktop can be assigned to the blade. The Web interface of the Advanced Management Module provides this capability.

As Table 3-1 shows, these two different drives can be identified by their location paths.

Table 3-1 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.

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.

Changing the assignment of physical optical drives

To change the assignment of physical optical drives, click the **Optical Devices** tab. Figure 3-41 on page 110 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 called Deimos.

The list of physical devices may vary, depending on the media tray assignment and the usage of remote media. Figure 3-41 on page 110 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 in the BladeCenter chassis.

Note: The remote control function for the IBM BladeCenter JS22 or JS12 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.

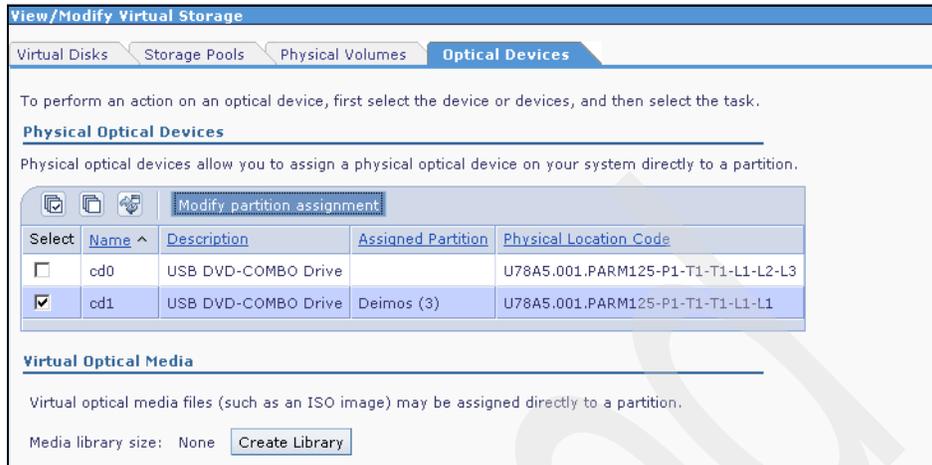


Figure 3-41 Physical optical devices in IVM

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 3-42.



Figure 3-42 Change physical optical device assignment

Virtual optical devices

Virtual optical devices are a new feature of the 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, a storage pool must be available. Refer to 3.6.2, "Storage pools" on page 101, for an explanation about how to set up a storage pool.

To set up a media library, follow these steps:

Click the **Optical Devices** tab in the View/Modify Virtual Storage window to create a media library. Then click the **Create Library** button, as shown in Figure 3-43.



Figure 3-43 Create media library

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** as shown in Figure 3-44.

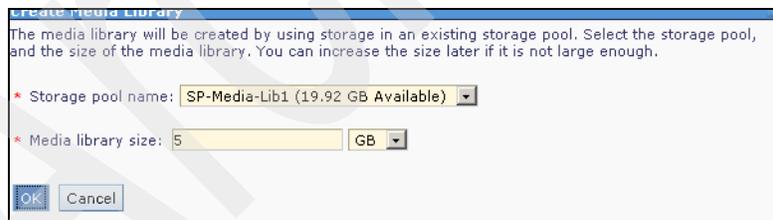


Figure 3-44 Media library size and storage pool

Depending on the size of the media library, the creation time will 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 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 need 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 onto the virtual optical device.

Note: Currently, the only way to perform a backup from an IBM i partition is to a virtual optical device assigned to the IBM i LPAR.

To add new media in the media library, click **Add Media...** as shown in Figure 3-45.

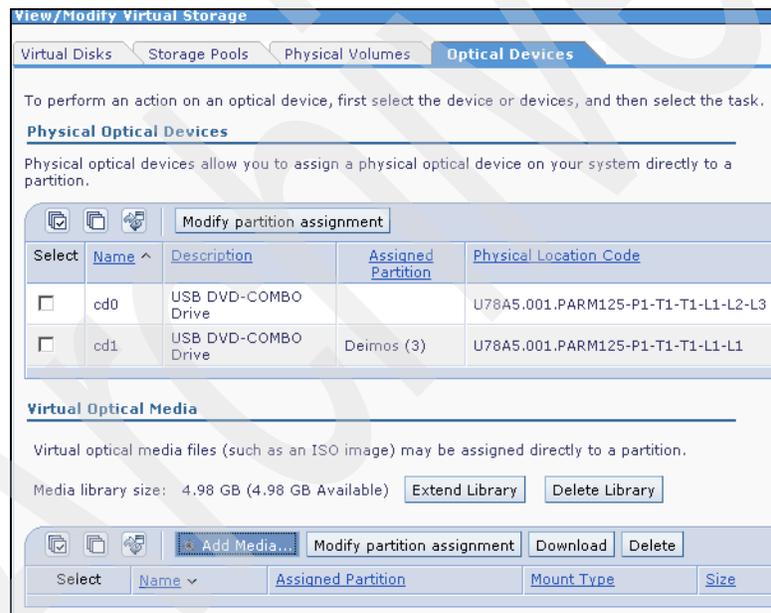


Figure 3-45 Add media to media library

There are four options to create new media:

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

The Upload media option allows you to transfer files or ISO images from a workstation directly to the media library. There is a limitation in the file size of 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 written to from an LPAR.

Figure 3-46 on page 114 shows as example that uses a physical optical device from the IBM BladeCenter 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 will return an error message stating there are too many parameters.

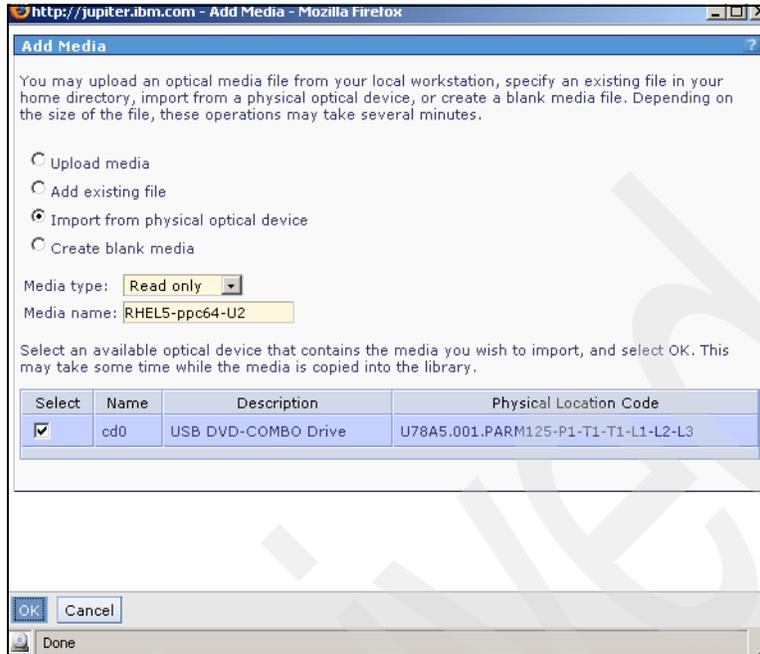


Figure 3-46 Add new media - settings

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 3-47.



Figure 3-47 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 3-48 on page 115.

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 will change to Read only.

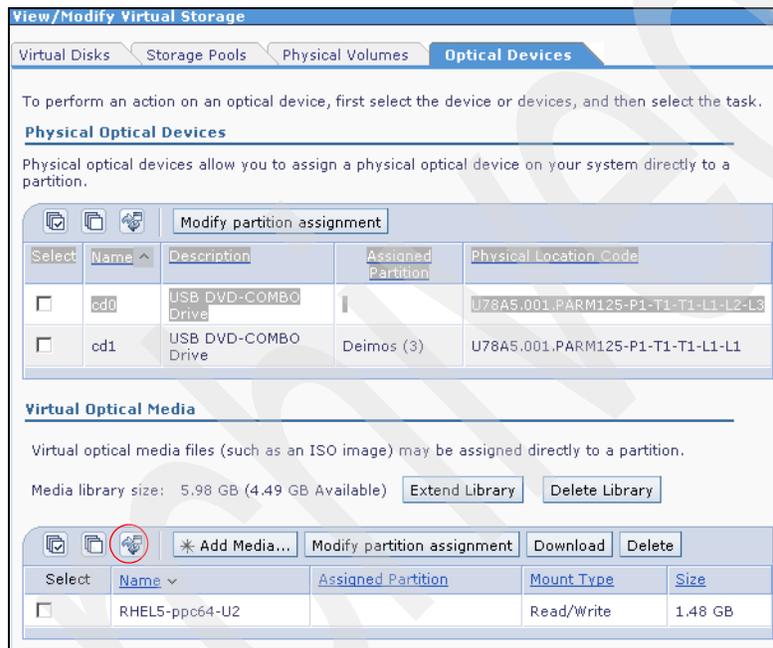


Figure 3-48 New created media with the copy operation in progress

Modifying media assignment to virtual optical devices in logical partitions

Media can be assigned from the Optical Devices tab in the View/Modify Virtual Storage window, when using the Create Partition wizard or from the Partition Properties window. The next step will be 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 desired media name then click **Modify partition assignment** under the Optical Devices tab from the View/Modify Virtual Storage window.

As shown in Figure 3-49, the media SLES10-ppc4-SP2 is already assigned to the LPAR Deimos. This media is again selected and will also be assigned to the LPAR Janus.

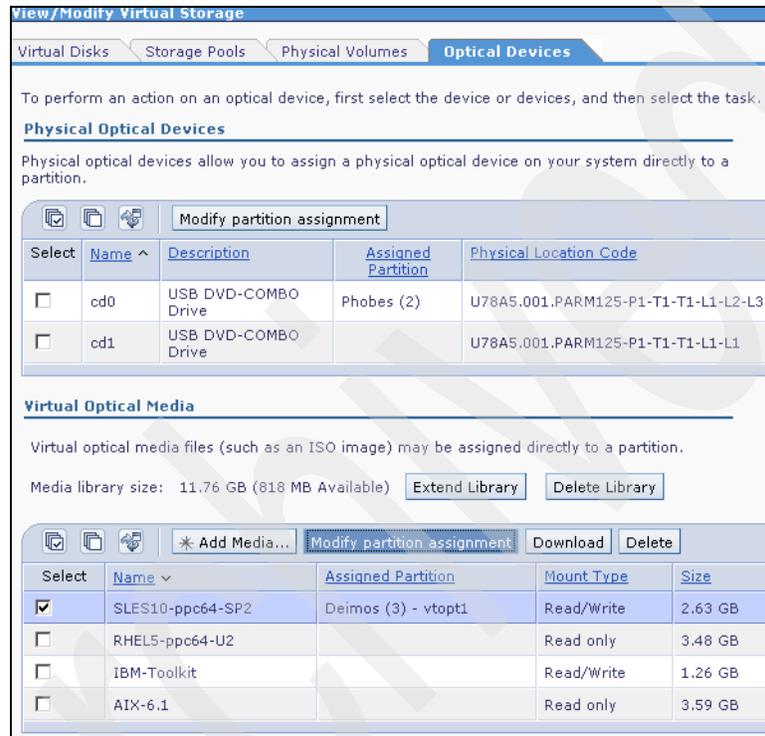


Figure 3-49 Modify partition assignment

Select the logical partition that needs the media assigned. 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 3-50 on page 117 shows that the LPARs Deimos and Janus are selected. The LPAR Deimos already has the media SLES10-ppc64-SP2 assigned. Next, LPAR Janus will be assigned the same media.

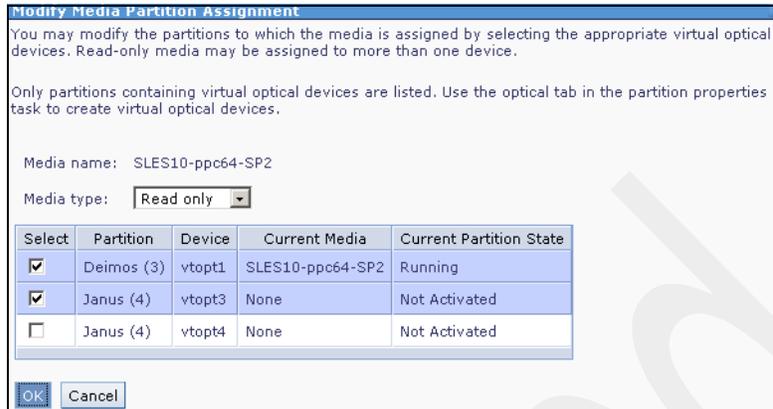


Figure 3-50 Modify media partition assignment

Click **OK** to return to the view of the optical devices. Notice that the updated table shown in Figure 3-51 now contains the LPARs Deimos and Janus in the Assigned Partition column as assigned partitions for the media SLES10-ppc64-SP2.

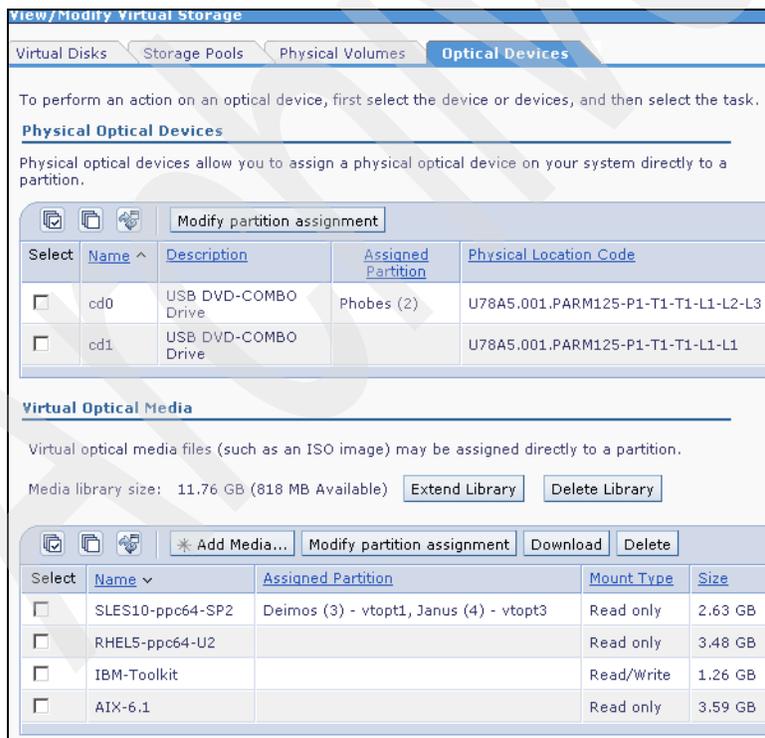


Figure 3-51 New assigned media to partition Janus

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

3.7 Partition configuration for Virtual I/O Client (VIOC)

With networking and storage defined, you can now create additional VIOC LPARs for the installation of additional supported operating systems.

3.7.1 Partition Mobility considerations

If Partition Mobility is going to be considered for a VIOC in logical partitions, you need to review the following planning considerations (currently unavailable on IBM i).

- ▶ VIOS running on source and target IBM BladeCenter JS22/JS12:
 - VIOS should be at the latest fixpack.
 - IBM BladeCenter JS22/JS12 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.
- ▶ Shared storage between the two Virtual I/O Servers.
- ▶ Only IVM- to-IVM managed systems are allowed (no HMC-to-IVM or IVM-to-MC is allowed)

Chapter 8, “Performing Live Partition Mobility” on page 329, provides setup and configuration details for Partition Mobility with a JS12 or JS22. The Redbooks publication *PowerVM Live Partition Mobility on IBM System p*, SG24-7460, provides an overview of Partition Mobility and explains how it applies to other System p configurations.

3.7.2 Partition Name and Environment

The create partition wizard is started with the **Create Partition** button on the **View/Modify Partitions** link in the navigation area of the IVM UI, as shown in Figure 3-52 on page 119.

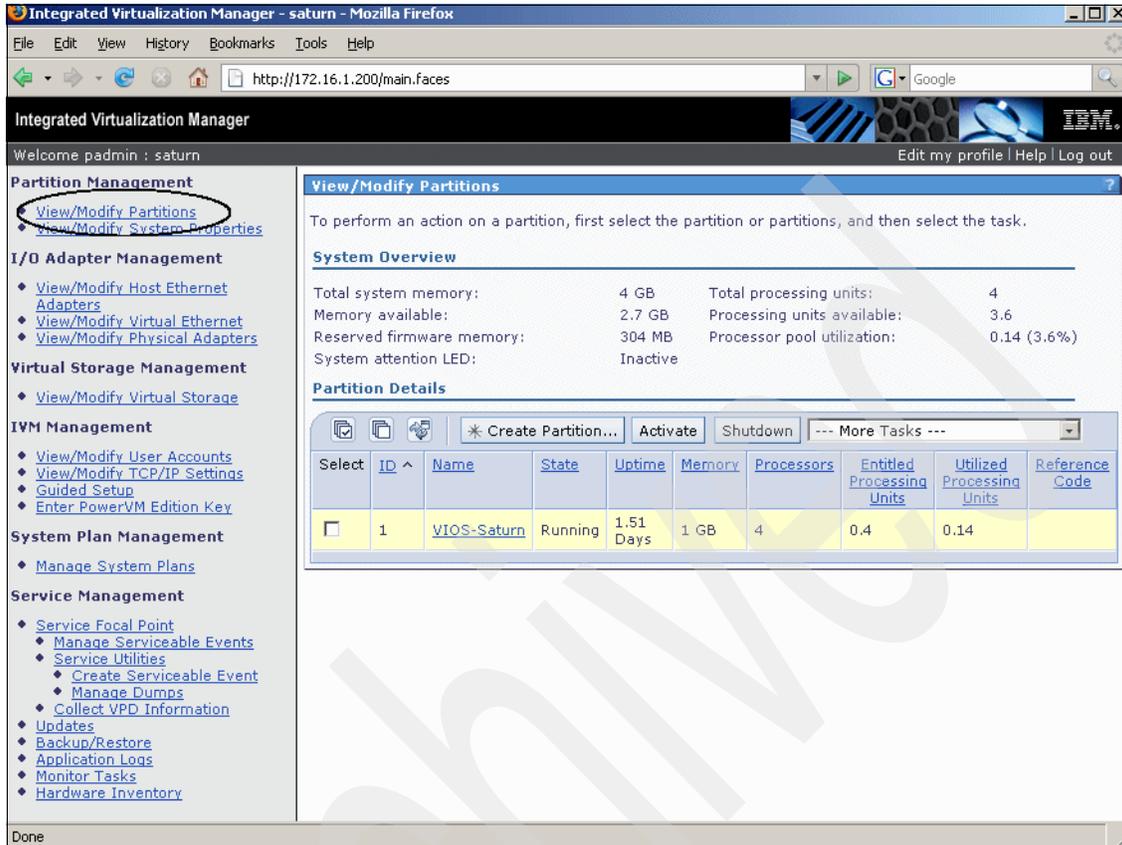


Figure 3-52 View/Modify Partition

3.7.3 Partition Name

When the wizard starts, a new window will open as shown in Figure 3-53 on page 120. This gives you the opportunity to change the Partition ID number, provide a Partition name, and select an operating system environment. Select the **Next** button for the memory step.

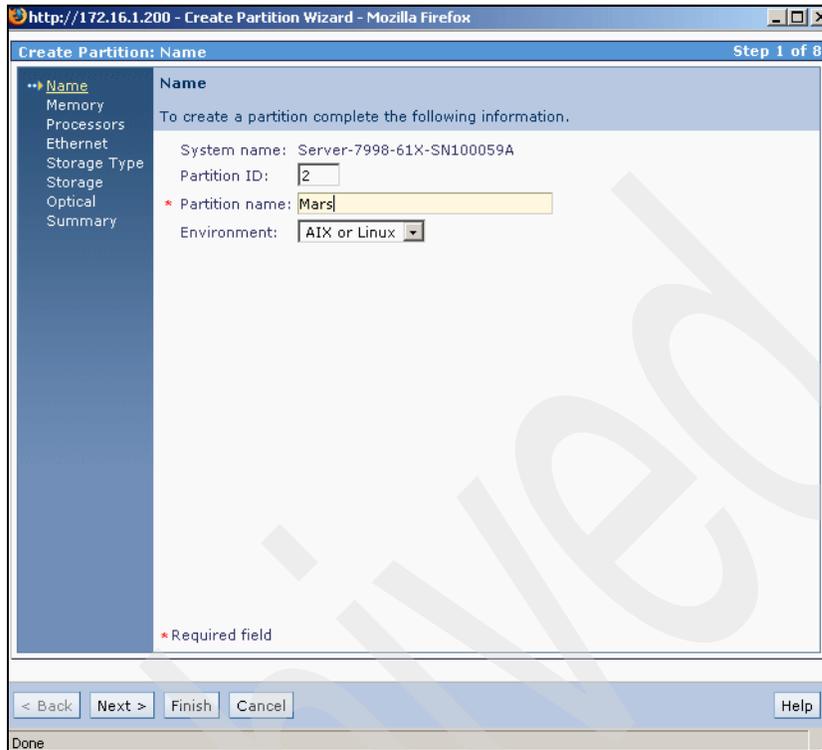


Figure 3-53 Create Partition: Name

3.7.4 Partition Memory

Figure 3-54 on page 121 shows how to assign memory to the partition. Total system memory and the current memory available for a new partition is summarized at the top of the window. After you enter the desired amount of memory, select the **Next** button.

Note: IVM does not allow you to over-commit memory or processor resources.

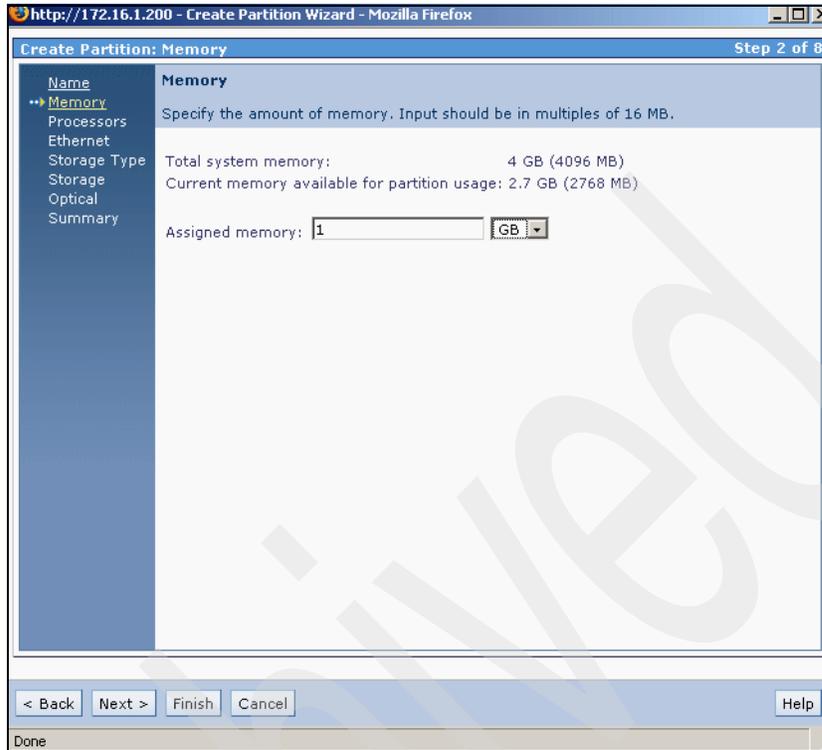


Figure 3-54 Create Partitions: Memory

3.7.5 Partition Processors

On the Create Partition: Processors window you have the option of assigning dedicated or shared processors. In shared mode, each virtual processor uses 0.1 physical processors (by default; more if available). In dedicated mode, each assigned processor uses one physical processor.

Available processor resources are displayed on the window and, as with memory resources, they cannot be over-committed. Figure 3-55 on page 122 shows a selection of shared mode and two assigned processors for this example. After you make your selections, click the **Next** button.

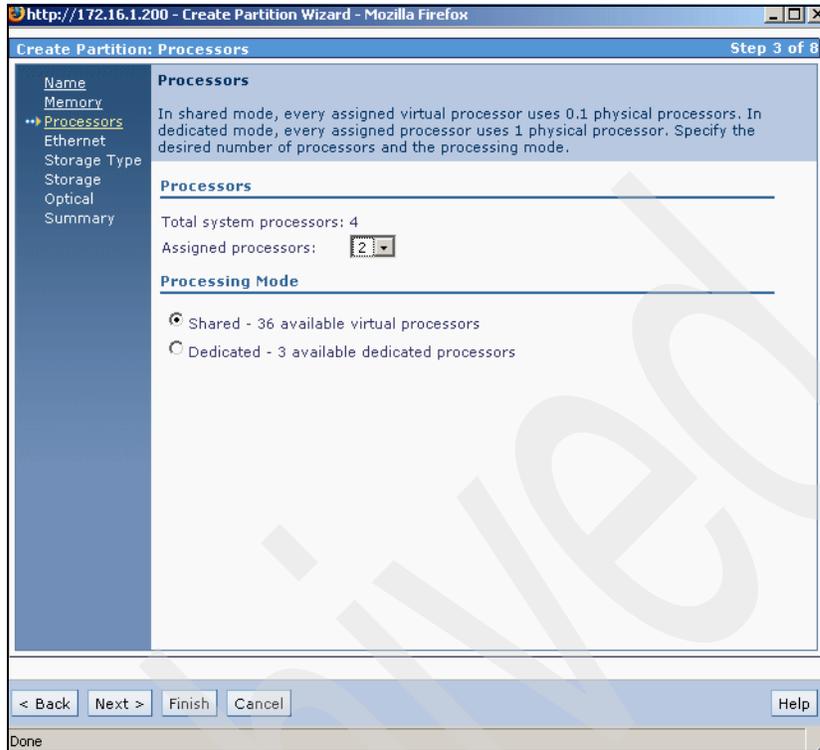


Figure 3-55 Create Partition: Processors

Note: After an LPAR is created, the processor mode cannot be changed from shared to dedicated or dedicated to shared.

3.7.6 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.

Figure 3-56 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 as this time to create additional virtual Ethernet adapters for the logical partition.

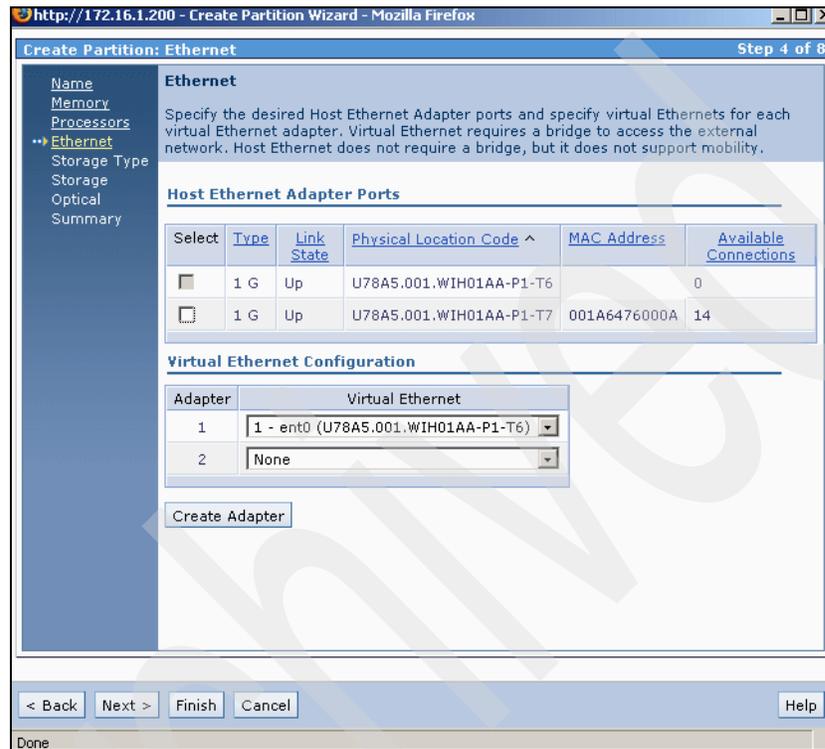


Figure 3-56 Create Partition: Ethernet

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

3.7.7 Partition Storage Types and Assignments

Logical partition disk storage can be virtual disks from a storage pool or physical volumes. Figure 3-57 on page 124 shows these choices and also provides the option to create additional virtual disks from a storage pool that has already been defined. Refer to 3.6.2, “Storage pools” on page 101 for an explanation about how to create storage pools.

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

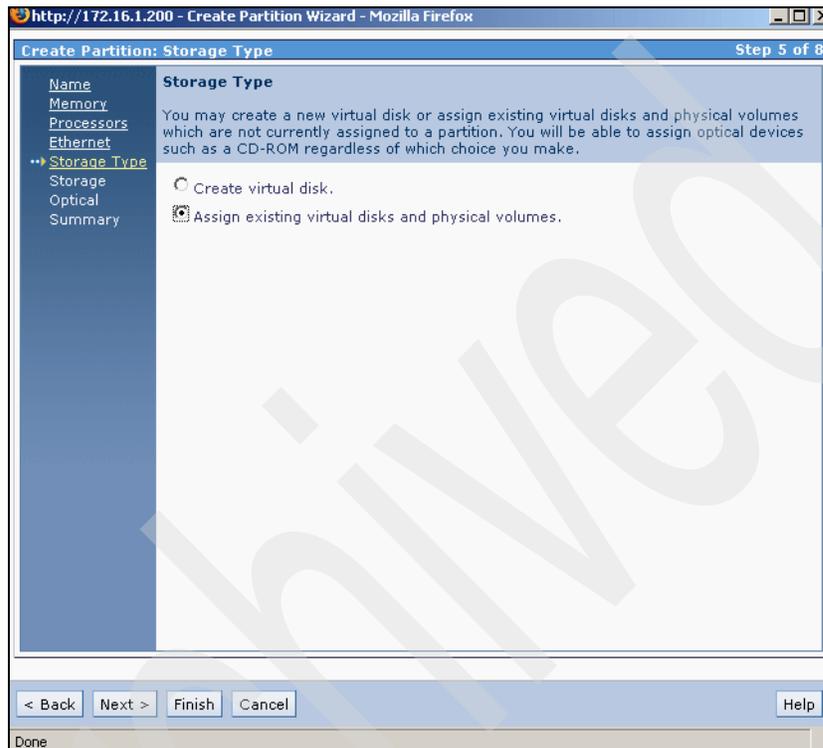


Figure 3-57 Create Partition: Storage Type

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

Figure 3-58 on page 125 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 the **Next** button.

Note: For initial LPAR setup and OS install, we recommend that you only select the install target disk at this time.

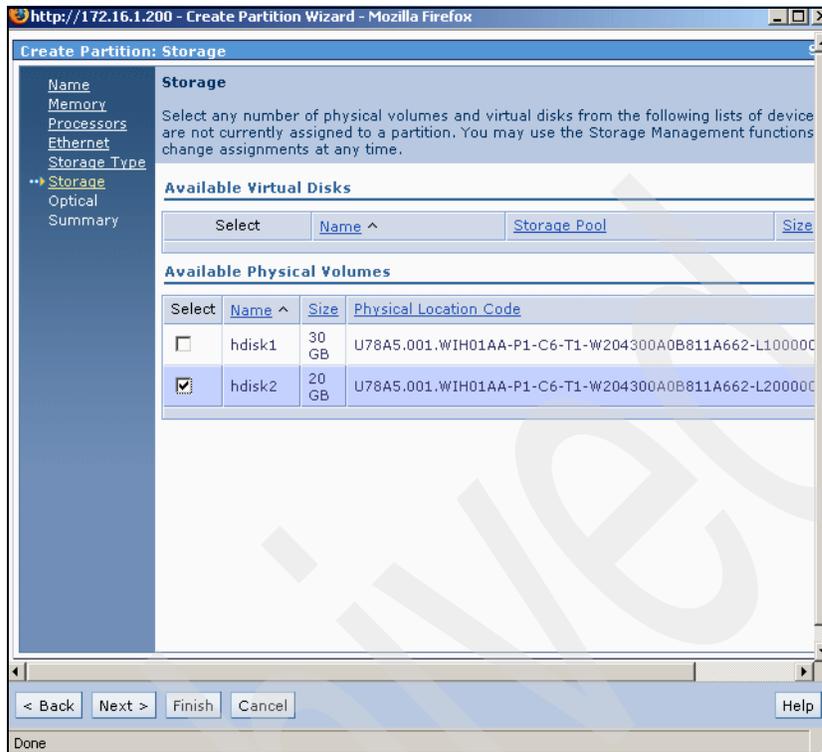


Figure 3-58 Logical Partition: Storage

3.7.8 Optical devices

Optical devices, both physical and virtual, can be assigned to an LPAR. With an IBM BladeCenter JS22 or JS12, the physical 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 110 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 optical device is required for backup for IBM i.

Figure 3-59 shows the optical device selection window. For this example, the physical optical device will not be required and no virtual optical devices have been created.

If your configuration would require a physical optical device for a media install, for instance, then you would select cd0 in the check box.

Click the **Next** button to proceed to the Physical Adapters window (if available) or the summary window.

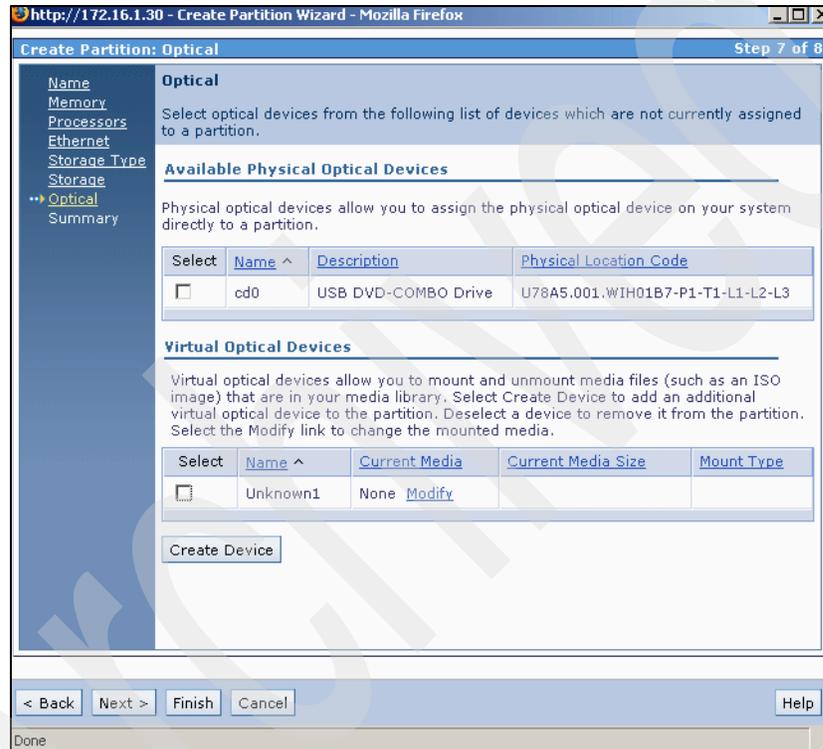


Figure 3-59 Create Partition: Optical

3.7.9 Physical adapters

The physical adapter step will only be shown if I/O adapters are available for assignment to an LPAR. Figure 3-60 on page 127 shows availability of an Ethernet adapter that could be assigned to the example LPAR.

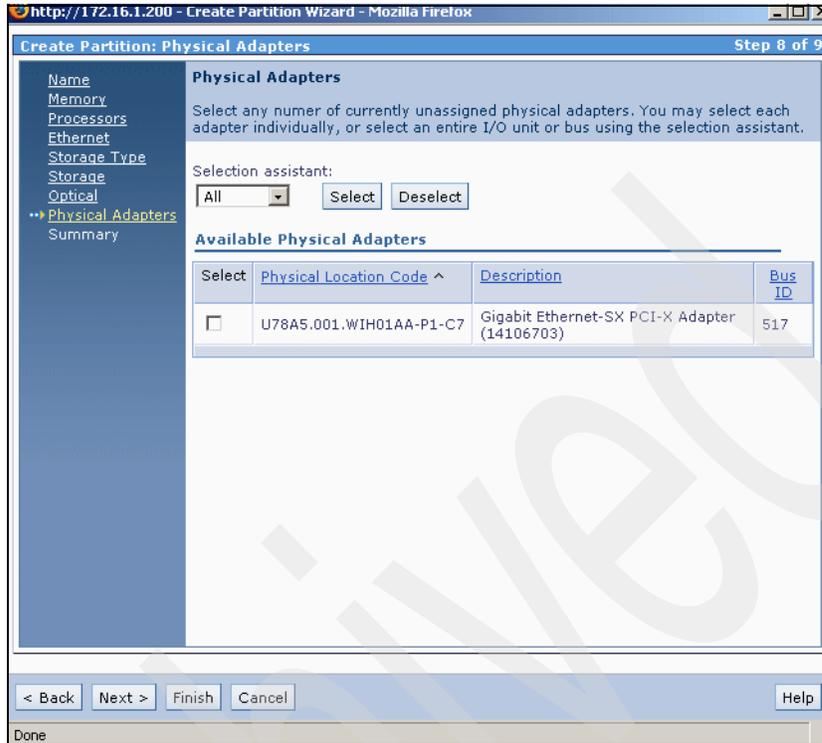


Figure 3-60 Create Partition: Physical Adapters

3.7.10 Partition Summary

The final window of the Create Partition wizard is the summary, as shown in Figure 3-61 on page 128. All of the previous selections can be reviewed on this window and edited if required by using the **Back** button.

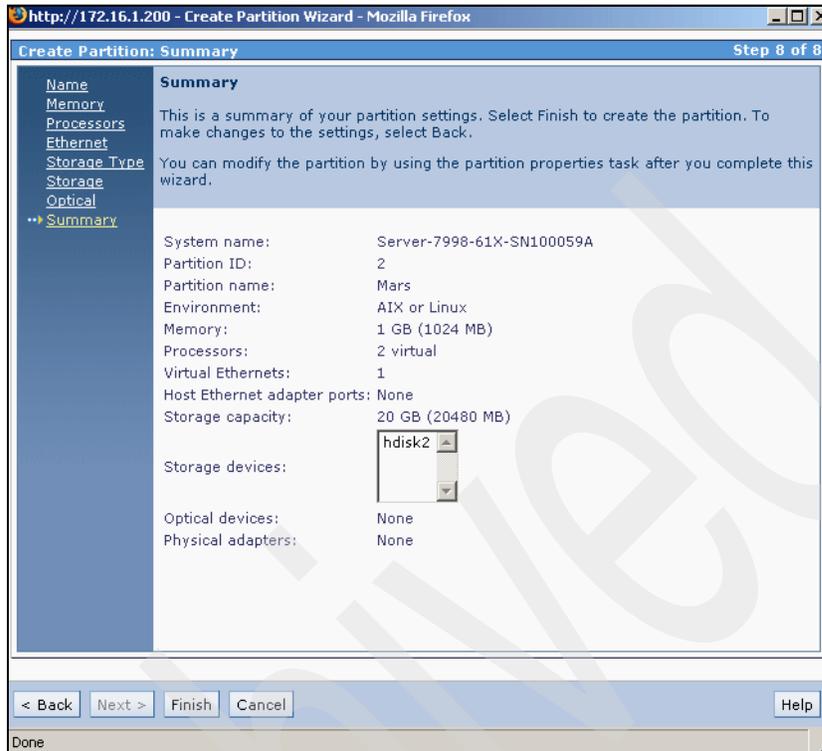


Figure 3-61 Create Partition: Summary

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

Figure 3-62 on page 129 of the View/Modify Partitions window shows the new logical partition that was created.

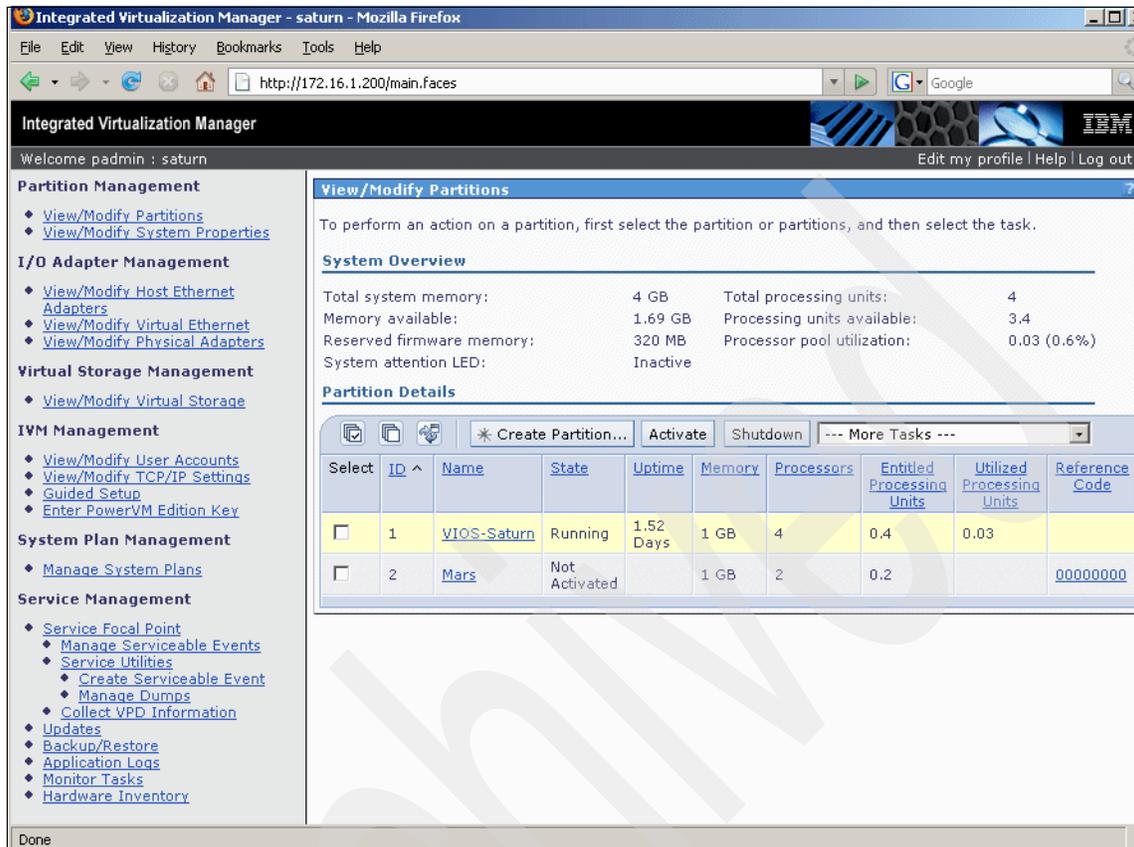


Figure 3-62 View/Modify Partition showing new partition

3.7.11 Partition Properties Changes and DLPAR operations

The IVM UI provides quick access to change an LPAR's properties and perform Dynamic LPAR (DLPAR) operations. The IBM BladeCenter JS22 and JS12 have the capability to perform DLPAR operations on memory, processors, and I/O adapters.

Partition properties changes and DLPAR operations are accessed by clicking the name of the partition to be modified when viewed from View/Modify Partitions view. When a partition name is selected, a new window will open to the General tab as shown in Figure 3-63 on page 130.

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.

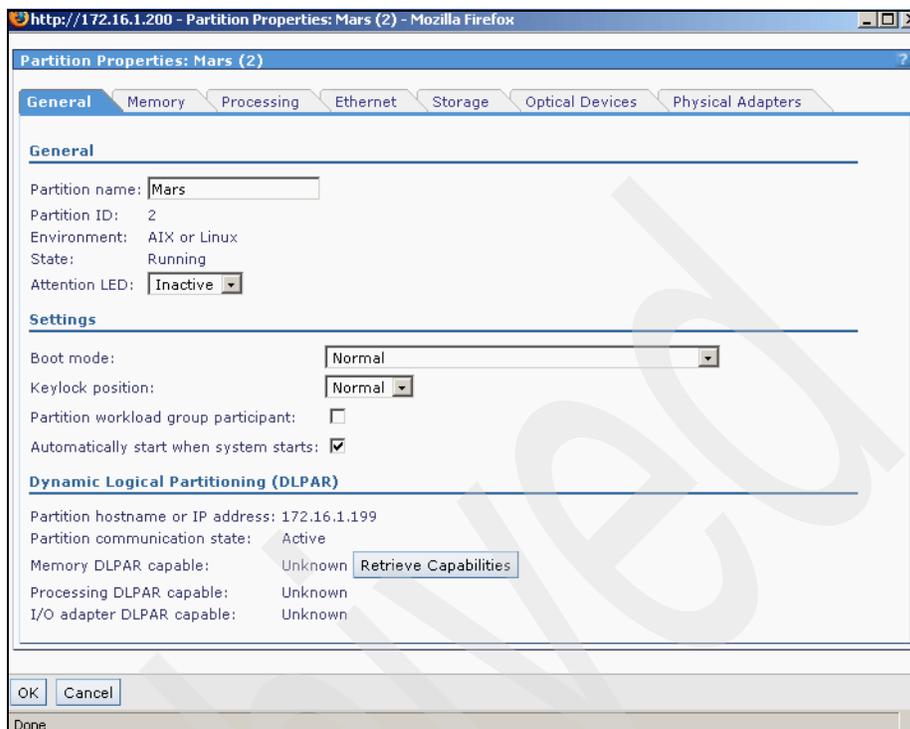


Figure 3-63 Partition Properties General tab

DLPAR capabilities can be retrieved by clicking the **Retrieve Capabilities** button. Figure 3-64 on page 131 shows the DLPAR capabilities of the IBM BladeCenter JS22 and JS12.

IBM i LPARs have a different Partition Properties General tab view, as shown in Figure 5-31 on page 249. They do not have a Physical Adapters tab.

Note: The LPAR must have a network interface configured and have an active Resource Monitoring and Control (RMC) connection with the management partition you can retrieve capabilities or perform DLPAR operations.

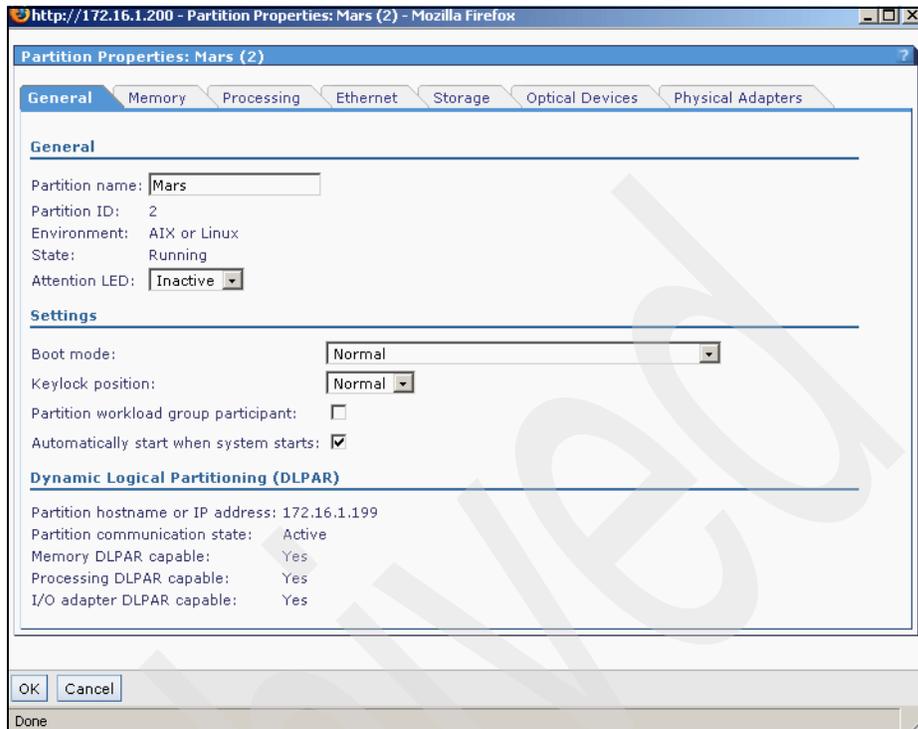


Figure 3-64 DLPAR retrieved capabilities

Selecting the **Memory** tab will display current and pending memory values for the LPAR, as shown in Figure 3-65 on page 132. 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 only be changed when the LPAR is *not* activated.

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

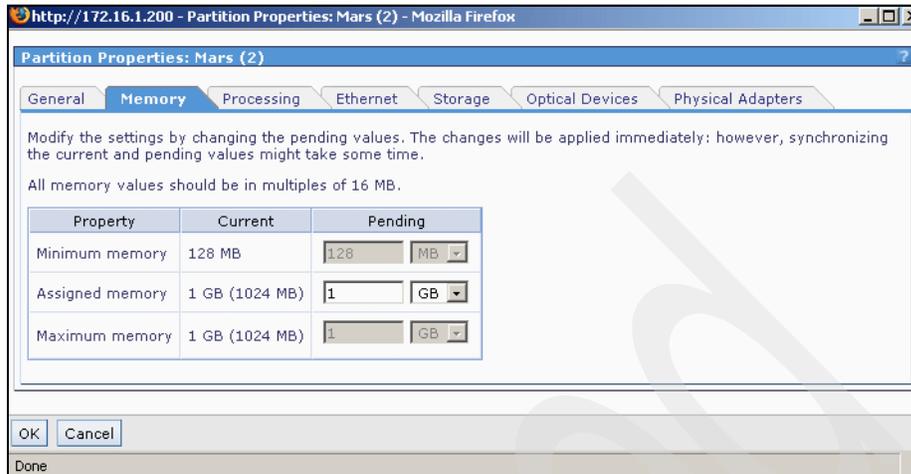


Figure 3-65 Partition Properties Memory tab

The **Processing** tab is used to change the processing units, virtual processors, and partition priority weighting for LPARS and the VIOS using the shared processor pool, as shown in Figure 3-66.

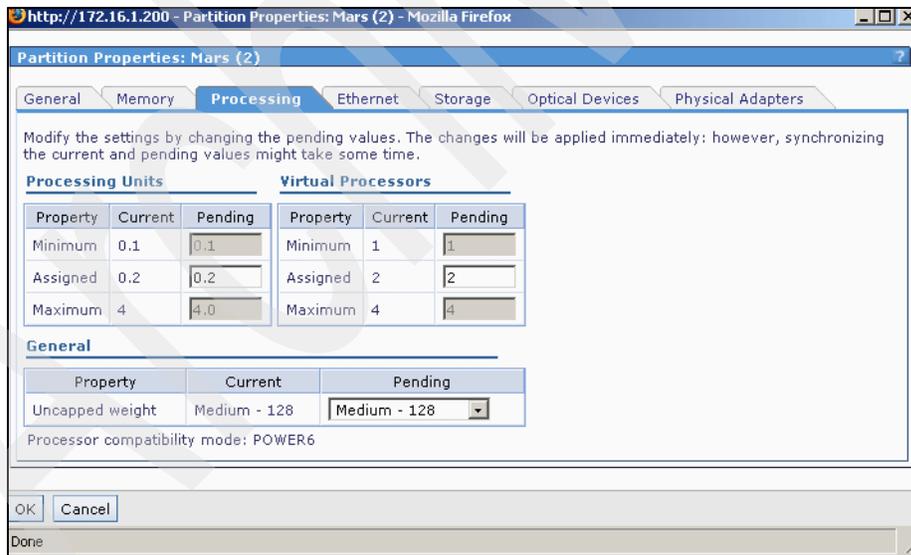


Figure 3-66 Partition Properties, Processing tab for shared pool

Partitions using dedicated processors will display the window as shown in Figure 3-67 on page 133. This example shows the LPAR in a not activated

state 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 allows 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

The excess capacity that is shared is utilized by the shared pool.

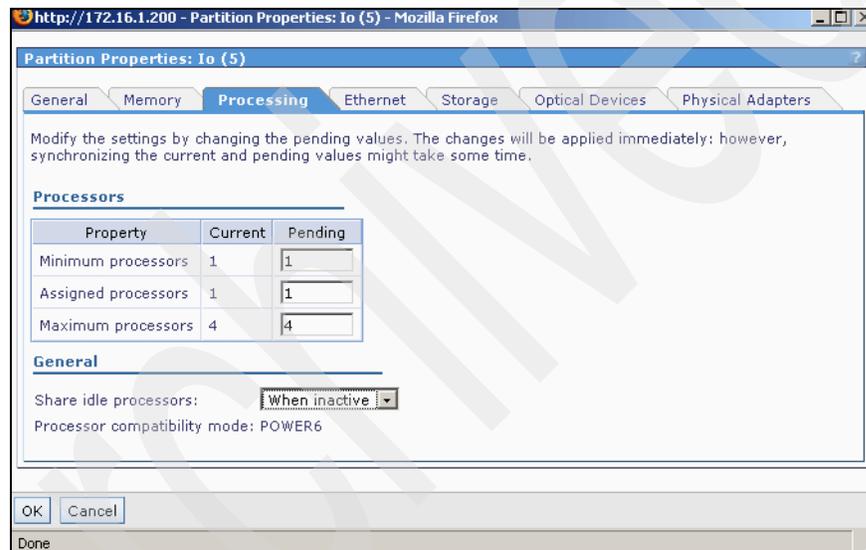


Figure 3-67 Partition Properties, Processing tab for dedicated processors

The **Ethernet** tab in Partition Properties allows the addition or removal of Ethernet adapters, as shown in Figure 3-68 on page 134.

Note: Before you can DLPAR remove Ethernet adapters from an active AIX LPAR, first use the `rmdev` command to removed 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 by the `lsdev` command. For example:

```
# lsdev -Cl ent1 -F parent
pci1
```

The alternate method is to try the DLPAR command; in that case IVM will return an error message with details about the parent and child devices that must be removed.

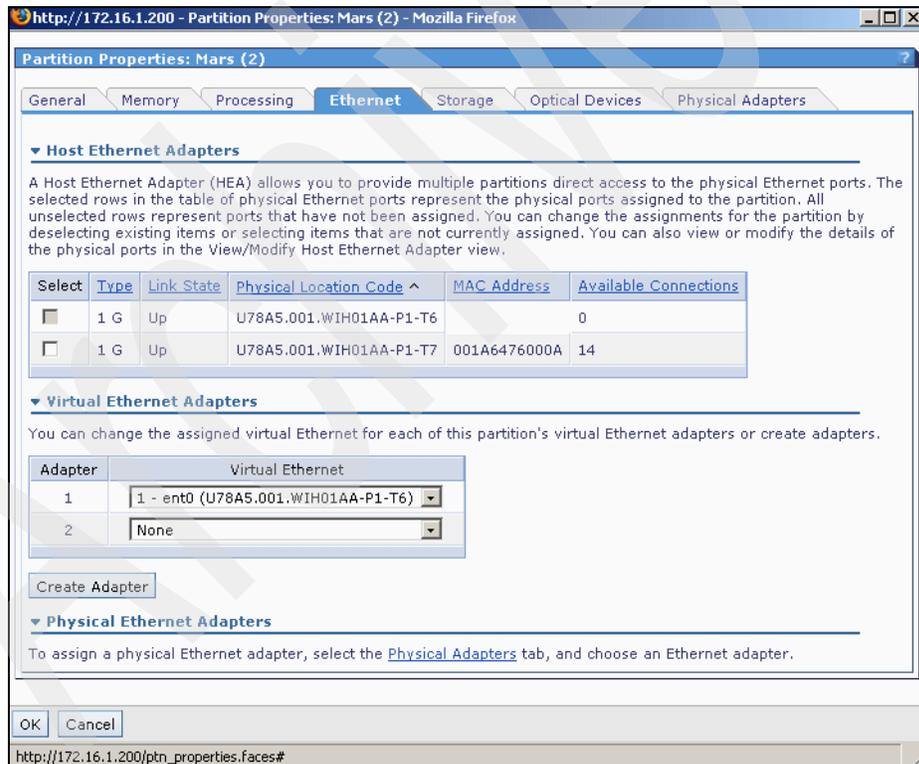


Figure 3-68 Partition Properties, Ethernet tab

Note: IBM i partitions cannot own HEAs. Therefore, the Host Ethernet Adapter section of this window will not be shown when viewing the properties of an IBM i partition. IBM i uses an HEA when configured as an SEA.

The **Storage** tab can be used to add or remove storage devices, either physical volumes or virtual disks, as shown in Figure 3-69.

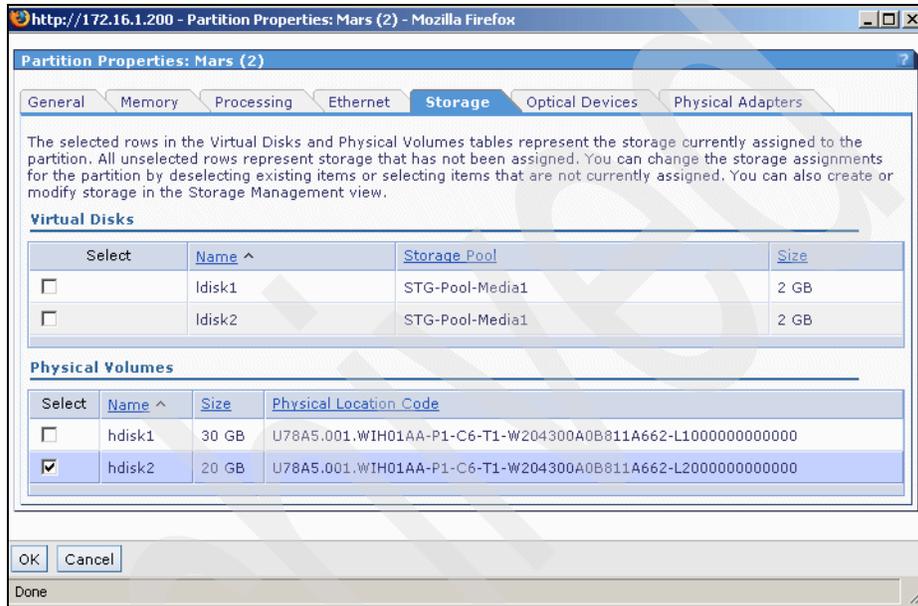


Figure 3-69 Partition Properties, Storage tab

Note: When removing storage from an active partition, IVM will require that you verify that a forced device removal from a running partition is requested.

Optical device assignments, both physical and virtual, can be managed from the **Optical Devices** tab shown in Figure 3-70 on page 136.

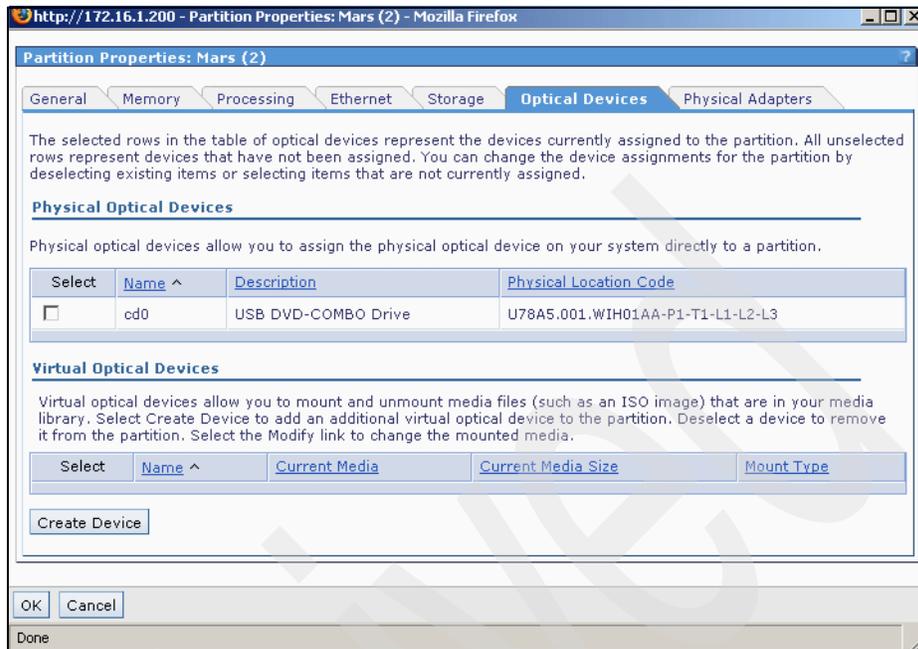


Figure 3-70 Partition Properties Optical Devices tab

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 110.

Figure 3-71 on page 137 shows an existing virtual optical device, vtopt1, being modified by changing the current media.

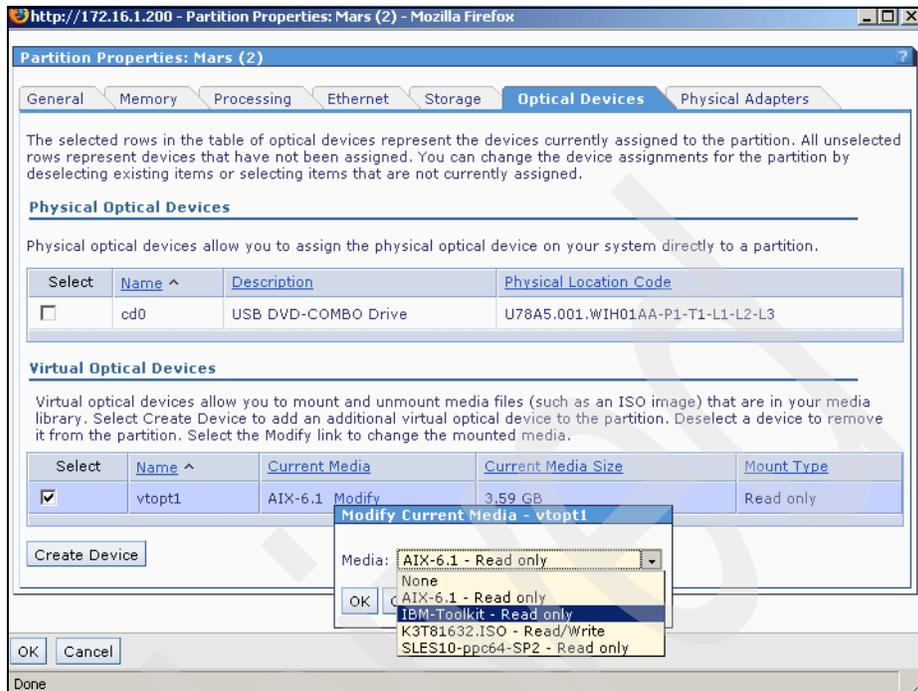


Figure 3-71 Partition Properties, changing the current mounted media

Physical adapters that are not assigned to an LPAR or any physical adapters that are already assigned to the selected LPAR will be displayed when the **Physical Adapters** tab is clicked.

Figure 3-72 on page 138 shows a PCI Express 4Gb FC Adapter available for assignment to this LPAR, and a Gigabit Ethernet-SX PCI-X Adapter that is already assigned.

Note: IBM i partitions cannot use physical adapters. Therefore, the Physical Adapter tab of this window will not be shown when viewing the properties of an i partition.

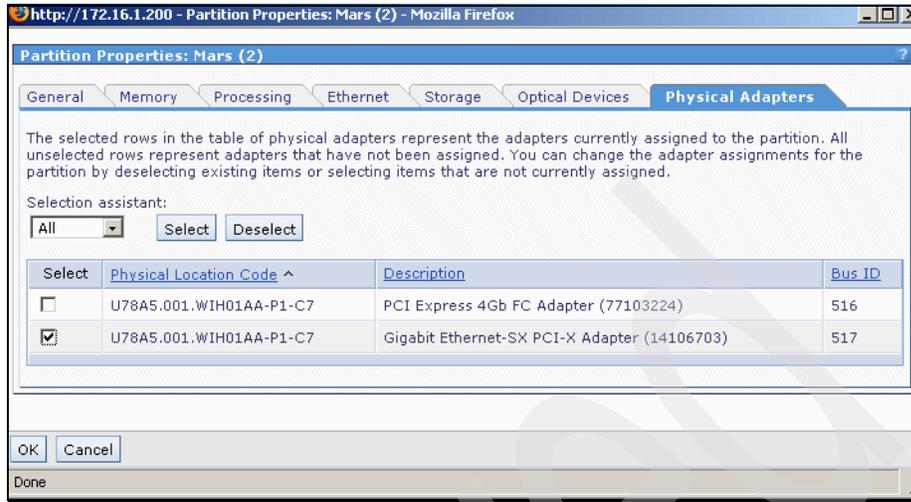


Figure 3-72 Partition Properties, Physical Adapters tab

Note: The Partition Properties window for the VIOS partition does not have the **Storage** and **Optical Devices** tabs.

3.8 Console access and activating a partition

The following sections discuss basic access to a partition and partition management functions.

3.8.1 Opening a virtual terminal

Accessing a partition virtual terminal from the VIOS can be done in two different 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 UI is shown in Figure 3-73 on page 139. From the View/Modify Partitions view, select the check box for the desired LPAR. From the drop-down box, select **Open terminal window**.

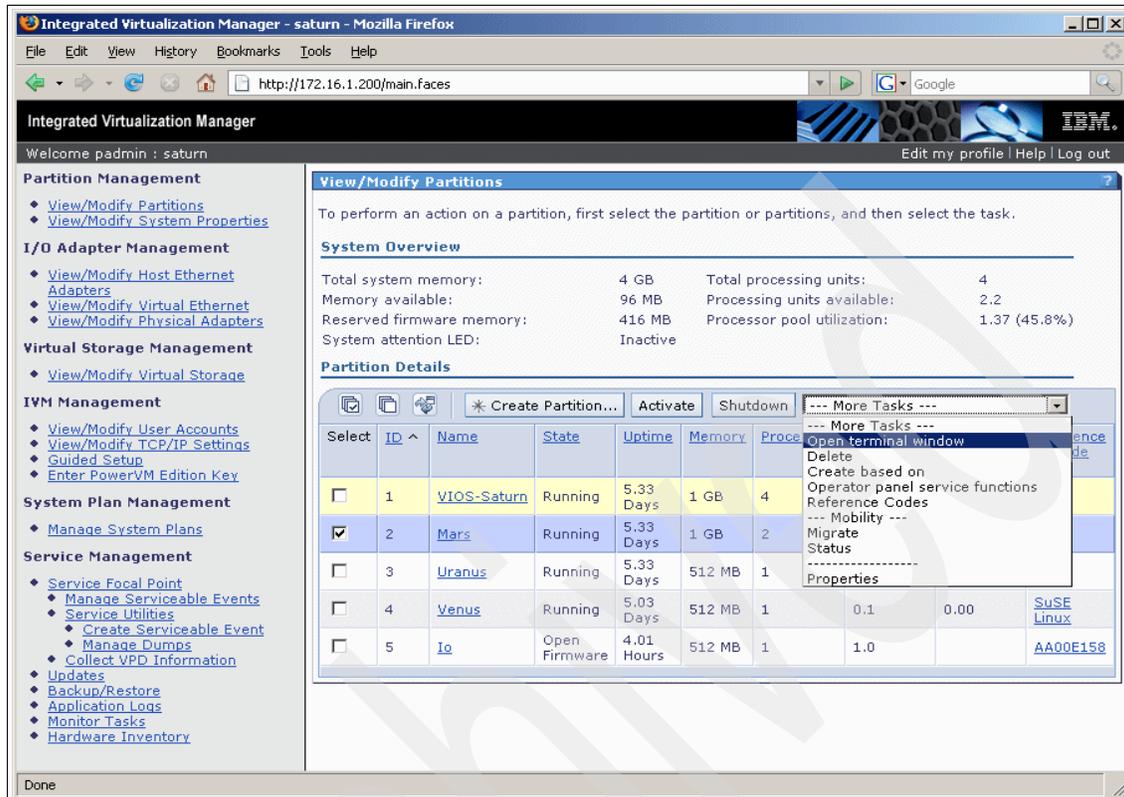


Figure 3-73 Start virtual terminal for an LPAR

A separate applet will start and open the virtual terminal window. This new window will prompt for the password of the login ID of the current IVM session.

Figure 3-74 shows a successful connection to the LPAR virtual terminal.



Figure 3-74 Virtual Terminal started from IVM UI

The second method to start a virtual terminal is from the VIOS command line. From the command line prompt, issue the `mkvt` command as shown in Example 3-14.

Example 3-14 Creating a virtual terminal from the command line

```
$ mkvt -id 2
```

Specify the partition number that you want to connect after the `-id` flag. Your command line session to the VIOS will now become the virtual terminal for the specified partition.

Note: The key sequence `<cr>~.` or enter key, tilde, period entered from the virtual terminal will allow you to break out of the `mkvt` command or close the virtual terminal applet.

3.8.2 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 UI

Activating a partition from the IVM UI starts from View/Modify Partitions. In this view, select the partition to be activated and click the **Activate** button as shown in Figure 3-75.

The screenshot displays the IVM 'View/Modify Partitions' page. The left sidebar contains navigation links for Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, and Service Management. The main content area includes a 'System Overview' section with system statistics and a 'Partition Details' table. The table lists five partitions, with partition 2 ('Mars') selected and its state set to 'Not Activated'. The 'Activate' button in the task bar is highlighted.

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	VIOS-Saturn	Running	5.39 Days	1 GB	4	0.4	1.44	
<input checked="" type="checkbox"/>	2	Mars	Not Activated		1 GB	2	0.2		00000000
<input type="checkbox"/>	3	Uranus	Running	5.39 Days	512 MB	1	0.1	0.01	
<input type="checkbox"/>	4	Venus	Running	5.09 Days	512 MB	1	0.1	0.04	SuSE Linux
<input type="checkbox"/>	5	Io	Open Firmware	5.38 Hours	512 MB	1	1.0		AA00E158

Figure 3-75 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 3-76 on page 142.

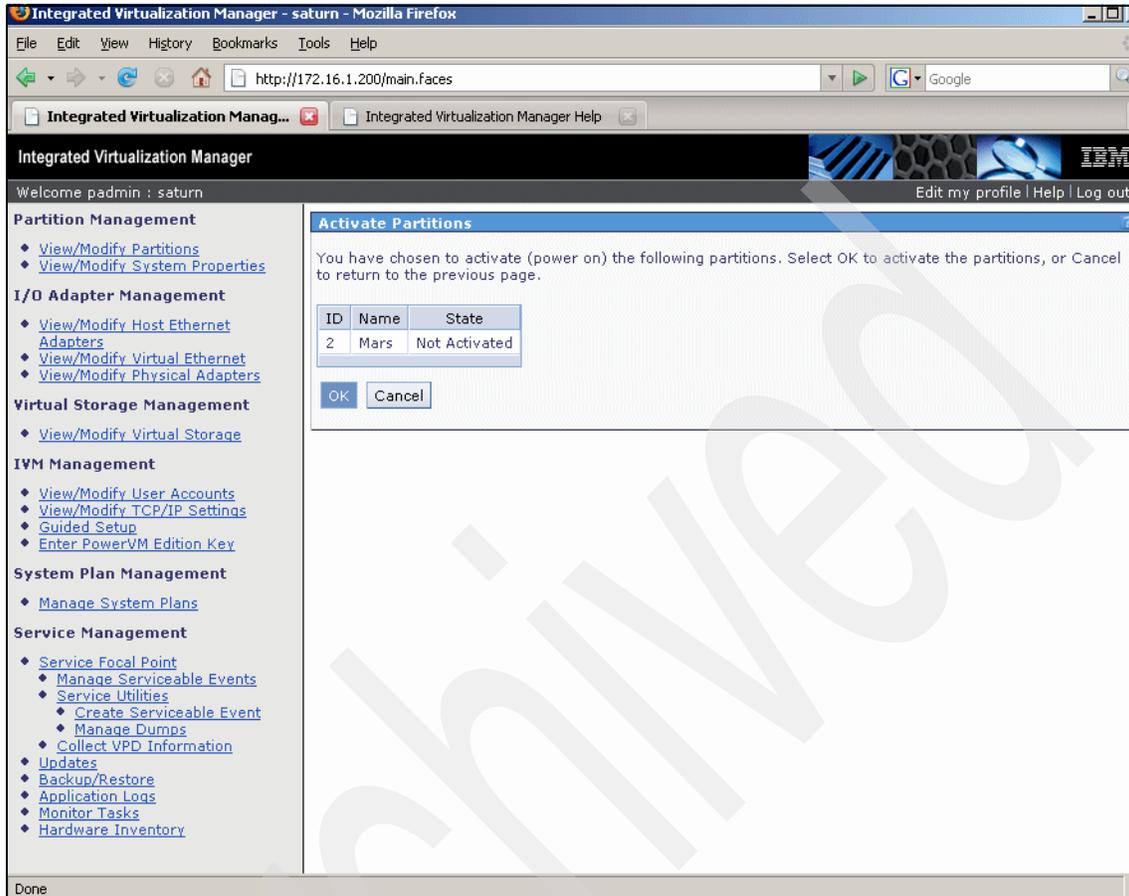


Figure 3-76 Confirm partition activation

When the LPAR activation starts, the message Performing Task - Please Wait will briefly appear, then the IVM UI will return to the View/Modify Partitions window.

Activating from the CLI

The **chsysstate** command is used to start a partition from the command line by either the LPAR number or name. Example 3-15 shows LPAR 2 being activated from the CLI.

Example 3-15 Using the CLI to activate a partition

```
$ chsysstate -o on -r lpar --id 2
```

or

```
$ chsysstate -o on -r lpar -n Mars
```

The **lsrefcode** command can be used to monitor the status codes as the LPAR becomes active. Example 3-16 shows the lsrefcode being used with both LPAR number and name.

Example 3-16 Using the CLI to monitor partition activation status codes

```
$ lsrefcode -r lpar --filter lpar_names=Mars -F refcode
00000000
$ lsrefcode -r lpar --filter lpar_names=Mars -F refcode
CA00E1F1
$ lsrefcode -r lpar --filter lpar_names=Mars -F refcode
AA00E158
```

or

```
$ lsrefcode -r lpar --filter lpar_ids=2 -F refcode
00000000
$ lsrefcode -r lpar --filter lpar_ids=2 -F refcode
CA00E1F1
$ lsrefcode -r lpar --filter lpar_ids=2 -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 3-17.

Example 3-17 The lssyscfg command used to display the LPAR state

```
$ lssyscfg -r lpar -F name,state
VIOS-Saturn,Running
Janus - AIX 6.1,Running
Uranus,Not Activated
Io,Not Activated
```

or

```
$ lssyscfg -r lpar -F lpar_id,state
1,Running
2,Running
3,Not Activated
5,Not Activated
```

3.8.3 Shutdown a partition

The shutdown of a partition can be initiated from the UI 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 options are available for a partition shutdown

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

The Operating System shutdown option is available only if the RMC connection is active. It is the recommended method. The Delayed option is the equivalent of pushing the white control-panel power button. AIX partitions will 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 are recommended to be shut down by using the IBM i session commands `SHTDWNSYS`, `ENDSBS`, or `PWRDWNSYS`.

Shutdown from the UI

From the View/Modify Partitions view, select the check box for the LPAR to be shut down and then click the **Shutdown** button as shown in Figure 3-77 on page 145.

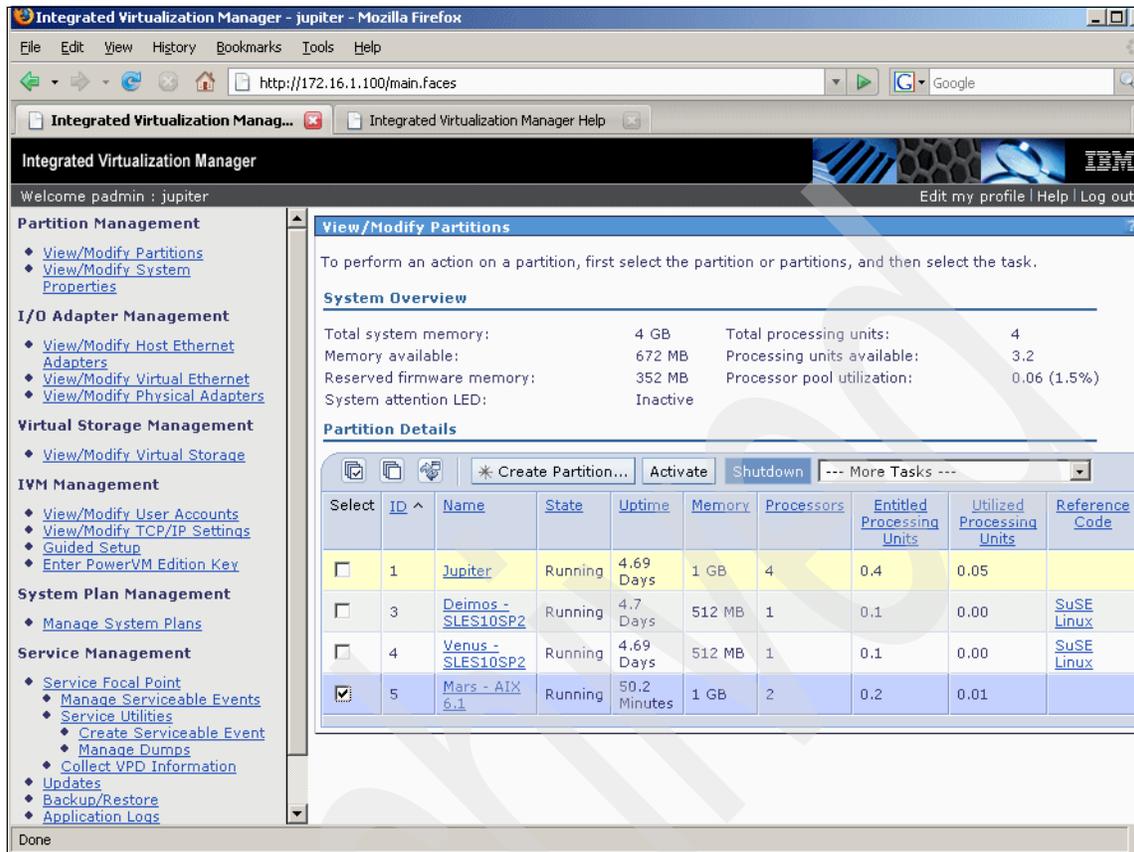


Figure 3-77 Shutdown an LPAR

The Shutdown partitions window, as shown in Figure 3-78 on page 146, will be displayed.

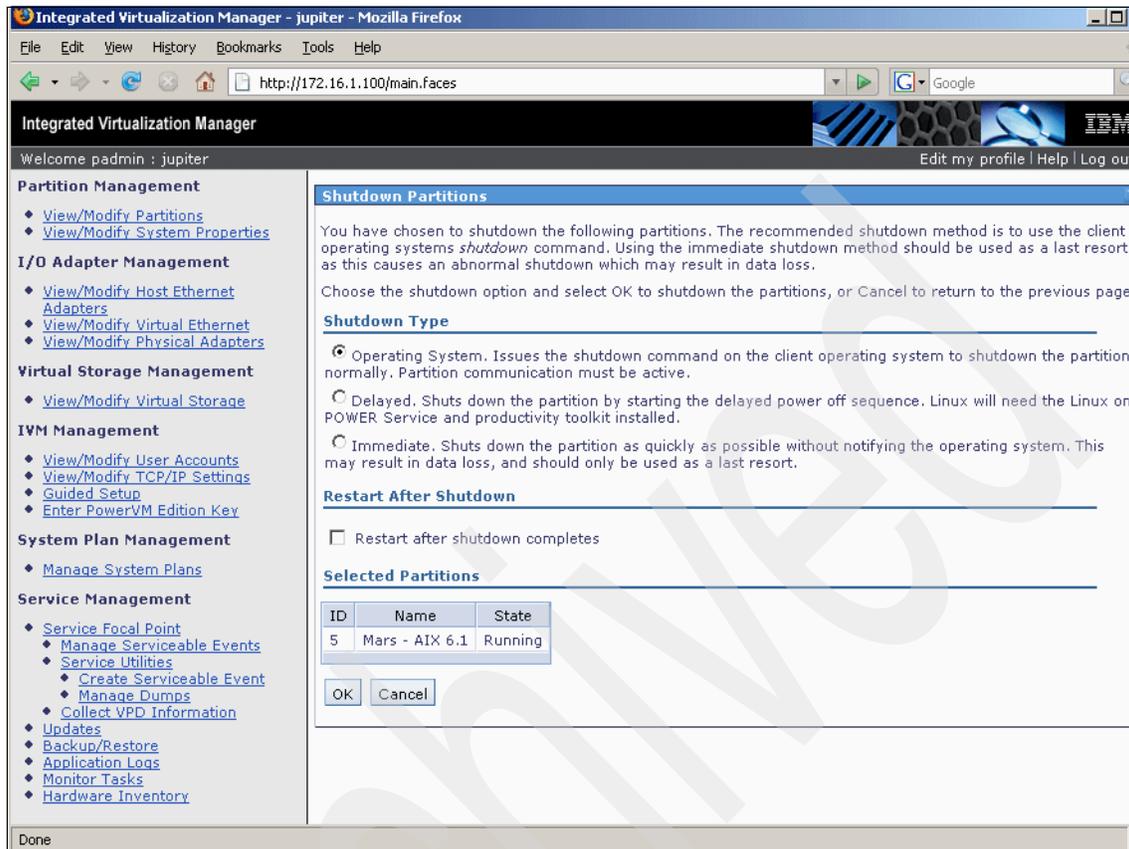


Figure 3-78 Partition shutdown options

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 will be shut down.

Note: The Operating System option will be disabled if RMC is not active between the LPAR and VIOS. The Delayed option will be selected by default.

When selecting the VIOS partition to be shut down, a warning is presented stating 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 3-79 on page 147 shows the VIOS Shutdown Partitions window.

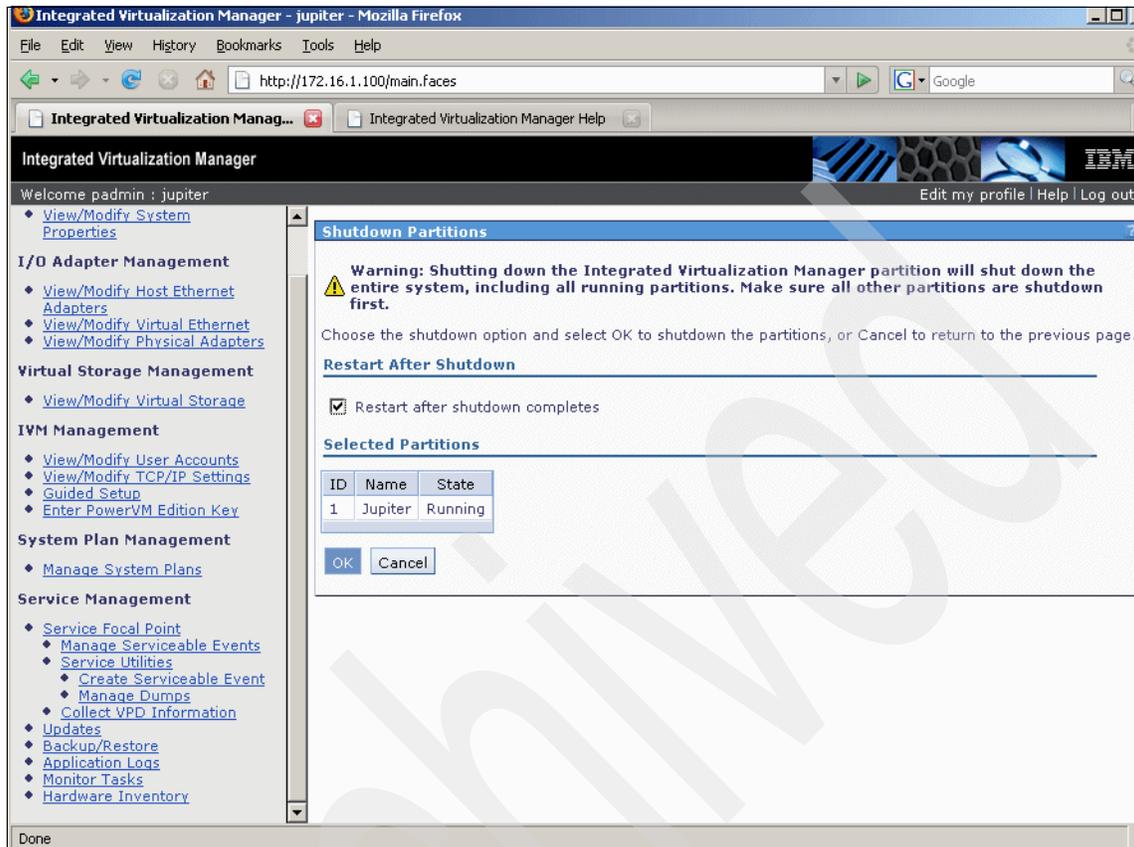


Figure 3-79 VIOS partition shutdown options

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 3-18. This example uses the operating system shutdown option on LPAR 5.

Example 3-18 CLI shutdown of a logical partition

```
$ chsysstate -r lpar -o ossshutdown --id 5
```

or

```
$ chsysstate -r lpar -o ossshutdown -n "Mars - AIX 6.1"
```

The corresponding CLI shutdown options to use with the `-o` flag are:

- ▶ `osshutdown` (Operating System)
- ▶ `shutdown` (Delayed, white button shutdown)
- ▶ `shutdown --immed` (Immediate)



IBM AIX Installation

IBM AIX can be installed native on IBM BladeCenter JS22 or in a client partition of IBM PowerVM. This chapter describes both methods. See 4.1, “Install IBM AIX 6.1 in a PowerVM client partition” on page 150 or 4.2, “Install IBM AIX 6.1 on IBM BladeCenter JS12 or JS22 using a iSCSI target” on page 177.

4.1 Install IBM AIX 6.1 in a PowerVM client partition

This section assumes that you have already installed VIOS 1.5.2.1 or any later version on the blade and performed the initial configuration. In case this was not done, go to Chapter 3, “System configuration using VIOS and IVM” on page 55.

To install IBM AIX 6.1 in a client partition it is necessary to first create the client partition with the IVM before you can start with the installation of AIX. This chapter is divided into the following parts:

- ▶ “Create a virtual media library” on page 150
- ▶ “Prepare the PowerVM client partition” on page 156
- ▶ “Install AIX 6.1 in a logical partition of the Virtual IO Server” on page 169

4.1.1 Create a virtual media library

A virtual media library is a new feature that was introduced in the Virtual IO Server version 1.5. It allows to store 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 install 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. Afterwards we will describe media library created and a DVD image from the AIX DVD.

Perform the following steps to set up 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 4-1.

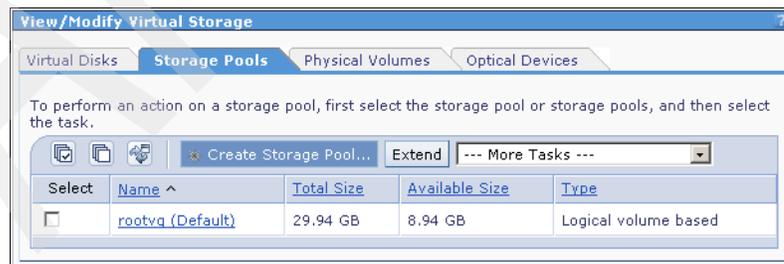


Figure 4-1 Media library - create storage pool

- Specify the name of the storage pool and select the physical volumes that will be assigned to this storage pool. Figure 4-2 shows that we used the name STG-Pool-Media1. The type of the volume group is *logical volume base*. This allows to increase the space of the media library when needed. Physical volume hdsik3 is assigned to this pool. Click **OK**.

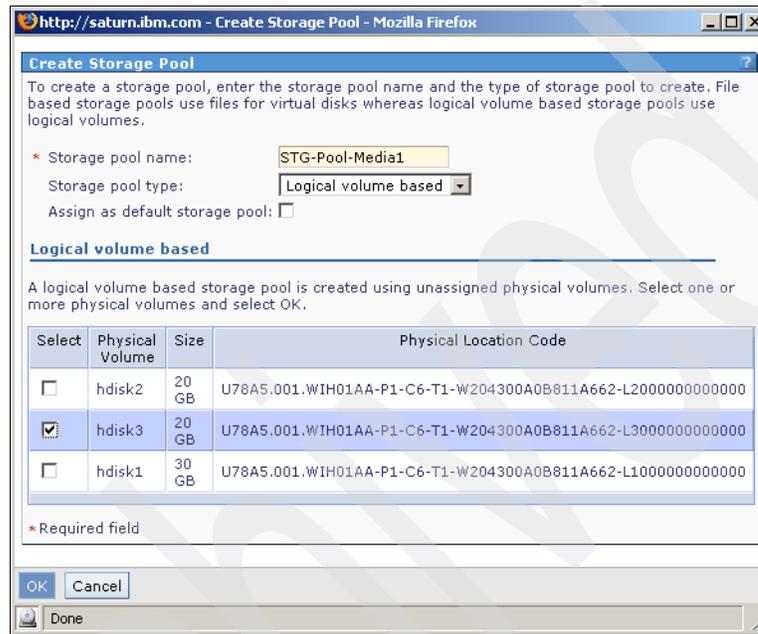


Figure 4-2 Media library - select the physical volume

- The storage pool was created. Now select the **Optical Devices** register card. See Figure 4-3.

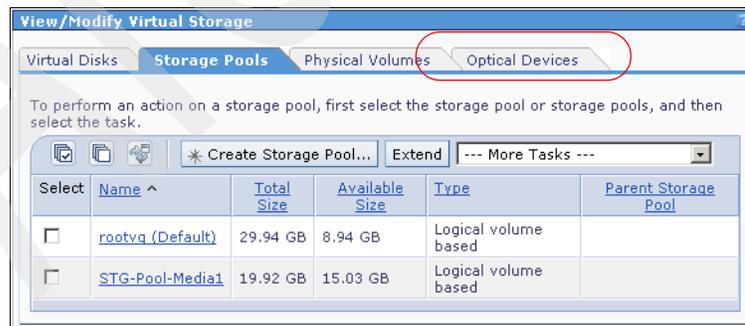


Figure 4-3 Media library - new storage pool

4. Click **Create Library**. See Figure 4-4.

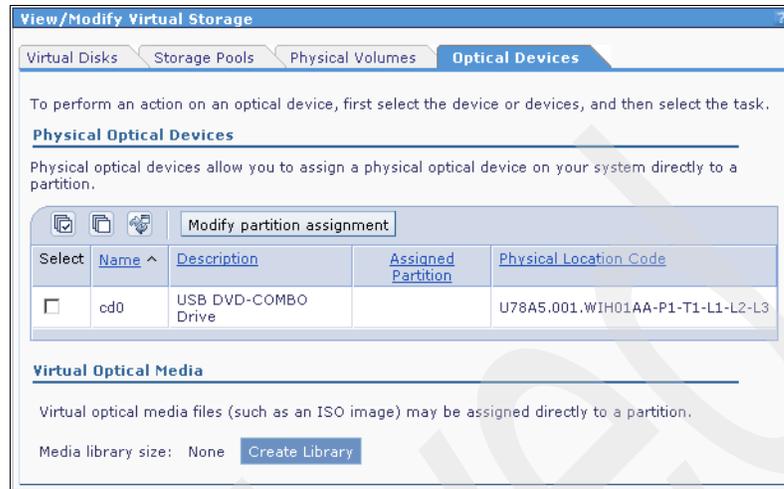


Figure 4-4 Media library - create library

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 150. The initial size was set to hold the AIX 6.1 DVD with a size of approximately 3.6 GB. See Figure 4-5. Click **OK**.

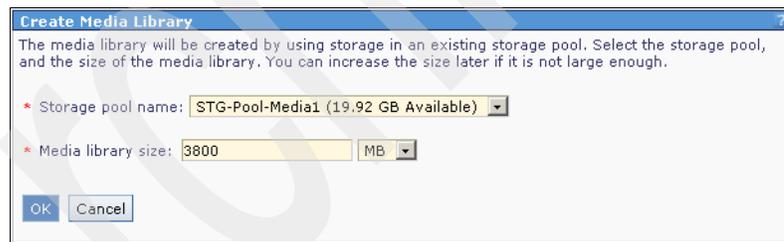


Figure 4-5 Media library - specify library size and volume group

- It takes a moment to create the library volume and file. After that is done, return to the panel shown in Figure 4-6. Click **Add Media** to create an image from the AIX DVD.



Figure 4-6 Media library - add media

- The add media dialog starts and guides you through the process of adding medias 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 will be used as 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 will lead to an error message like the one shown in Figure 4-7.

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



Too many parameters.

Usage: mkvopt -name FileName [-size FileSize | -dev SourceDevice
-file SourceFile] [-ro]

Figure 4-7 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 4-8 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.

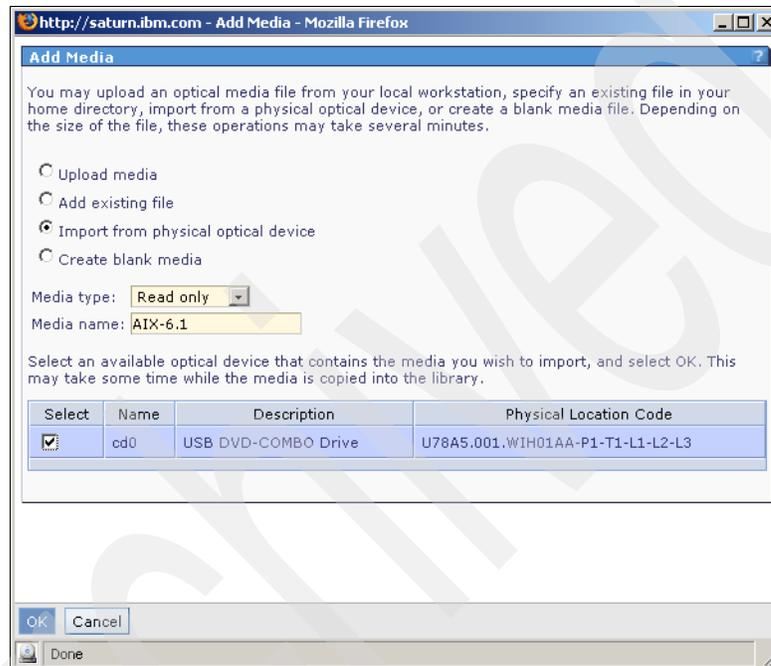


Figure 4-8 Media library - specify details about media creation

8. The copy task will take 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 left-hand navigation under **Service Management** → **Monitor Task**. See Figure 4-9.

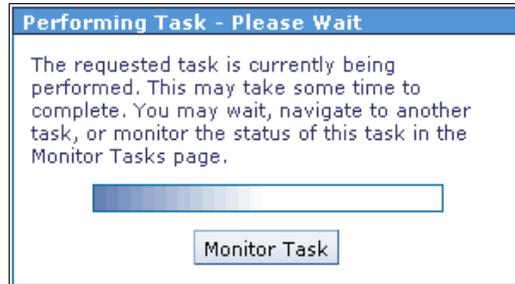


Figure 4-9 Media library - performing task

9. After closing the browser window of the add media dialog, you return to the view shown in Figure 4-10. The new media is already listed here. Clicking **Refresh** updates the size information during the copy operation.

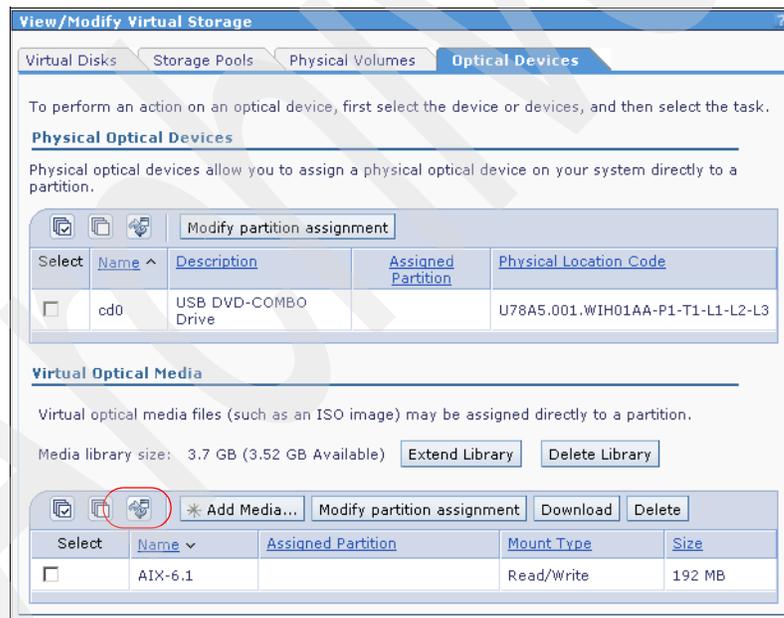


Figure 4-10 Media library - newly created media

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.

4.1.2 Prepare the PowerVM client partition

Perform the following steps 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 3.3.3, “Initial network setup” on page 74.

A Web page comes up that allows you to log in. Use the default account that was created during setup when you had not yet created your own account. The default account is called *padmin* and uses the password *padmin*. Figure 4-11 shows the logon window of the IVM. Click **Login**.

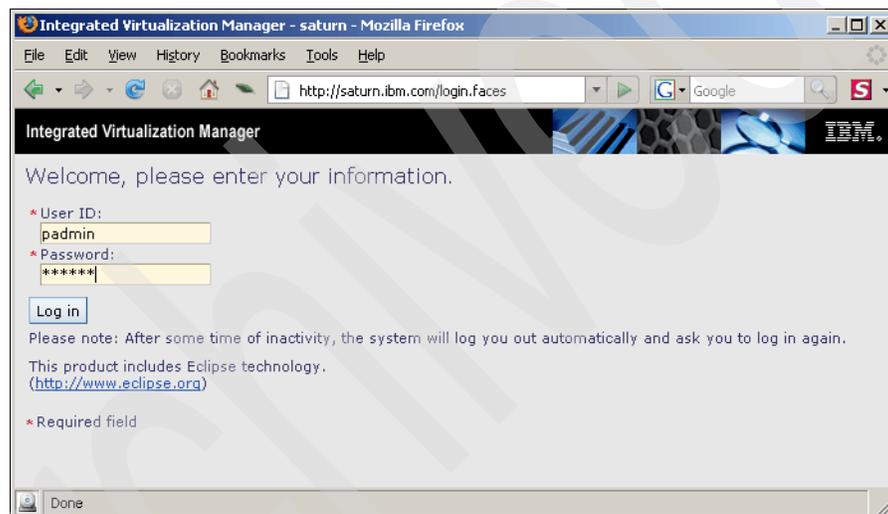


Figure 4-11 IVM login window

- Depending on the setup of your IVM, you will be at the Guided Setup or on the View/Modify Partitions page. Figure 4-12 shows the usual page that you see after logon when the IVM is fully configured.

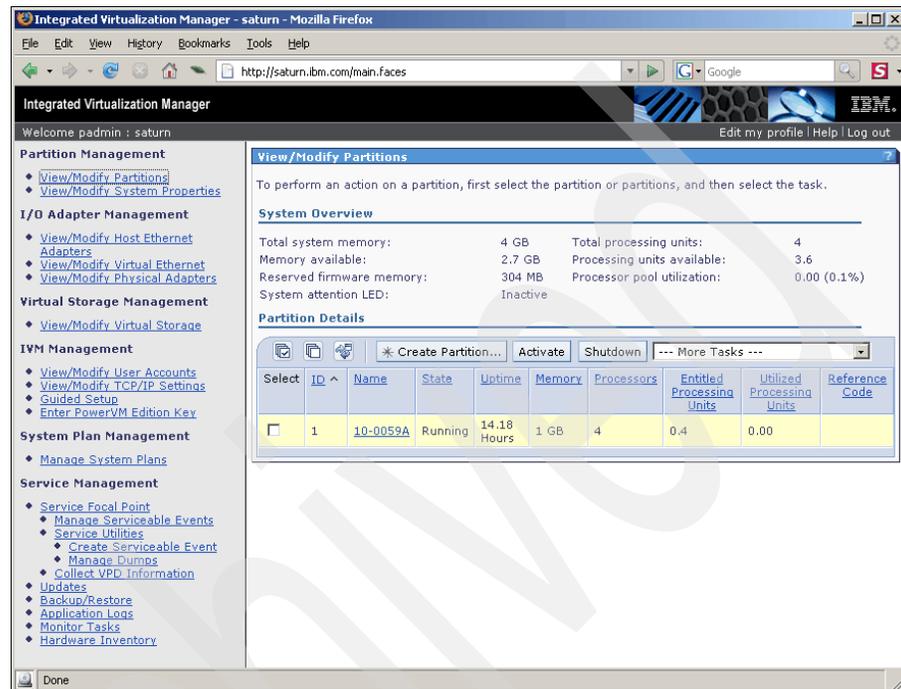


Figure 4-12 View/Modify Partitions page after logon

3. Verify that you have your storage available to the VIOS. Click **View/Modify Virtual Storage** in the left menu under Virtual Storage Management. See Figure 4-13.

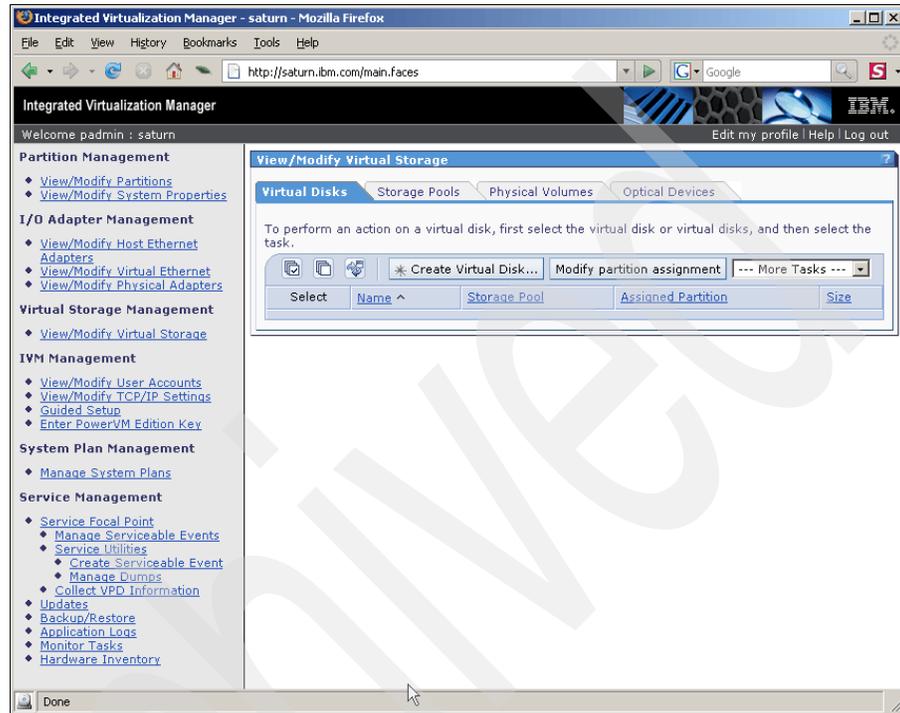


Figure 4-13 View/Modify Storage

- Click the **View/Modify Storage** page on the Physical Volumes tab to see a list of available hard drives to the VIOS. Verify that the expected drives are available. See Figure 4-14.

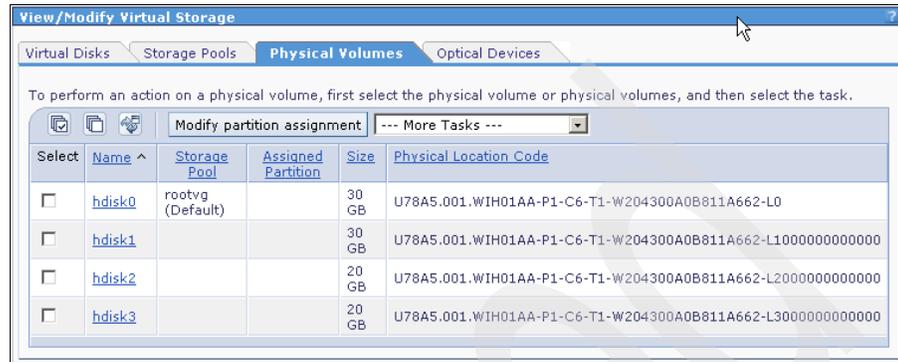


Figure 4-14 Available physical volumes

Figure 4-15 shows that there are four physical volumes available. They are all located on a DS4800. HDISK0 and HDISK1 are used for the VIOS itself. 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 4-15.

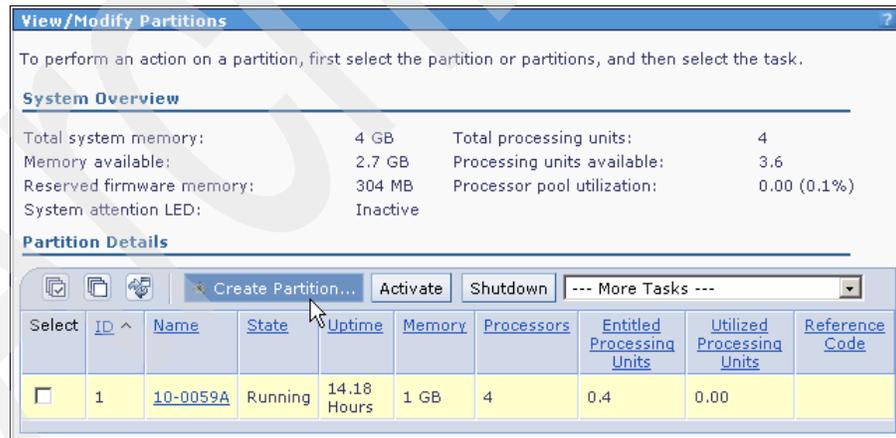


Figure 4-15 View/Modify Partition - Create Partition

A dialog will open that 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 4-16 we chose the host name as partition name. The type can be either AIX/Linux or i5/OS. Choose the type according to the OS you plan to install. We chose AIX/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 4-16 Create partition - define name

6. Define the amount of memory that will be assigned to the partition. In Figure 4-17 we chose 1 GB. Click **Next** to proceed.

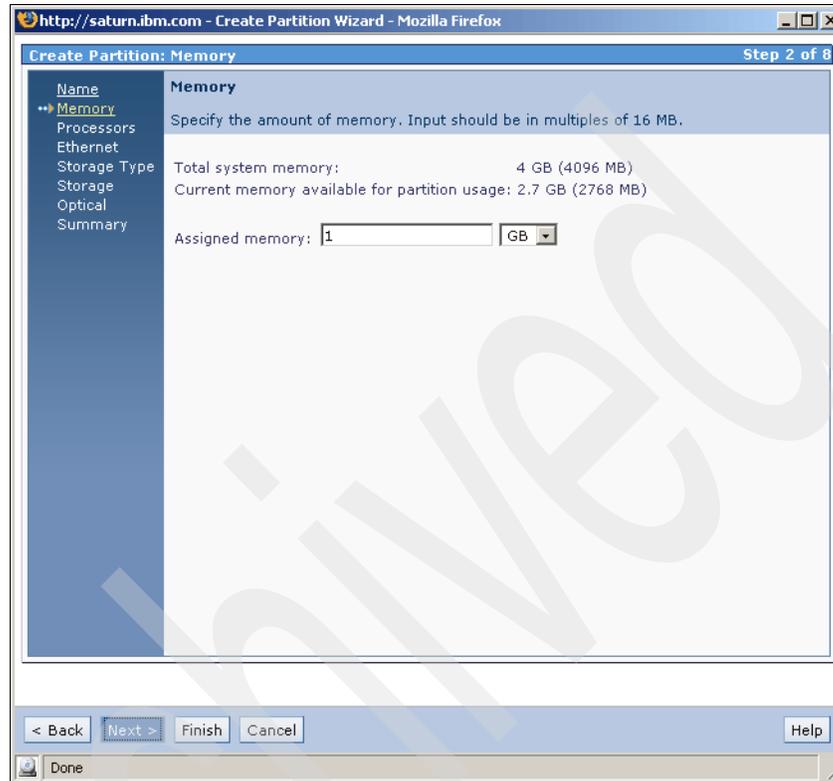


Figure 4-17 Create partition - configure the amount of memory

7. Choose the number of CPUs that will be used by the partition. You have to decide whether to use dedicated or shared CPUs. When a dedicated CPU is used, no load can be moved to other currently free CPUs because this may lead to a performance issue. In Figure 4-18 you see that we configured two CPUs and shared processor mode. Click **Next** to proceed.

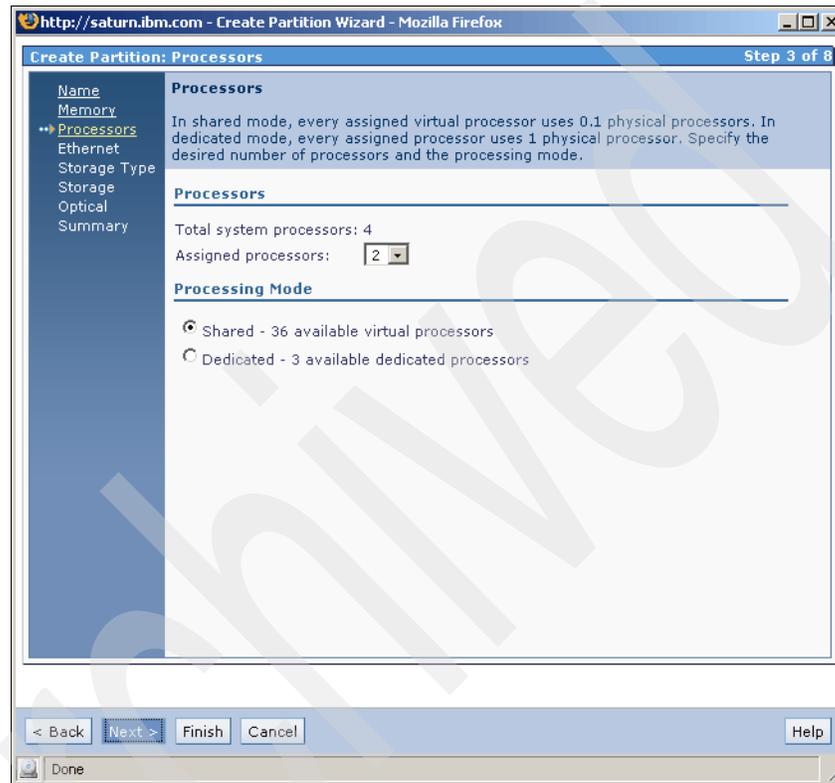


Figure 4-18 Create partition - CPU configuration

8. Depending on the setup of the network in the BladeCenter chassis and in the Virtual IO Server, you may have different settings. When you define host Ethernet adapters to the partition then 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 it is a requirement that the partition be configured with virtual Ethernet adapters. The Ethernet bridge is configured in the Virtual IO Server. For more details about this see 3.5.2, “Virtual Ethernet Adapters and SEA” on page 89.

As shown in Figure 4-19, we chose one virtual Ethernet adapter. Click **Next** to proceed.

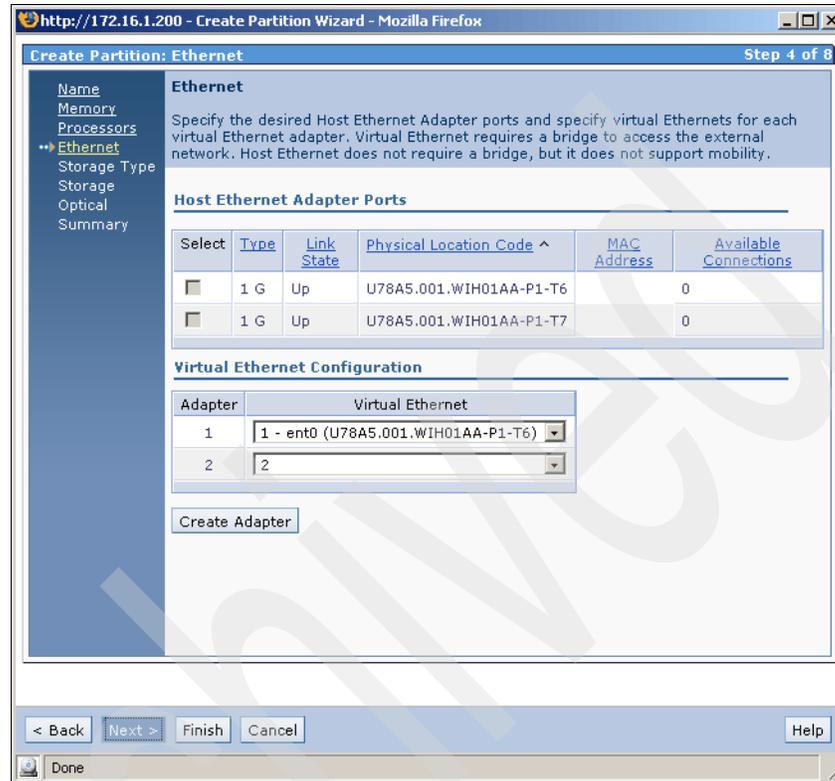


Figure 4-19 Create partition - ethernet network

9. Set up the storage type you plan to use. There are three different options available. You may use volume group or file-based storage. In addition there is an option to use a dedicated physical volume for the partition. Depending on the type of storage subsystem, there are limitations about the maximum number of physical volumes you may have per host adapter or host adapter group. In 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 4-20, select **Assign existing virtual disks and physical volumes**. Click **Next** to proceed.

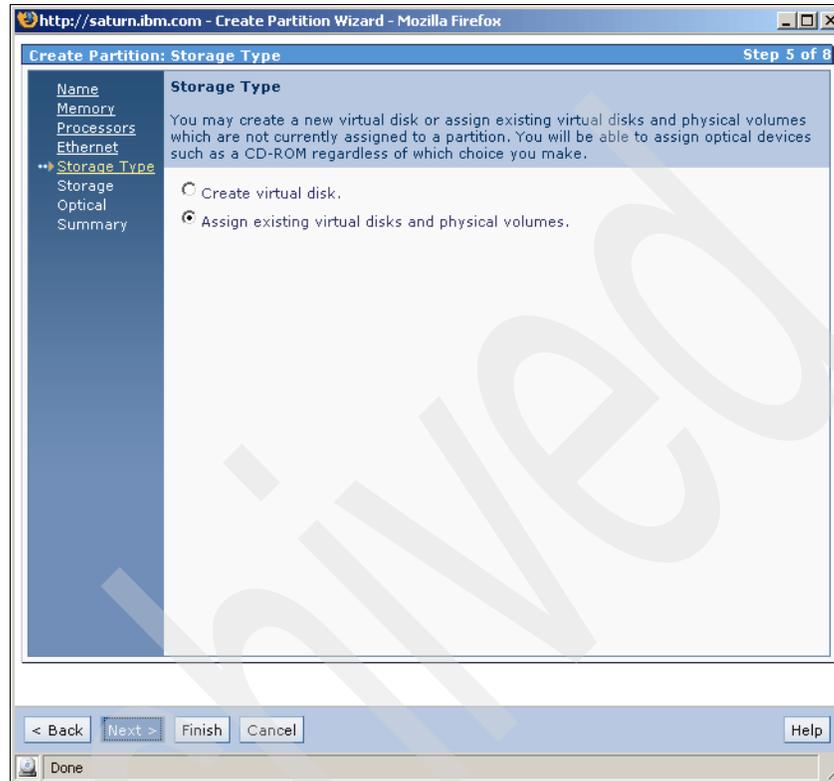


Figure 4-20 Create partition - storage type

10. Select the physical volume or volumes that need to be available to the partition. Figure 4-21 shows the section of hdisk1. Click **Next** to proceed.

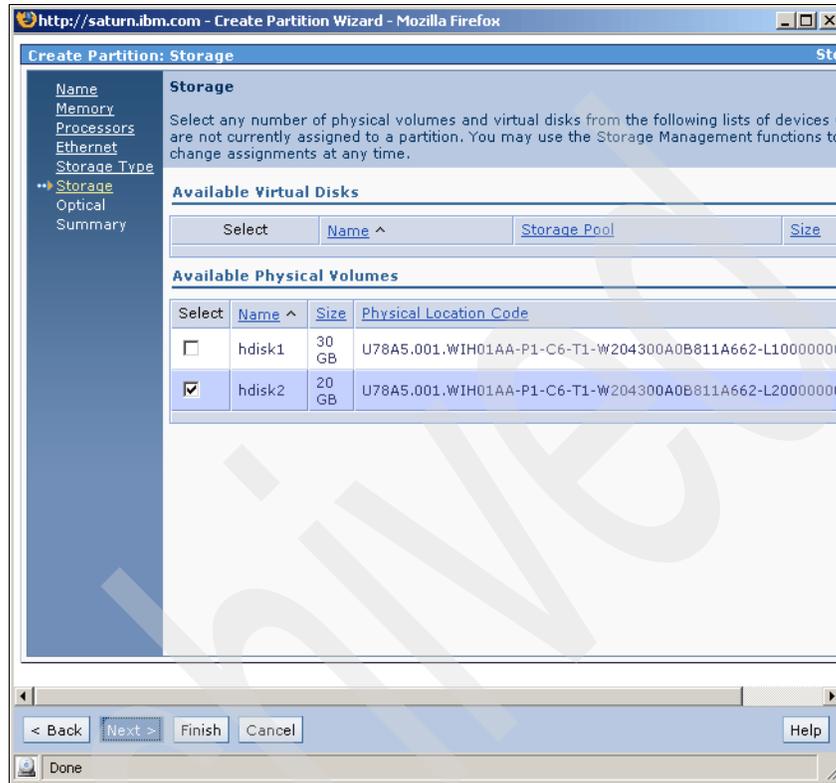


Figure 4-21 Create partition - select physical volumes

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

There might be multiple physical CDROM drives available. Use the location code to differentiate between the CDROM drives:

- U78A5.001.WIH01AA-P1-T1-L1-L2-L3 - CDROM drive in the media tray
- U78A5.001.WIH01AA-P1-T1-L1-L1 - Remote media CDROM 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 CDROM drive that has the location code U78A5.001.WIH01AA-P1-T1-L1-L1

is the CDROM drive that is provided via 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 may have to issue `cfgdev` on the command line of the Virtual IO Server to get it recognized by VIOS.

Virtual CDROM drives are used to mount CDs that are placed in the media library. See 3.6.2, “Storage pools” on page 101 and 3.6.4, “Optical devices” on page 108.

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 4-22.



Figure 4-22 Create partition - modify virtual optical device

12. Change the selected media from *none* to AIX-6.1 and click **OK**. See Figure 4-23.

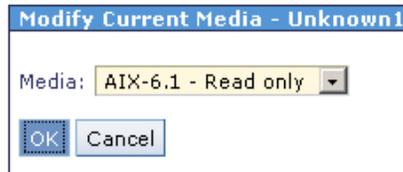


Figure 4-23 Create partition - modify current media of virtual optical device

13. Click **Next** to see an overview of the setting of the new partition. See Figure 4-24.

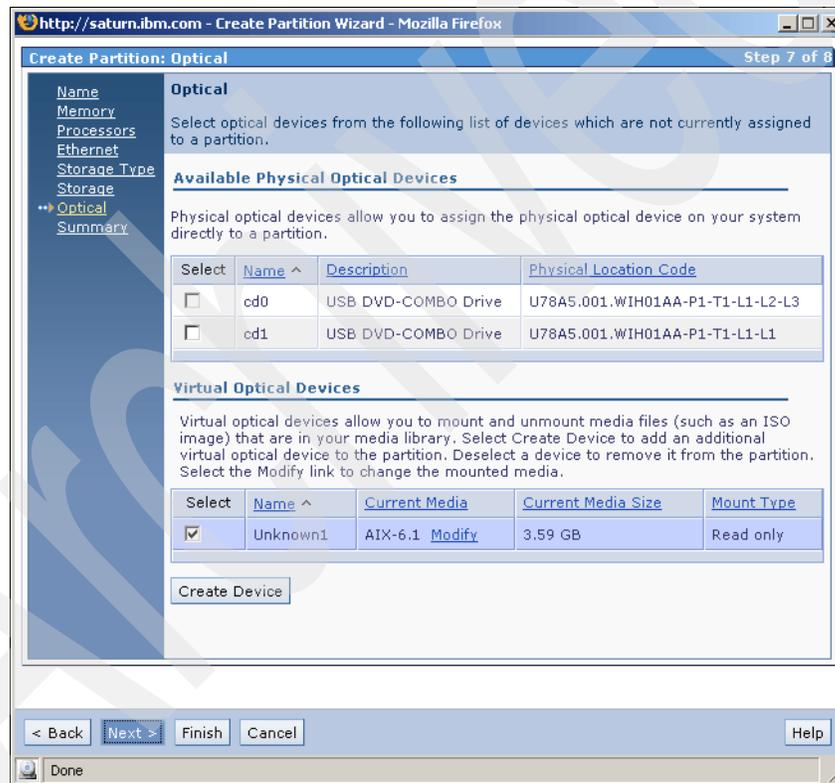


Figure 4-24 Create partition - virtual optical device

14. Verify your setting and click **Finish** to create a partition with the settings you defined. See Figure 4-25.

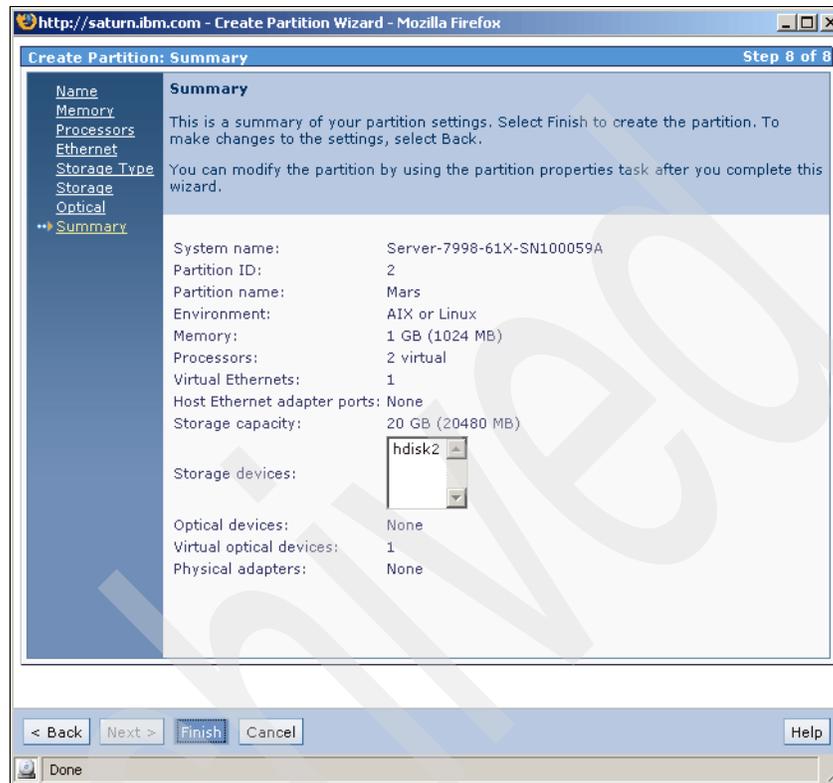


Figure 4-25 Create partition - summary

15. The new partition will be listed under View/Modify Partitions, as shown in Figure 4-26.

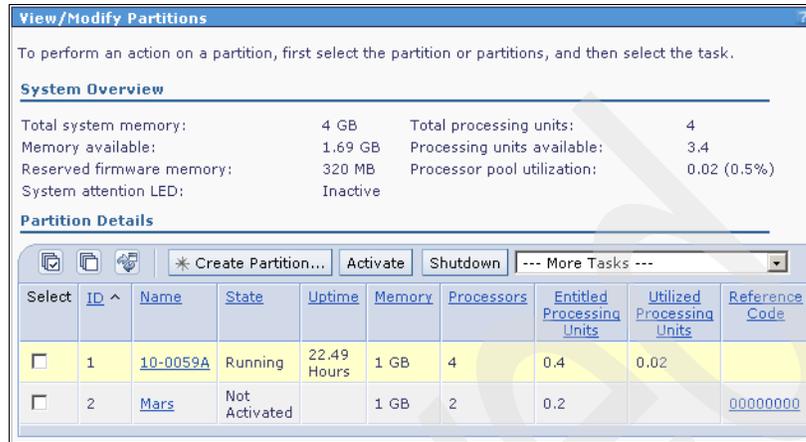


Figure 4-26 Newly created AIX/Linux partition

The preparation of the partition is done. Proceed with the installation of AIX in the newly created partition.

4.1.3 Install AIX 6.1 in a logical partition of the Virtual IO Server

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. Follow the outlined steps:

1. To activate the partition, click the check box of the partition and click **Activate**. See Figure 4-27.

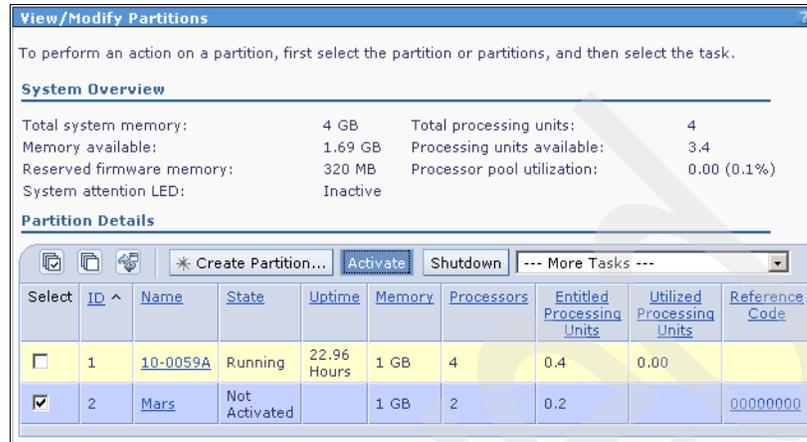


Figure 4-27 Activate a partition

2. Confirm the activation of the partition with **OK** as shown in Figure 4-28.

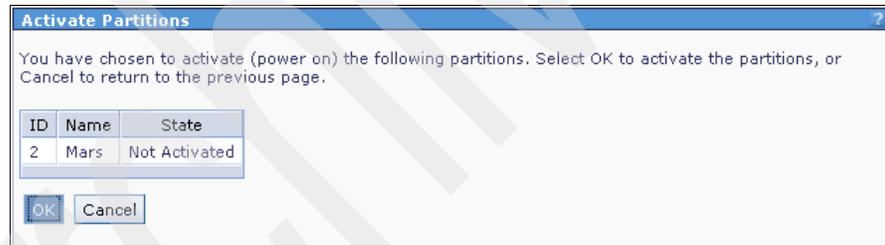


Figure 4-28 Confirm partition activation

- The status of the partition has changed to *running*. Select **Open Terminal** from the More Tasks drop-down list box to open a terminal connected to the selected partition. See Figure 4-29.

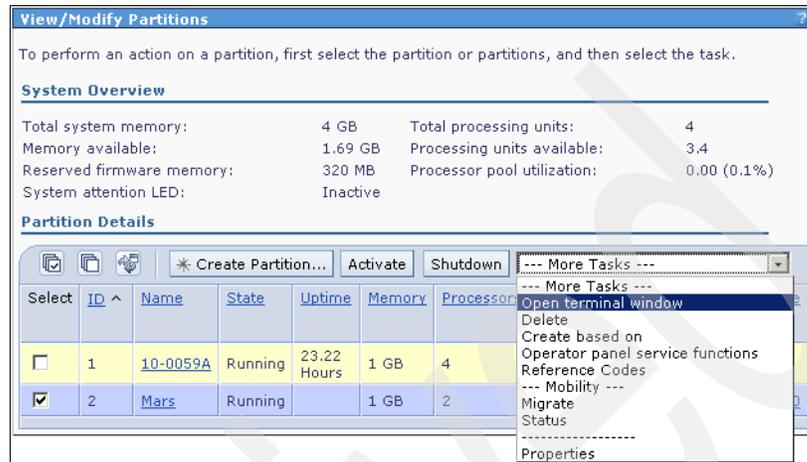


Figure 4-29 Open a virtual terminal to the partition

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 in case 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 4-30.

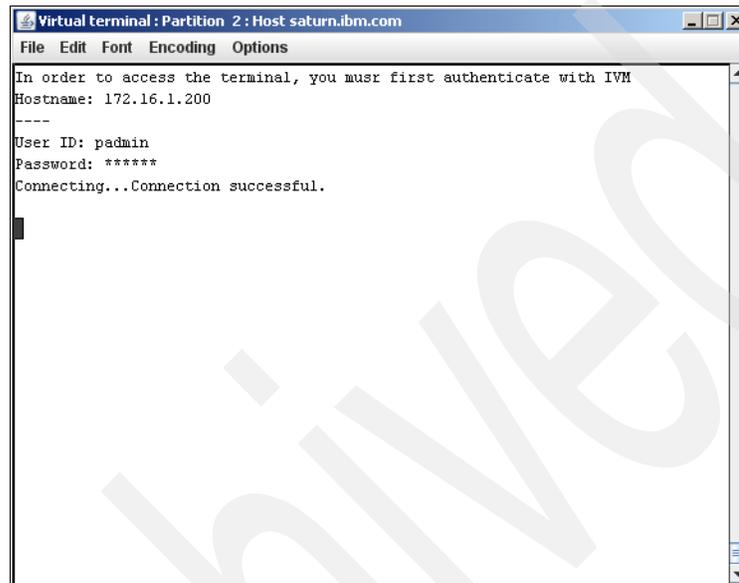


Figure 4-30 Virtual terminal connection

On the virtual terminal you will see the POST of the partition with the possibility to enter the SMS menu. There is no change 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 a few messages, as shown in Example 4-1.

Example 4-1 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-scsi@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 4-2. Depending on the console you are using, you may need to also enter **F1** or **2**.

Example 4-2 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 1 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 would like to use during the installation of IBM AIX 6.1. You can define the language of the operating system that now gets installed in a later step. Click Enter to proceed; see Example 4-3.

Example 4-3 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 proceed with the installation by entering **1** followed by Enter, as shown in Example 4-4.

Example 4-4 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]:

9. The copy process starts after you click **1** followed by the Enter key.
10. After the installation is done, a reboot of the partition is performed. Then select your terminal type, as shown in Example 4-5.

Example 4-5 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 []:
```

-
11. Select **Show Installed License Agreements** and click Enter to read the license agreement; see Example 4-6.

Example 4-6 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
```

12. Select the software package from which you would like to read the license agreements. The default is to show all license agreements. Click Enter to start showing the license text; see Example 4-7.

Example 4-7 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
```

13. Navigate through the licenses. When you have finished reading, click F3 twice. You are returning to the Software License Agreements panel. Select Accept License Agreements and click Enter; see Example 4-8.

Example 4-8 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	

14. Click Tab followed by Enter to accept the license and change the *no* to a *yes*, as shown in Example 4-9.

Example 4-9 Accept license agreements

Accept License Agreements

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

			[Entry Fields]	
ACCEPT	Installed License Agreements		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	

15. After the status of the command has changed to OK, click F10 or ESC+0 to leave the license agreement. The keys you need depend on the terminal you are using; see Example 4-10.

Example 4-10 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

16. The installation assistant will guide 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 get to a login prompt, use ESC+0 or F10. You may start this installation assistant at any time again by using the command `install_assist` after login as `root`. The installation assistant is shown in Example 4-11.

Example 4-11 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
Esc+9=Shell

F2=Refresh
Esc+0=Exit

F3=Cancel
Enter=Do

Esc+8=Image

The installation of AIX 6.1 in the logical partition has completed. You may now start with configuring your AIX or installing your applications.

4.2 Install IBM AIX 6.1 on IBM BladeCenter JS12 or JS22 using a iSCSI target

A native installation of IBM AIX 6.1 is not different from an installation into a logical partition of the Virtual IO Server. Instead of using virtual resources, physical resources are required.

A difference is seen when iSCSI is used as boot device. The iSCSI host bus adapter that is supported in an IBM BladeCenter JS 12 or JS22 is a Standard Form Factor I/O expansion card that requires the removal of the hard drive and the hard drive tray next to the backplane connectors of the blade. In the JS22 there can only be one SAS hard drive installed. Thus it is necessary to boot from a network resource via the iSCSI host bus adapter when this adapter is installed in the JS22. JS12 allows the installation of two SAS hard drives where one is still available when the iSCSI host bus adapter is installed.

An iSCSI storage subsystem is required to boot JS12 or JS22 over iSCSI. Total Storage N series models currently have limited support for supporting an iSCSI AIX boot. Linux on Power is supported by the Total Storage DS3300, but only with a software initiator.

An alternative to iSCSI storage subsystems are the iSCSI software targets. iSCSI software targets are applications that share a file, logical volume, or disks as physical disk over iSCSI to hosts using iSCSI initiators to access those disks. Products are available that provide iSCSI software target functionality. IBM AIX and SuSE Linux Enterprise Server 10 also contain a iSCSI software target implementation.

This section explains the installation of AIX V6.1 on a iSCSI target that is provided by the iSCSI Software Target, included in AIX V6.1. Tests with the software target included with SuSE Linux Enterprise Server 10 Service Pack 1 have shown that this target cannot be used to boot from with the iSCSI HBA in JS12 or JS22.

4.2.1 Prepare the iSCSI target

The iSCSI Software target is not included in a standard installation. It must be installed from the AIX Expansion Pack DVD and then configured.

Install the iSCSI Software Target file set

The iSCSI software target driver code for AIX can be found on the AIX Expansion Pack DVD. To install the target driver, place this DVD in the DVD drive and run **smitty install_latest**. Select the DVDROM drive that contains the AIX Expansion Pack DVD. Typically this is /dev/cd0. A list of available drives is shown using F4, shown in Example 4-12.

Example 4-12 Select the input device for the installation

Install Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

* INPUT device / directory for software [Entry Fields] +

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	

Select the **Software to install** field and click F4, as shown in Example 4-13. A list of available file sets is shown. Not all file sets are required.

Example 4-13 Install software - choose file sets

Install Software

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

```

                                     [Entry Fields]
* INPUT device / directory for software      /dev/cd0
* SOFTWARE to install                        [_all_latest]      +
  PREVIEW only? (install operation will NOT occur)  no      +
  COMMIT software updates?                     yes      +
  SAVE replaced files?                         no      +
  AUTOMATICALLY install requisite software?      yes      +
  EXTEND file systems if space needed?          yes      +
  OVERWRITE same or newer versions?            no      +
  VERIFY install and check file sizes?         no      +
  Include corresponding LANGUAGE filesets?      yes      +
  DETAILED output?                             no      +
  Process multiple volumes?                   yes      +
  ACCEPT new license agreements?              no      +
  Preview new LICENSE agreements?             no      +

```

```

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

```

In the list of all file sets on the installation medium, search for the iSCSI software target file set by clicking /. This will open a dialog to enter a search pattern, as shown in Example 4-14.

Example 4-14 List of all file sets on the installation medium

Install Software

```

Ty+-----+
Pr|                SOFTWARE to install
  |
  | Move cursor to desired item and press F7. Use arrow keys to scroll.
  | * ONE OR MORE items can be selected.
  | * Press Enter AFTER making all selections.
  |
  | [TOP]
  |
  | Java14.ext                                ALL
  | + 1.4.2.0 Java SDK 32-bit Comm API Extension
  | + 1.4.2.0 Java SDK 32-bit Java3D
  |-----+

```



```

Move cursor to desired item and press F7. Use arrow keys to scroll.
* ONE OR MORE items can be selected.
* Press Enter AFTER making all selections.

[MORE...78]
des ALL
+ 5.3.0.0 Data Encryption Standard Library Routines

devices.tmiscsw ALL
+ 6.1.0.0 iSCSI Target Software Device Driver

gskak ALL
+ 6.0.5.41 AIX Certificate and SSL Base Runtime ACME Toolkit
[MORE...135]

Esc+1=Help      Esc+2=Refresh      Esc+3=Cancel
Es| F7=Select     F8=Image          F10=Exit
Es| Enter=Do     /=Find            n=Find Next
F9+-----+

```

The file set `devices.tmiscsw` is listed in the Software to install field. Click Enter; see Example 4-17.

Example 4-17 Install software with specified file set to be installed

Install Software

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

```

[Entry Fields]
* INPUT device / directory for software /dev/cd0
* SOFTWARE to install [devices.tmiscsw > +
PREVIEW only? (install operation will NOT occur) no +
COMMIT software updates? yes +
SAVE replaced files? no +
AUTOMATICALLY install requisite software? yes +
EXTEND file systems if space needed? yes +
OVERWRITE same or newer versions? no +
VERIFY install and check file sizes? no +
Include corresponding LANGUAGE filesets? yes +
DETAILED output? no +
Process multiple volumes? yes +
ACCEPT new license agreements? no +
Preview new LICENSE agreements? no +

```

```

Esc+1=Help      Esc+2=Refresh      Esc+3=Cancel      Esc+4=List
Esc+5=Reset     F6=Command         F7=Edit           F8=Image

```

Confirm that you are sure about the installation with by clicking Enter; see Example 4-18.

Example 4-18 Confirm the installation of the selected file set

Install Software

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

```

                                [Entry Fields]
* INPUT device / directory for software /dev/cd0
* SOFTWARE to install [devices.tmiscsw > +
  PREVIEW only? (install operation will NOT occur) no +
  COMMIT software updates? yes +
  SAVE replaced files? no +
  AUTOMATICALLY install requisite software? yes +
+-----+
|                                     |
|                ARE YOU SURE?      |
|                                     |
|  Continuing may delete information |
|  you may want to keep. This is   |
|  your last chance to stop         |
|  before continuing.               |
|  Press Enter to continue.         |
|  Press Cancel to return to the    |
|  application.                     |
|                                     |
|  Esc+1=Help      Esc+2=Refresh     |
|  Esc+3=Cancel   |                                     |
|  F8=Image       | F10=Exit         |
|                                     |
+-----+

```

Verify the output of the install command and leave smitty with F10 when the installation was successful; see Example 4-19.

Example 4-19 Successful installed file set devices.tmiscsw

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

[MORE...63]

Installation Summary

Name	Level	Part	Event	Result
devices.tmiscsw.rte	6.1.0.0	USR	APPLY	SUCCESS

```
devices.tmiscsw.rte      6.1.0.0      ROOT      APPLY      SUCCESS
```

File /etc/inittab has been modified.

One or more of the files listed in /etc/check_config.files have changed.
See /var/adm/ras/config.diff for details.

[BOTTOM]

```
Esc+1=Help      Esc+2=Refresh      Esc+3=Cancel      F6=Command  
F8=Image        F9=Shell           F10=Exit          /=Find  
n=Find Next
```

The iSCSI software target is now installed and needs to be configured.

Configure the iSCSI software target

After the file set for the iSCSI software target has been installed, proceed with the configuration of the iSCSI software target. It is assumed that you have a working network configuration and that your host running the iSCSI software target is accessible over the network.

Run smitty with the fast path tmiscsi using the command **smitty tmiscsi**. Select **iSCSI Target Protocol Device** and click Enter; see Example 4-20.

Example 4-20 iSCSI target device menu

iSCSI Target Device

Move cursor to desired item and press Enter.

```
iSCSI Target Protocol Device  
iSCSI Targets  
iSCSI Target LUNs
```

```
F1=Help          F2=Refresh        F3=Cancel         F8=Image  
F9=Shell         F10=Exit          Enter=Do
```

Select **Add an iSCSI Target Protocol Device** and click Enter (Example 4-21).

Example 4-21 iSCSI target protocol device

iSCSI Target Protocol Device

Move cursor to desired item and press Enter.

```
List All iSCSI Target Protocol Devices  
Add an iSCSI Target Protocol Device  
Change / Show Characteristics of an iSCSI Target Protocol Device
```

Remove an iSCSI Target Protocol Device

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

Create a single instance of the iSCSI target driver. This instance acts as a container for the other iSCSI objects. Click Enter to create the instance (Example 4-22). You can configure only one instance.

Example 4-22 iSCSI target protocol instance

Add an iSCSI Target Protocol Device

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

```

                                     [Entry Fields]
iSCSI Target Protocol Device Name    tmsw0
Maximum number of sessions           [200]                                +#
Maximum transfer size in one PDU     [262144]                              +#
Access Control enabled                [off]                                  +
Apply change to DATABASE only        no                                     +
```

F1=Help F2=Refresh F3=Cancel F4=List
F5=Reset F6=Command F7=Edit F8=Image
F9=Shell F10=Exit Enter=Do

Verify that the iSCSI target driver instance creation was successful. Click F3 three times to return to the iSCSI target main menu (Example 4-23).

Example 4-23 Successful creation of an iSCSI target protocol instance

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

tmsw0 Available

Esc+1=Help Esc+2=Refresh Esc+3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next

Select **iSCSI Targets** and click Enter (Example 4-24).

Example 4-24 iSCSI target device menu

iSCSI Target Device

Move cursor to desired item and press Enter.

iSCSI Target Protocol Device
iSCSI Targets
iSCSI Target LUNs

Esc+1=Help	Esc+2=Refresh	Esc+3=Cancel	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Select **Add an iSCSI Target** and click Enter (Example 4-25). An iSCSI target should be configured for each host that will access LUNs over iSCSI.

Example 4-25 Add an iSCSI target

iSCSI Targets

Move cursor to desired item and press Enter.

List All iSCSI Targets
Add an iSCSI Target
Change / Show Characteristics of an iSCSI Target
Remove an iSCSI Target

Esc+1=Help	Esc+2=Refresh	Esc+3=Cancel	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Create one iSCSI target device for each virtual iSCSI target that is allocated by the iSCSI target driver; see Example 4-26. The target followed by a number will be used as logical name for the SCSI targets when no other name is specified. The iSCSI target name is the iSCSI qualified name of the iSCSI target. A host will use this IQN to access the LUNs available on this target. The example uses `iqn.2008-07.com.ibm:storage.sys1`. This value will be used later when the iSCSI connection during the AIX or Virtual I/O Server installation is configured. Click Enter to create the target.

Example 4-26 iSCSI target setting

Add an iSCSI Target

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

[Entry Fields]

```

iSCSI Target Protocol Device Name      tmsw0
iSCSI Target Logical Name              []
* iSCSI Target Name                    [iqn.2008-07.com.ibm:s>
Maximum Target LUNs allowed            [512]          +#
Paths                                  []
Force disconnection of iSCSI clients on remove [no]          +
Initiator Reconnection delay           []            +#
Addressing method                       []            +#
Apply change to DATABASE only          no            +

```

```

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

```

Verify that the iSCSI target was successfully created. Click F3 three times to return to the iSCSI Target Device menu; see Example 4-27.

Example 4-27 Successfully created iSCSI target target0

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

target0 Available

```

Esc+1=Help      Esc+2=Refresh      Esc+3=Cancel      F6=Command
F8=Image        F9=Shell          F10=Exit          /=Find
n=Find Next

```

Select **iSCSI target LUNs** and click Enter (Example 4-28).

Example 4-28 iSCSI target LUN's

iSCSI Target Device

Move cursor to desired item and press Enter.

```

iSCSI Target Protocol Device
iSCSI Targets
iSCSI Target LUNs

```

```

Esc+1=Help      Esc+2=Refresh      Esc+3=Cancel      F8=Image
F9=Shell        F10=Exit          Enter=Do

```

Choose the Backing Device Type disk or logical volume. Then select the disk or logical volume that will be accessible via the iSCSI software target. Click Enter to create this new LUN of the iSCSI target; see Example 4-31. The default logical unit name is lu followed by consecutive numbers. It is advisable to choose a name that also reflects the target to which the LUN will be assigned. The order of the LUNs available over the iSCSI target is the order in which they were assigned to the target. Click Enter when the disk space has been selected to create the new LUN of the selected iSCSI target.

Example 4-31 Create new LUN on an iSCSI target

Add an iSCSI Target Logical Unit

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

			[Entry Fields]
iSCSI Parent Target Name			target0
Logical Unit Name			[]
* Backing Device Type			[disk] +
* Backing Device Name			[hdisk1] +
Maximum Number of Commands to Queue to Device			[3] +#
Apply change to DATABASE only			no +

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	

Verify that the disk or logical volume was defined as an iSCSI target LUN. Click F10 to leave smitty (Example 4-32).

Example 4-32 Successfully created iSCSI target LUN

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

lu0 Available

F1=Help	F2=Refresh	F3=Cancel	F6=Command
F8=Image	F9=Shell	F10=Exit	/=Find
n=Find Next			

The setup of the iSCSI target is done. Multiple iSCSI targets with multiple LUNs can be defined.

4.2.2 Perform the installation of AIX V6.1 from DVD onto a iSCSI target

This method may not be officially supported going forward. It is included here for example purposes only.

1. You can perform an iSCSI boot with the integrated network adapter or with an optional QLogic iSCSI HBA when you use IBM AIX V6.1. The Virtual I/O server does not yet support the iSCSI boot over regular network cards. So there must be an iSCSI host bus adapter present in the blade.
2. The iSCSI qualified name (IQN) needs to be defined during the installation together with all network settings that are required. You need the following configuration values:
 - IP Address of the iSCSI host bus adapter port or Ethernet network interface
 - Subnet mask
 - Gateway that might be required to access the target
 - iSCSI qualified name of the host to access the target
 - iSCSI qualified name of the iSCSI target
 - IP Address of the iSCSI target

Currently there is no support for authentication, such as CHAP.

3. Use a console as discussed in “Consoles of the IBM BladeCenter JS12 and JS22” on page 392 and power on the blade. Configure the blade to boot from the AIX CD or DVD. You may choose remote media or the CD or DVD drive in the media tray of the BladeCenter chassis.
4. Boot the blade from the AIX installation media. Example 4-33 shows the console over SOL. When you use the graphical console or Remote Control you need to press F1 to select this console as system console. To use the SOL console as system console, click 2.

Note: The console can be changed after the installation is finished.

Example 4-33 AIX system console selection

Preserving 22646 bytes of symbol table [/usr/lib/drivers/lft_loadpin]

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

Type a 2 and press Enter to use this terminal as the system console.

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

5. Select the language that is used during install. The default language is English. Click Enter; see Example 4-34.

Example 4-34 AIX installation language selection

```
>>> 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]:
```

6. An error message appears stating that there are no disks available in the system. So far this is a correct message. The iSCSI host bus adapter is not yet configured to have disks available. Click 1 and Enter to proceed with the installation as shown in Example 4-35.

Example 4-35 AIX Installation with no disk error message

Error Warning

There are no disks available on this system.

To reboot the system, press reset.

```
>>> 1 Continue with Install
```

88 Help ?

>>> Choice [1]:

7. Select **Configure Network Disks (iSCSI)** by clicking 4 followed by Enter. See Example 4-36. This will start a dialog that allows you to configure the iSCSI host bus adapter.

Example 4-36 Installation main menu

Welcome to Base Operating System

Installation and Maintenance

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

>>> 1 Start Install Now with Default Settings

2 Change/Show Installation Settings and Install

3 Start Maintenance Mode for System Recovery

4 Configure Network Disks (iSCSI)

88 Help ?

99 Previous Menu

>>> Choice [1]:

8. Select **Configure iSCSI** and click Enter. The Network Utilities allow you to ping a iSCSI target, but use the on-board network interfaces. Therefore, this option is not the best to test the network connectivity. See Example 4-37.

Example 4-37 Configure iSCSI

Configure iSCSI

Move cursor to desired item and press Enter.

Configure iSCSI

Network Utilities

Esc+1=Help

Esc+2=Refresh

Esc+3=Cancel

F8=Image

F9=Shell

F10=Exit

Enter=Do

the domain name of the initiator or target vendor. The domain name is written reverse to clearly differentiate between IQNs and domain names.

Then follows an identifier of the host or target or additional values for identification. The maximum length of an IQN is 223 alphanumerical characters. The minimum length is 16 characters. The IQN used in this setup for the iSCSI target is `iqn.2008-07.com.ibm:storage.sys1`, as configured in “Configure the iSCSI software target” on page 183.

When you create your own IQNs, develop your own naming schema that allows the identification of the iSCSI port of each host and target and that allows to structure them. When all values are defined, click Enter.

Example 4-39 Configure the iSCSI host bus adapter port

iSCSI Configuration -- TOE Adapter

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

	[Entry Fields]
* iSCSI Adapter	ics0
* IP Address of iSCSI Adapter	[172.16.1.44]
* Network Mask of iSCSI Adapter	[255.255.255.0]
* Gateway	[172.16.1.1]
* iSCSI Target Name	[iqn.2008-07.com.ibm:s>
* IP Address of iSCSI Target	[172.16.1.195]
* Port Number of iSCSI Target	[3260]
* iSCSI Initiator Name	<7.com.ibm:blade5port1]

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	

11. Example 4-40 shows a successful login to an iSCSI target. When the line showing the hdisk does not show up, something went wrong. The virtual I/O Server will return a message that no child devices are found in this case. Not much troubleshooting is possible. During the setup it is helpful to make a network trace and trace the communication between the iSCSI initiator and target. This provides a clear indication where your problem lies.

Example 4-40 Established iSCSI connection to a disk

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

```
ics0 changed
sh: /usr/sbin/savebase: not found
mkiscsi: 0321-007 Failed to execute savebase command.
hdisk0 Available 03-09-01 N/A
```

```
Esc+1=Help      Esc+2=Refresh   Esc+3=Cancel    F6=Command
F8=Image        F9=Shell        F10=Exit        /=Find
n=Find Next
```

12. When the iSCSI disk is found, click F10 or ESC+0 to leave the iSCSI configuration. You return to the panel where you started the iSCSI configuration; see Example 4-41. Choose 2 to verify the details of your iSCSI disk, and define the new disk to be used for the installation.

Example 4-41 Installation main menu

```
Welcome to Base Operating System
                Installation and Maintenance

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

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

      88 Help ?
      99 Previous Menu

>>> Choice [1]:
```

13. Click 1 for System Settings to select the disk onto which IBM AIX will be installed; see Example 4-42.

Example 4-42 IBM AIX installation settings

```
Installation and Settings

Either type 0 and press Enter to install with current settings, or type the
number of the setting you want to change and press Enter.
```

```
1 System Settings:
```

Method of Installation.....New and Complete Overwrite
Disk Where You Want to Install.....

- 2 Primary Language Environment Settings (AFTER Install):
 - Cultural Convention.....English (United States)
 - LanguageEnglish (United States)
 - KeyboardEnglish (United States)
 - Keyboard Type.....Default
- 3 Security Model.....Default
- 4 More Options (Software install options)

>>> 0 Install with the current settings listed above.

```
88 Help ? | +-----+
99 Previous Menu | | WARNING: Base Operating System Installation will
                | | destroy or impair recovery of ALL data on the
                | | destination disk .
```

>>> Choice [0]:

14. Verify that you have the expected disk. Select the disk with **1** and leave this screen with **0**; see Example 4-43.

Example 4-43 Select the disk to install

Change Disk(s) Where You Want to Install

Type one or more numbers for the disk(s) to be used for installation and press Enter. To cancel a choice, type the corresponding number and Press Enter. At least one bootable disk must be selected. The current choice is indicated by >>>.

	Name	Location Code	Size(MB)	VG Status	Bootable
>>> 1	hdisk0	03-09-01	4766	none	Yes No

- >>> 0 Continue with choices indicated above
- 55 More Disk Options
- 66 Devices not known to Base Operating System Installation
- 77 Display More Disk Information
- 88 Help ?
- 99 Previous Menu

>>> Choice [0]:

15. Define your cultural conventions, language, and keyboard settings. Choose the security model and additional software settings. Proceed with the installation with **0**; see Example 4-44.

Example 4-44 IBM IAIX installation settings

Installation and Settings

Either type 0 and press Enter to install with current settings, or type the number of the setting you want to change and press Enter.

```
1 System Settings:
  Method of Installation.....New and Complete Overwrite
  Disk Where You Want to Install.....hdisk0

2 Primary Language Environment Settings (AFTER Install):
  Cultural Convention.....English (United States)
  Language .....English (United States)
  Keyboard .....English (United States)
  Keyboard Type.....Default

3 Security Model.....Default
4 More Options (Software install options)

>>> 0 Install with the current settings listed above.

      88 Help ?           | +-----+
      99 Previous Menu  | | WARNING: Base Operating System Installation will
                        | | destroy or impair recovery of ALL data on the
                        | | destination disk .
>>> Choice [0]:
```

16. Verify your settings and start the installation with **1**; see Example 4-45.

Example 4-45 IBM 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 .
>>> Choice [1]:

```

17. Installation will start as shown in Example 4-46.

Example 4-46 IBM AIX 6.1 Installation in progress

Installing Base Operating System

Please wait...

Approximate % tasks complete	Elapsed time (in minutes)	
13	5	Over mounting /.

18. When the installation is done, a reboot is performed. During the reboot, shortly before the kernel is loaded, you will see the configuration details of the iSCSI connection, as shown in Example 4-47.

Example 4-47 Configuration details of the iSCSI connection

IBM
1000Mbps full duplex link up.

```

iSCSI BOOT -----
Server IP.....172.16.1.197
Client IP.....172.16.1.50
Gateway IP.....172.16.1.1
Subnet Mask.....255.255.255.0
iSCSI Initiator.....iqn.2008-07.com.ibm:blade5port1
iSCSI Target.....iqn.2008-07.com.ibm:storage.sys1
Target Port.....3260
Target LUN.....0

```

19. After the installation is done, a reboot of the partition is performed. Select your terminal type; see Example 4-48.

Example 4-48 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-----

88 Help ? | If the next screen is unreadable, press Break (Ctrl-c)
| to return to this screen.

>>> Choice []:

-
20. Select **Show Installed License Agreements** and click Enter to read the license agreement; see Example 4-49.

Example 4-49 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

-
21. Select the software package from which you would like to read the license agreements. The default is to show all license agreements. Click Enter to start showing the license text; see Example 4-50.

Example 4-50 Select a software license

Show Installed License Agreements

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

* SOFTWARE name	[Entry Fields]	
SHOW license agreement text?	[a11]	+
	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

22. Navigate through the licenses. After finishing reading click F3 twice. You return to the Software License Agreements panel. Select Accept License Agreements and click Enter; see Example 4-51.

Example 4-51 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

23. Click Tab followed by Enter to accept the license and change the no into a yes; see Example 4-52.

Example 4-52 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

24. After the status of the command has changed to OK, click F10 or ESC+0 to leave the license agreement. The keys you need depend on the terminal you are using; see Example 4-53.

Example 4-53 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

25. The installation assistant will guide 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 get to a login prompt, use ESC+0 or F10. You may start this installation assistant at any time again by using the command `install_assist` after login as `root`. The installation assistant is shown in Example 4-54.

Example 4-54 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 using a iSCSI disk to boot from has completed. You may now start to configure your AIX or install your applications.

4.3 Boot from iSCSI LUN

When you want to set up the iSCSI boot using an existing and already installed iSCSI LUN, as you need to do after the planar of the blade was replaced, or you need to set up the iSCSI initiator using the SMS menu, perform the following steps:

1. Use a console to power on or restart the blade and enter the SMS menu.
2. Select **Setup Remote IPL (Initial Program Load)** and click Enter; see Example 4-55.

Example 4-55 Select Remote IPL

Version EA320_031
SMS 1.7 (c) Copyright IBM Corp. 2000,2007 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:

3. Choose the device you want to use for the iSCSI boot. Example 4-56 on page 202 shows the two on-board network interface ports and four QLogic Dual Port 1000Base-T entries. These four ports belong to the same dual port HBA. The iSCSI HBA has two different functions per port. The first one is the regular Ethernet functionality. There is no iSCSI or TCP off-load with this function. The HBA behaves as any other regular network card.

The second function is like a SCSI adapter. SCSI commands are sent from the OS over the driver to this port. The TCP off-load engine encapsulates the SCSI commands in TCP packages which are then sent over the Ethernet to the iSCSI target. Figure 4-31 on page 202 shows a high-level overview of this iSCSI HBA.

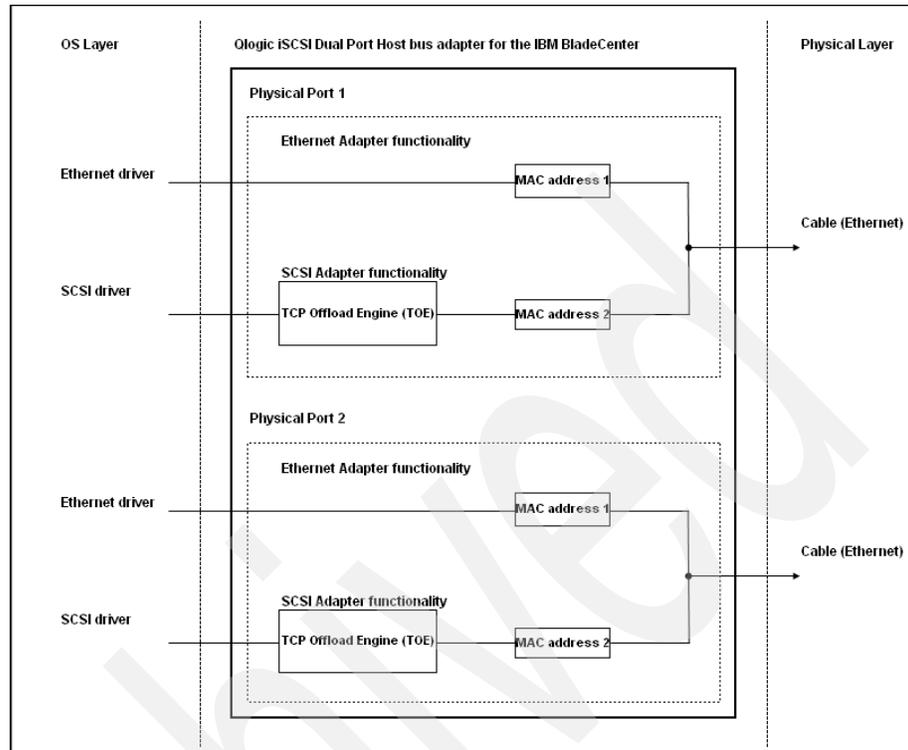


Figure 4-31 iSCSI HBA high-level overview

4. In Example 4-56 the selection of the device for the iSCSI boot is made.

Example 4-56 Choose the device to be used for the iSCSI boot

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

Version EA330_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.

NIC Adapters	Device	Location Code	Hardware Address
1.	PORT - 1 IBM Host Ethernet Ada	U78A5.001.WIH01AA-P1-T6	001a64760008
2.	PORT - 2 IBM Host Ethernet Ada	U78A5.001.WIH01AA-P1-T7	001a64760009
3.	QLogic Dual Port 1000Base-T PC	U78A5.001.WIH01AA-P1-C6-T1	00c0dd03ee29
4.	QLogic Dual Port 1000Base-T PC	U78A5.001.WIH01AA-P1-C6-T1	00c0dd03ee2a
5.	QLogic Dual Port 1000Base-T PC	U78A5.001.WIH01AA-P1-C6-T2	00c0dd03ee2b
6.	QLogic Dual Port 1000Base-T PC	U78A5.001.WIH01AA-P1-C6-T2	00c0dd03ee2c

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:

5. Choose the protocol used for remote IPL. Choose **iSCSI** by clicking **2** and click Enter (Example 4-57).

Example 4-57 Choose the protocol used for remote IPL

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

Select Network Service.

No alias : PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

1. BOOTP
 2. iSCSI
-

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:

6. Select the alias to cofeature. Up to two aliases might be configured. These aliases are used later to specify the boot device. You will find them again in the boot device list (Example 4-58). Choose the first alias with **1** and click Enter.

Example 4-58 Choose iSCSI alias

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

iSCSI Alias Selection:

No alias : PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

1. iSCSI_DISK1
2. iSCSI_DISK2

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:

7. Click **1** to specify all iSCSI and networking parameters manually. There is DHCP support available here (Example 4-59).

Example 4-59 Enter iSCSI and network parameters statically

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

iSCSI Discovery Selection.

iscsi_disk1: PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

1. Enter Static Parameters.
-

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:

8. This form will be empty and does not reflect the setting you made during the installation of AIX. To see the settings you made during the installation you need to enter the Open Firmware prompt and use the **printenv** command to show the settings that are used to boot from the iSCSI target. This command must be issued on a system that actually was used to install the iSCSI LUN and must still have the boot configuration set up. See “Open Firmware interface” on page 407 for more details.

Example 4-60 shows the empty menu that needs to be filled out with the network and iSCSI settings. Click **1** to define the IP address of the network interface. Again, this can be an iSCSI host bus adapter or a regular network card interface and which one is supported depends on the operating system. Example 4-60 shows the usage of a regular network adapter port.

When you click **1** you can enter the IP address of the iscsi initiator in the last line in Example 4-60. To cancel the input, use **x** and Enter to save the new address. Proceed with the gateway and the subnet mask used for the iSCSI initiator as you did with the IP address. Then click **4** to define the iSCSI Qualified name or initiator name.

Example 4-60 iSCSI initiator settings

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

iSCSI Initiator Parameters:

iscsi_disk1: PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

1. Client IP Address
 2. Gateway IP Address
 3. Subnet Mask
 4. Set Initiator Name
 5. Enter iSCSI Target Parameters
-

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:

9. To specify the iSCSI qualified name, enter **2** to edit the current empty field. The last line will be used to enter the IQN. Use ESC to cancel the input and Enter to save the IQN, as shown in Example 4-61. An already defined IQN may be cleared with **1**.

IQNs start with iqn. Then follows—separated by a dash—the registration year and month of the domain name of the initiator or target vendor. The domain name is written reverse to clearly differentiate between IQNs and domain names. Then follows an identifier of the host or target or additional values for identification. The maximum length of an IQN is 223 alphanumeric characters. The minimum length is 16 alphanumeric characters.

Example 4-61 iSCSI qualified name of the initiator

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

Enter iSCSI Initiator Name:
iscsi_disk1: PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

Must be at least 4 characters.
Press Backspace key to delete.
Press Enter key to save.
Press Esc key to exit without saving name.

1. Clear Current Name.
 2. Edit Current Name.
- []

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:

10. Click 5 to specify the iSCSI target settings (Example 4-62).

Example 4-62 iSCSI initiator settings

Version EA320_031
SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

iSCSI Initiator Parameters:
iscsi_disk1: PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

1. Client IP Address [172.16.1.171]
2. Gateway IP Address [255.255.255.0]
3. Subnet Mask [255.255.255.0]
4. Set Initiator Name
5. Enter iSCSI Target Parameters

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:

11. Click **1** to enter the iSCSI target portal IP address. Enter the IP address of the iSCSI target that you want to connect to. Save the address with Enter or cancel the input with X. Verify the defined iSCSI port. The default port is 3260 but your target may use a different port. The default LUN 0 should be the correct setting for most setups. It can be changed here. Click **34** to define the iSCSI qualified name of the iSCSI target; see Example 4-63.

Example 4-63 iSCSI target settings

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

iSCSI Target Parameters:

iscsi_disk1: PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

1. Target IP Address
2. iSCSI PORT
3. Target lun
4. Set iSCSI Target Name

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:

12. There is no iSCSI qualified name defined as default here. Click **2** to edit the currently empty field and add the iQN from the iSCSI target. An existing value might be deleted with **1**. When you are done, click **ESC** to return to the iSCSI target setting menu, as shown in Example 4-64.

Example 4-64 iSCSI target name

Version EA320_031

SMS 1.7 (c) Copyright IBM Corp. 2000,2007 All rights reserved.

Enter iSCSI Target Name:

iscsi_disk1: PORT - 1 IBM Host Ethernet Adapter: U78A5.001.WIH0188-P1-T6

Must be at least 4 characters.

Press Backspace key to delete.

Press Enter key to save.

Press Esc key to exit without saving name.

1. Clear Current Name.
2. Edit Current Name.

[]

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:

13. Click ESC again to leave the iSCSI target settings menu and return to the iSCSI initiator settings menu.
14. Click ESC again to save the iSCSI settings you chose.
15. Now verify your boot device settings and be sure that the iSCSI alias you configured is listed as boot device. The alias is called iscsi-disk1 or iscsi-disk2. After this verification you can leave the System Maintenance Services menu and boot the blade from the iSCSI target.

IBM i 6.1 installation

This chapter explains the installation process of the IBM i 6.1 Operating System on an IBM BladeCenter JS12 Express server installed in a BladeCenter S chassis using the disks provided in the disk storage modules.

For the IBM BladeCenter JS22 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.

For a technical overview and complete information, as well as latest updates on IBM i on Power blades, refer to the readme file available at the following address:

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

This chapter discusses the following base topics:

- ▶ Preparing for installation
- ▶ IBM System Access for Windows V6R1
- ▶ Creating an IBM i 6.1 partition
- ▶ Installing Licensed Internal Code (LIC)
- ▶ Installing the IBM i 6.1 Operating System
- ▶ IBM i 6.1 Backup/Restore

5.1 Preparing for installation

There are important considerations for setting up and using IBM i 6.1 client logical partitions on IBM Power servers or the IBM BladeCenter JS12 or JS22 Express server. On Power blades, you use IBM to manage IBM i partitions.

A *client logical partition* is a partition that uses some of the I/O resources of another partition.

When the IBM i 6.1 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.

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 allows you to optimize the amount of hardware used by the host partition.

5.1.1 Software installation process

The IBM i 6.1 installation process involves three phases:

- ▶ Pre-installation activities
- ▶ Main installation
- ▶ Post-installation activities

Figure 5-1 on page 211 illustrates one type of software installation of the IBM i 6.1 Operating System.

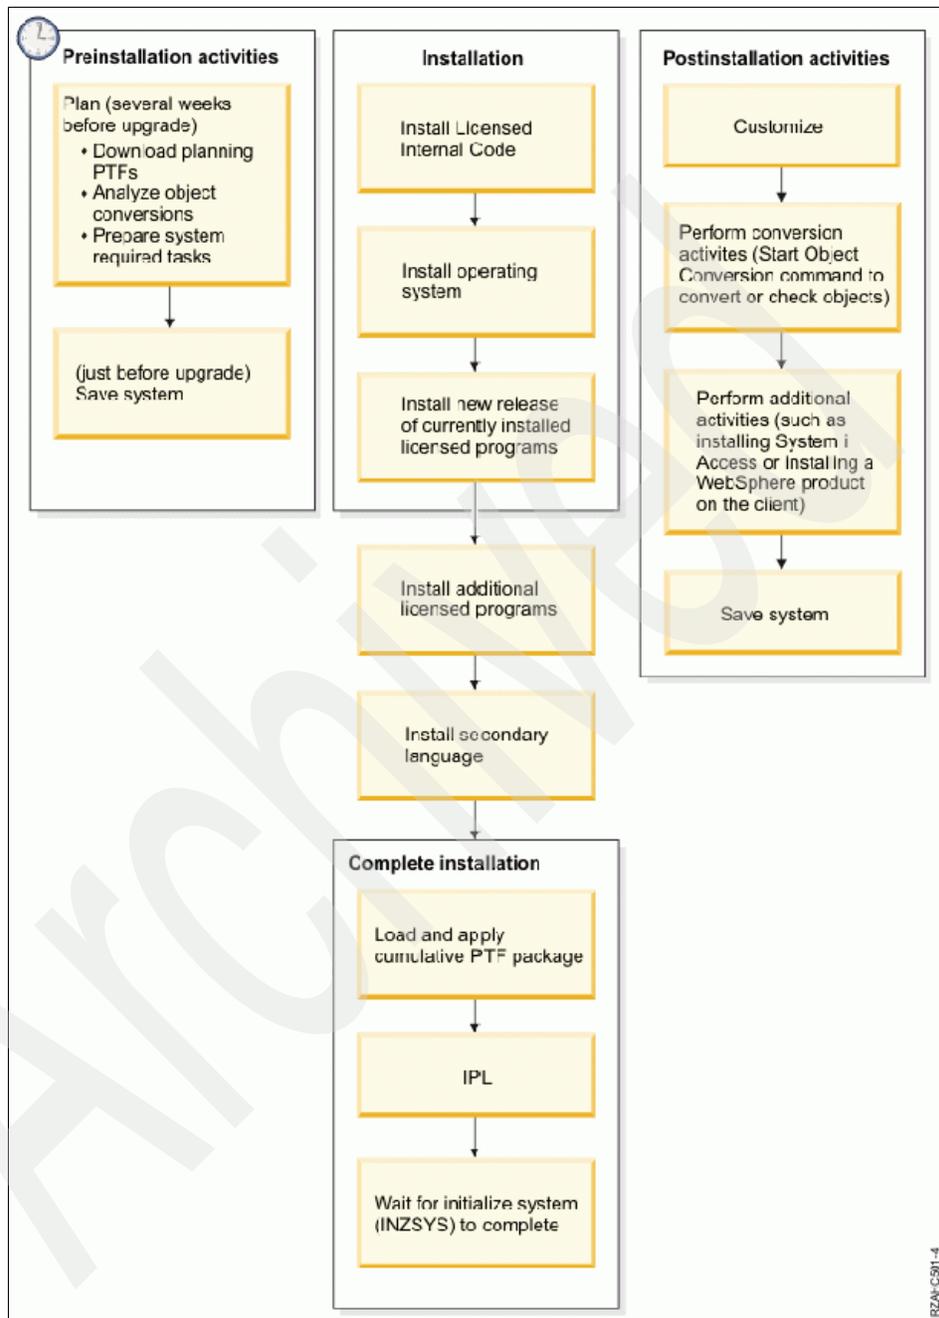


Figure 5-1 IBM i 6.1 installation process

5.1.2 Hardware environments

This section describes an example IBM BladeCenter chassis and IBM BladeCenter JS12 Express server configuration with recommended firmware levels.

Note: The disk configurations are dependent on the I/O requirements. For example, two SAS disk drives will not be enough with mirroring and backup to the media library.

A combination of BladeCenter S chassis and IBM BladeCenter JS22 or IBM BladeCenter JS12 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 JS22 Express

- ▶ 1 JS22 Express server
- ▶ 2 GB memory
- ▶ SAS Expansion Card (CFFv) for IBM BladeCenter
- ▶ 1 SAS disk drive

IBM BladeCenter JS12 Express

- ▶ 1 JS12 Express server
- ▶ 2 GB memory
- ▶ SAS Expansion Card (CFFv) for IBM BladeCenter
- ▶ 1 SAS disk drive

A combination of BladeCenter H chassis and BladeCenter JS22 blade 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 JS22 Express

- ▶ 1 JS22 Express server
- ▶ 2 GB memory
- ▶ 1 QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)
- ▶ 1 SAS disk drive

Table 5-1 lists the minimum and required features required to manage an IBM BladeCenter JS22 Express system with the IBM i 6.1 Operating System.

Table 5-1 BladeCenter JS22 feature list

Feature	Description	Notes
8400	IBM BladeCenter JS22 Express	
8401	Processor Entitlement (4-core)	Four processor Entitlement required
8401	Express Configuration	
8399	Processor Entitlement (2-core) Zero Priced Proc. (2-core)	
8232	2 GB DDE2 667 MHz DIMMs	One required, up to two supported
8233	4 GB DDE2 667 MHz DIMMs	
8234	8 GB DDE2 667 MHz DIMMs	
8235	16 GB DDE2 533 MHz DIMMs	
8237	73 GB SAS 10k SFF disk	One disk is required
8236	146 GB SAS 10k SFF disk	
8252	QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)	- Required for SAN connection in BladeCenter H - Not supported in BladeCenter S
8250	SAS Expansion Card (CFFv) for IBM BladeCenter	- Optional for SAS Tape Connection in BladeCenter H - Required for SAS Disk and Tape in BladeCenter S

Table 5-2 on page 213 lists the minimum and required features needed to manage an IBM BladeCenter JS12 Express system with the IBM i 6.1 Operating System.

Table 5-2 BladeCenter JS12 feature list

Feature	Description	Notes
8442	IBM BladeCenter JS12 Express	

Feature	Description	Notes
8444	Processor Entitlement (2-core)	Two processor Entitlement required
8444	Express Configuration	
8443	Processor Entitlement (1-core) Zero Priced Proc. (1-core)	
8220	2 GB DDE2 667MHz DIMMs	One required, up to four supported
8229	4 GB DDE2 667MHz DIMMs	
8239	8 GB DDE2 667MHz DIMMs	
8245	16 GB DDE2 400 MHz DIMMs	
8237	73 GB SAS 10k SFF disk	One disk is required
8236	146 GB SAS 10k SFF disk	
8250	SAS Expansion Card (CFFv) for IBM BladeCenter	- Required for SAS Disk and Tape in BladeCenter S

For more information on supported devices on a BladeCenter JS12 server, refer to the following sites:

<http://www.ibm.com/systems/i/os/i5os/v6r1/blades/config.html>

http://www-03.ibm.com/systems/power/hardware/blades/i_on_blade_readme.pdf

Note: At the time of writing, the IBM BladeCenter JS12 in combination with the IBM i 6.1 Operating System is supported in the BladeCenter S chassis using the CFFv cards only.

5.1.3 BladeCenter hardware preparation

For a smooth installation process of IBM i 6.1, all installed hardware components should be at the latest firmware levels.

The following components should be checked and updated if necessary:

- ▶ IBM JS12 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 Channel CFFh expansion cards and SAS expansion cards.)

For more information about supported components and about procedures to install updates to the IBM BladeCenter JS22 or IBM BladeCenter JS12 Express server, BladeCenter S, or BladeCenter H chassis, refer to:

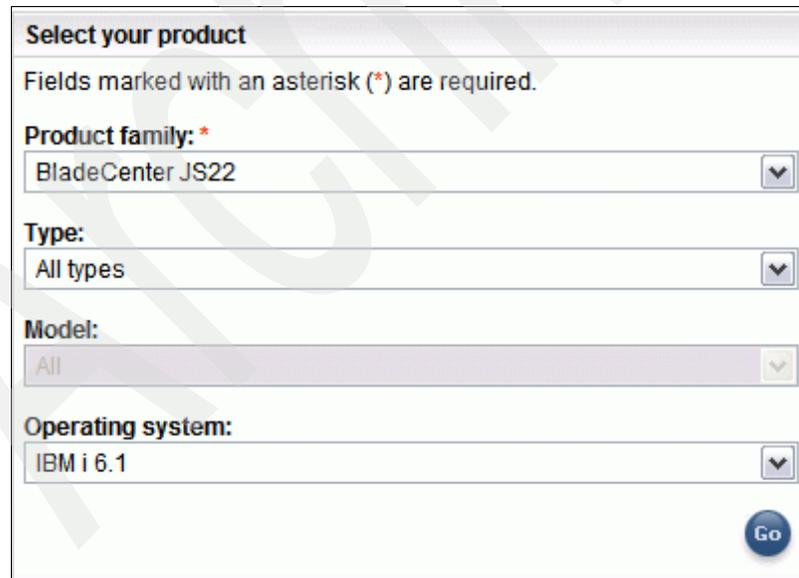
<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

For firmware updates for the IBM BladeCenter JS22 or IBM BladeCenter JS12 Express, BladeCenter S, and BladeCenter H chassis and integrated components, refer to:

<http://www.ibm.com/system/support/supportsite.wss/brandmain?brandind=5000020>

When you reach that window, follow these steps:

1. Click the down arrow in the Product family box and select the corresponding product: IBM BladeCenter JS22, BladeCenter JS12, BladeCenter S, or BladeCenter H.
2. Click the down arrow button in the Operating system box and select **IBM i 6.1**, as shown in Figure 5-2. Then click the **Go** button to activate the search.



The screenshot shows a web form titled "Select your product". Below the title is a note: "Fields marked with an asterisk (*) are required." The form contains 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 5-2 Firmware information and download

Click **Go** to continue.

3. On the next window, scroll down to view all the available updates and drivers related to the selected product and operating system for downloading.

5.1.4 VIO Server software environments

VIO Server is part of IBM PowerVM Editions (formerly Advanced POWER Virtualization). It is required in the IBM i 6.1 for IBM BladeCenter JS12 Express environment.

Work with your local sales channel to ensure that PowerVM (Standard or Enterprise Edition) and the latest fix pack are part of the BladeCenter JS12 order. Consult the supported environments page to verify that you have the minimum supported release of VIOS:

<http://www.ibm.com/systems/i/os/i5os/v6r1/blades/config.html>

To get the latest recommended fixes and information refer to:

http://www-912.ibm.com/s_dir/slkbases.nsf/slkbases

For detailed information about using the Integrated Virtualization Manager and VIO Server, refer to Chapter 3, “System configuration using VIOS and IVM” on page 55.

5.1.5 Network considerations

Before the initial installation process starts, a set of IP addresses should be reserved for setting up the minimum environment.

AMM

The Advanced Management Module address is a physical LAN IP address. It is used to manage the BladeCenter chassis and IBM BladeCenter JS12 Express server remotely.

Note: We suggest that you 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 is used to connect 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 is used to connect 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 is used to allow the 5250 console to connect to the VIOS using the IBM System i Access for Windows® software.
6.1 production interface	This IP address on the external LAN is used to provide 5250 production network access. This address will be configured after 6.1 is installed using LAN console. It is recommended that the 6.1 LAN console and production network interface 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.

5.1.6 Storage consideration BladeCenter H

An IBM BladeCenter JS12 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 JS12 to the 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 JS22 or JS12 blade installed in a BladeCenter H chassis using LUN attached disks to install IBM i 6.1, refer to:

http://www.ibm.com/systems/power/hardware/blades/i_on_blade_readme.pdf

5.1.7 Disk consideration in BladeCenter S

This section describes the configuration environment used in this chapter to install IBM i 6.1 on an IBM BladeCenter JS12 Express server in a BladeCenter S chassis.

IBM i 6.1 running on an IBM BladeCenter JS12 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 JS12 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 (CFFv) also must be installed in each IBM BladeCenter JS22 or IBM BladeCenter JS12 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 to the 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 JS12, 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 6.1.

Note: SATA drives installed in the Disk Storage Modules (DSM) in the BladeCenter S chassis are not supported for IBM i 6.1.

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 JS12 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 RAS, use the IBM i 6.1 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 5.1.8, “Disk configuration in BladeCenter S” on page 219, helps you to visualize or select a predefined DSM configuration.

5.1.8 Disk configuration in BladeCenter S

To use a pre-defined configuration to a BladeCenter JS12 server, you must establish a connection to the SAS Module as shown in Figure 5-3 on page 219. using a browser window directly connected to the SAS Module.

An alternative, that is more intuitive for clients is the SCM GUI.

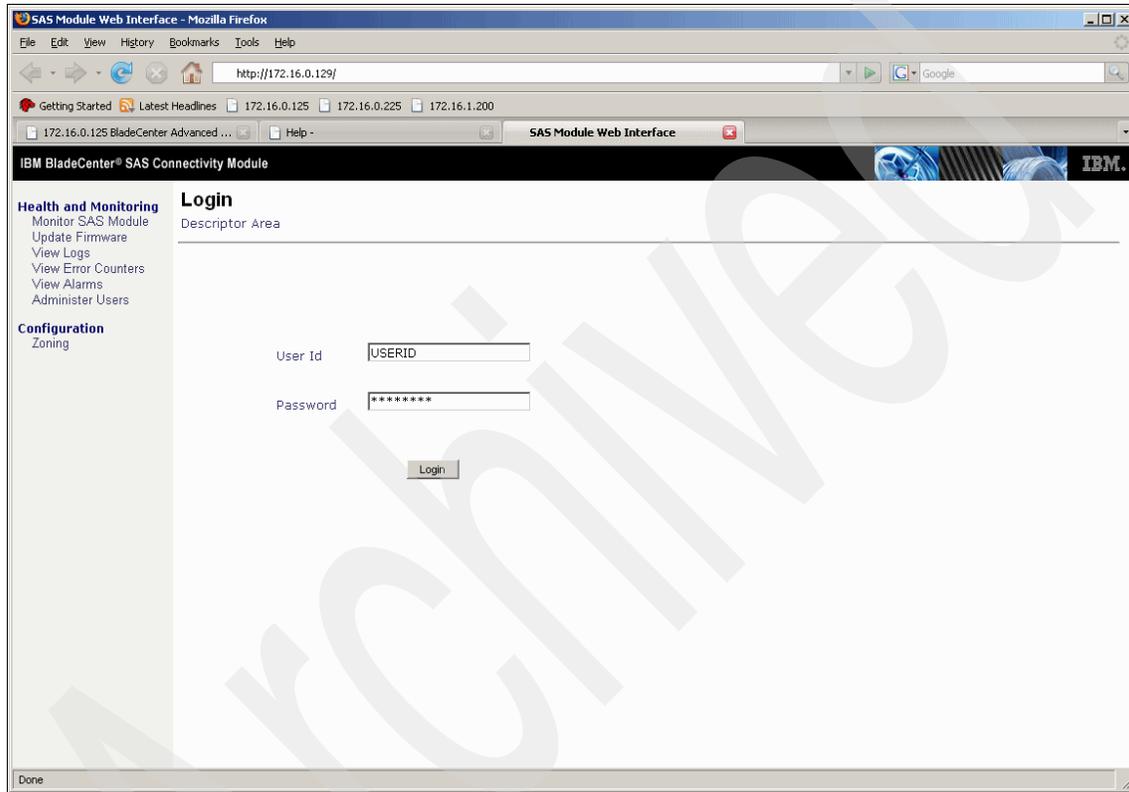


Figure 5-3 SAS Connection module login

Enter **User ID** and **Password** of the account that has access to the SAS module and click **Login**

Next screen that appears in the SAS module welcome screen provided in Figure 5-4 on page 220.

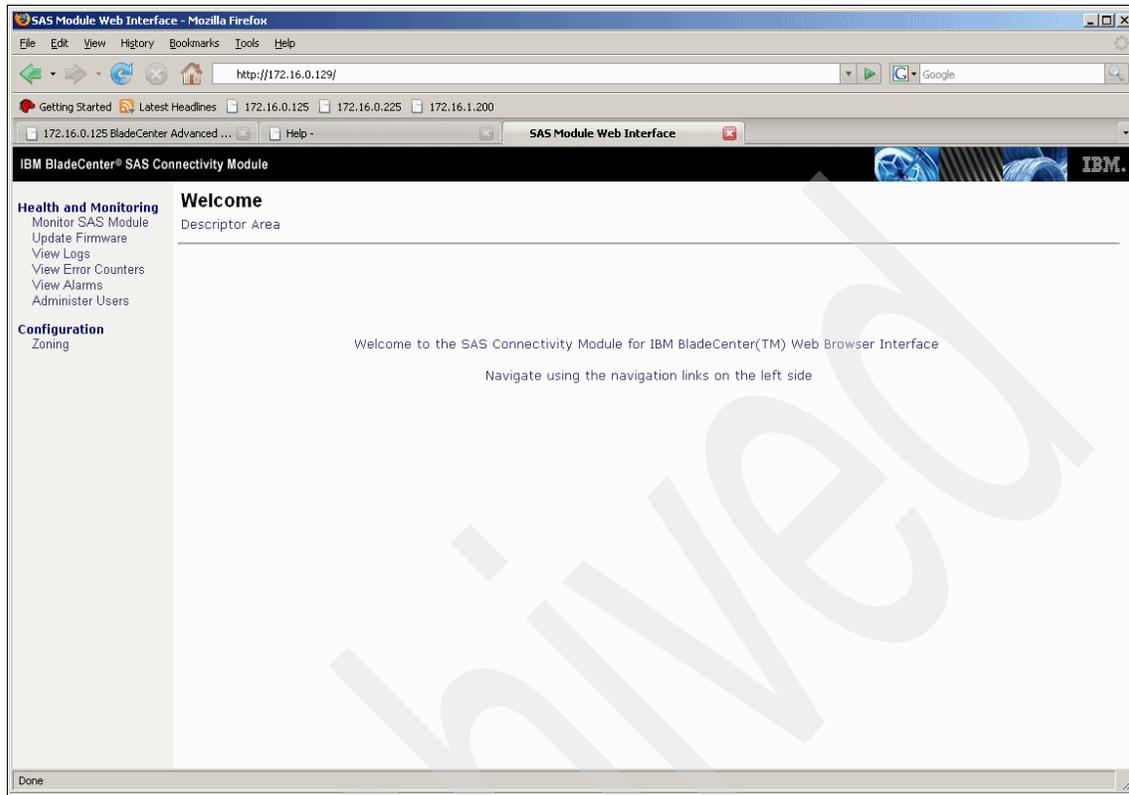


Figure 5-4 SAS connection module welcome

Select **Zoning**.

In the example shown in Figure 5-5 on page 221, **Predefined Config 09** is selected and active. Notice that our BladeCenter JS12 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.

Click **Basic Zone Permission Table**.

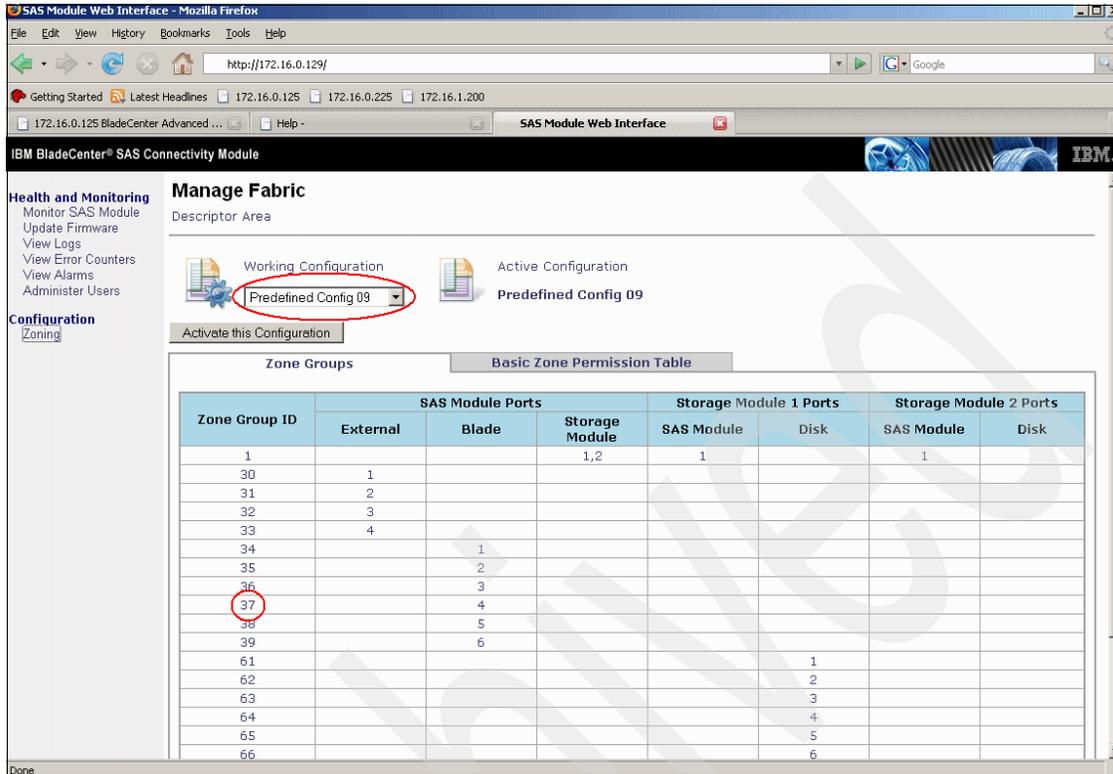


Figure 5-5 SAS connection module zone groups

Figure 5-6 on page 222 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 detailed information about this topic, refer to: *Implementing the IBM BladeCenter S Chassis, REDP-4357*.

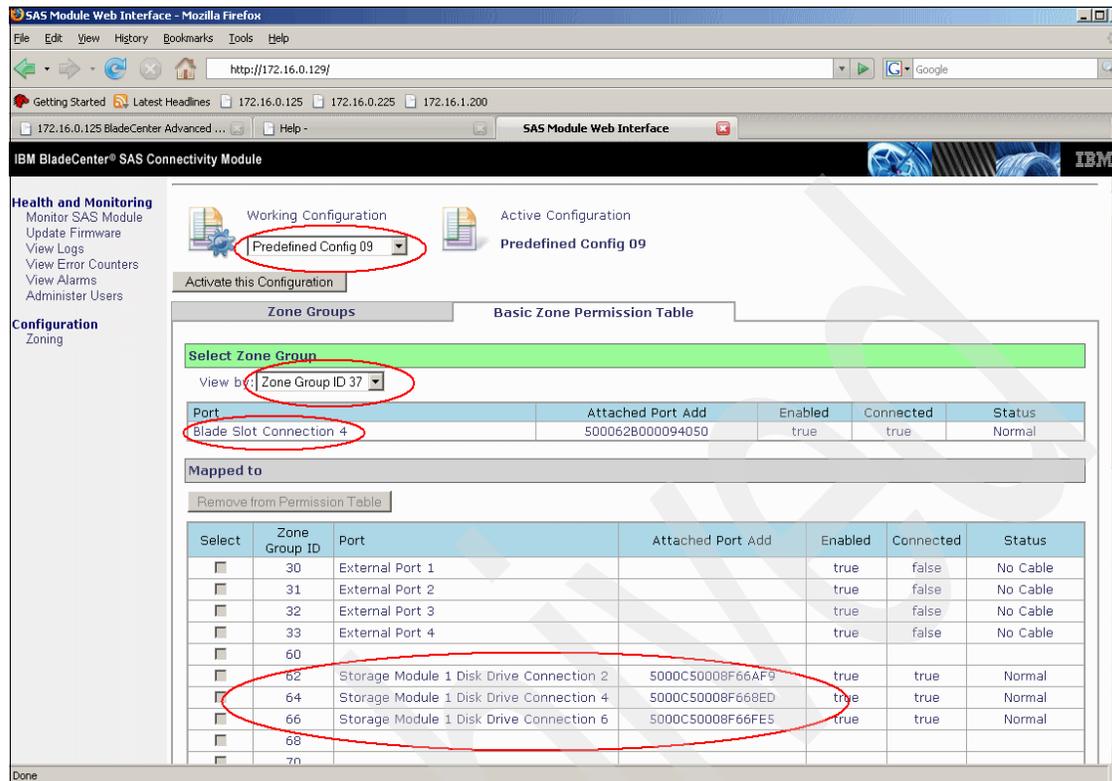


Figure 5-6 SAS connection module zoning

To verify the configuration in the SAS module configuration menus, logon to the IBM BladeCenter Advanced Management Module. Under Storage Tasks, select **Configuration** as shown in Figure 5-7 on page 223.

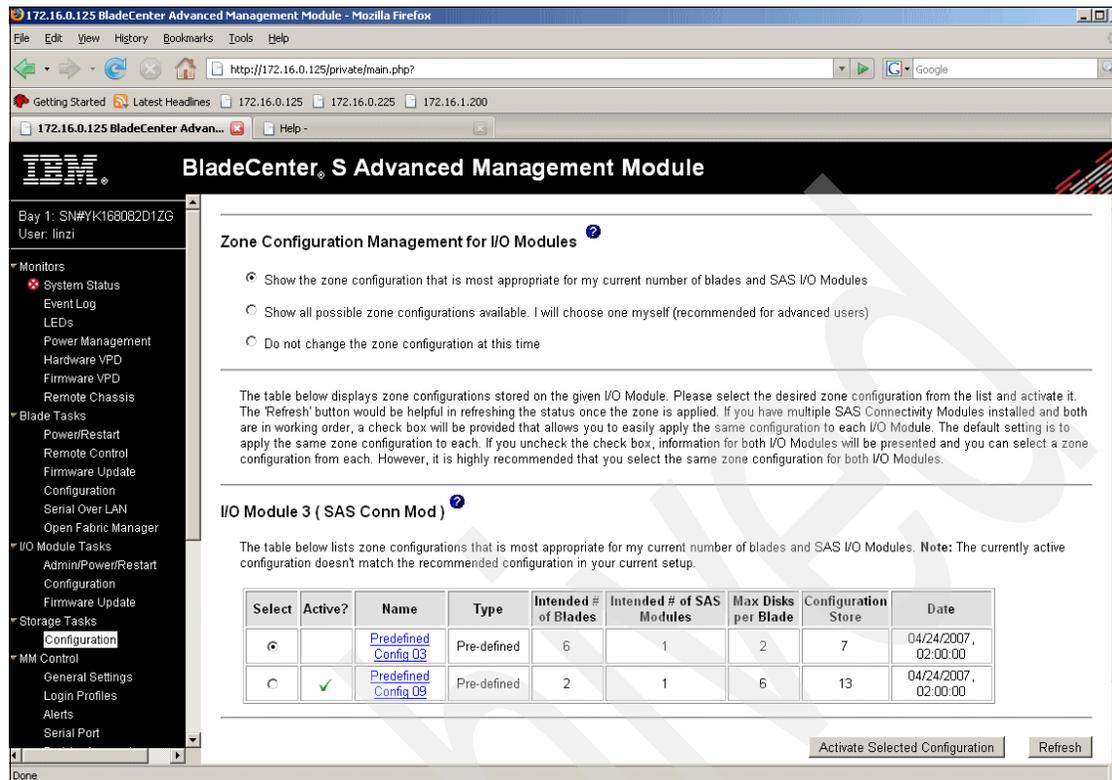


Figure 5-7 AMM SAS configuration zone

Click **Predefined Config 09** to proceed.

Figure 5-8 on page 224 shows the current configuration. Select the blade in the upper rectangle to highlight the assigned disks to that blade.

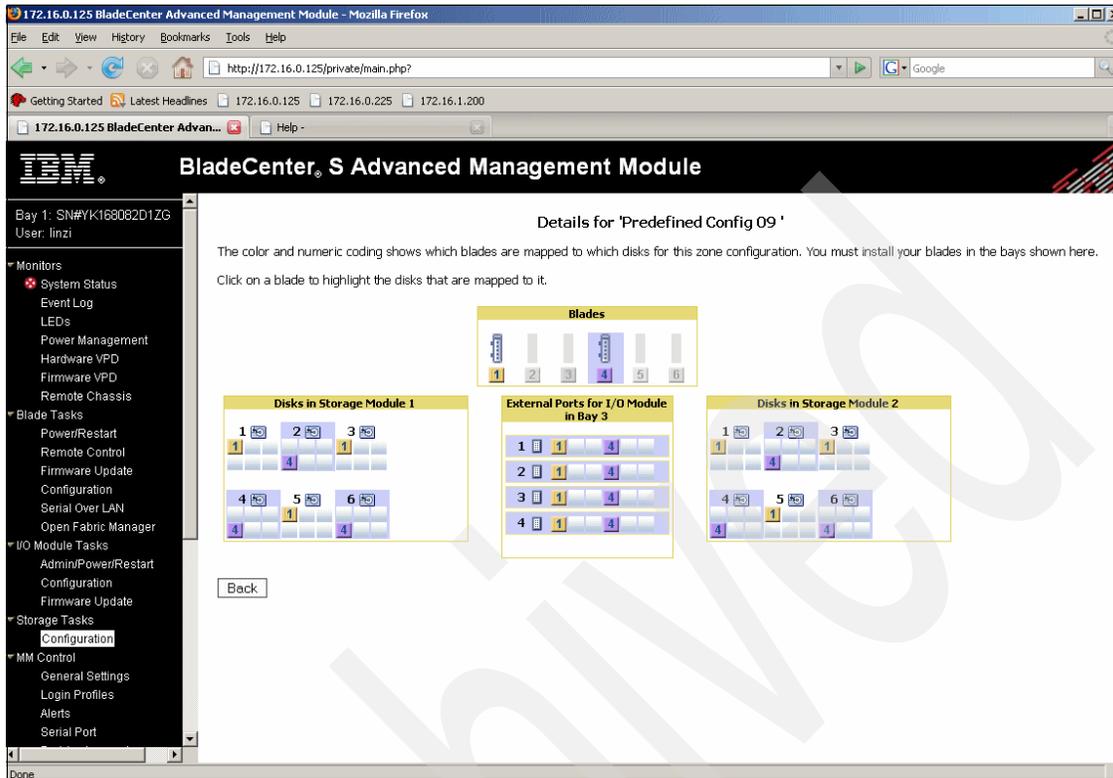


Figure 5-8 AMM SAS configuration zone 9

For detailed information, refer to *Implementing the IBM BladeCenter S Chassis*, REDP-4357 and *IBM BladeCenter Products and Technology*, SG24-7523.

5.1.9 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) may 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>

5.2 IBM System Access for Windows V6R1

System i Access for Windows fully integrates the power of the IBM i 6.1 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.

Prior to 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 Socket Layer (SSL)
- ▶ Printer Driver
- ▶ Display and Printer Emulation
- ▶ Operations Console

5.2.1 Preparation System i Access for Windows

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.

To meet all the System i Access for Windows hardware requirements, follow the instructions described in:

<http://publib.boulder.ibm.com/infocenter/systems/scopt/i5os/topic/rzajr/rzajrhardwarereq.htm#hardwarereq>

To meet all the System i Access for Windows software requirements, follow the instructions described in:

5.2.2 Installation System i Access for Windows

This section describes the sequence to install System i Access for Windows.

1. After you insert the System i Access for Windows V6R1 DVD, the window shown in Figure 5-9 appears. Select the language you require, and click **OK** to continue.

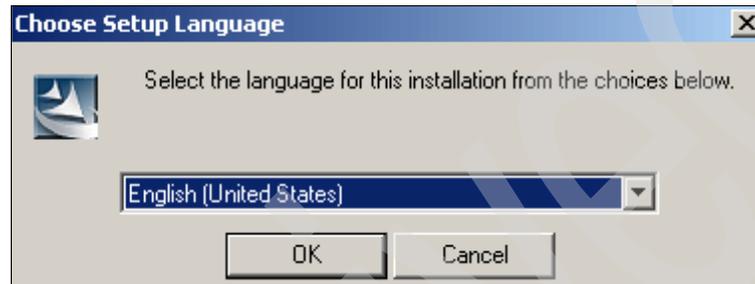


Figure 5-9 Setup LAN Console

2. The InstallShield Wizard Welcome window shown in Figure 5-10 on page 227 appears.

Click Next> to continue.



Figure 5-10 IBM System i Access for Windows welcome screen

3. The License Agreement shown in Figure 5-11 on page 228 appears. You can select **I accept the terms in the license agreement**. Click **Next>** to continue.

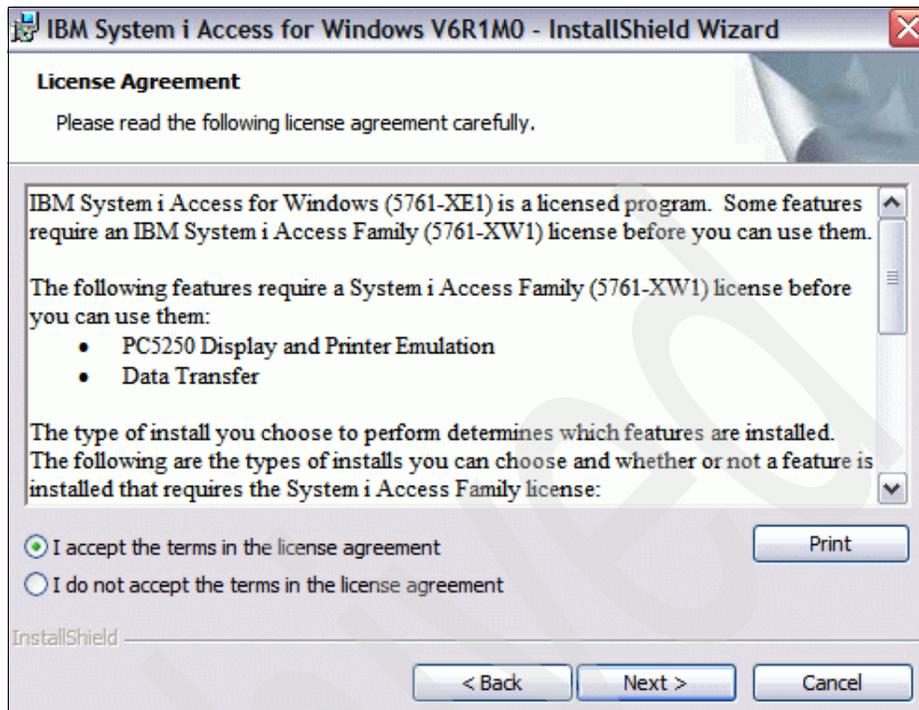


Figure 5-11 IBM System i Access for Windows Licence Agreement

4. IBM System i Access for Windows can be installed at a different location, as shown in Figure 5-12 on page 229.

To store the software at a different location, click **Change...** and choose a new location.

Or, accept the predefined path and click **Next>** to continue.

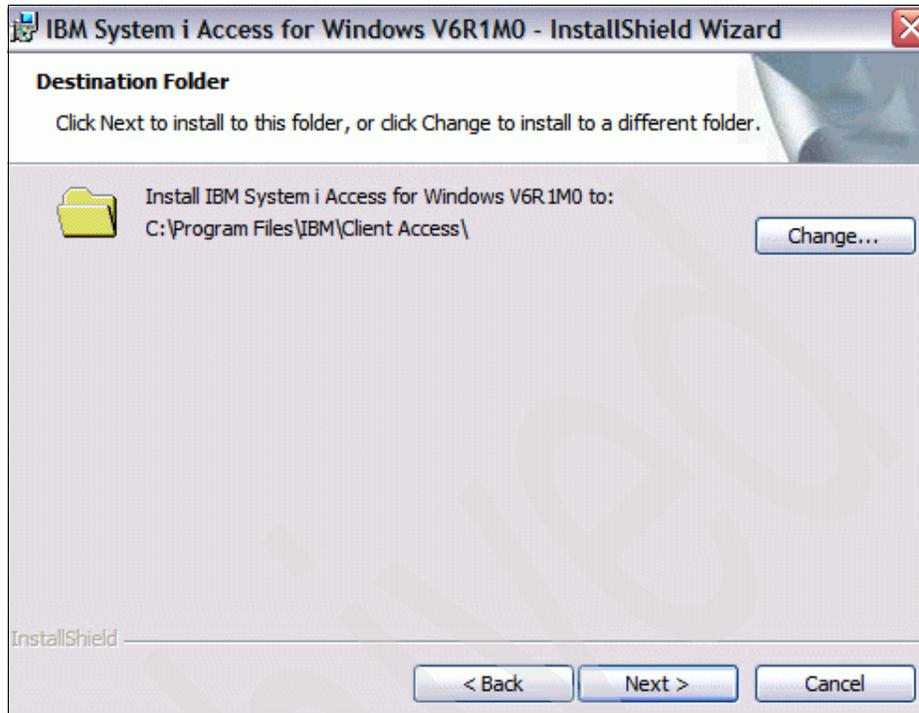


Figure 5-12 IBM System i Access for Windows install location

5. Depending on the native language, a selection can be made in the following window as shown in Figure 5-13 on page 230. Normally the same language will be chosen as the language for the IBM i 6.1 Operating System.

Click **Next>** to continue.



Figure 5-13 IBM System i Access for Windows Primary language

6. Depending on the complexity of functions, several choices are available as shown in Figure 5-14 on page 231. 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.

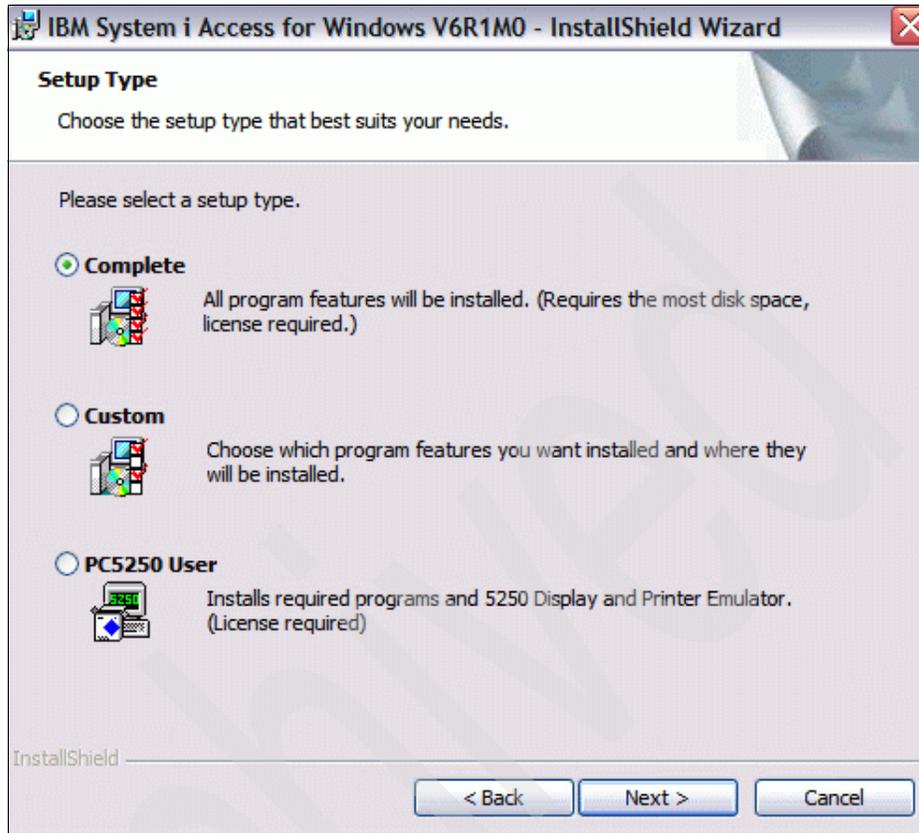


Figure 5-14 IBM System i Access for Windows Setup Type

7. Select **Complete** and click **Next>**
8. Some features require a license agreement to use their functionality, as shown in Figure 5-15 on page 232. Ask your service representative to receive a valid license key.
Click **Next>** to continue.

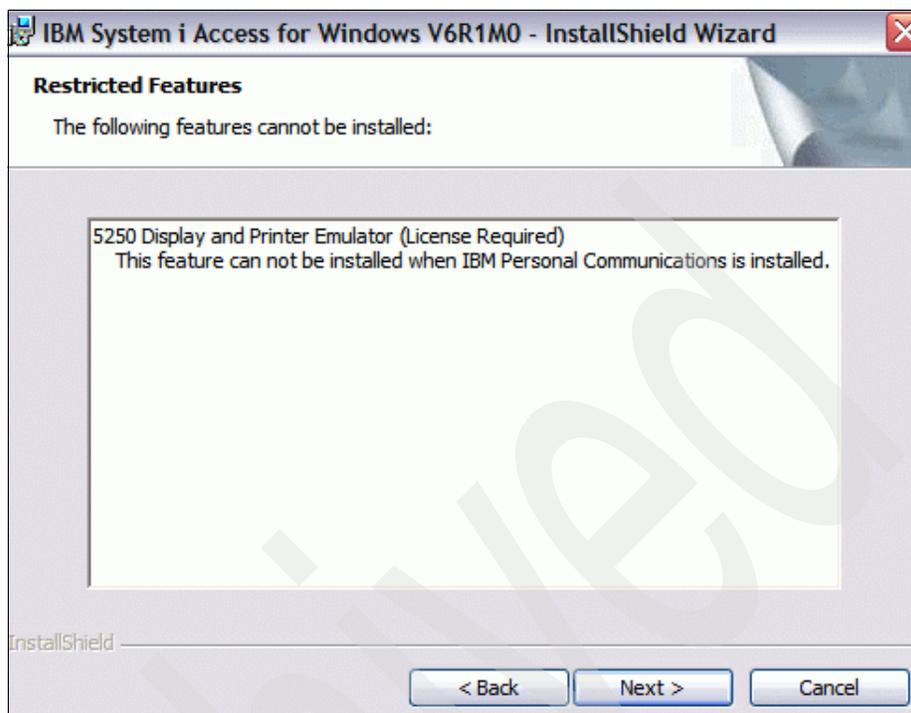


Figure 5-15 IBM System i Access for Windows Restricted Features

9. The installation starts automatically after you select **Next>** in the previous menu. Figure 5-16 on page 233 shows the progress of the installation process.



Figure 5-16 IBM System i Access for Windows installation progress

10. Figure 5-17 on page 234 indicates the installation process was successful. Click **Finish** to continue.

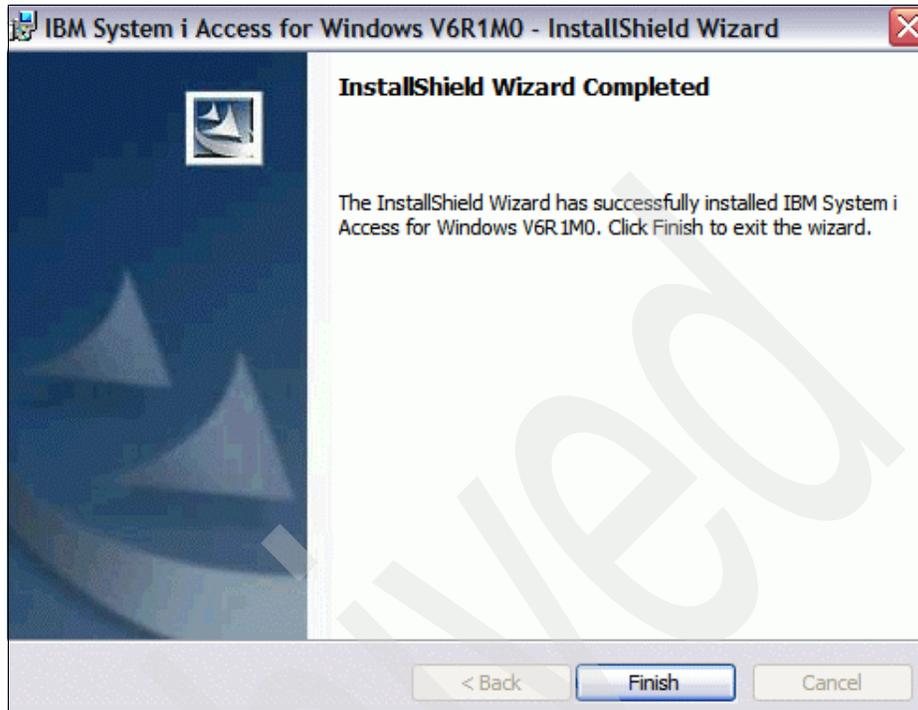


Figure 5-17 IBM System i Access for Windows installation completed

11. To finalize the IBM i Access for Windows installation a reboot is required, as indicated in Figure 5-18.

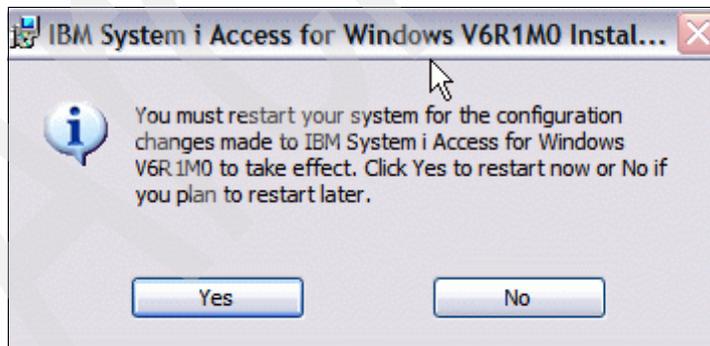


Figure 5-18 IBM System i Access for Windows Reboot

Click **Yes** to reboot the system.

After the console PC is successfully rebooted, the information screen shown in Figure 5-19 is displayed. The Welcome window provides additional information about the software just installed. (For some information, the administration PC needs a connection to the Internet.)

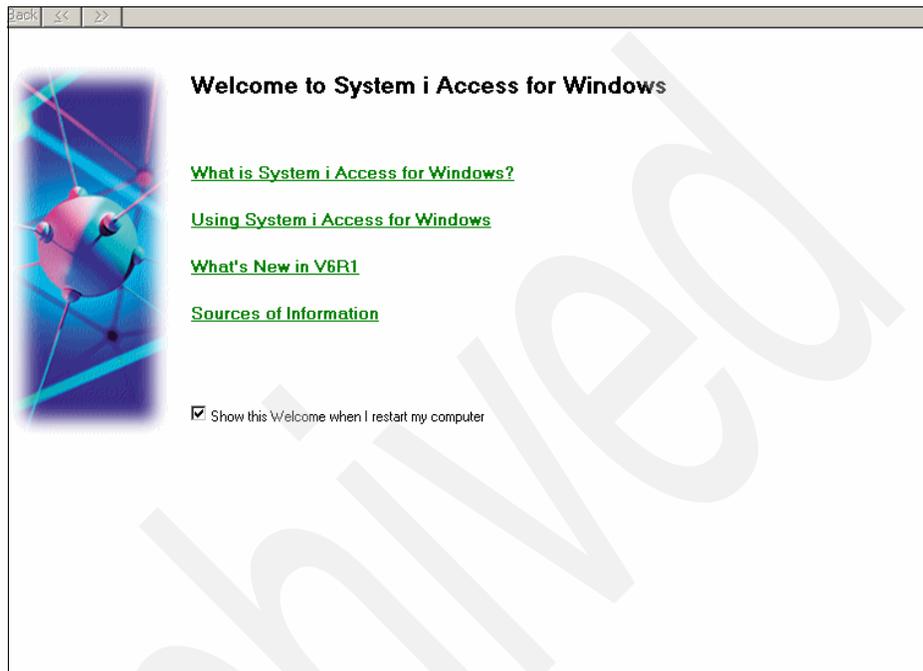


Figure 5-19 IBM System i Access for Windows Welcome

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 the following address:

<http://www.ibm.com/systems/i/software/access/caorder.html>

After the System i Access for Windows product is installed, you must meet two conditions 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 6.1 Operating System.

The following section explains these tasks.

5.2.3 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 previously mentioned, to establish a connection to the IBM i 6.1 partition in the BladeCenter S chassis, the PC must be in the same TCP/IP subnet as the BladeCenter JS12 blade.

If you experience problems establishing a connection, check the settings of the installed firewall. The ports used by the System i LAN console may have been disabled by the firewall.

To establish a connection to an IBM i 6.1 created partition, follow these steps:

1. In the active Windows session, select **Start -> All Programs -> IBM System i Access for Windows -> Operations Console**. The window shown in Figure 5-20 should appear.

Select the **Connection** drop-down menu to continue.

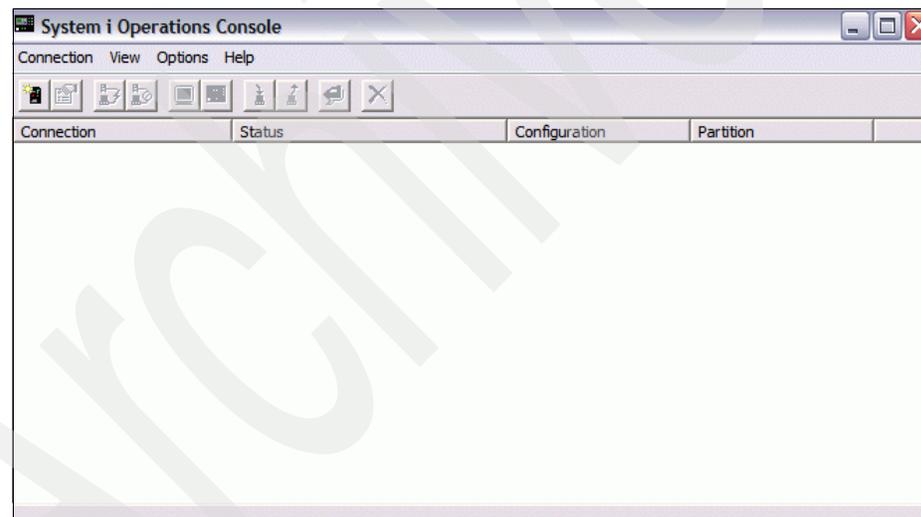


Figure 5-20 IBM System i Operations Console

You will reach the window shown in Figure 5-21 on page 237.

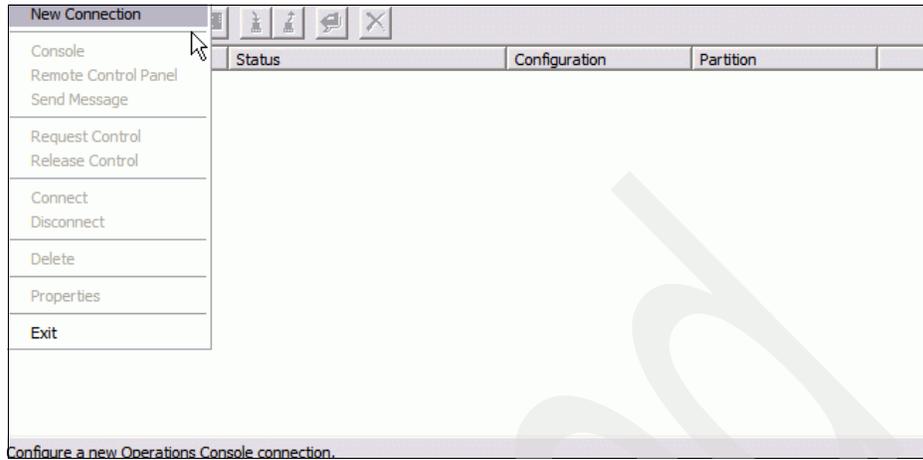


Figure 5-21 IBM System I operator console

2. Click **New Connection** to continue.

You reach the Operation Console Configuration wizard Welcome screen, as shown in Figure 5-22 on page 238.

A connection to the Internet is required to reach the InfoCenter services.

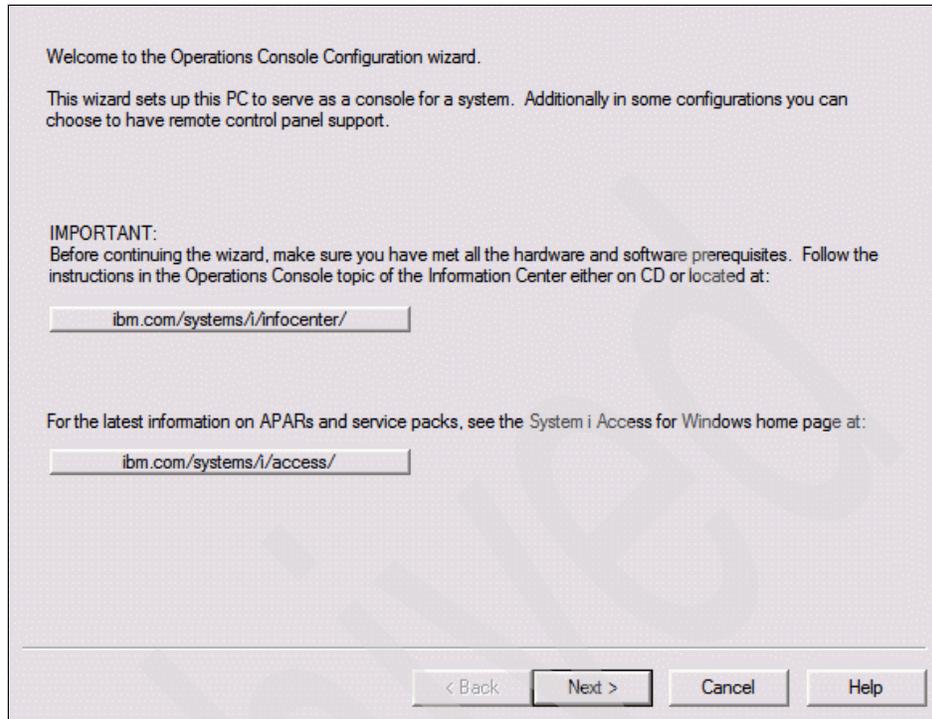


Figure 5-22 IBM System i Operations Console Welcome

Click **Next>** to continue.

3. 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 5-23 on page 239.

In our case we selected the method **Local console on a network (LAN)**.

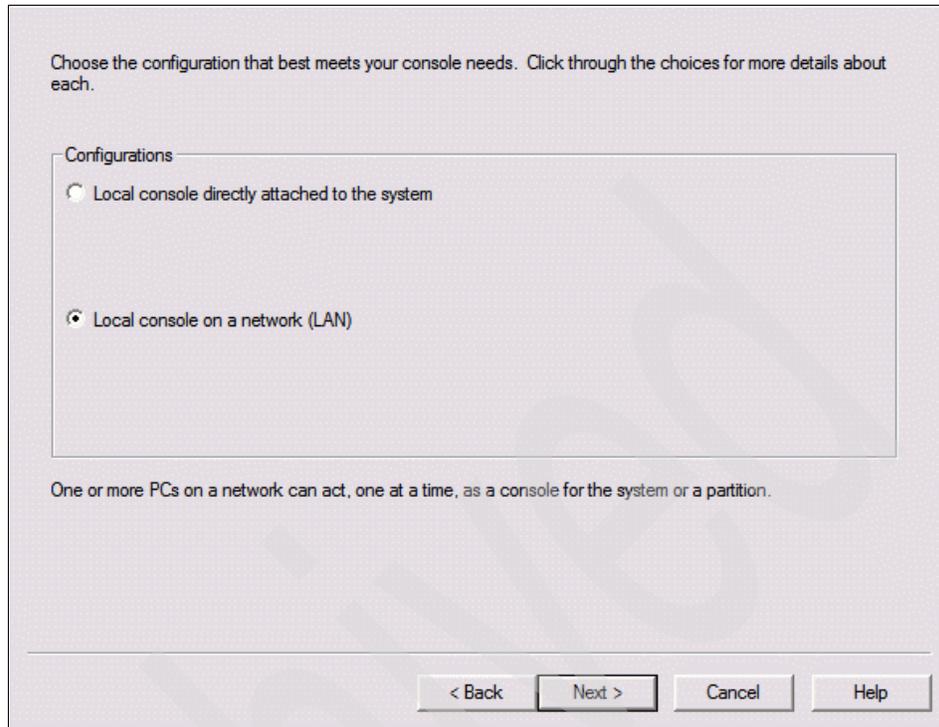


Figure 5-23 IBM System i Operations Console - choose a configuration

Click **Next>** to continue.

4. The System i service hostname must be defined first to establish a connection to the BladeCenter JS22 blade; see Figure 5-24 on page 240.

To configure an Operations Console connection, you must specify the System i service host name. If the connection has never been assigned or configured before, you will be assigning the name and network data in this wizard.

For an existing configuration, match the existing service host name defined in Dedicated Service Tools (DST). To find the service host name, go into DST and use the Configuration Service Tools Adapter display.

Note: The service host name is the name that identifies the card used for Operations Console. It may not be the TCP/IP address of any existing configurations in System i/OS or name of the PC.

What is the service host name of the system you are connecting to?

Service host name:

Service TCP/IP Address:

< Back Next > Cancel Help

Figure 5-24 IBM System i Operations Console - enter the Service host name

Enter the service host name and click **Next>**.

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. This is assigned by your system or network administrator and must be resolved through DNS.

You need 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 may have a non-networked console, having a remote control panel to a secondary partition may be desirable.

5. If this is a first-time connection, the console will respond as shown in Figure 5-25 on page 241.

If the Service host name is already defined in the host file of the console PC, then the window will resolve the IP address automatically as shown in Figure 5-26 on page 242.

Note: Choose a service host name that is related to the IBM i 6.1 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 in 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

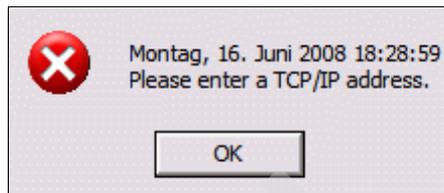


Figure 5-25 IBM System i operator console

Click **OK** to continue.

6. Enter the service TCP/IP address in the field shown in Figure 5-26 on page 242.

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.

To configure an Operations Console connection, you must specify the System i service host name. If the connection has never been assigned or configured before, you will be assigning the name and network data in this wizard.

For an existing configuration, match the existing service host name defined in Dedicated Service Tools (DST). To find the service host name, go into DST and use the Configuration Service Tools Adapter display.

Note: The service host name is the name that identifies the card used for Operations Console. It may not be the TCP/IP address of any existing configurations in System i/OS or name of the PC.

What is the service host name of the system you are connecting to?

Service host name:

Service TCP/IP Address:

< Back Next > Cancel Help

Figure 5-26 IBM System i Operations Console - enter the Service TCP/IP Address

Click **Next>** to continue.

7. 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, as shown in Figure 5-27 on page 243.

System serial number This is the BladeCenter JS12 unique system number. To find the System serial number, use the Integrated Virtualization Management (IVM) console and look under System Properties.

Target partition This is the Target partition is the partition ID of the IBM i 6.1 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 6.1 partition ID is 2.

Enter the appropriate values and click **Next>** to continue.

What is the TCP/IP information and serial number of the system to which you are making a service connection?

Service host name:	V6R1M0
Service TCP/IP Address:	172.16.1.61
Service subnet mask:	<input type="text" value="255.255.255.0"/>
Service gateway address 1:	<input type="text" value="172.16.1.1"/>
Service gateway address 2 (optional):	<input type="text"/>
System serial number:	<input type="text" value="1002D1A"/>
Target partition:	<input type="text" value="2"/>

< Back Next > Cancel Help

Figure 5-27 IBM System i Operations Console - enter System serial number

Enter values and click **Next>** to continue.

8. 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 5-28 on page 244.

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.

Service tools user IDs are created through DST or SST and are separate from IBM i 6.1 user profiles. 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>

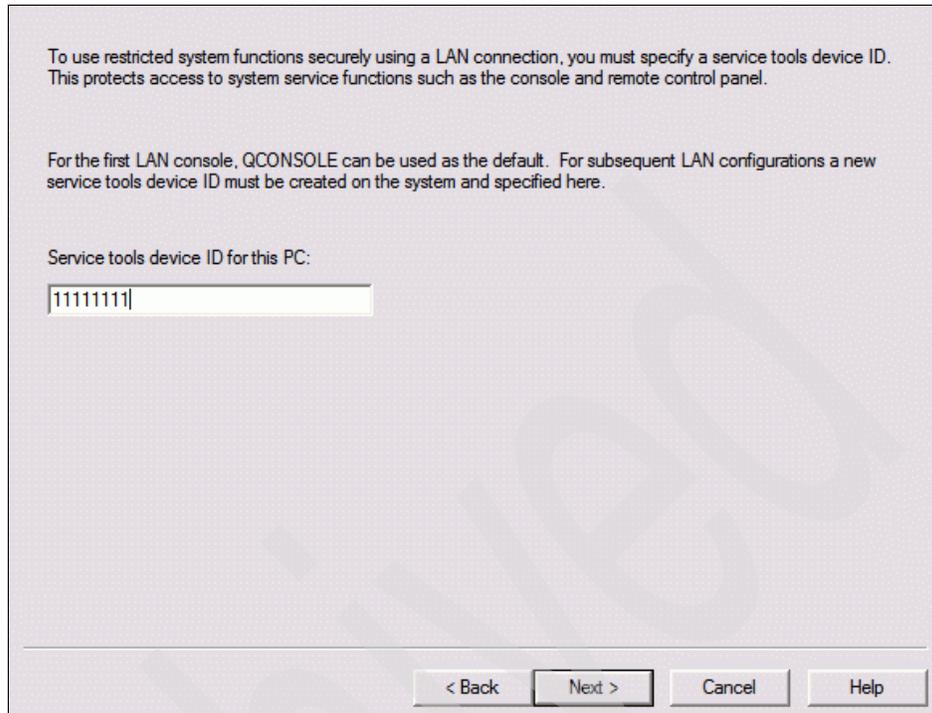


Figure 5-28 IBM System i Operations Console - enter Service tools device ID

Enter the **Service tool device ID** and click **Next>** to continue.

9. Figure 5-29 on page 245 shows the final window that is displayed after you define the recommended information for an IBM System i Operations Console.

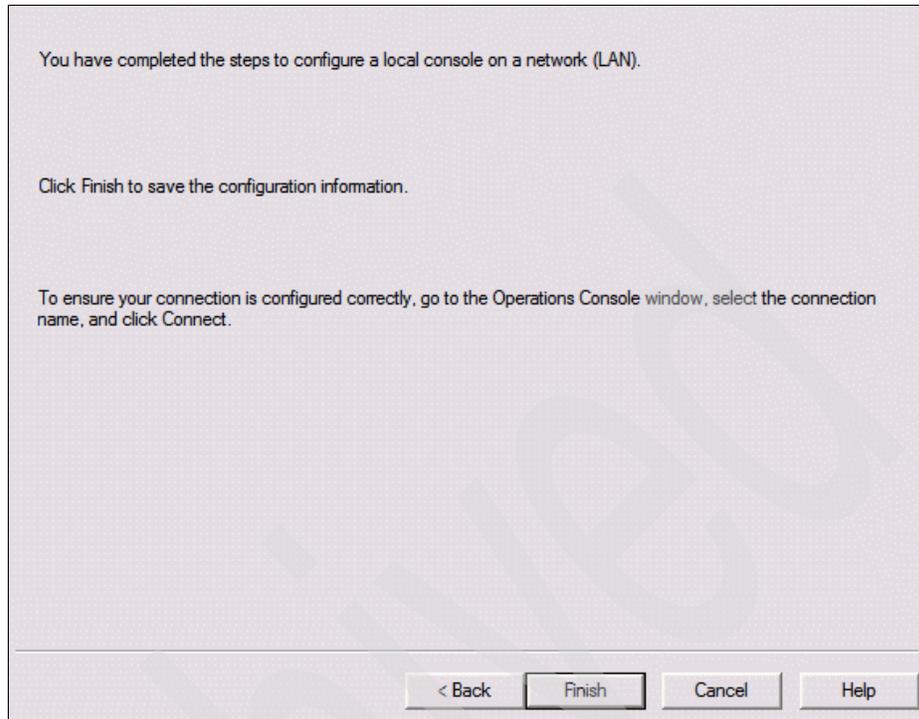


Figure 5-29 IBM System i Operations Console - finalizing the setup

Click **Finish** to save the configuration information.

The configuration window will close immediately and you will return to the initial window with the predefined console definitions for a BladeCenter JS12 blade, as shown in Figure 5-30 on page 246.

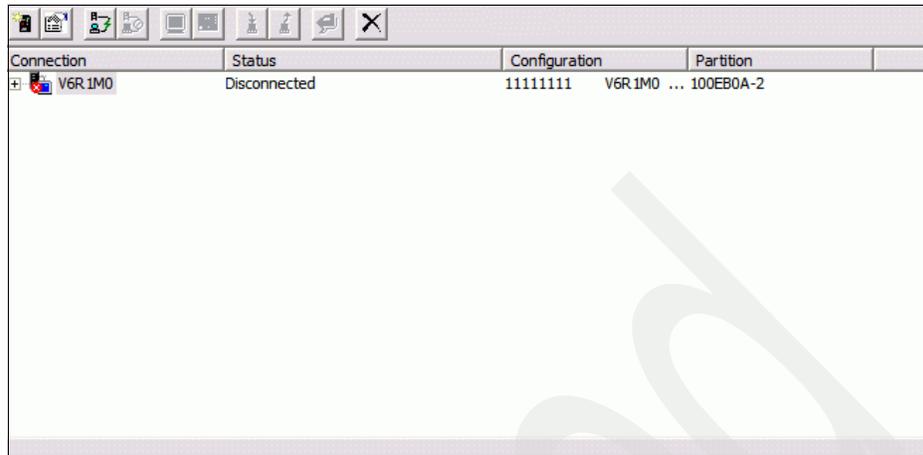


Figure 5-30 IBM System i Operations Console

To connect the IBM System i Operations Console to the IBM i 6.1 partition, complete the tasks described in the following section.

5.3 Creating an IBM i 6.1 partition

Using Integrated Virtualization Manager (IVM) to create an IBM i 6.1 partition is similar to using the HMC. IVM uses a number of defaults that simplify partition creation. For example, because IBM i 6.1 partitions cannot own physical hardware on an IVM-managed system such as a BladeCenter JS12, those screens are omitted from the creation wizard.

IVM defaults the load source and alternate IPL adapters to the Virtual SCSI client adapter in the IBM i 6.1 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.

5.3.1 IBM i 6.1 minimum requirements

The minimum recommended amount of memory for an IBM i 6.1 client partition on the BladeCenter JS12 blade is 1 GB. If you running two or more IBM i 6.1 partitions on the same BladeCenter JS12 blade, consider increasing the VIO Server memory to 2 GB.

The actual memory and CPU values should be sized individually for each IBM i 6.1 workload using the IBM Workload Estimator, which is available at:

<http://www.ibm.com/systems/support/tools/estimator/index.html>

Note: Consider using 512 MB of free memory for the system firmware or POWER Hypervisor for the BladeCenter JS12 blade.

5.3.2 VIO Server configuration

For a detailed explanation of how to create a partition using the Integrated Virtualization Manager (IVM), refer to 3.7, “Partition configuration for Virtual I/O Client (VIOC)” on page 118.

This section provides a brief explanation of how to create an IBM i 6.1 partition.

5.3.3 Creating an IBM i partition

To create an IBM i 6.1 partition using the Integrated Virtualization Manager (IVM), follow these steps:

1. Click **View/Modify Partitions**.
2. Click **Create Partition**. The next available Partition ID is preselected.
 - a. Enter a name for the partition in the Partition name field.
 - b. Select **IBM i** or **i5/OS** in the Environment field.
 - c. Click **Next** to continue.
3. Define the Assigned memory value. Available memory for the partition will be shown above the Assigned memory field.
 - a. Click **Next** to continue.
4. Select the desired processor configuration.
 - a. Click **Next** to continue.
5. Specify the VLAN used for bridging.
 - a. Click **Next>** to continue.
6. Select **Assign existing virtual disks and physical volumes**.
 - a. Click **Next** to continue.
7. Depending on the BladeCenter chassis configuration in the storage window, you can select either:
 - a. A LUN that is already configured for the IBM i 6.1 partition from the list of **Available Physical Volumes**.

- b. Or an entire disk provided from the VIO Server to the IBM i 6.1 partition from the list **Available Virtual Disks**.
 - c. Click **Next** to continue.
8. Also depending on the installation preparation in the Optical devices menu, you can select either:
 - a. The USB DVD drive to the IBM i 6.1 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 6.1 partition on the same BladeCenter JS12 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.
9. Review the summary of your definition.
 - a. Click **Finish** to create the IBM i 6.1 partition.

Load Source and Console Identification

During the creation phase of an IBM i 6.1 partition, a different section will be visible after selecting i5OS as the target operating system in the Environment field.

Figure 5-31 on page 249 shows the section Load Source and Console Identification in the General tab view of the IBM i partition. These are predefined fields for the load-source input/output processors (IOPs), the load-source I/O adapter, and for either System Service Tool (SST) display or the Dedicated Device Tools (DST) display.

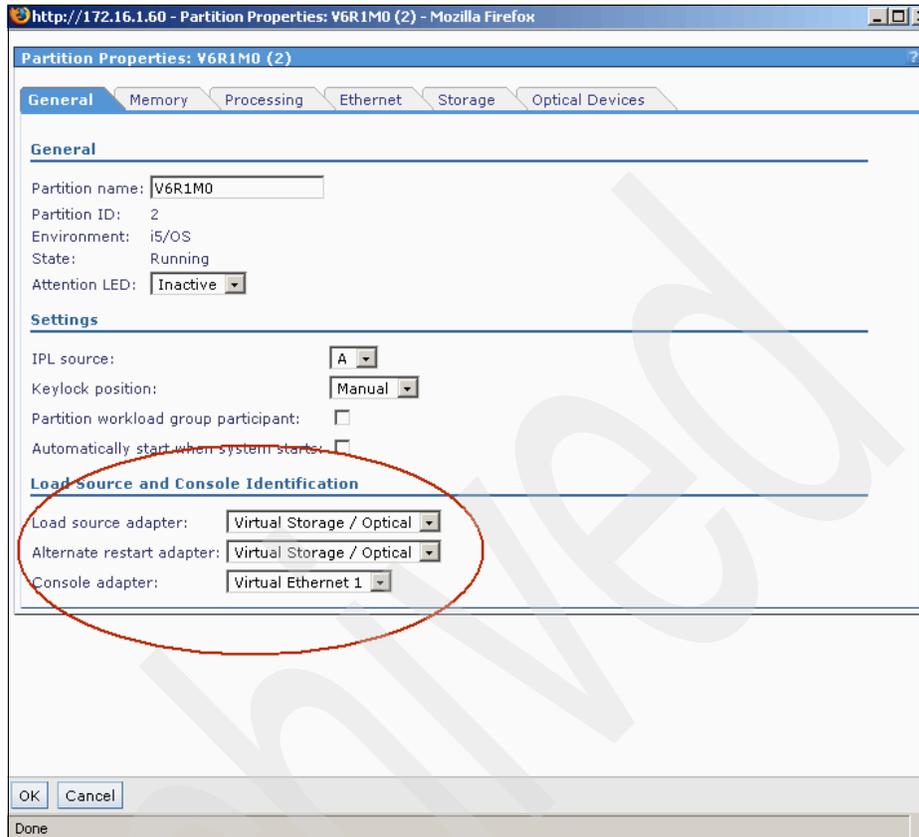


Figure 5-31 Load Source and Console Identification fields

For a first-time installation of IBM i 6.1, to ensure a complete and successful installation do not change the parameters.

5.3.4 IBM i 6.1 install media preparation

There are two general methods for installing IBM i Licensed Internal Code (LIC) and the 6.1 operating system on a BladeCenter JS12 blade in an IBM BladeCenter chassis. You can use the CD/DVD drive in the IBM BladeCenter chassis Media module attached to the IBM i 6.1 partition, or you can 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 need 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 3.6.4, “Optical devices” on page 108.

5.3.5 IBM i 6.1 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 <i>not</i> 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 6.1 partition, use IPL Type D to install and set up the environment. After the LIC is successfully installed, the installation process will automatically change the IPL type to A.

After the prerequisites are completed, the steps required to install 6.1 on a BladeCenter JS22 are essentially the same as on any other supported system:

1. Place the IBM i 6.1 installation media in the DVD drive in the BladeCenter media tray, which at this point should be assigned to your BladeCenter JS12. Or use the previous created Virtual Optical Media device described in 5.3.4, “IBM i 6.1 install media preparation” on page 249.
2. In IVM, select the IBM i 6.1 partition and click **Activate**.

3. After connecting the IBM System i for Windows LAN console, the screen shown in Figure 5-32 appears.

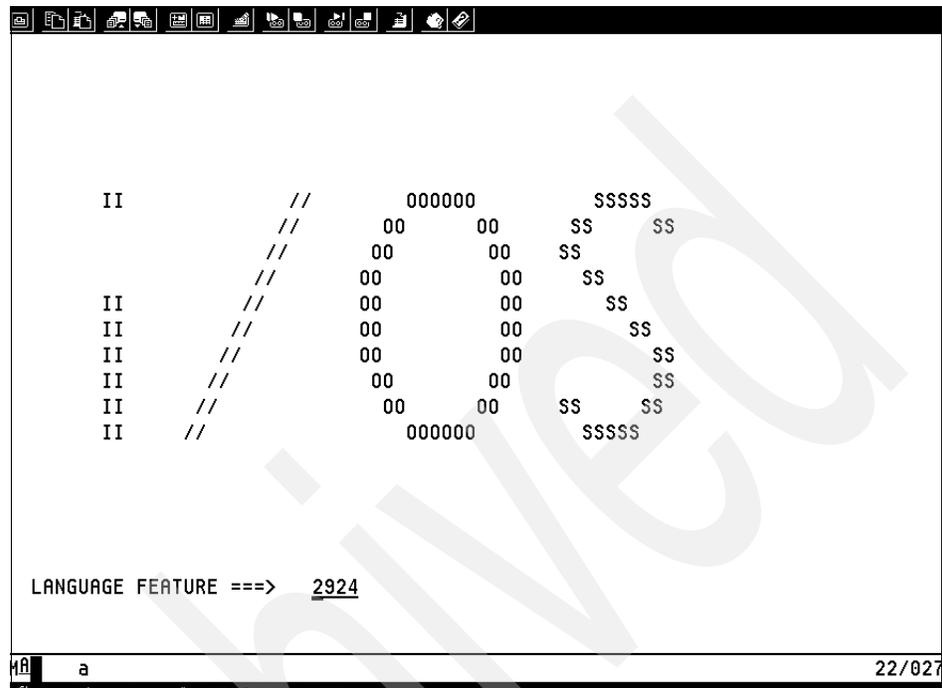


Figure 5-32 LIC initial installation screen

Press **Enter** to continue.

4. Depending on the native language, a selection can be made in the following screen as shown in Figure 5-33 on page 252.

Normally the same language will be chosen as the language for the IBM i 6.1 operating system. Language feature 2924 enables the English environment.

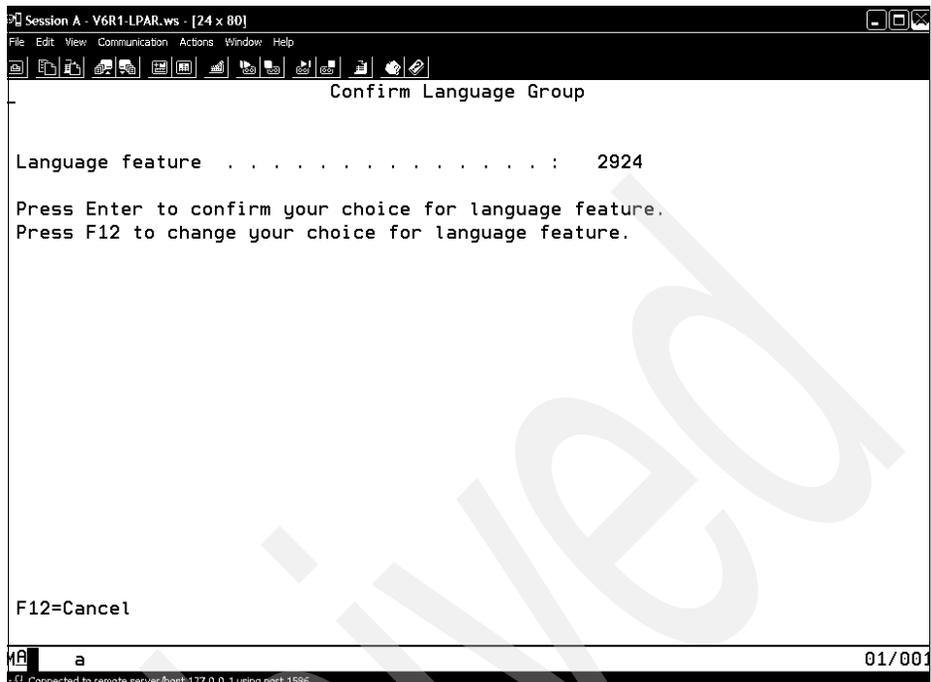


Figure 5-33 Confirm Language setup

Press **Enter** to continue.

The next screen displays several options, as shown in Figure 5-34 on page 253. To install the Licensed Internal Code, type 1 and press Enter.

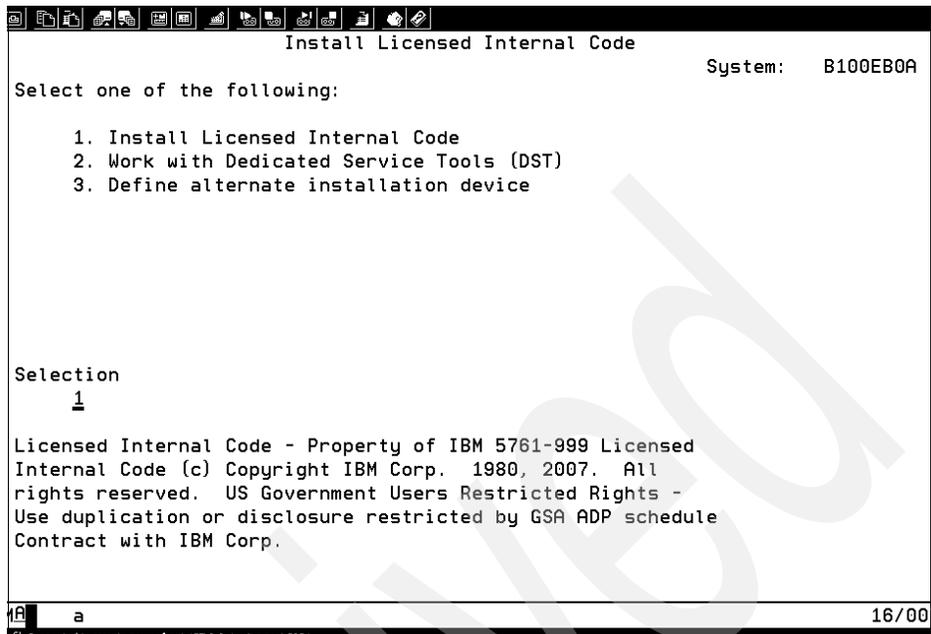


Figure 5-34 Install LIC

5. Now select the target install device. Move the cursor to the target device, type 1 and press **Enter**; see Figure 5-35 on page 254.

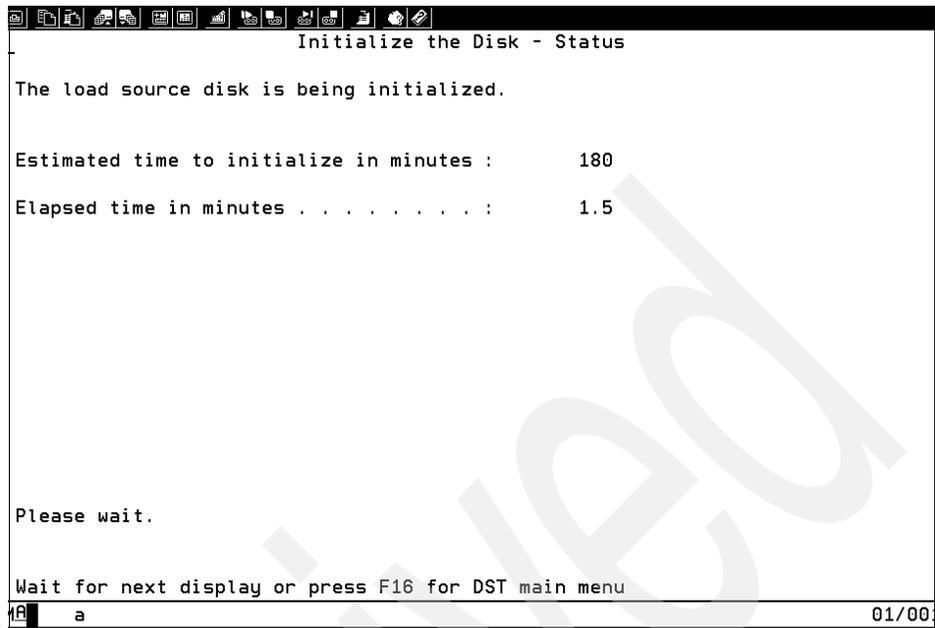


Figure 5-39 Initialize disk

9. Next, the Install Licensed Internal Code status display appears on the console as shown in Figure 5-40 on page 259. It will remain on the console for approximately 30 minutes.

The logical partition is automatically restarted at this time to complete the Licensed Internal Code installation.

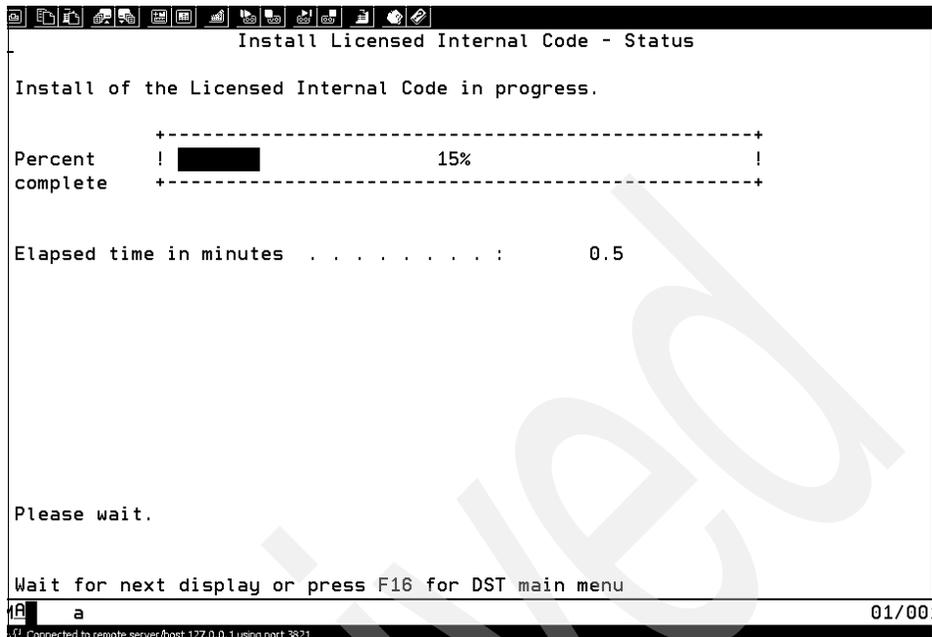


Figure 5-40 Install LIC status

10. The Disk Configuration Attention Report display might appear on the console. Figure 5-41 on page 260 shows the report for a new disc 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 InfoCenter:

<https://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzajy/rzajoverview.html>

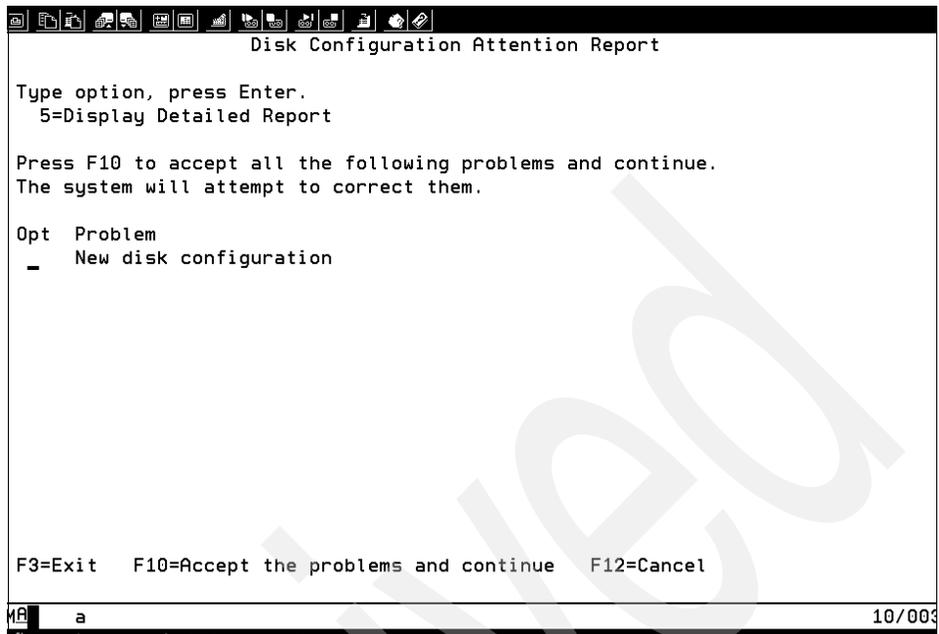


Figure 5-41 Attention Report

After the Licensed internal Code installation is complete, you will see the screen shown in Figure 5-42 on page 261. To continue the installation process for IBM i, you can select option 2 to install the IBM i operating system after an IPL.

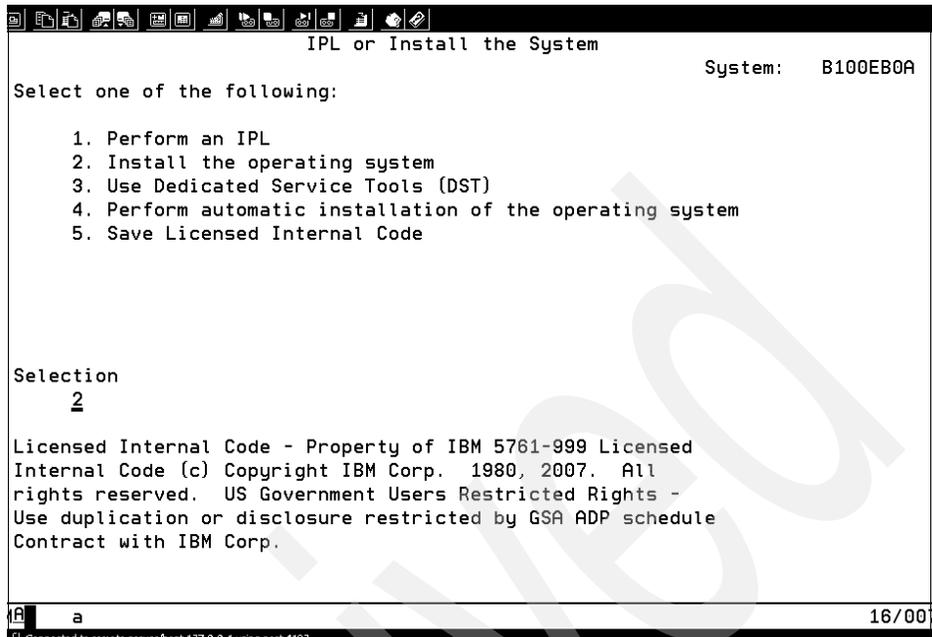


Figure 5-42 Install the operating system

5.4 Installing the IBM i 6.1 Operating System

From the IPL or Install the System screen, 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 6.1 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.

1. On the IPL or Install the System screen on the console, type the number that corresponds to the correct device type as shown in Figure 5-43 on page 262. 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.

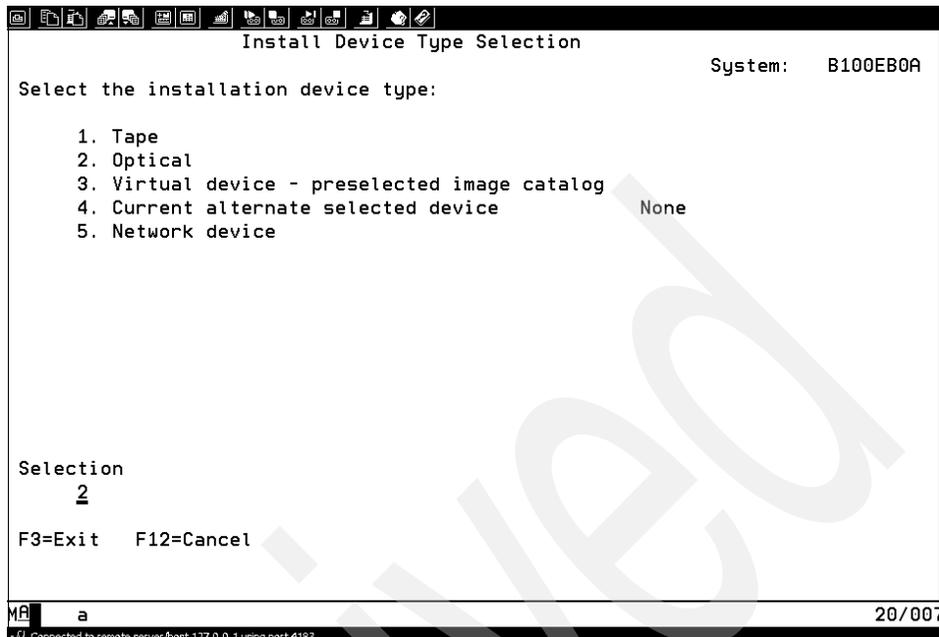


Figure 5-43 Select install device

Type **2** and press **Enter** to continue.

2. The Confirm Install of the Operating System screen is displayed on the console screen, as shown in Figure 5-44 on page 263.

Press **Enter** to continue the installation process.

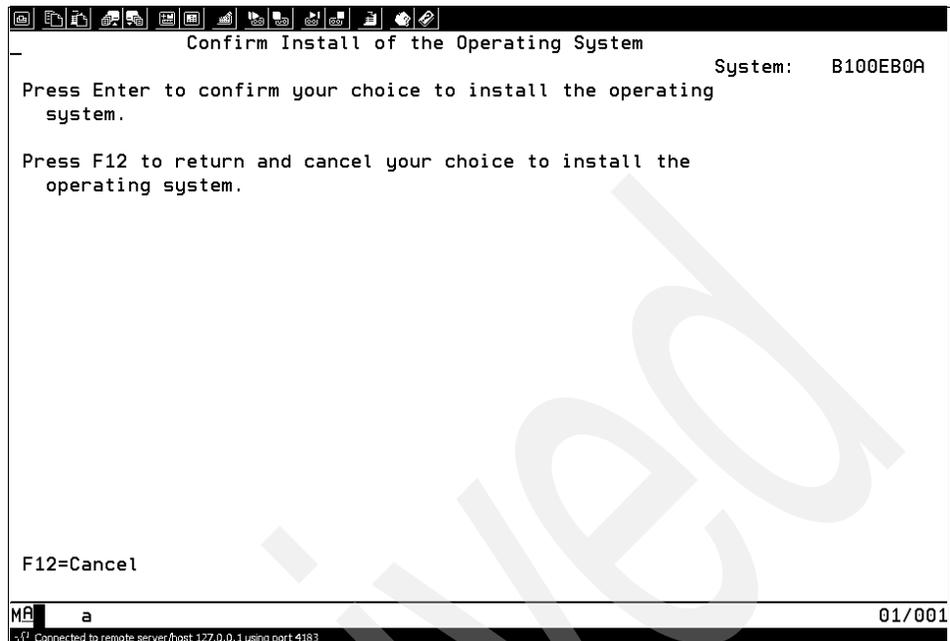


Figure 5-44 Confirm installation

3. The Select a Language Group screen displays the primary language preselection, as shown in Figure 5-45 on page 264. This value should match the Language feature number that is printed in the installation media.

The following URL provides the Language feature codes:

<http://publib.boulder.ibm.com/infocenter/scope/i5os/topic/rzahc/rzahcnlvfeaturecodes.htm#rzhcnlvfeaturecodes>

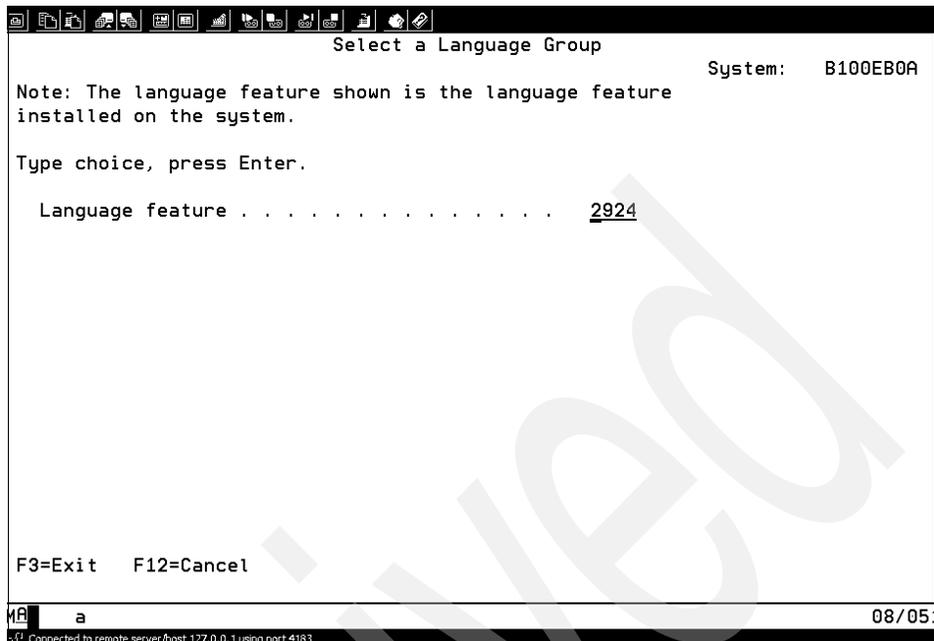


Figure 5-45 Select language feature

4. Type choice and press **Enter** to continue.

The Confirm Language Feature Selection appears on the console, as shown in Figure 5-46 on page 265.

Press **Enter** to confirm and continue.

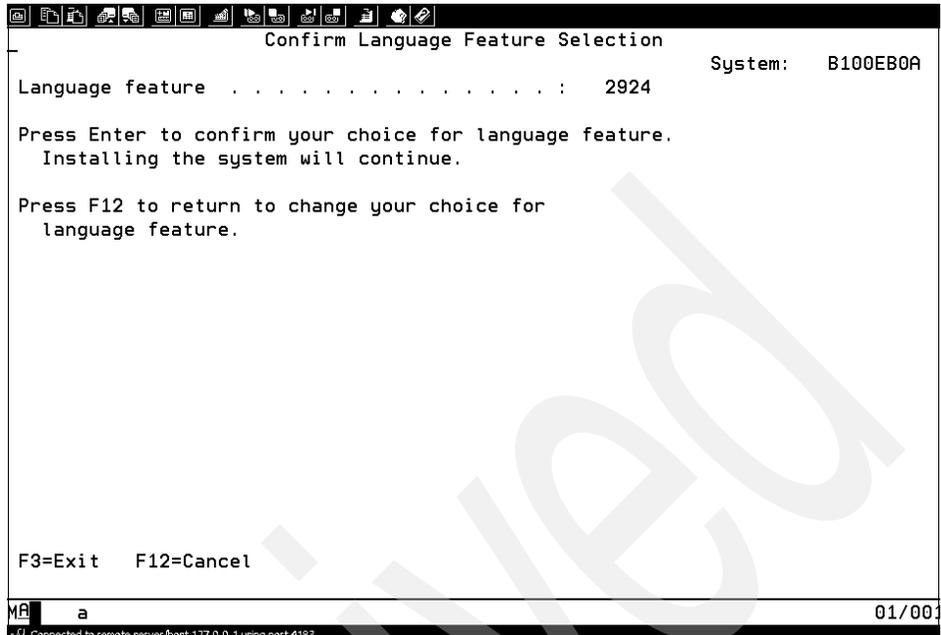


Figure 5-46 LIC install confirm language

5. The Licensed Internal Code IPL in Progress screen appears on the console, as shown in Figure 5-47 on page 266. 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

```

Figure 5-47 IPL in progress

The Install the Operating System screen appears on the console, as shown in Figure 5-48 on page 267.

6. Change the date and time values to the appropriate settings. You must use the 24-hour clock format to set the current time.

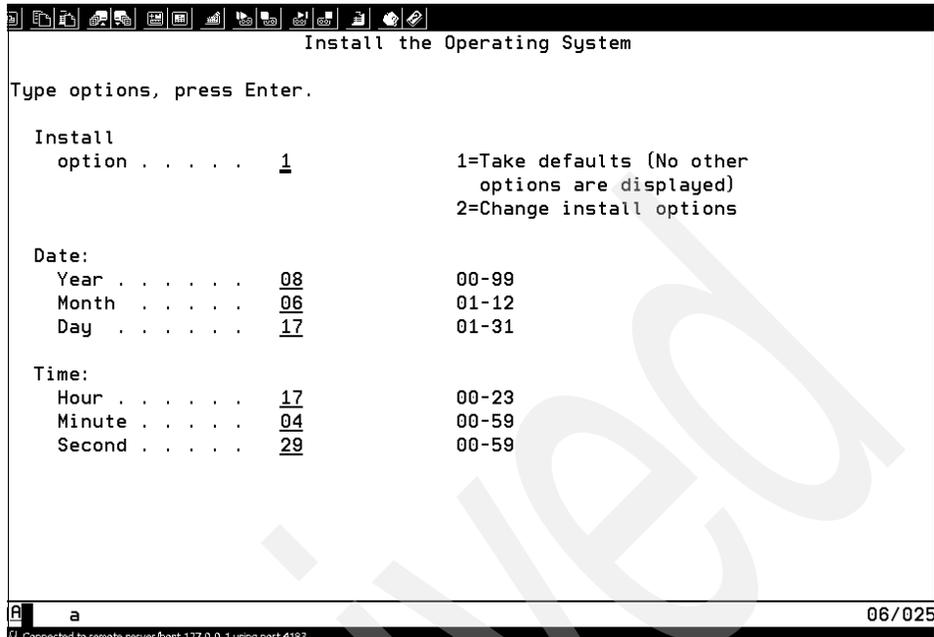


Figure 5-48 Set date and time

7. Figure 5-49 on page 268 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 a while between Installation Stage 4 and 5.

```

Session A - LIVIO.ws - [24 x 80]
File Edit View Communication Actions Window Help
Message ID . . . : CPI2070                               i5/OS Installation Status
Stage 2 ! ██████████ 17% !
Installation
Stage          Completed      Objects Restored
  1 Creating needed profiles and libraries . . . . . : X
>> 2 Restoring programs to library QSYS . . . . . :      02272
  3 Restoring language objects to library QSYS . . . :
  4 Updating program table . . . . . :
  5 Installing database files . . . . . :
  6 Installing base directory objects. . . . . :
i5/OS is a trademark of IBM in the United States and other countries.
17/068

```

Figure 5-49 Installation status

8. When the Sign On screen is displayed, as shown in Figure 5-50 on page 269, the base installation of the IBM i 6.1 Operating System is finished.

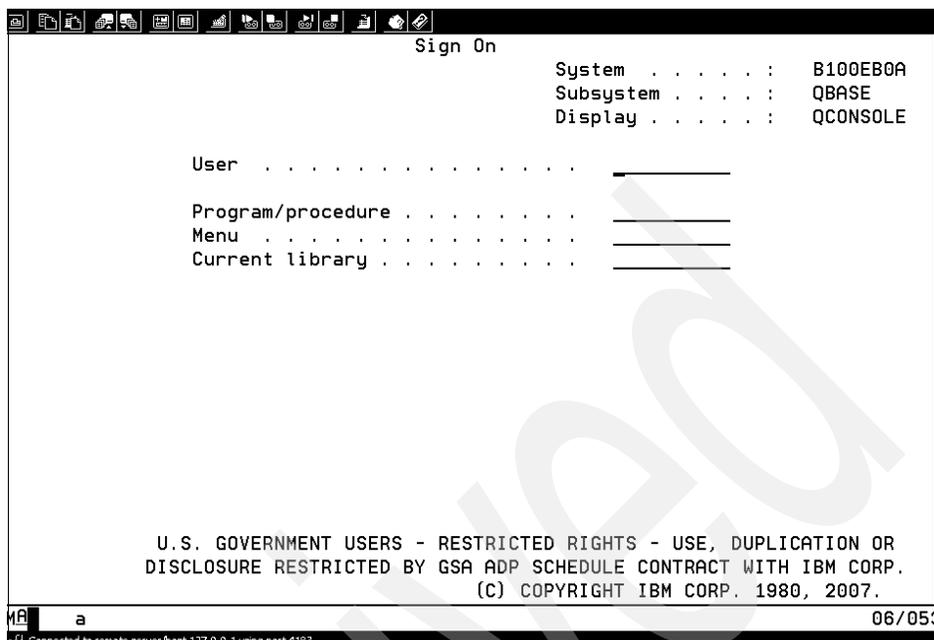


Figure 5-50 Sign On screen

At this stage, the IBM i 6.1 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.bolder.ibm.com/infocenter/iserics/v5r3/topic/rzahc/rzahcs_winstallprocess.htm

5.4.1 IBM i 6.1 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 recommended PTFs:

http://www-912.ibm.com/s_dir/slkbases/recommendedfixes

Refer to the primary Web site for downloading fixes for all operating systems and applications:

<http://www-912.ibm.com/eserver/support/fixes>

For IBM i 6.1, IBM i5/OS, or OS/400® Operating Systems, fixes are available. To Obtain an IBM i 6.1 fix overview for downloading:

- ▶ Select **System i** in the Product family field.
- ▶ Select **IBM i, i5/OS, and OS/400** in the Product field.
- ▶ Select one the following options in the Ordering option field:
 - Groups, Hyper, Cumulative fixes
 - Individual fixes
 - Search for fixes
- ▶ 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 need a valid IBM ID and password.

5.4.2 IBM i 6.1 TCP/IP setup

Ethernet on a BladeCenter JS12 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 JS12 server. Refer to the Ethernet on System i 6.1 topic in the IBM Information Center to configure the IBM i 6.1 communication:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzajy/rzajyoverview.html>

Note: IBM i 6.1 installed on a IBM BladeCenter JS12 server will communicate to the external LAN using Virtual Ethernet Adapter only, and only those with the Virtual Ethernet bridge enabled.

5.4.3 IBM i 6.1 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 6.1 management and will 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 JS12 server or Power Systems, as shown in Figure 5-51.

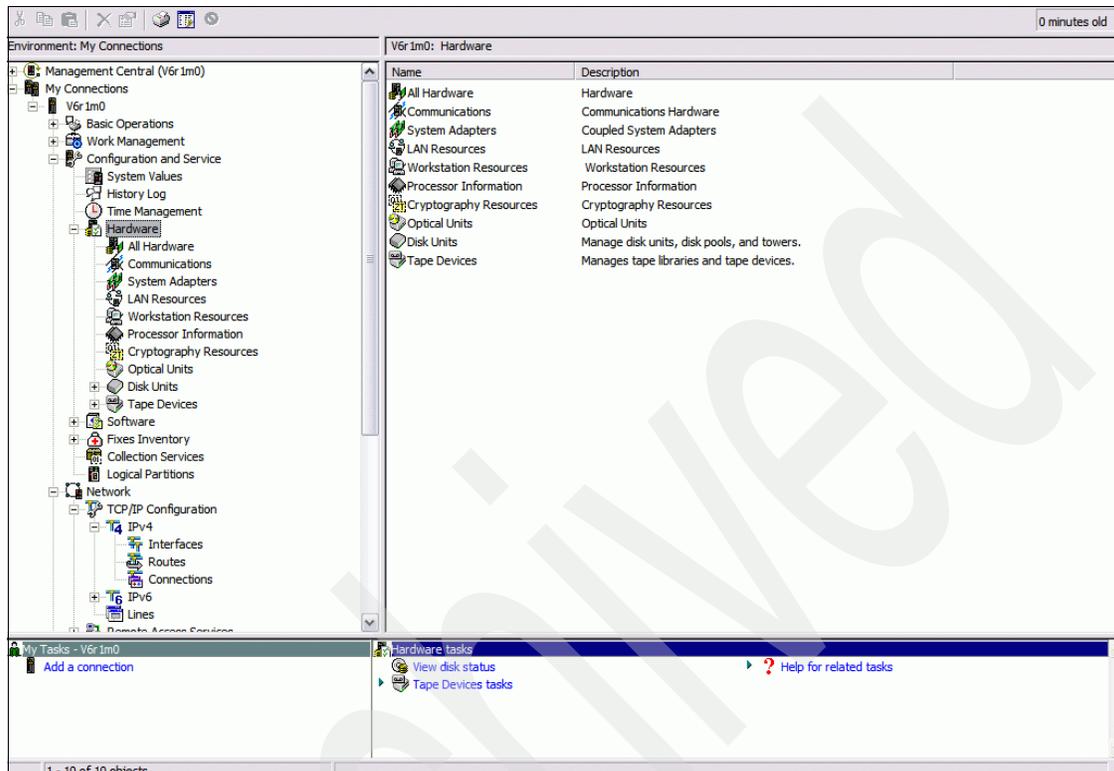


Figure 5-51 IBM Systems Navigator for i

More detailed information to the IBM Systems Director Navigator for i functionality can be found at:

<http://www.ibm.com/systems/i/software/navigator/index.html>

or in *Managing OS/400 with Operations Navigator V5R1 Volume 1: Overview and More*, SG24-6226.

5.5 IBM i 6.1 Backup/Restore

IBM i 6.1 installed on a BladeCenter JS12 server in a partition in VIOS does not have direct access to a physical device, like a tape drive for backup and restore. This changes the save/restore process as a whole, although the steps performed within 6.1 remain the same for the most part.

To back up and restore an IBM i 6.1 partition on a IBM BladeCenter JS12 blade as described here, you need the following hardware:

- ▶ IBM SAS LTO tape drive attached to a SAS Switch Module

Note: For further information about using the IBM SAS LTO tape library attached to a BladeCenter S chassis, refer to:

http://www.ibm.com/systems/power/hardware/blades/i_on_blade_readme.html

IBM i 6.1 backup

Performing an IBM i 6.1 Operating System backup is a 2-stage process on the IBM BladeCenter JS12 blade:

1. A standard 6.1 **save** command or 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.
2. The virtual media image, containing the 6.1 save, is written out to a SAS attached tape drive using a VIOS command.

IBM i 6.1 restore

Performing a restore follows the same 2-stage process in reverse:

1. The virtual media image file is restored from the SAS tape drive onto VIOS disk using a VIOS command. The image file is then mounted on the correct virtual optical drive assigned to the IBM i 6.1 partition and becomes available as a volume from which to restore.
2. A standard IBM i 6.1 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.

Important: The virtualized DVD-ROM drive in the chassis *cannot* be used for IBM i 6.1 backups, because it is not writable.

5.5.1 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.

1. Logon the Integrated Virtualization Manager on the BladeCenter JS22.
2. Select **View/Modify Virtual Storage**, the **Storage Pool** tab, and then click **Create Storage Pool**. The window displayed in Figure 5-52 is shown.

Note: We recommend using a separate storage pool. However, rootvg can be used.

http://172.16.1.60 - Create Storage Pool - Mozilla Firefox

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 5-52 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 click the **Optical Devices** Tab and select **Create Library**.
5. Select the name of the new storage pool and enter an appropriate size for the media library. Select **OK** to continue.

- To add a new virtual optical device to the media library, select **Add Media** in the Virtual Optical Media section. The window displayed in Figure 5-53 will appear.

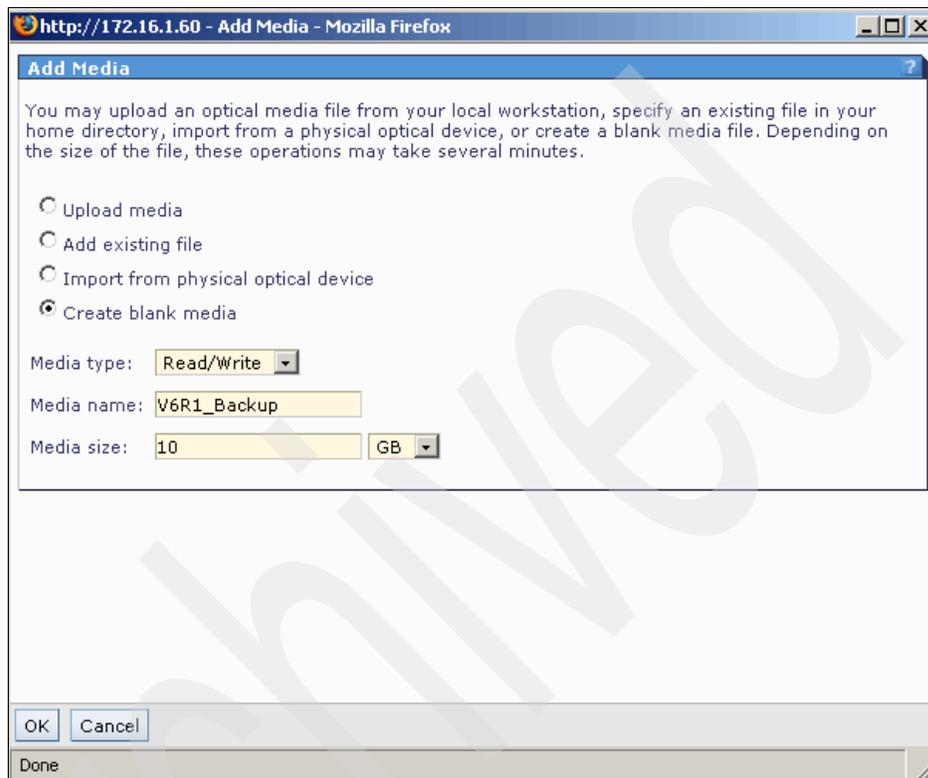


Figure 5-53 IVM Create blank media

- Select **Create blank media** and enter a meaningful Media Device name and an appropriate size for the new volume. Click **OK** to continue.
- The new virtual optical device should be listed in the Virtual Optical device list, as shown in Figure 5-54 on page 275.

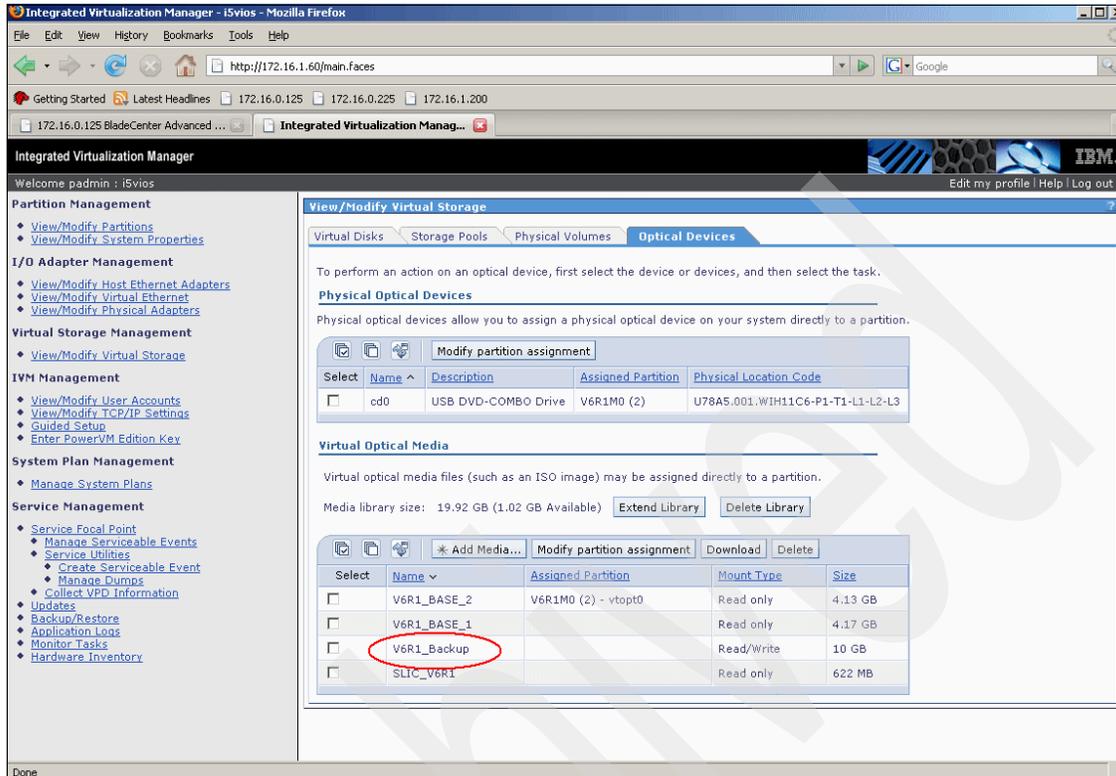


Figure 5-54 IVM Virtual optical device created

To assign the new created virtual optical device to the IBM i 6.1 partition, select the virtual optical device and click **Modify partition assignment** as shown in Figure 5-55 on page 276.

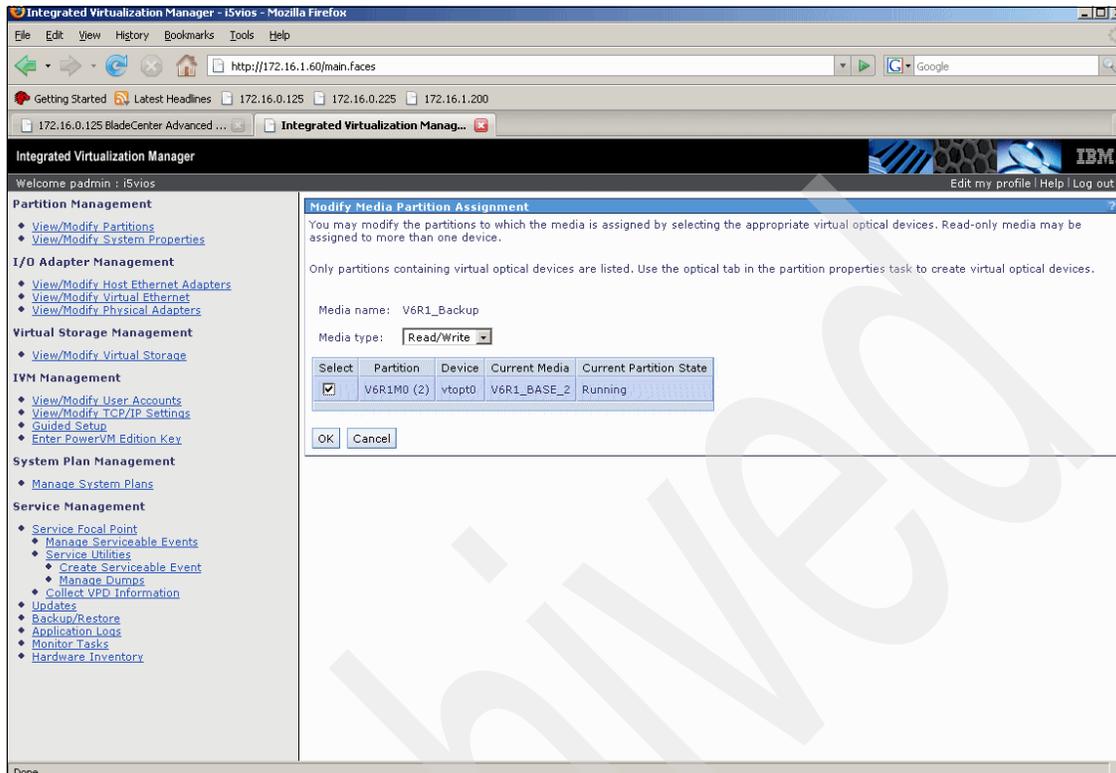


Figure 5-55 Virtual optical device assign to partition

9. Select the IBM i 6.1 partition and click **OK** to continue.

Figure 5-56 on page 277 shows the IVM Virtual Storage Management window with the current assignment of the virtual optical device to the partition.

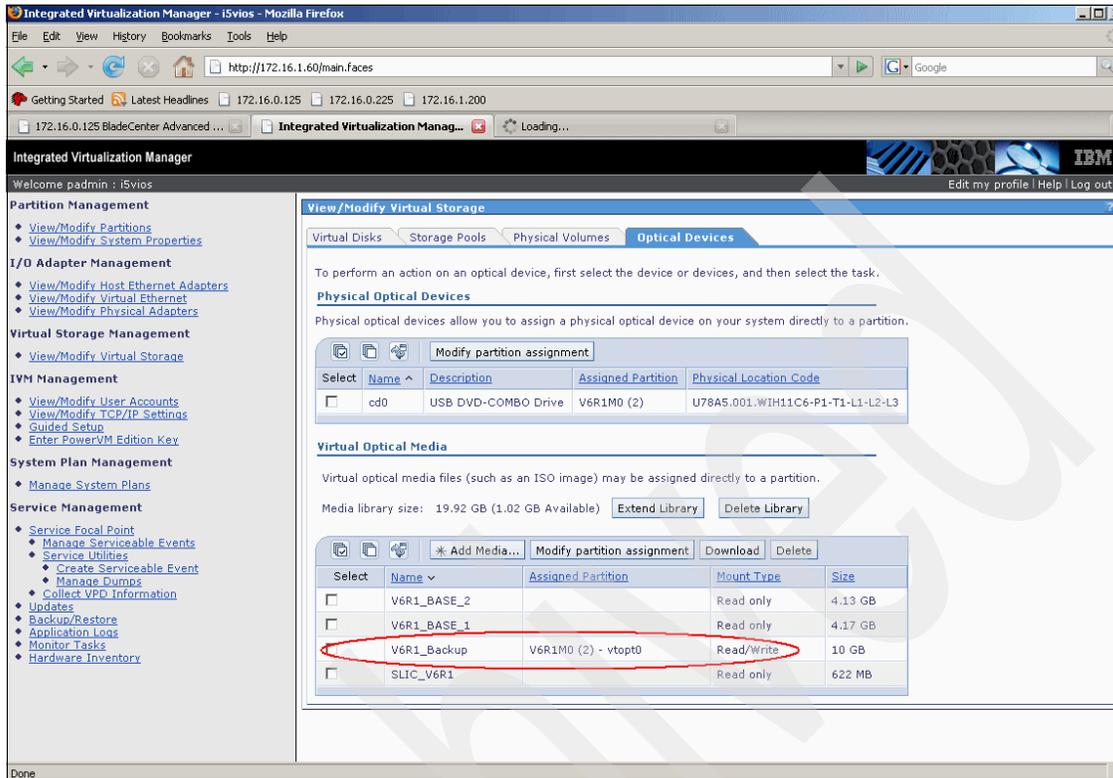


Figure 5-56 IVM Virtual optical device assignment done

After the virtual optical device is mounted to the correct virtual optical device, it will become available in the IBM i 6.1 partition. The IBM i 6.1 Operating System will not use the device name of the virtual optical device given in Integrated Virtualization Manager.

An IBM i 6.1 screen execute command **WRKOPTVOL** and the screen shown in Figure 5-57 on page 278 should appear. The virtual optical device will be identified with a time stamp volume ID.

```

Work with Optical Volumes
System: B100EB0A
Device . . . . . *ALL
Side information . . . . . : *ALL
Type options, press Enter.
  1=Add  2=Change  3=Copy  4=Remove  5=Display  8=Work with directories
 10=Initialize  11=Work with object links  12=Duplicate ...

Opt  Volume      Device      Volume      Media      Authorization
   Volume      Device      Type        Type        List
---  -
   080619122216  OPT02      *UNKNOWN    *DVD-RAM    QOPTSEC

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve  F11=View 2
F12=Cancel  F14=Show extended information  F24=More keys

MA a MW 11/00

```

Figure 5-57 Virtual optical device check device

5.5.2 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 different operating environments, including the VIO Server, on a variety of different 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 install images by default. To ensure the IBM Tivoli Storage Manager client is installed on the VIO Server, run the command shown in Example 5-1.

Example 5-1 TSM client check

```

ls|pp -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 the IBM i 6.1 save 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 *PowerVM Virtualization on IBM System p Managing and Monitoring*, SG24-7590.

5.5.3 IBM i 6.1 shutdown and restart

This section describes the shutdown and restart procedure on an IBM i 6.1 partition.

Shut down an IBM i 6.1 partition

Before you shut down an IBM i 6.1 logical partition, ensure that all jobs are completed and all applications are ended.

The correct way to shut down an IBM i 6.1 logical partition is:

- ▶ 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.

To shut down the IBM i 6.1 partition, follow these steps:

1. Open a console connection to the IBM i partition.
2. Sign in with an user ID that has the privilege to shut down the partition.
3. Execute the command **PWRDWSYS** in the command line, as shown in Figure 5-58 on page 280.

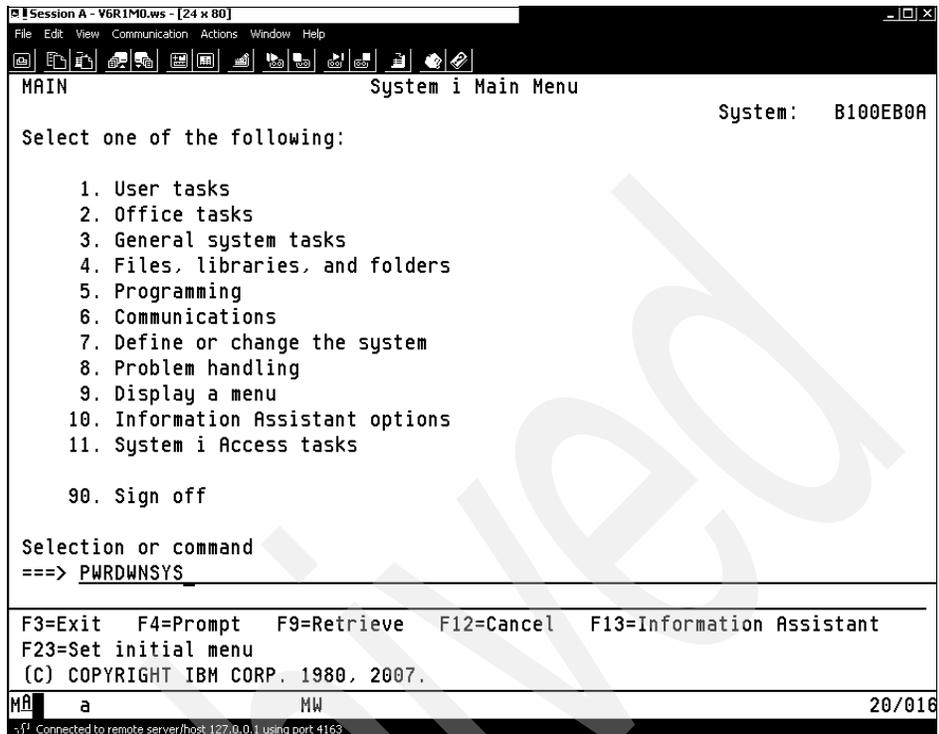


Figure 5-58 IBM i power down partition

4. Confirm the shutdown action by pressing **F16**.
5. This process can take a while. Check the Integrated Virtualization Manager (IVM) window for the message Not Activated in the State column of the IBM i partition.

Start an IBM i 6.1 partition

To activate an IBM i 6.1 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. During the boot phase, the status information in the **Reference Code** column for the IBM i partition changes from time to time. If status code C6004031 appears in the Reference code column, you can connect the System i Access for Windows operating console to the IBM i partition to finally IPL the IBM i partition. Figure 5-59 on page 281 shows the detailed Reference code to connect the operation console.

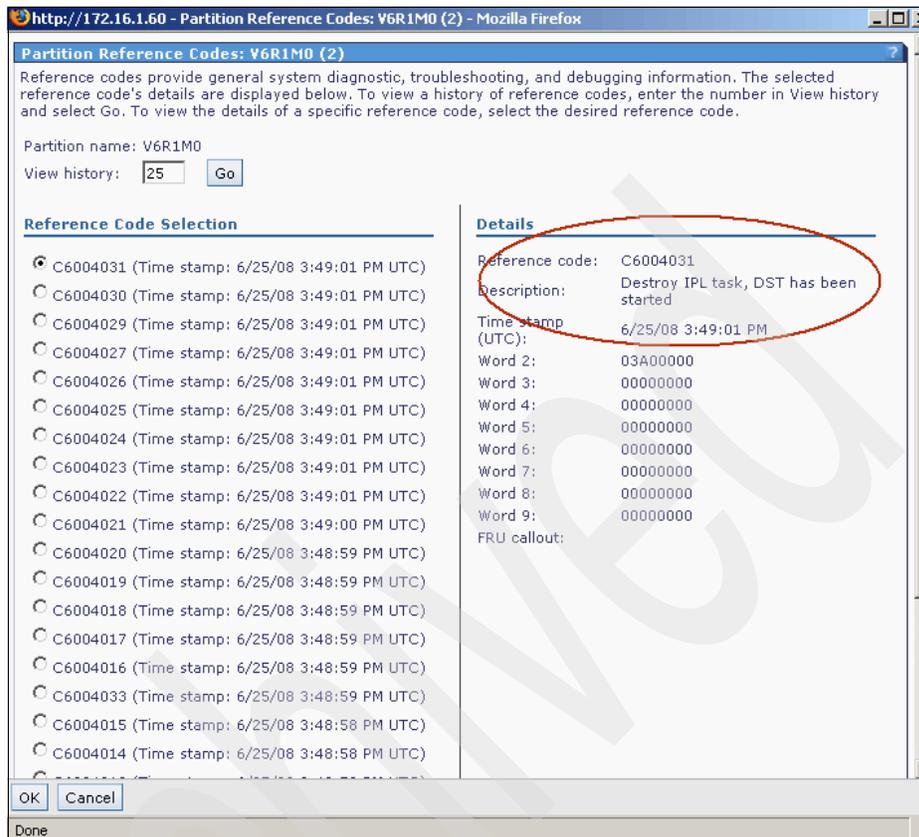


Figure 5-59 Boot reference code for console connection

After you successfully connect to the IBM i 6.1, perform the appropriate tasks to IPL the IBM i 6.1 partition.

Archived



Red Hat Enterprise Linux installation

This chapter describes the procedures to install Red Hat Enterprise Linux on a JS22 BladeCenter.

We discuss the following topics:

- ▶ Supported Red Hat operating systems
- ▶ Linux LPAR installation using DVD
- ▶ Linux network installation (detailed)
- ▶ Native Red Hat Enterprise Linux 5.2 installation
- ▶ Red Hat Enterprise Linux 5.2 automated installation
- ▶ IBM service and productivity tools

6.1 Supported Red Hat operating systems

Red Hat Enterprise Linux for POWER Version 4.6 or later and Red Hat Enterprise Linux for POWER Version 5.1 or later support installation on a JS22.

This chapter specifically covers installing Red Hat Enterprise Linux for POWER Version 5.2 with a DVD and over the network on a PowerVM logical partition (LPAR).

6.1.1 Considerations and prerequisites

There are some system configuration considerations and prerequisites prior to installing Red Hat Enterprise Linux 5.2 on a JS22 partition. They are covered here.

PowerVM LPAR considerations

By the characteristics of the virtualization features in the System p JS22 BladeCenter, the operating system and applications do not know they are running in a micro-partitioned or virtualized I/O environment. This 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 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 on virtualization and LPAR setup is available in Chapter 3, “System configuration using VIOS and IVM” on page 55.

Red Hat Enterprise Linux 5.2 installation considerations

The minimum RAM required to install Red Hat Enterprise Linux 5.2 is 1 GB, while the recommended RAM is 2 GB.

Note: Systems with less than 1 GB of RAM may experience installation process hangs or other unexpected failures.

IBM JS22 BladeCenters with 1 GB of RAM experience significant performance degradation under certain RAM-intensive workloads, so we recommend 4 GB of RAM for these machines.

In addition, ensure there is enough unpartitioned disk space or 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/5.2/html/Installation_Guide/ch11s03.html

6.2 Linux LPAR installation using DVD

With PowerVM installed and the system partitioned into LPARs using the PowerVM LPAR considerations and Red Hat Enterprise Linux 5.2 prerequisites, we are ready to install Linux on the JS22 BladeCenter LPAR.

Start an LPAR installation of Red Hat Enterprise Linux 5.2 using a DVD

The steps to start the installation of Red Hat Enterprise Linux 5.2 from a DVD are as follows:

1. On the Advanced Management Module (AMM) select **Blade Tasks** → **Remote Control**.
2. In the right side window, look for Media Tray Owner and click the drop-down arrow to display the various blade bays. Select your blade bay location with the mouse as shown in Figure 6-1.

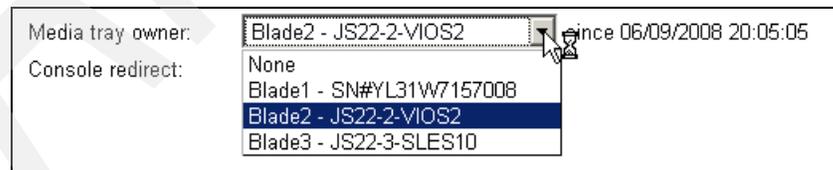


Figure 6-1 Start Remote Control window

3. Press **Refresh** on the right side.

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 pressing this button. Their MT light is on if the media tray is assigned to them.

4. Double-check that your blade bay owns the media tray by opening the AMM window and selecting **Monitors** → **System Status**. The right window will show an X in the MT column of your blade bay location. Figure 6-2 gives one example of this assignment.

Bay	Status	Name	Pwr	Owner**	
				KVM	MT*
1	●	JS22-1-VIOS1	On		X
2	●	JS22-2-VIOS2	On		
3		No blade present			

Figure 6-2 The X indicates the Media Tray (MT) is owned by the blade in bay 1

5. Place the DVD into the BladeCenter media tray.
6. Log in to the Integrated Virtualization Manager (IVM) via a Web browser.
 - a. Select one of the available PowerVM LPARs to install with Red Hat Enterprise Linux 5.2 by placing a checkmark in the Select box.
 - b. Click **Activate**, as shown in Figure 6-3.

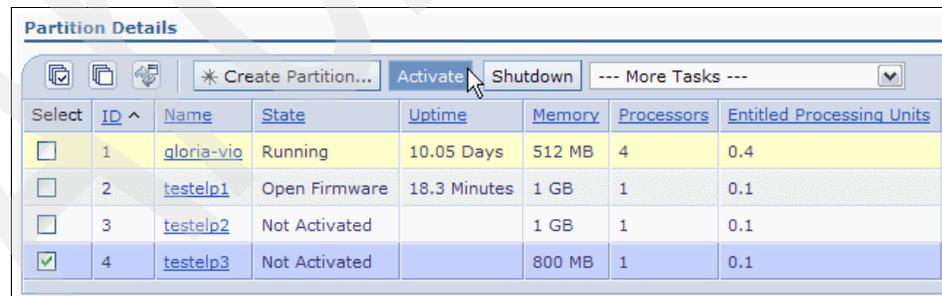


Figure 6-3 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field and select **Open terminal window**.

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 1.4.2_08, or higher.

Note: Even though this section covers installation via the Integrated Virtualization Manager (IVM) console, there are other console options available on the JS22. They are covered in Appendix A, “Consoles, SMS, and Open Firmware” on page 391.

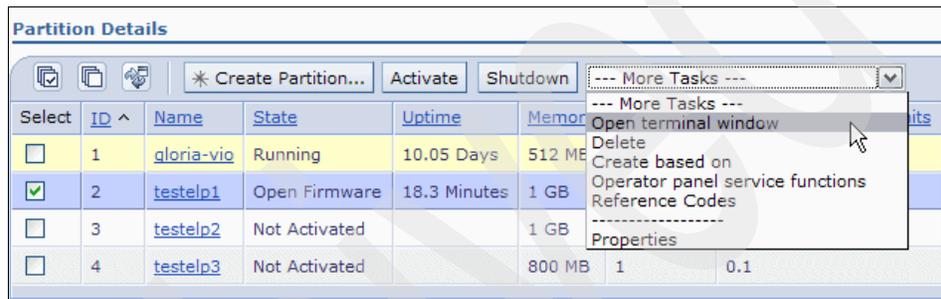


Figure 6-4 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

7. The SMS menu appears in the IVM terminal; see Figure 6-5 on page 288.

Tip: The SMS menu appears very quickly after activating the box, 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.

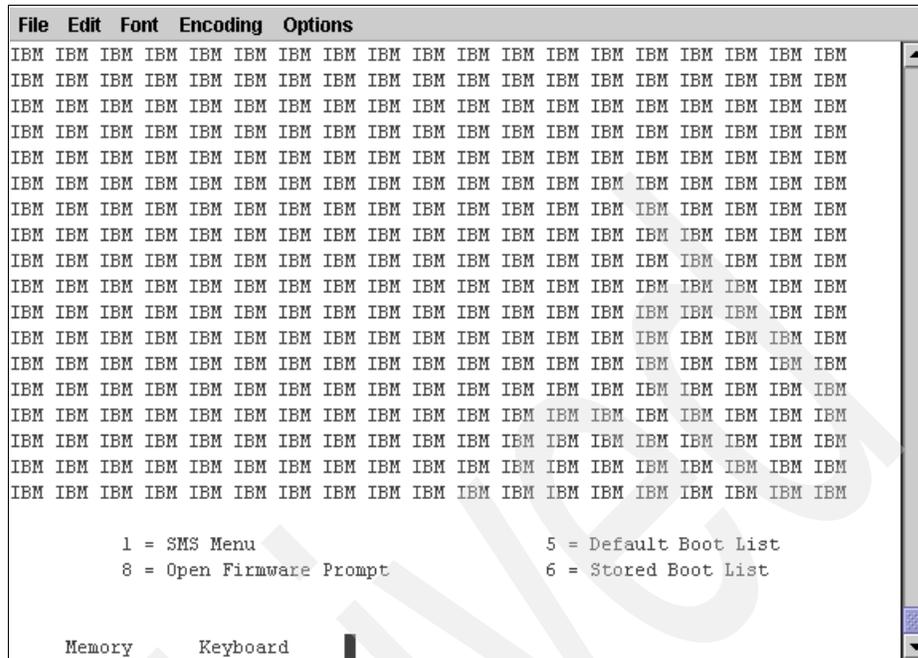


Figure 6-5 SMS menu

- a. Select 1 = SMS Menu by pressing the number 1 on the keyboard.

Tip: Press the number next to the desired system function to select and navigate through the SMS menu.

- b. Select option **5. Select Boot Options.**
- c. Choose option **1. Select Install/Boot Device.**
- d. Pick **3. CD/DVD.**
- e. Select **6. USB.**
- f. Finally, select **1. USB CD-ROM.**
- g. Choose **2. Normal Mode Boot.**
- h. Pick **1.Yes** to exit the SMS menu.
- i. At the boot : prompt press the Enter key.

At this point the system will begin reading off the DVD, which can take a few minutes.

8. After successfully booting from the DVD, the Red Hat Installer will ask if you want to perform a media check. If this is your first time installing using this

media, we highly recommend running the media check. Once the media check is complete, the Red Hat Anaconda Installer will assist with the completion of the install. More detailed installation instructions are available here:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/pt-install-info-ppc.html

6.3 Linux network installation (detailed)

This section describes a Network File System (NFS) installation on an external Storage Area Network (SAN) device which is part of a PowerVM LPAR. Even though this example uses a specific storage device and network option to complete the installation, this process is generic enough that it will work for all supported internal or external storage devices on JS22 and only a few slight modifications are needed in the Red Hat Anaconda installer setup if a different network installation option is selected. The goal of this section is to show the external SAN disk install 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 *Installing Linux using the network - General remarks of Appendix D, “Additional Linux installation configuration options”* on page 449.

This section assumes that you have already set up the NFS server properly and have read “PowerVM LPAR considerations” on page 284 and followed the installation prerequisites in “Red Hat Enterprise Linux 5.2 installation considerations” on page 284.

Start an LPAR installation of Red Hat Enterprise Linux 5.2 over the network

To start an LPAR installation of Red Hat Enterprise Linux 5.2 over the network, use the following steps:

1. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
 - a. Select one of the available PowerVM LPARs to install with Red Hat Enterprise Linux 5.2 by placing a checkmark in the Select box.

- b. Click **Activate**, as shown in Figure 6-6.

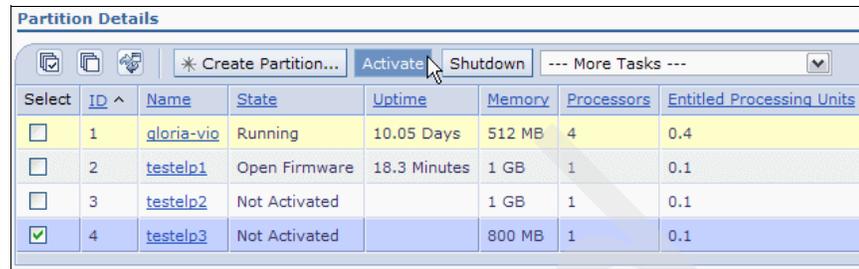


Figure 6-6 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field (Figure 6-7) and select **Open terminal window**.

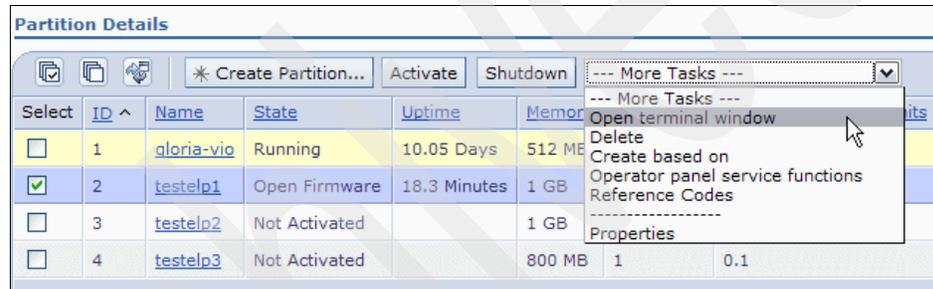


Figure 6-7 Opening a terminal window from the IVM

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 6-8 on page 291).

Tip: The SMS menu appears very quickly after activating the LPAR, 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 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 shut down.

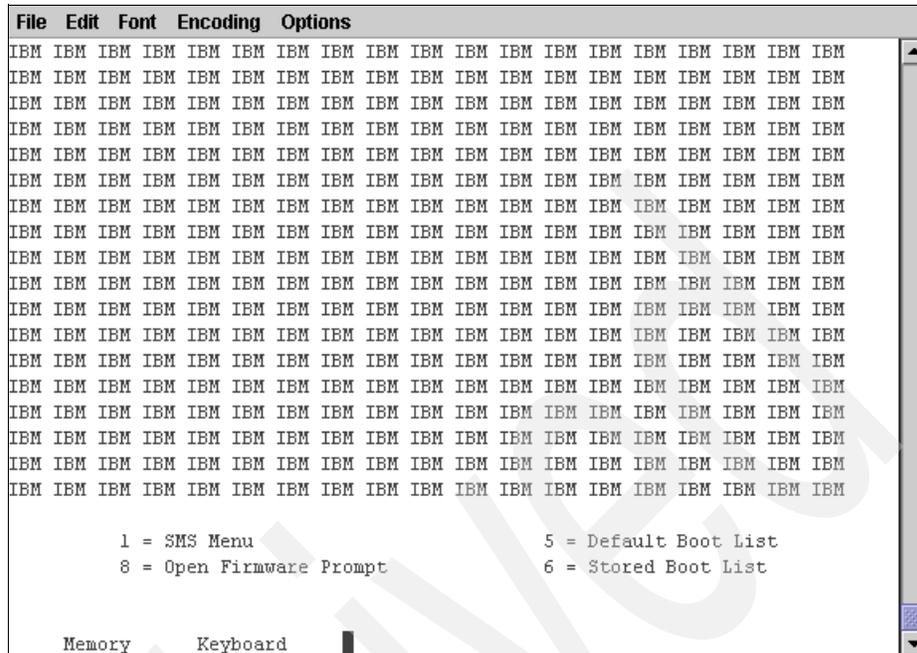


Figure 6-8 SMS menu

- 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. Choose **5. Select Boot Options.**
- c. Choose **1. Select Install/Boot Device.**
- d. Choose **6. Network.**
- e. Pick a specific network port.
- f. Choose **2. Normal Mode Boot**
- g. Pick **1.Yes** to exit the SMS menu.

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.

Note: Review Appendix D, “Additional Linux installation configuration options” on page 449 prior to starting an NFS installation.

You will notice the packet count value increasing. After the complete boot image is uploaded, the system boots off of it to show the Red Hat Enterprise Linux 5.2 welcome screen, shown in Figure 6-9.

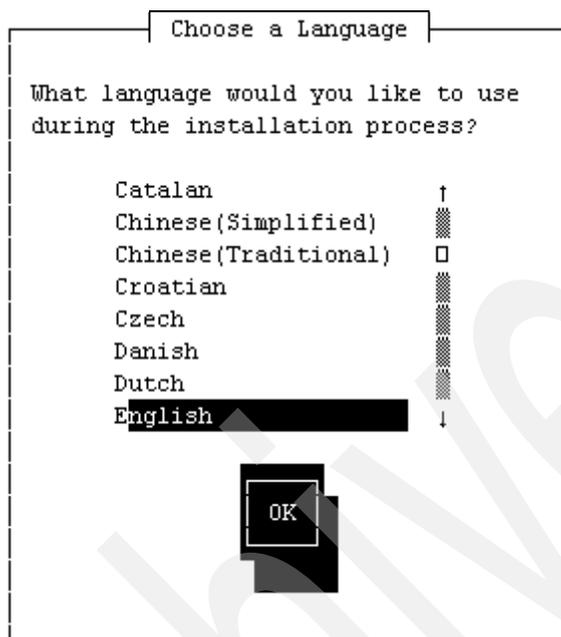


Figure 6-9 Red Hat Enterprise Linux 5.2 welcome screen

3. Select the language to use during the install process. In this example we are using English. Then press the Tab key to move to the **OK** button and then press the Space bar to confirm.

Tip: Use the Tab key to navigate between selections, the Up/Down Arrow keys to move within a list, and the Space bar to select a specific option in the Red Hat 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 navigate to the **OK** button and press the Space bar to confirm.
5. Select the networking device to install the image from.
6. To configure DHCP, select either IPv4 or IPv6 support and then Dynamic IP configuration (DHCP) from the TCP/IP window. Then select **OK**. See Figure 6-10 on page 293 for more details and skip the next two steps.

As an example, if you wish to configure static IP parameters with IPv4 support, configure the TCP/IP parameters and **Enable IPv4** support, select **Manual configuration** and disable **Enable IPv6 support** as shown in Figure 6-10, and configure the device as shown in Figure 6-11 on page 294.

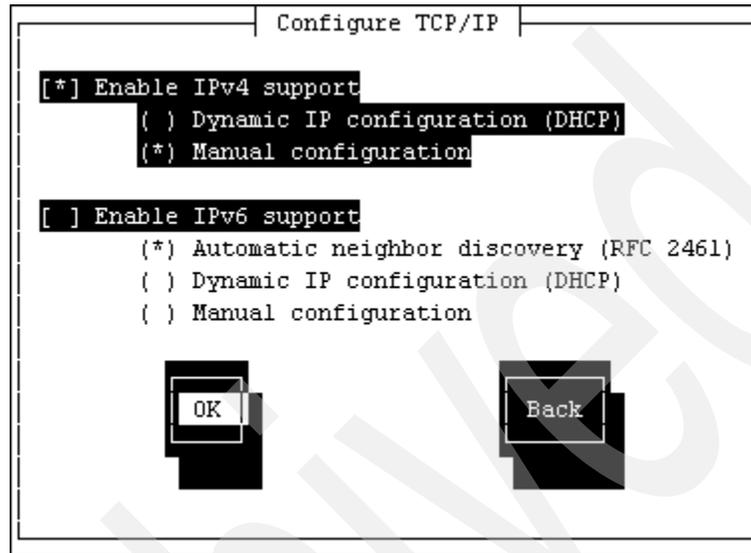


Figure 6-10 Configure TCP/IP

7. In the next panel, configure the LPAR's IPv4 address, subnet mask, gateway, and name server. The IP address, subnet mask, and gateway are exactly the same as the values entered in Figure 6-11 on page 294.

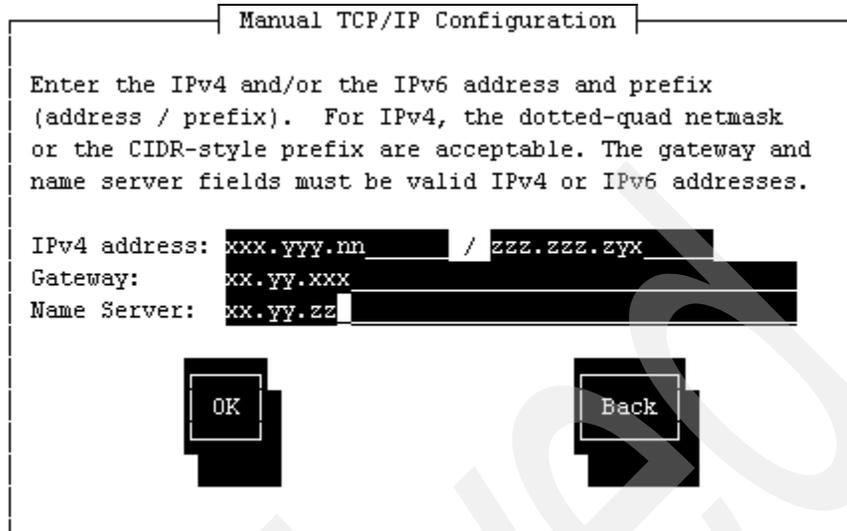


Figure 6-11 TCP/IP configuration of IP address, gateway, and name server

8. In the NFS Setup window in Figure 6-12, enter the IP address of the NFS server and in the field directly below that, enter the NFS directory that contains the Red Hat Enterprise Linux 5.2 install image.

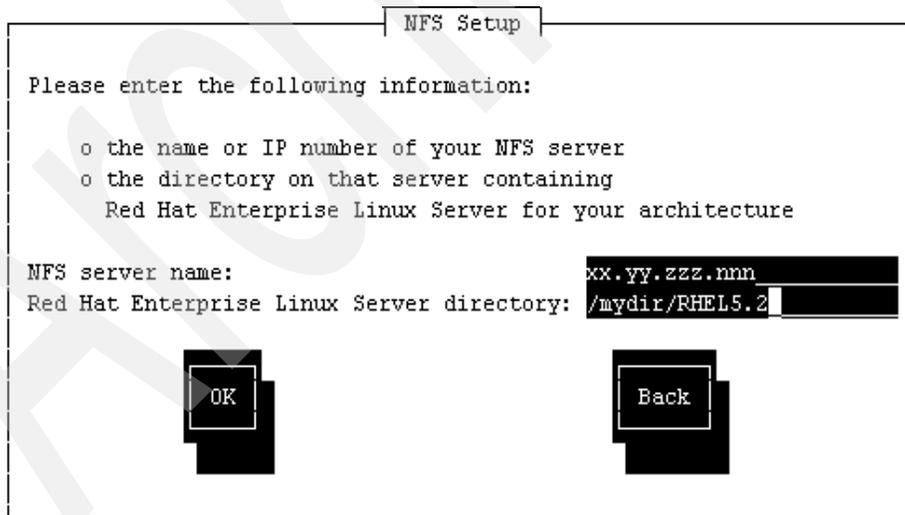


Figure 6-12 NFS server configuration window

9. Approximately one minute later the Welcome to Red Hat Enterprise Linux Server message panel appears. Select **OK**.

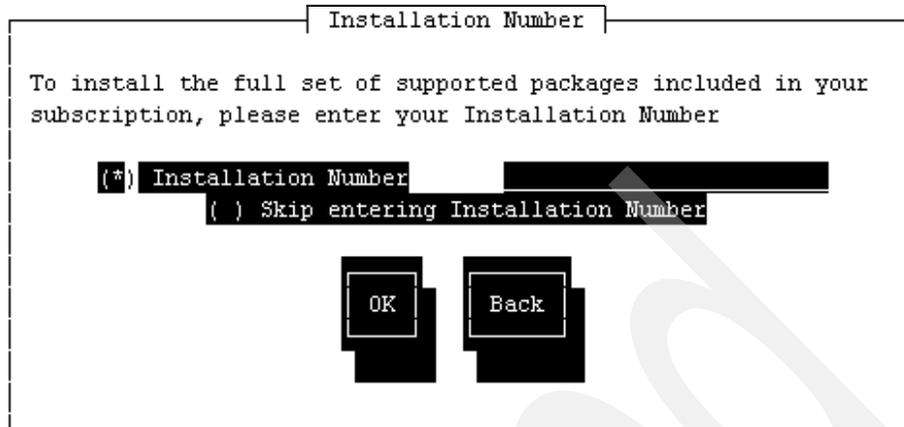


Figure 6-13 Installation number page

10. Enter the installation number in Figure 6-13. Click **OK** to confirm.

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. If you skip entering this number, you will only have the basic packages to select from later on.

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 6-14 on page 296 as an example.

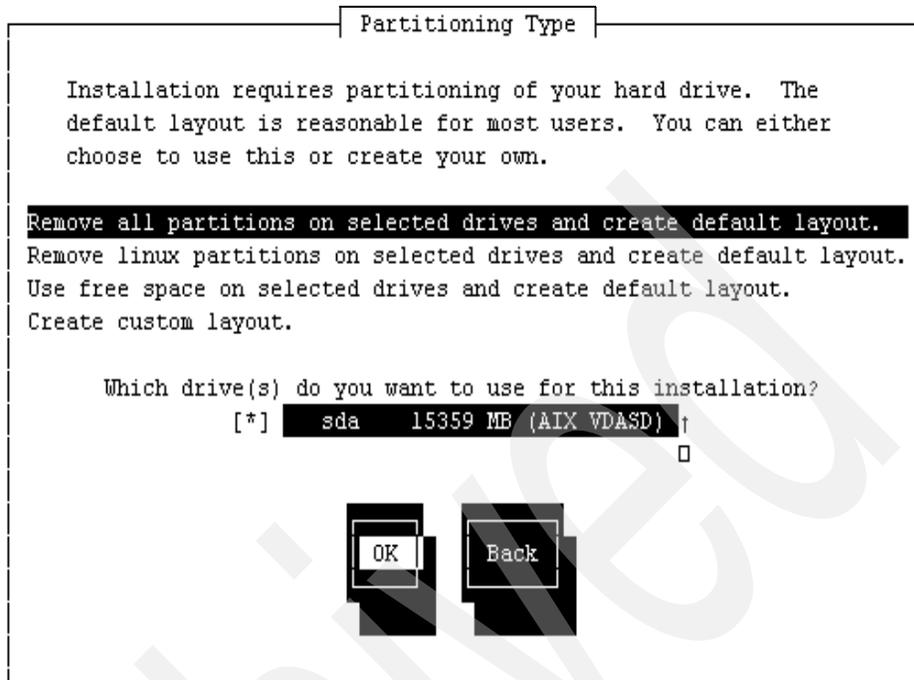


Figure 6-14 Select Partitioning Type panel

12. A warning appears asking 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 6-15 on page 297. Press **OK** to confirm.

Device	Start	End	Size	Type	Mount Point
VG VolGroup00			15232M	VolGroup	
LW LogVol01			1984M	swap	
LW LogVol00			13248M	ext3	/
/dev/sda					
sda1	1	1	0M	PPC PReP B	
Free space	1	1	7M	Free space	
sda2	2	13	94M	ext3	/boot
sda3	14	1958	15257M	physical v	

↑

↓

New Edit Delete RAID OK Back

Figure 6-15 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. Press **OK** on the Network Configuration panel. The default is fine because this was already set up in Figure 6-11 on page 294.
16. Press **OK** for the Miscellaneous Network Setting window. The gateway and primary DNS are already configured.
17. Press **OK** after editing/confirming the hostname.
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 Space bar on the **OK** button to confirm.
19. Enter the root password.
20. Select any additional software applications with the Space bar and press **OK** when complete.

Note: These packages can be installed later using **yum** from the command line if you skip this step during the installation.

21. Press **OK** to allow the installation to begin.
22. The next window has two progress bars: One for the package currently being installed and another detailing the overall progress of the installation.

```
Package Installation

Name   : glibc-common-2.5-24-ppc
Size   : 65636k
Summary: Common binaries and locale data for glibc

100%

Total   :           Packages      Bytes      Time
Completed:           16          24M    0:00:04
Remaining:           864         1870M   0:06:05

1%
```

Figure 6-16 Installation progress window

23. Press **Reboot** after the Install Complete window appears.

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**.

24. During boot the Setup Agent window appears (Figure 6-17). You can modify any of the fields if desired or press **Exit** to finish booting the LPAR.

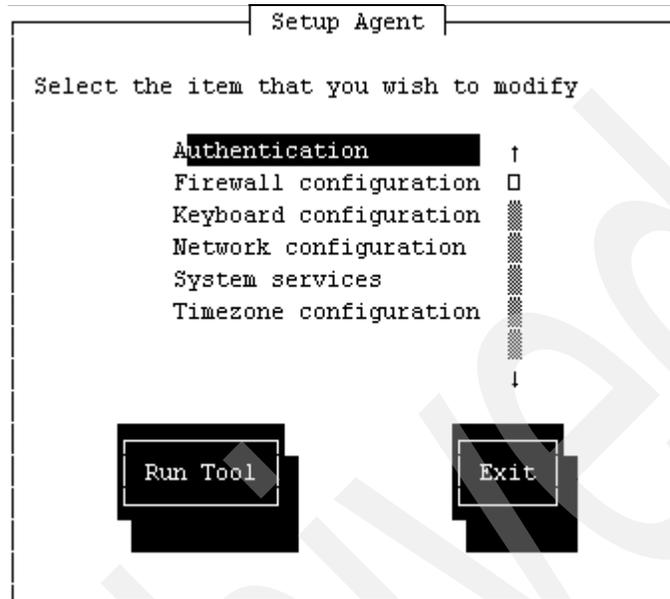


Figure 6-17 Setup Agent panel

25. The Red Hat Enterprise Linux 5.2 login prompt appears, as shown in Figure 6-18. The installation is complete.



Figure 6-18 Finished Red Hat Enterprise Linux 5.2 installation

6.4 Native Red Hat Enterprise Linux 5.2 installation

A native Red Hat Enterprise Linux 5.2 installation of a JS22 blade follows a similar process to the one given in the previous VIOS LPAR installation sections. However, there are some key differences:

- ▶ 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 as an alternative. See Appendix A, “Consoles, SMS, and Open Firmware” on page 391 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.

6.5 Red Hat Enterprise Linux 5.2 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 resides on a single server system and multiple clients can read it during installation. There are three methods to create a Kickstart file:

1. Create the file using the Red Hat Kickstart Configurator tool. This tool and the steps to create this file are described in detail in Appendix B, “Red Hat Enterprise Linux Kickstart” on page 419.
2. Have Red Hat Anaconda Installer generate this file for you during a manual installation. After the installation completes, the Kickstart file resides as `/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 6-1.
3. Generate the file yourself using a text editor.

Example 6-1 Basic Kickstart configuration file created during manual installation

```
# Kickstart file automatically generated by anaconda.

install
nfs --server=9.3.80.16 --dir=/install/linuxRHEL5.2
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
js22-turkey-3-lp1.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 of the partition information is commented out with a # symbol. This section needs to be uncommented and edited to support the partition schemes of systems that will use the automated Kickstart install process. The automated Kickstart process will not work without these edits. Use your favorite text editor to make edits to the Kickstart configuration file so that it looks like the one given in Example 6-2. This file must reside on a server that is accessible to all systems that will use this as their Kickstart configuration file.

Example 6-2 Complete Kickstart file ks.cfg

```
install
nfs --server=9.3.80.16 --dir=/install/linuxRHEL5.2
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
js22-turkey-3-lp1.austin.ibm.com
rootpw --iscrypted $1$0VjuAHA9$6cqrT/Nfnh6pRYFBIBECy/
firewall --enabled --port=22:tcp
authconfig --enablshadow --enablemd5
selinux --enforcing
timezone --utc America/New_York
bootloader --location=partition --driveorder=sda --append="console=hvc0
rghb quiet"
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
@dialog
```

```
@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
```

Performing a Kickstart installation

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 6-19.

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

Memory   Keyboard   Network   SCSI   Speaker   ok
0 > _
```

Figure 6-19 Open Firmware prompt

2. Type the following command in the Open Firmware prompt to start automated installation. For example, if the configuration file is served using NFS:

```
boot net ks=nfs://192.168.1.254/ks.cfg ksdevice=eth1 ip=dhcp
```

Press the Enter key and the process will begin. The automated Red Hat Enterprise Linux installation is now complete.

Tip: *Red Hat Enterprise Linux 5.2 Installation Guide* is a good source of information regarding all Kickstart file options. It is also available at:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/ch-kickstart2.html

6.6 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. They are IBM owned and distributed.

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 F, “Service and productivity tools for Linux” on page 465 for more information on installing the service and productivity tools specific to your system’s configuration.

SUSE Linux Enterprise Server installation

This chapter describes the procedures to install SUSE Linux Enterprise Server (SLES) on a JS22 BladeCenter.

We discuss the following topics:

- ▶ Supported SUSE operating systems
- ▶ Considerations and prerequisites
- ▶ Linux LPAR installation using DVD
- ▶ Linux network installation (detailed)
- ▶ Native SLES10 SP2 installation
- ▶ SLES 10 SP2 automated installation
- ▶ IBM service and productivity tools

7.1 Supported SUSE operating systems

SUSE Linux Enterprise Server 10 for POWER Service Pack 1 (SLES 10 SP1) or later supports installation on a JS22.

This chapter specifically covers installing SUSE Linux Enterprise Server 10 Service Pack 2 (SLES 10 SP2) for POWER with a DVD and over the network on a PowerVM LPAR.

7.2 Considerations and prerequisites

There are some system configuration considerations and prerequisites prior to installing SLES 10 SP2 on a JS22 partition. They are covered here.

7.2.1 PowerVM LPAR considerations

By the characteristics of the virtualization features in the System p JS22 Blade Center, the operating system and applications do not know they are running in a micro-partitioned or virtualized I/O environment. This 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 3, "System configuration using VIOS and IVM" on page 55.

7.2.2 SUSE Linux Enterprise Server 10 SP2 installation considerations

The SUSE Linux Enterprise Server 10 SP2 document recommends at least 256 MB of RAM to install SUSE Linux Enterprise Server 10 SP2.

Tip: We recommend at least 1-2 GB per PowerVM LPAR and 4 GB on LPARs running CPU intensive loads.

In addition, the SLES10 installation guide suggests to have at least 1.8 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 recommend 10 GB or more *total* hard disk space for each PowerVM LPAR.

The Novell Web site has additional installation preparation information for SLES10 Sp2 available at:

<http://www.novell.com/documentation/sles10/index.html>

7.3 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 10 SP2 prerequisites, we are ready to install Linux on the JS22 Blade Center LPAR.

Note: Even though this section covers installation via the Integrated Virtualization Manager (IVM) console, there are other console options available on the JS22; these are covered in Appendix A, “Consoles, SMS, and Open Firmware” on page 391.

Start an LPAR installation of SLES 10 SP2 using a DVD

The following directions are for installing SLES 10 SP2 from a DVD:

1. On the Advanced Management Module (AMM) select **Blade Tasks** → **Remote Control**.
2. In the right panel look for Media Tray Owner and click the drop-down arrow to display the different blade bays. Select your blade bay location with the mouse as shown in Figure 7-1.

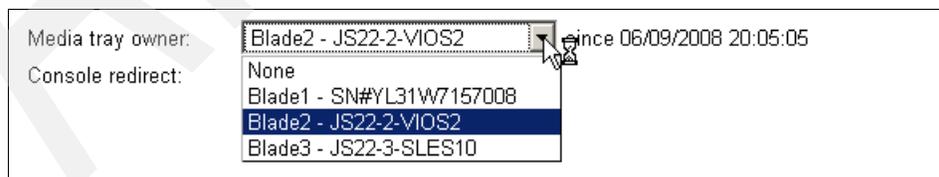


Figure 7-1 Start Remote Console panel

3. Press **Refresh**.

Note: The other option is to press the **MT** button on the blade to assign the media tray to the blade.

Important: Make sure no other blade in the Blade Center is using the media tray before pressing this button. The blade's MT light is on if the media tray is assigned to them.

4. 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 7-2 gives one example of this assignment.

Bay	Status	Name	Pwr	Owner**	
				KVM	MT†
1	●	JS22-1-VIOS1	On		X
2	●	JS22-2-VIOS2	On		
3		No blade present			

Figure 7-2 The X indicates that the Media Tray (MT) is owned by the blade in Bay 1

5. Place the DVD into the Blade Center media tray.
6. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
 - a. Select one of the available PowerVM LPARs to install with SLES 10 SP2 by placing a checkmark in the Select box.
 - b. Click **Activate** as shown in Figure 7-3

Partition Details							
Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units
<input type="checkbox"/>	1	gloria-vio	Running	10.05 Days	512 MB	4	0.4
<input type="checkbox"/>	2	testelp1	Open Firmware	18.3 Minutes	1 GB	1	0.1
<input type="checkbox"/>	3	testelp2	Not Activated		1 GB	1	0.1
<input checked="" type="checkbox"/>	4	testelp3	Not Activated		800 MB	1	0.1

Figure 7-3 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field and select **Open terminal window**.

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 1.4.2_08, or higher.

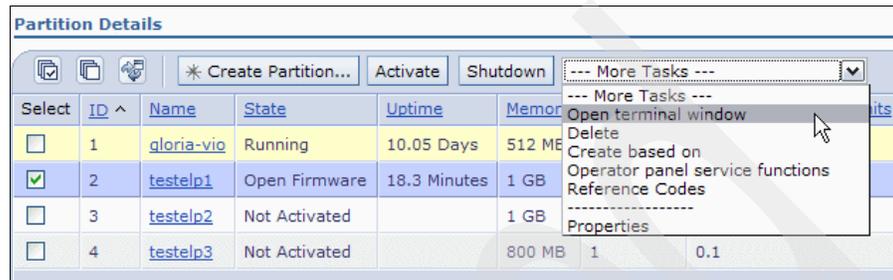


Figure 7-4 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

7. The SMS menu appears in the IVM terminal.

Note: The SMS menu appears very quickly after activating the box, 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.

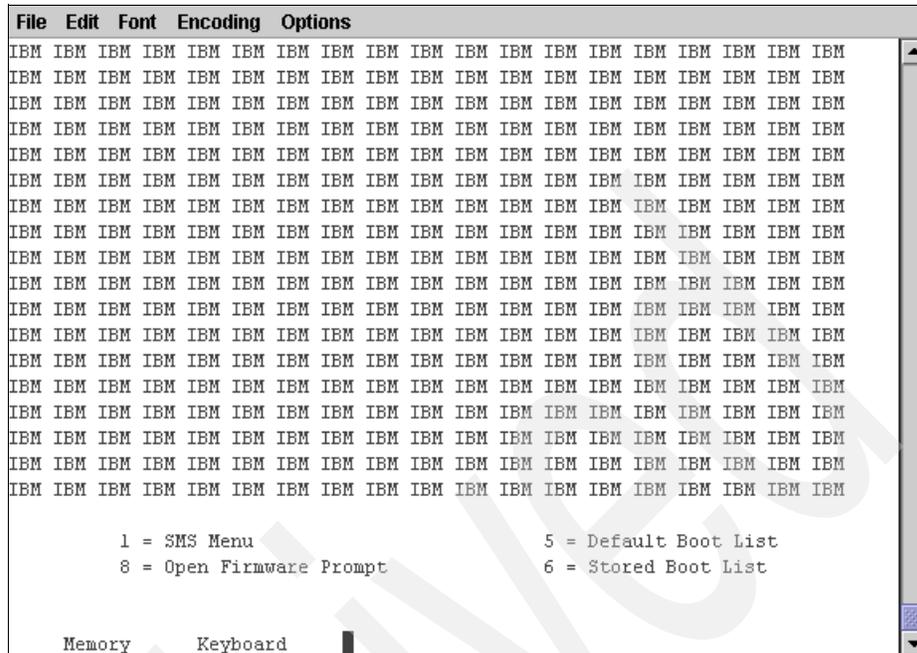


Figure 7-5 SMS Menu

- a. 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.

- b. Select option **5. Select Boot Options**.
 - c. Choose option **1. Select Install/Boot Device**.
 - d. Pick **3. CD/DVD**.
 - e. Select **6. USB**.
 - f. Finally, select **1. USB CD-ROM**.
 - g. Choose **2. Normal Mode Boot**.
 - h. Pick **1.Yes** to exit the SMS menu.
 - i. At the Linux **boot:** prompt, type `install`, then press Enter to confirm. The LPAR will start reading from the DVD, which can take a couple of minutes.
8. After successfully booting off the DVD, the SUSE language selection panel appears on the IVM terminal and Your awesome Setup Tool (YaST) will

assist with the completion of the install. More detailed installation instructions are available here:

http://www.novell.com/documentation/sles10/sles_admin/index.html?page=/documentation/sles10/sles_admin/data/sles_admin.html

7.4 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 will work for all supported internal or external storage devices on JS22, and only a few slight modifications are needed in the YaST installer setup if a different network installation option is selected. The goal of this section is to show the SAN disk install 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 *Installing Linux using the network - General remarks of Appendix D, “Additional Linux installation configuration options”* on page 449.

This section assumes that the user has already set up the NFS server properly and has read 7.2.1, “PowerVM LPAR considerations” on page 306 and followed the installation prerequisites in 7.2.2, “SUSE Linux Enterprise Server 10 SP2 installation considerations” on page 306.

Start an LPAR installation of SLES 10 SP2 over the network

The following instructions guide you through the installation of SLES 10 SP2 through a network.

1. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
 - a. Select one of the available PowerVM LPARs to install with SLES10 SP2 by placing a checkmark in the Select box.
 - b. Click **Activate** as shown in Figure 7-6 on page 312.

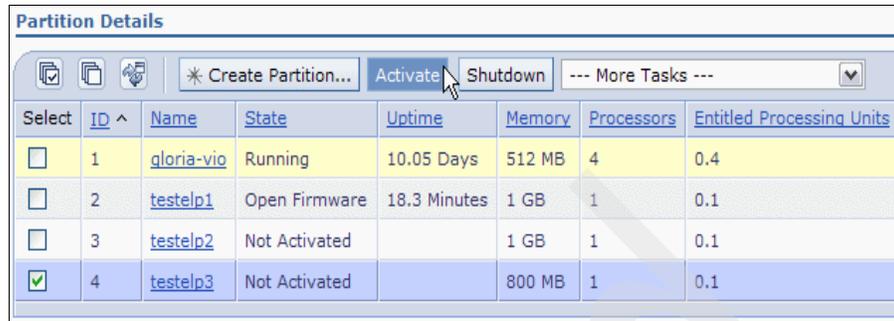


Figure 7-6 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field and select **Open terminal window** as shown in Figure 7-7.

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 1.4.2_08, or higher.

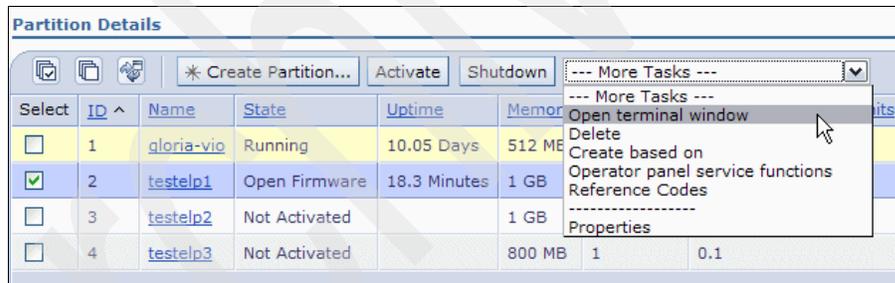


Figure 7-7 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

2. The firmware boot panel appears in the IVM terminal.

Note: The SMS menu option appears very quickly after activating the LPAR, 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 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.

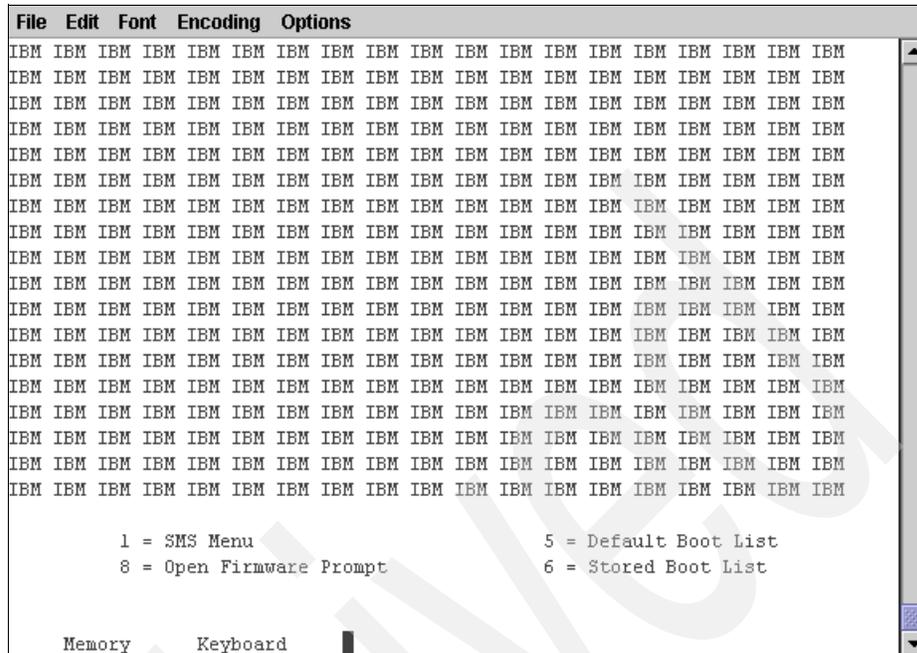


Figure 7-8 SMS Menu

- 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. Choose **5. Select Boot Options.**
- c. Choose **1. Select Install/Boot Device.**
- d. Choose **6. Network.**
- e. Pick **1. BOOTP.**
- f. Choose a network port.
- g. Choose **2. Normal Mode Boot.**
- h. Pick **1.Yes** to exit the SMS menu.

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 a couple of minutes the Main Menu panel appears as shown in Figure 7-9 on page 314 will.

1. Select **3) Kernel Modules (Hardware Drivers).**

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 will move you back to the previous option window.

```
Main Menu
1) Settings
2) System Information
3) Kernel Modules (Hardware Drivers)
4) Start Installation or System
5) Verify Installation CD-ROM/DVD
6) Eject CD
7) Exit or Reboot
8) Power off
>
```

Figure 7-9 Main Menu selection window

2. Choose **1) Load ppc Modules**.

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 up/down scroll bar on the IVM terminal to navigate the module list. The most commonly used modules are 5) e1000 : Intel PR0/1000 and 15) ehea : EHEA.

3. Press the **Enter** key after you have finished loading the modules to go back to the Main Menu.

4. Select **4) Start Installation or System**.

5. Choose **1) Start Installation or Update** as shown in Figure 7-10.

```
Start Installation or System

1) Start Installation or Update
2) Boot Installed System
3) Start Rescue System

>
```

Figure 7-10 Start installation or update option

6. Select **2) Network**.

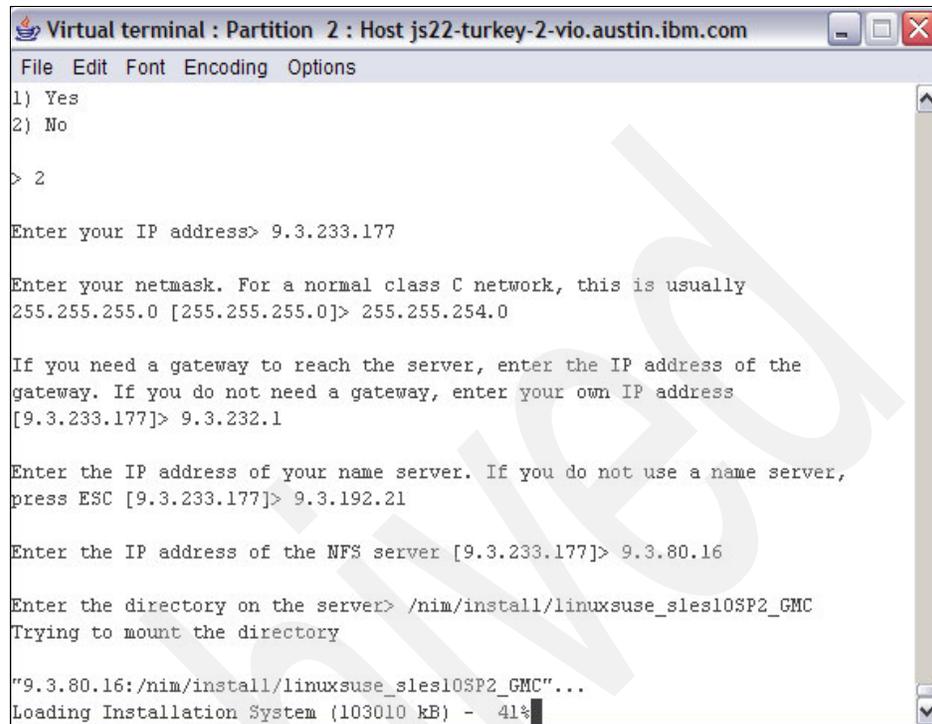
7. Select **3) NFS** as the network protocol.

8. Choose **1) Yes** or **2) No** for network configuration via DHCP.

a. If option **2) No** was selected, enter the static IP information:

- 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 SLES10 SP2 image.

See Figure 7-11, showing a sample configuration.



```
Virtual terminal : Partition 2 : Host js22-turkey-2-vio.austin.ibm.com
File Edit Font Encoding Options
1) Yes
2) No
> 2
Enter your IP address> 9.3.233.177
Enter your netmask. For a normal class C network, this is usually
255.255.255.0 [255.255.255.0]> 255.255.254.0
If you need a gateway to reach the server, enter the IP address of the
gateway. If you do not need a gateway, enter your own IP address
[9.3.233.177]> 9.3.232.1
Enter the IP address of your name server. If you do not use a name server,
press ESC [9.3.233.177]> 9.3.192.21
Enter the IP address of the NFS server [9.3.233.177]> 9.3.80.16
Enter the directory on the server> /nim/install/linuxsuse_sles10SP2_GMC
Trying to mount the directory
"9.3.80.16:/nim/install/linuxsuse_sles10SP2_GMC"...
Loading Installation System (103010 kB) - 41%
```

Figure 7-11 Static network configuration example

The LPAR begins reading from the SLES10 SP2 image directory and then displays the Your awesome Setup Tool (YaST) Welcome panel (Figure 7-12 on page 317).

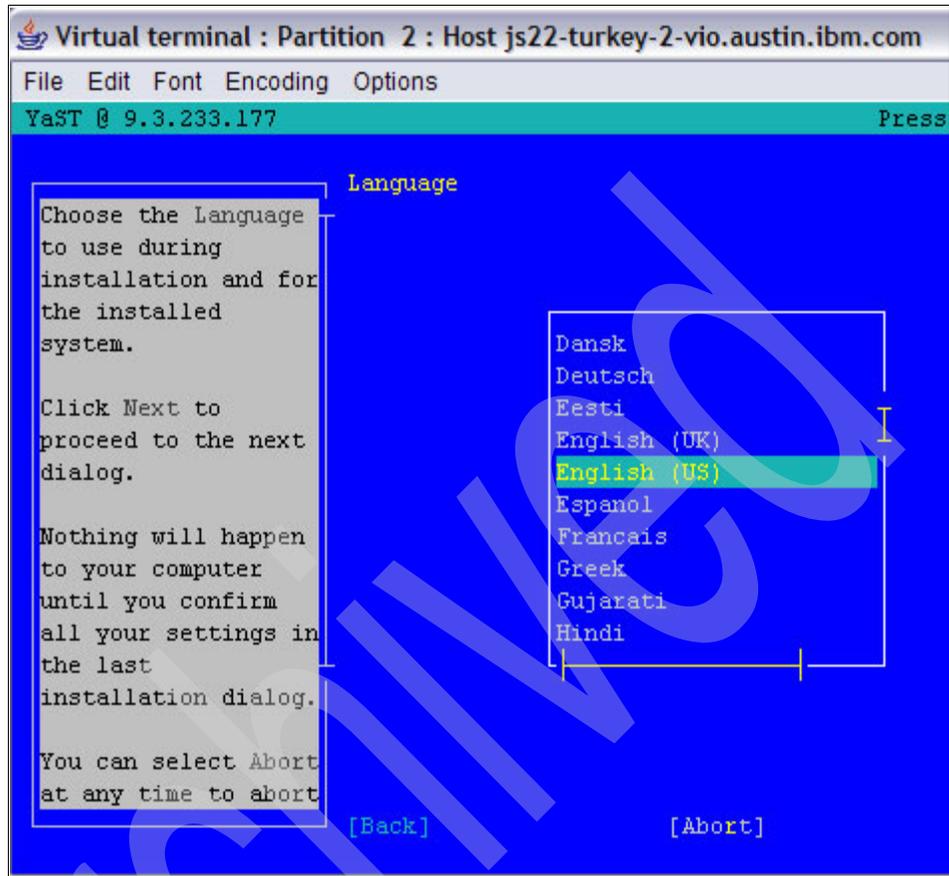


Figure 7-12 YaST Welcome panel

Tip: Navigate the YaST tool by using the Tab key to move between sections, the up/down arrow keys to move within a specific window section, the space bar to check a “()” entry with an “x,” the Enter key to confirm a selection with square brackets “[]” around it, and the Delete key to erase entries.

1. Highlight **English** in the Language selection panel. Press the Tab key to **[Next]** and then press Enter.
2. On the next panel, select **Yes, I Agree to the Licence Agreement** and then press the Tab key to **[Next]** and press Enter to confirm.
3. Select **New Installation** for the Installation Mode.
4. Configure your clock and time zone information.

5. 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.
6. Select **[I Agree]** to the AGFA Monotype Corporation License Agreement.
7. Choose **[Install]** to start the installation.

The YaST window refreshes to the installation progress bars as shown in Figure 7-13. 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.

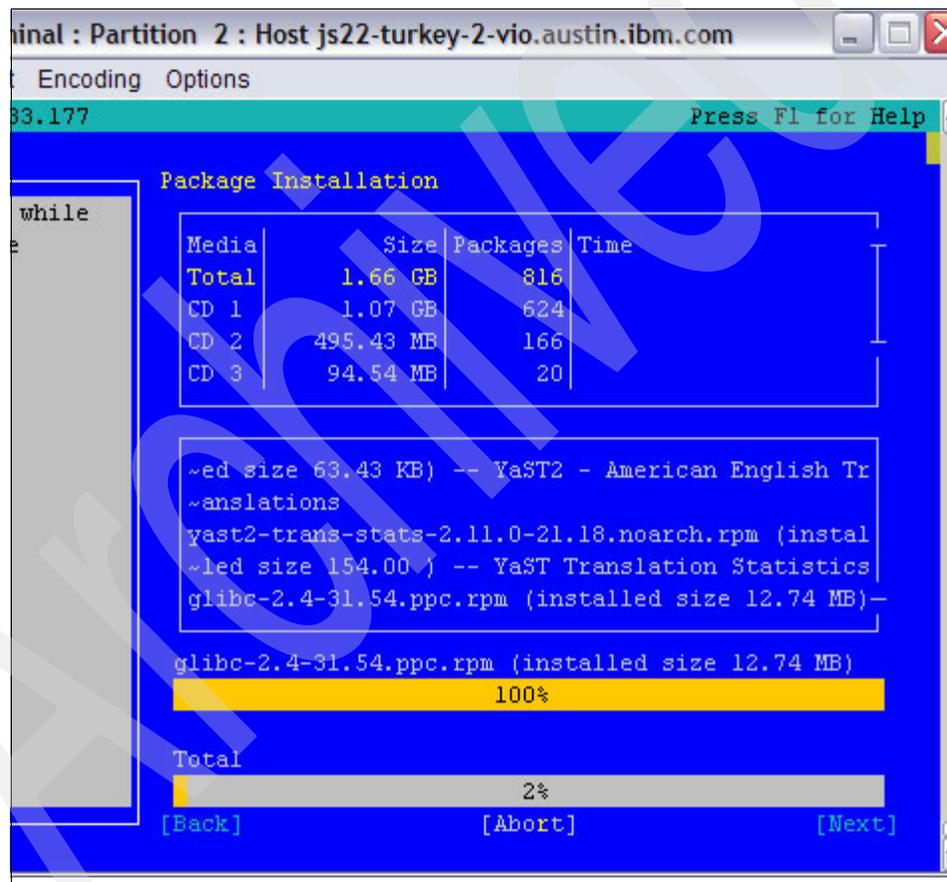


Figure 7-13 YaST installation progress window

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**.
- ▶ Choose the appropriate hard disk with the Linux image from the given list.
- ▶ Select **2. Normal Mode Boot** → **1. Yes**.

8. After booting from the appropriate storage device YaST will start again. Select **[Continue]** for each hardware device YaST detects, as shown in Figure 7-14 on page 320.

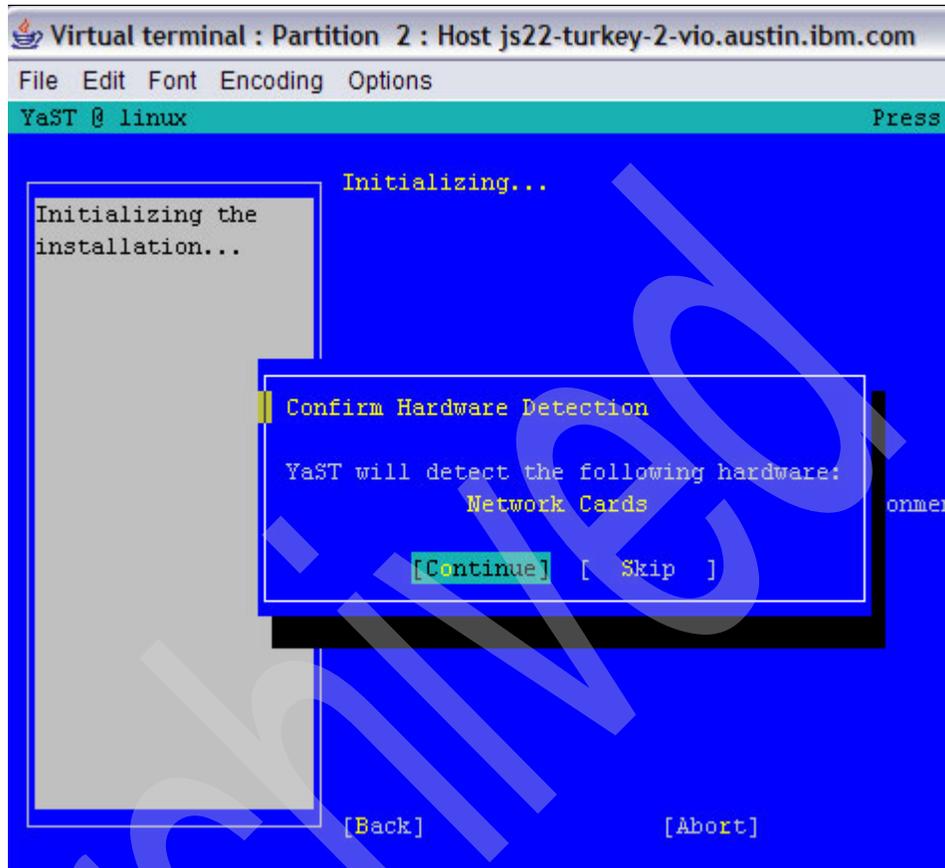


Figure 7-14 Confirm hardware detection window

9. Enter the root user's password. Press **[Next]** to confirm.
10. Provide the hostname and the domain. Press **[Next]** to confirm.
11. Select () **Use Following Configuration** in the Network Configuration window (Figure 7-15 on page 321) and verify that the Firewall is marked as **enabled**. Press the Tab key to **[Change....]** to change the SSH port settings to open.

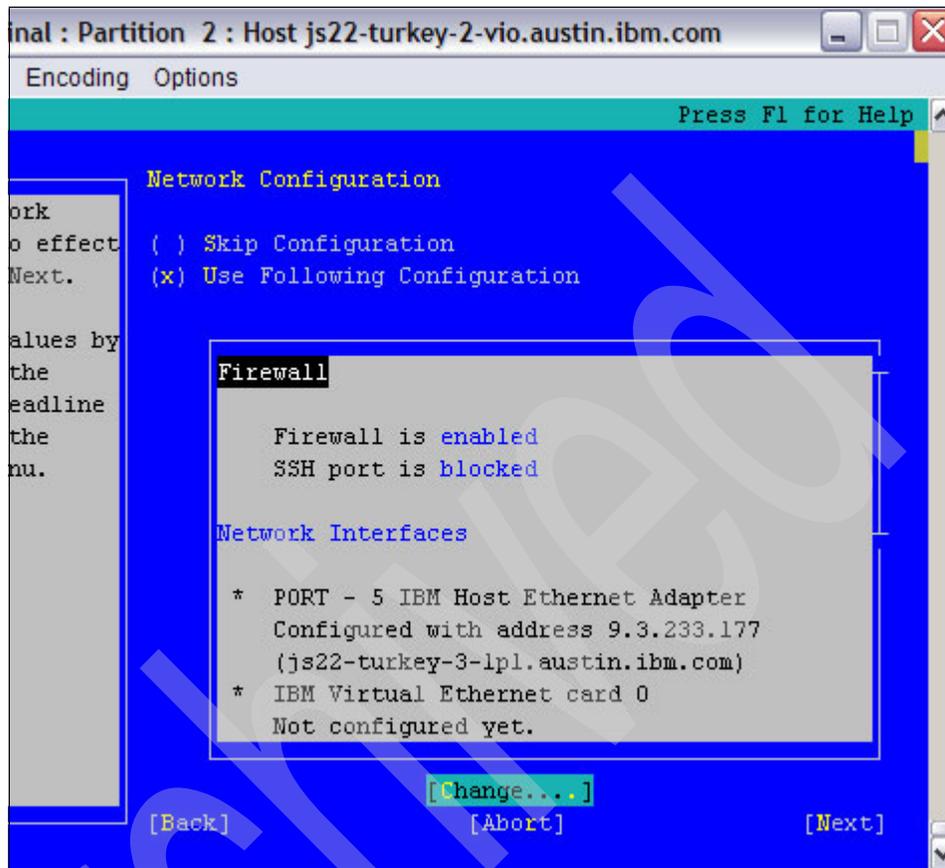


Figure 7-15 Change network configuration

- a. Select **Firewall** in Figure 7-16.

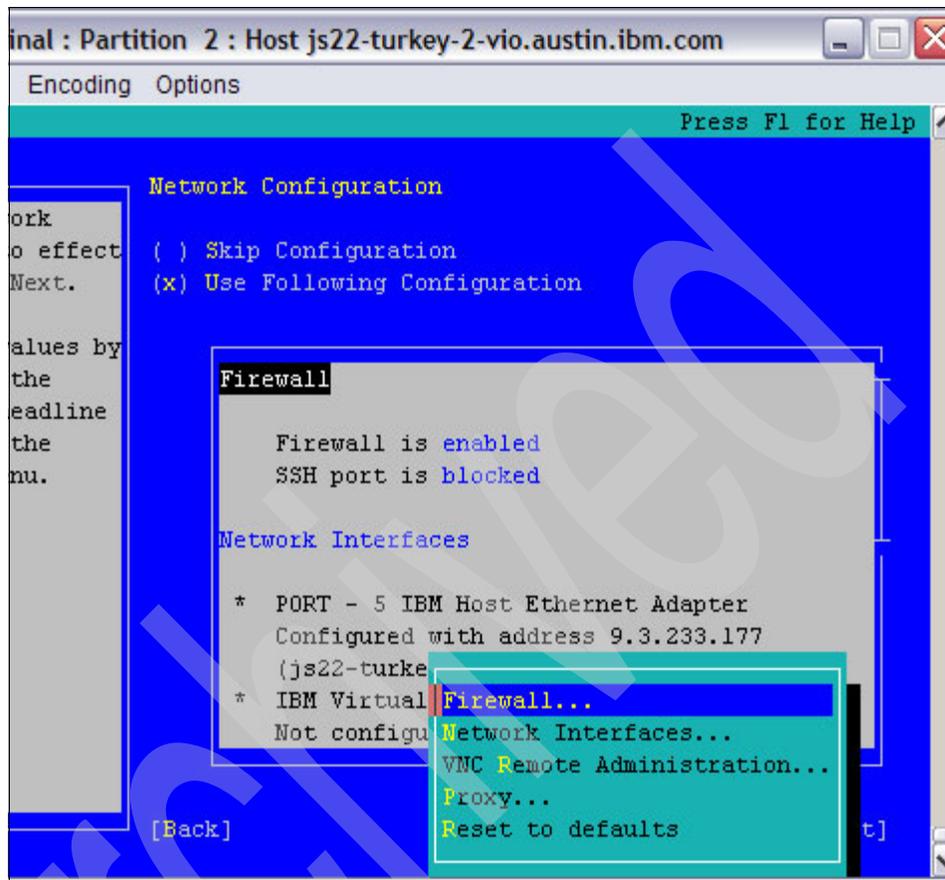


Figure 7-16 Select firewall

- b. Scroll to **Allowed Services**.
- c. Press the Tab key to highlight **DHCP Client** in the Service to Allow section, then press the down arrow key to bring up a new window. Find and highlight **SSH** in the new window, as shown in Figure 7-17 on page 323. Finally, press Enter to confirm.

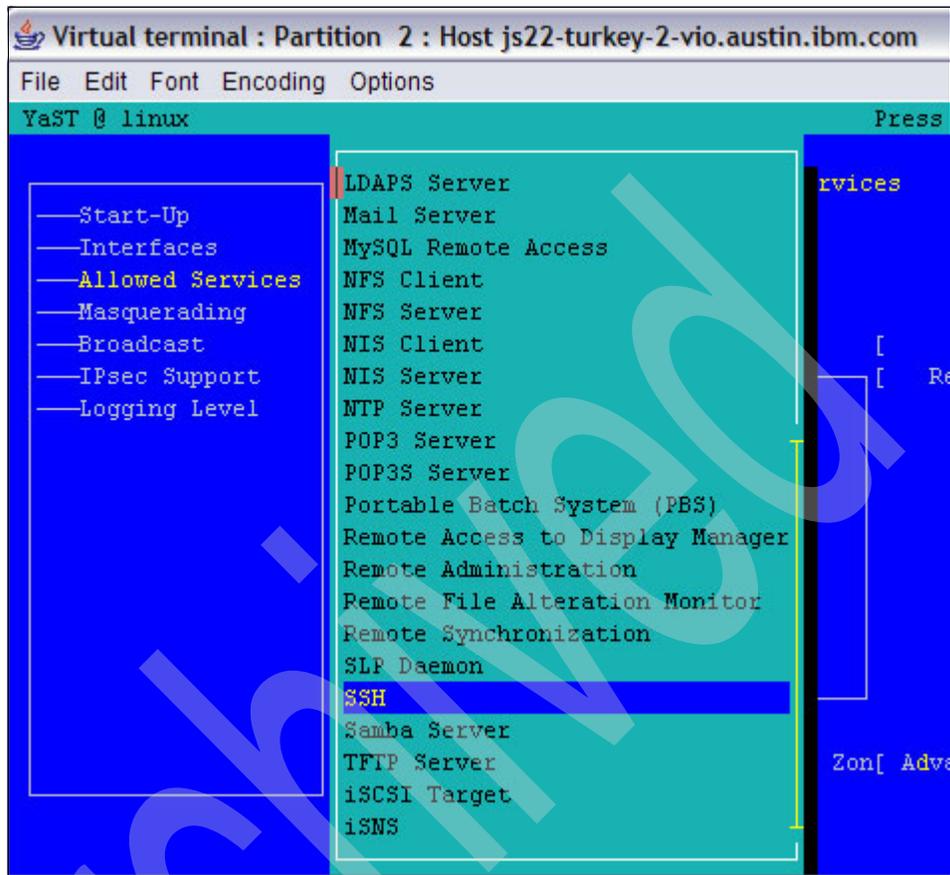


Figure 7-17 Services to allow list and selecting SSH service

- Press the Tab key to highlight **[Add]** and the press Enter to confirm.
- SSH will appear in the **Allowed Service** list. Press **[Next]** to confirm.

- f. Now the Firewall section of the Network Configuration window (Figure 7-18) shows “SSH port is open.”

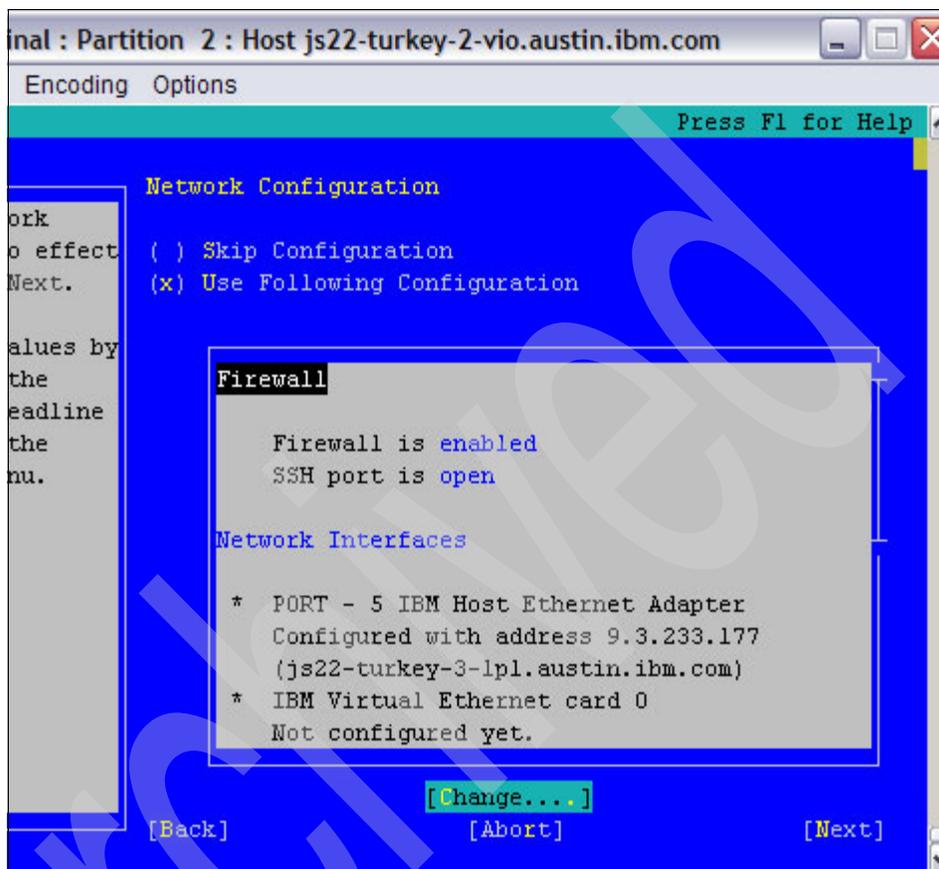


Figure 7-18 SSH port is open

1. Test the Internet connection, if desired.
2. Change the Certification Authority (CA) Installation setting, if desired. Select **[Next]** to confirm the changes.
3. Select the user authentication method appropriate for this LPAR and select **[Next]**.
4. Create a local user and select **[Next]**.
5. YaST will write the configuration settings and then display the release notes. Choose **[Next]** after reading the release notes.
6. Configure Hardware (Printers) if desired, then confirm the described configuration with **[Next]**.

- YaST displays the Installation Completed window (Figure 7-19). Select () **Clone This System for Autoyast** (see “SLES 10 SP2 automated installation” on page 326 for more information) if desired and then select **[Finish]**.

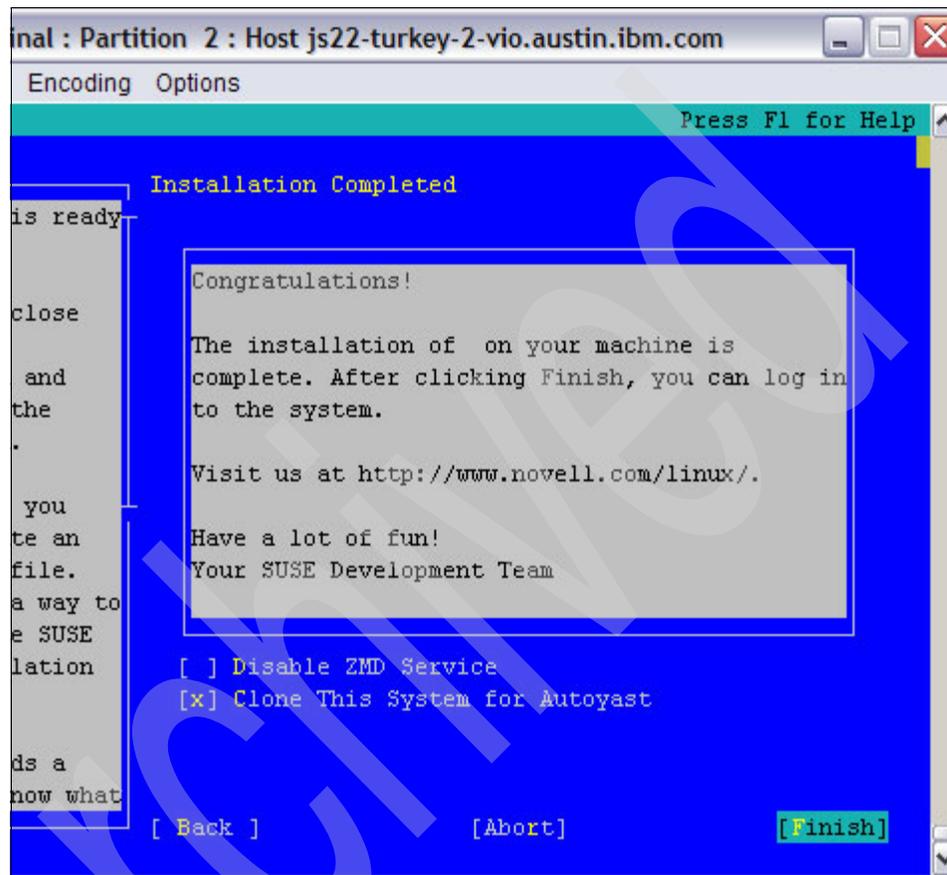


Figure 7-19 Installation completed window

7.5 Native SLES10 SP2 installation

A native SLES10 SP2 installation of a JS22 blade follows a similar process to those given in the VIOS LPAR installation sections. However, there are some key differences:

- ▶ 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 as an alternative. See Appendix A, “Consoles, SMS, and Open Firmware” on

page 391 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

7.6 SLES 10 SP2 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. There are multiple methods to create 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 C, “SUSE Linux Enterprise Server AutoYaST” on page 429 for more information on 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 7-19 on page 325), select () **Clone This System for Autoyast** and then select **[Finish]**. This creates an Autoyast profile in `/root/autoinst.xml` that is ready for immediate use.

Tip: *SUSE Linux Enterprise Server 10 Doc* is a good source of information regarding all Autoyast profile options. It is available at:

http://www.novell.com/documentation/sles10/sles_admin/index.html?page=/documentation/sles10/sles_admin/data/cha_deployment_autoinst.html

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.

4. Type 8 during the LPAR boot process to go to the Open Firmware prompt, as shown in Figure 7-20.

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

Memory      Keyboard      Network      SCSI      Speaker  ok
0 > _
```

Figure 7-20 Open Firmware prompt

5. 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/SLES10_SP2
```

Note: This automated install assumes there is a DHCP server that will provide the client with the correct IP address information. A system using static IP will require additional parameters for `sshpassword`, `hostip`, `netmask`, `gateway`, and `nameserver`.

Press the Enter key to start the process. The automated SUSE Linux Enterprise Linux installation is now complete.

7.7 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. They are IBM owned and distributed.

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 F, “Service and productivity tools for Linux” on page 465 for more information on installing the service and productivity tools specific to your system’s configuration.



Performing Live Partition Mobility

This chapter discusses the requirements and configuration procedures to perform Live Partition Mobility between two IBM BladeCenter JS22 blades. A complete technical discussion of Partition Mobility can be reviewed in *IBM System p Live Partition Mobility*, SG24-7460

8.1 Requirements

Partition mobility places certain demands on hardware, software, network and storage configurations. These considerations need to be reviewed early in the setup of an IBM BladeCenter JS22 or JS12 to avoid reconfiguration and rework.

8.1.1 Hardware

The IBM BladeCenter JS22 or JS12 requires a supported expansion card for SAN or iSCSI support. All storage assigned to a logical partition must have external backing devices and be visible to both the local and remote VIOS systems.

8.1.2 Firmware

In a new blade environment the recommendation would be to have the most current system firmware available 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 may be used for system maintenance reasons. The actual levels used are shown as displayed by the BladeCenter AMM UI.

- ▶ EA320_046 released 5/29/08
- ▶ EA330_031 released 5/20/08

IBM BladeCenter JS22 and JS12 firmware updates can be downloaded from <http://www.ibm.com>

Select the **Support & Downloads** link, then **BladeCenter** from the drop-down box to be directed to the Support for IBM BladeCenter page. Select **BladeCenterJS12** or **BladeCenterJS22** 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.

8.1.3 VIOS Version

Similarly to system firmware, the VIOS version and fixpack level should be the most recent. To display the current code level from the UI, click **Updates** from the navigation area. The Management Partition Updates view will open and the code level shown in Figure 8-1 on page 331.

If the workstation that you are using has Internet access, the link displayed on the Management Partition Updates page will take you to a download site for newer updates and fixes, if available. The link to the Virtual I/O Server is also available here:

<http://techsupport.services.ibm.com/server/vios/download>

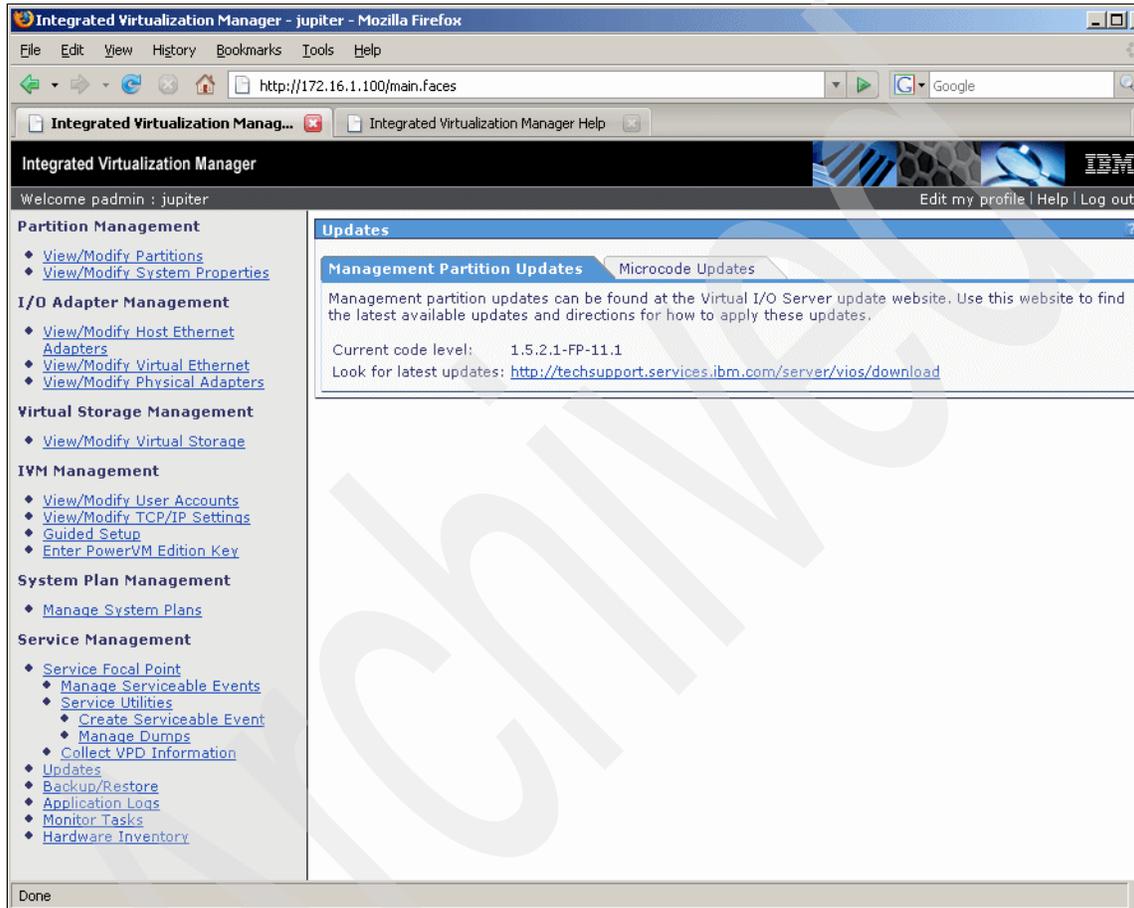


Figure 8-1 Management Partition Updates view

From the CLI use the `ioslevel` command to display the VIOS version and fixpack level, as shown in Example 8-1

Example 8-1 `ioslevel` command

```
$ ioslevel  
1.5.2.1-FP-11.1
```

8.1.4 PowerVM Enterprise

PowerVM Enterprise Edition is an optional feature on an IBM BladeCenter JS22 or JS12 and is required to enable Partition Mobility. The currently enabled features can be viewed using the `lsvet` command. Example 8-2 shows the `lsvet` command used to verify that Partition Mobility is enabled.

Example 8-2 lsvet command

```
$ lsvet -t hist | grep mobility
time_stamp=06/18/2008 22:52:15,entry=[VIO500042A-0341] Inactive
partition mobility enabled.
time_stamp=06/18/2008 22:52:15,entry=[VIO500042B-0342] Active
partition mobility enabled.
```

If 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 at this link:

<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 JS22 or JS12, it can be upgraded through the Miscellaneous Equipment Specification (MES) process.

Entering an enablement key through the IVM UI

The PowerVM key is entered from the UI by clicking the **Enter PowerVM Edition Key** link in the navigation area. The window that opens will have a box to enter the key number. Type in the key and select **Apply** as shown in Figure 8-2 on page 333. When PowerVM Enterprise is enabled, a Mobility section is added to the More Tasks drop-down box on the View/Modify Partitions view.

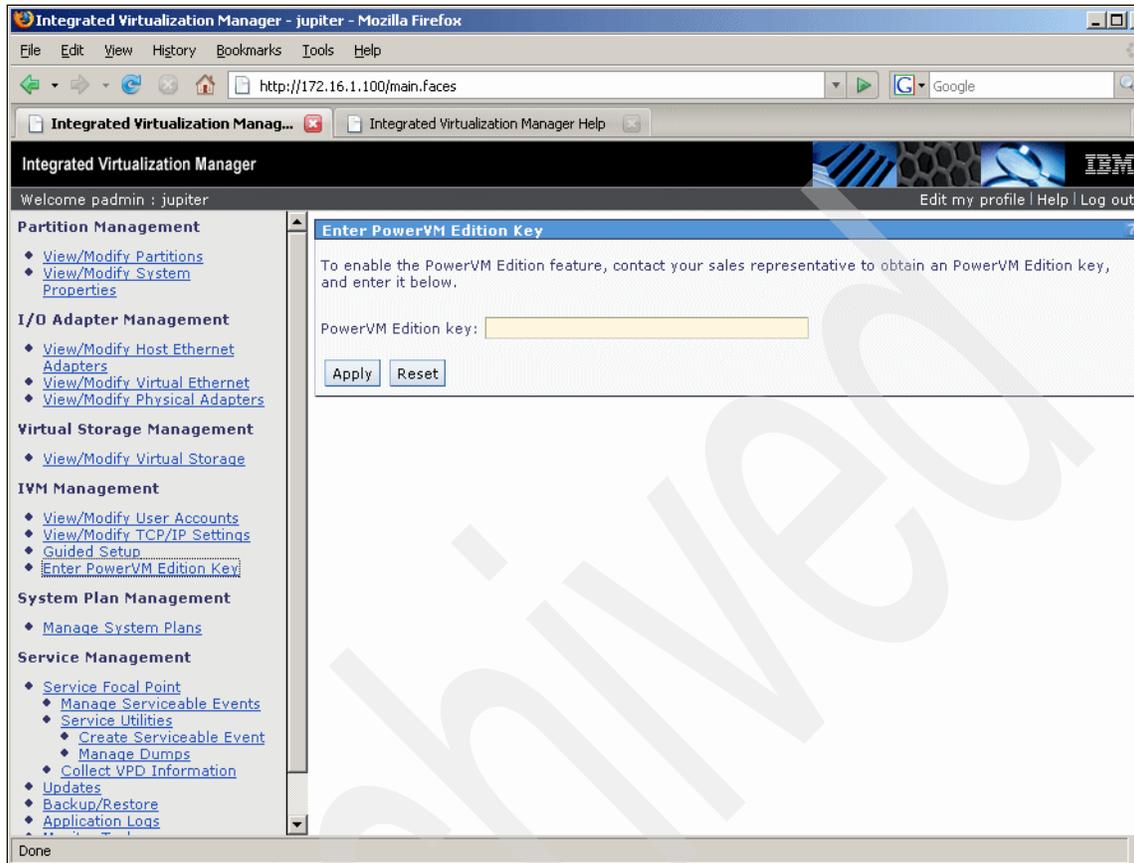


Figure 8-2 PowerVM Enterprise key entry

8.1.5 LPAR OS versions

The running operating system in the mobile partition must be AIX or Linux. The currently 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

8.2 Preparation

This section describes the settings and configurations that must be verified and possibly changed to prepare the local and remote VIOS systems for partition mobility.

8.2.1 VIOS (source and target)

Memory region size

The memory region size is the smallest block of memory that can be assigned to or changed in an LPAR. The current setting can be reviewed by clicking the **View/Modify System Properties** link from the Navigation area and then selecting the **Memory** tab. By default the value is set to 16 MB (automatic) as shown in Figure 8-3 on page 335. 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).

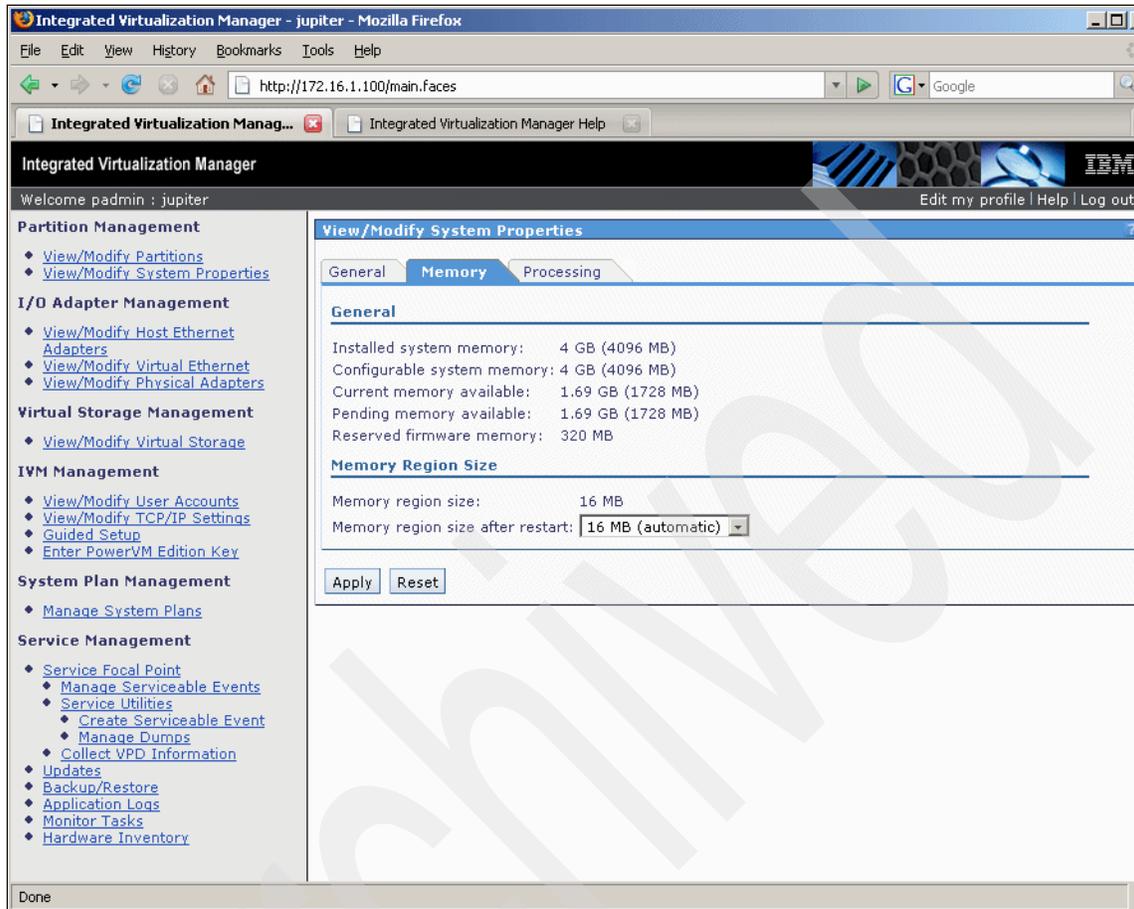


Figure 8-3 Memory region size

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. The `lsdev -dev hdisk# -attr | grep reserve` command can be used to review the current policy setting. The `chdev` command is used to change the attribute of the hdisk. Example 8-3 shows hdisk1 with the default reserve policy, the `chdev` command being issued, and the new reserve policy.

Example 8-3 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 will fail with the following message:

Some error messages may contain invalid information for the Virtual I/O Server environment.

```
Method error (/etc/methods/chgfcparry):
    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 UI 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 unknown, it is likely the reserve policy has not been changed and the validation process will fail. Figure 8-4 on page 337 shows hdisk0 and hdisk7 in the unknown size condition.

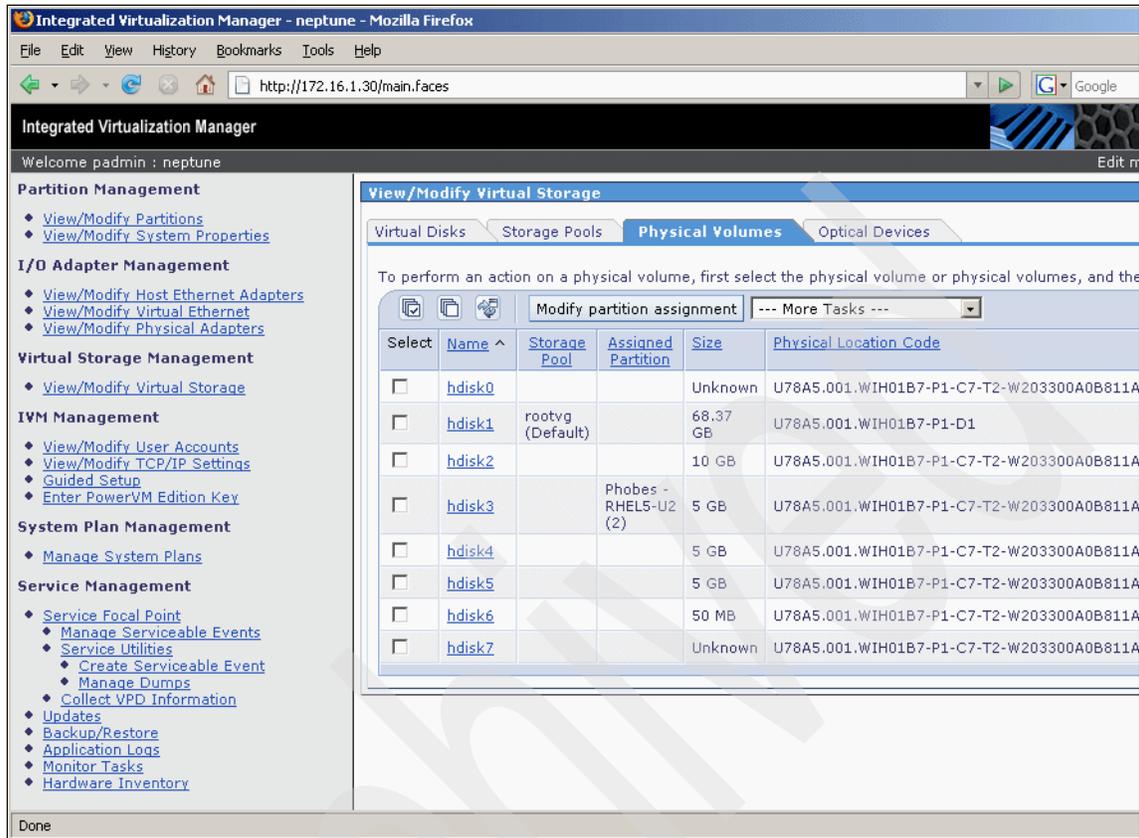


Figure 8-4 *hdisk reserve policy not set correctly*

When the validation process is run, an error message similar to Figure 8-5 on page 338 will be displayed.

This problem can be resolved by:

1. Shutting down the mobile LPAR on the local VIOS if running.
2. Modifying the mobile LPAR hdisk assignments on the local VIOS to *none*.
3. Using the `chdev` command to change the hdisk's reserve policy to *no_reserve*.
4. Modifying the mobile LPAR hdisk assignments to the original assignments.
5. Refreshing the View/Modify Virtual Storage view on the remote VIOS.

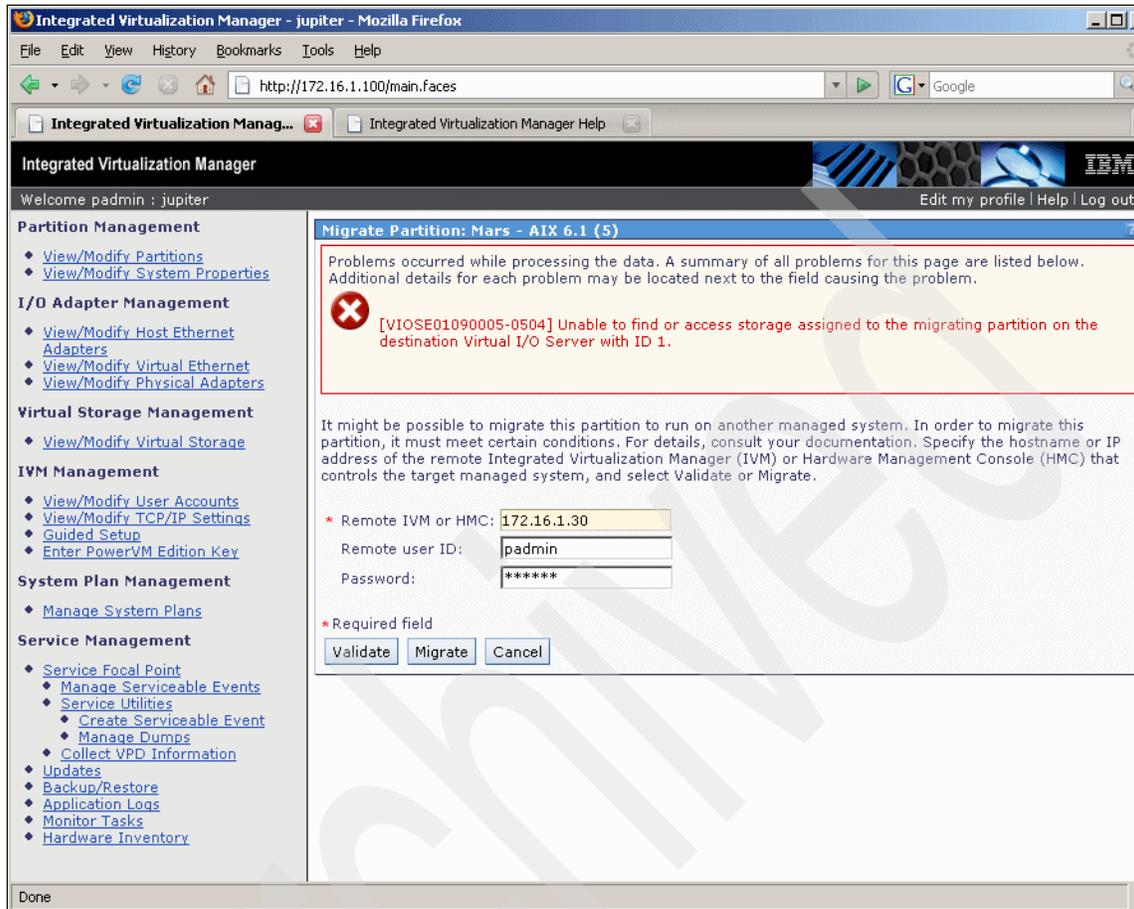


Figure 8-5 Partition Migration validation error message for target storage

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 8-6 on page 339 indicates that the virtual device `vtopt0` is still assigned to the mobile partition. The device can be removed by unchecking the box and clicking **OK**.

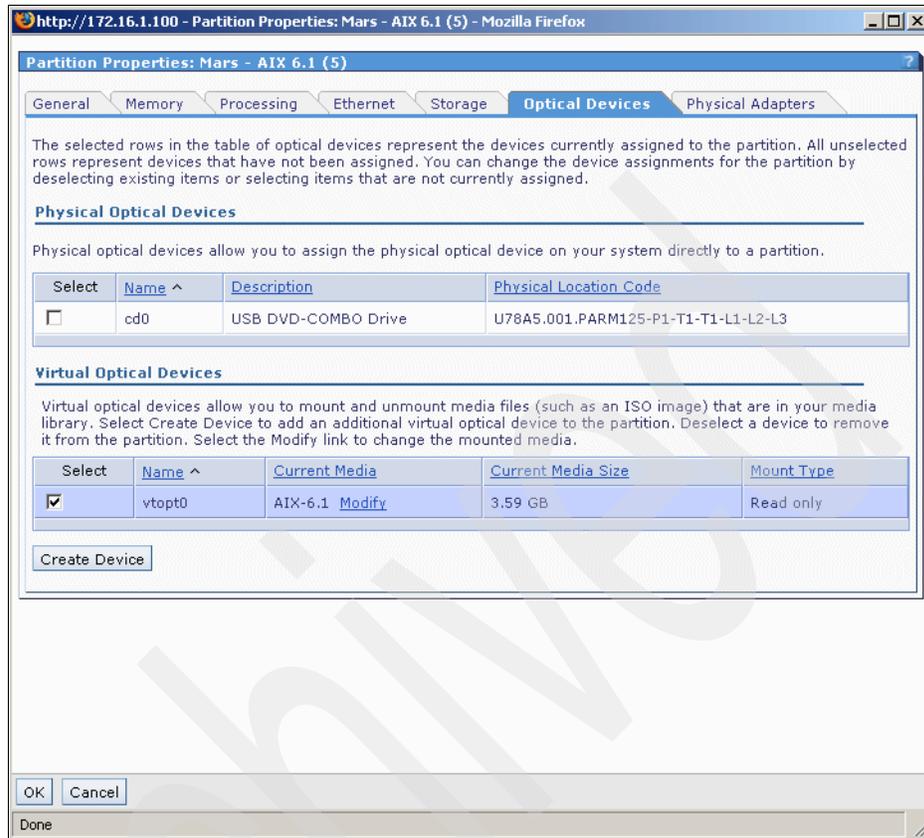


Figure 8-6 Virtual optical device to be removed

8.2.2 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 3.5.2, “Virtual Ethernet Adapters and SEA” on page 89.

The Resource Monitoring and Control (RMC) daemon must be active on the mobile partition or the validate and migration process will fail. The IVM UI can be used to verify the status as described in 3.7.11, “Partition Properties Changes and DLPAR operations” on page 129.

The CLI `lssyscfg` command can also be used to determine the RMC status as shown in Example 8-4.

Example 8-4 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/1opdiags/home.html>

Service and Productivity tools are discussed in Appendix F, “Service and productivity tools for Linux” on page 465.

8.3 Migrating the LPAR

The following sections describe how to use the IVM UI and CLI to validate, migrate, and check status on mobile LPAR.

8.3.1 Using the IVM UI

Validate

The migration process is started by first selecting **View/Modify Partitions** from the Navigation area. With this view open, use the check box for the desired mobile partition and select **Migrate** from the More Tasks drop-down box as shown in Figure 8-7 on page 341.

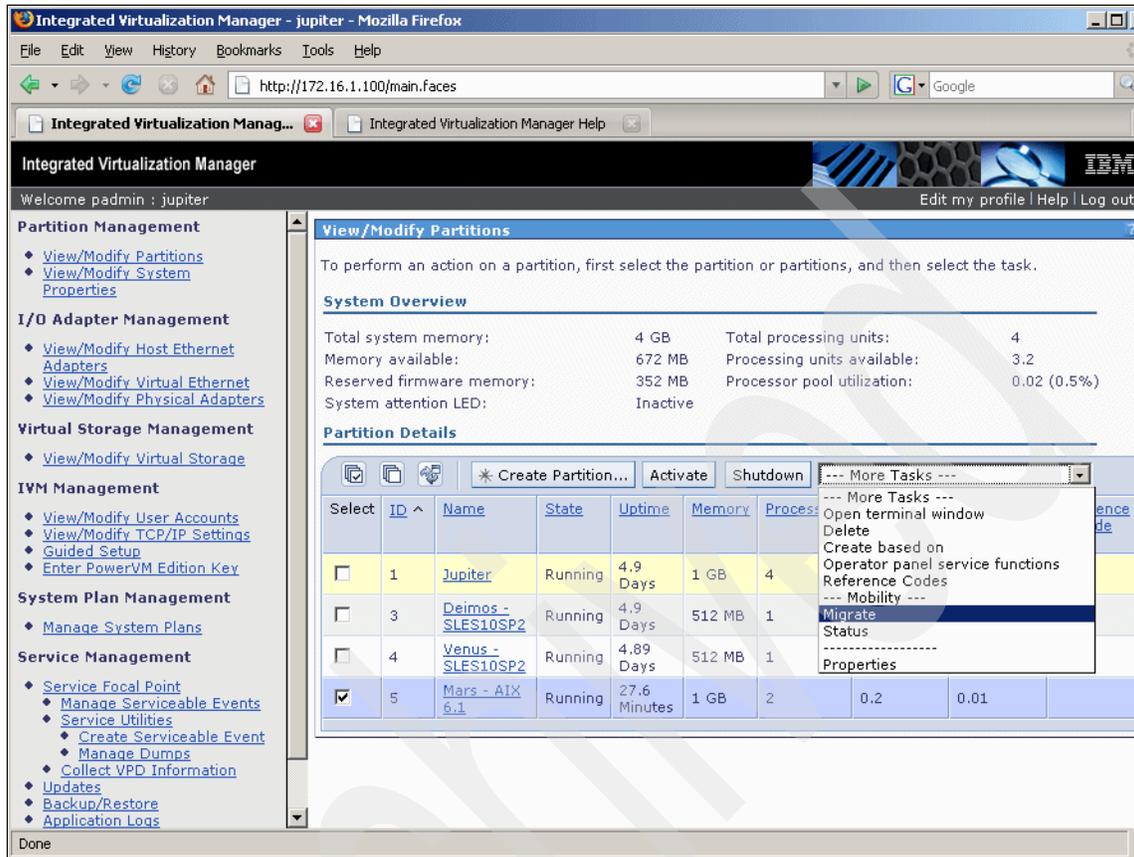


Figure 8-7 Partition Migrate option

The Migrate Partition view will open 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 8-8 on page 342. 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.

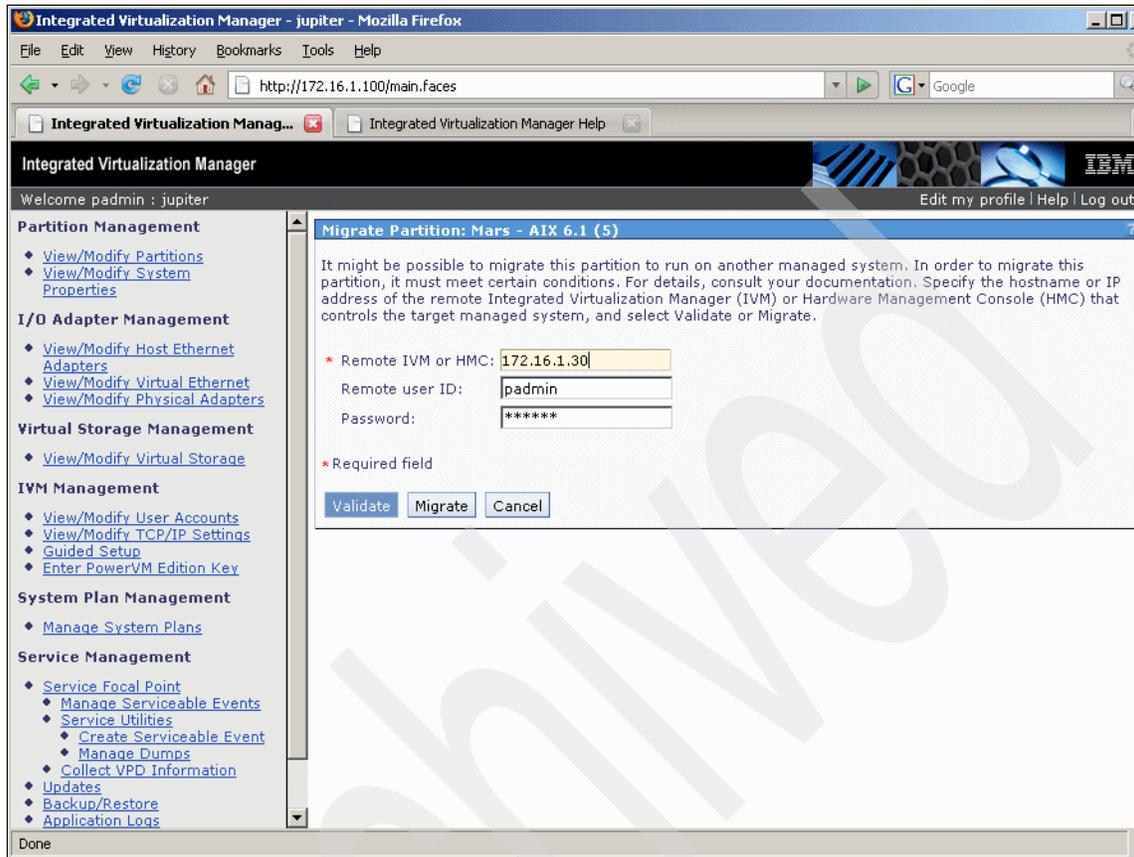


Figure 8-8 Partition Mobility validation

At the end of the successful validation process, the Migrate Partition window will be updated similar to Figure 8-9.

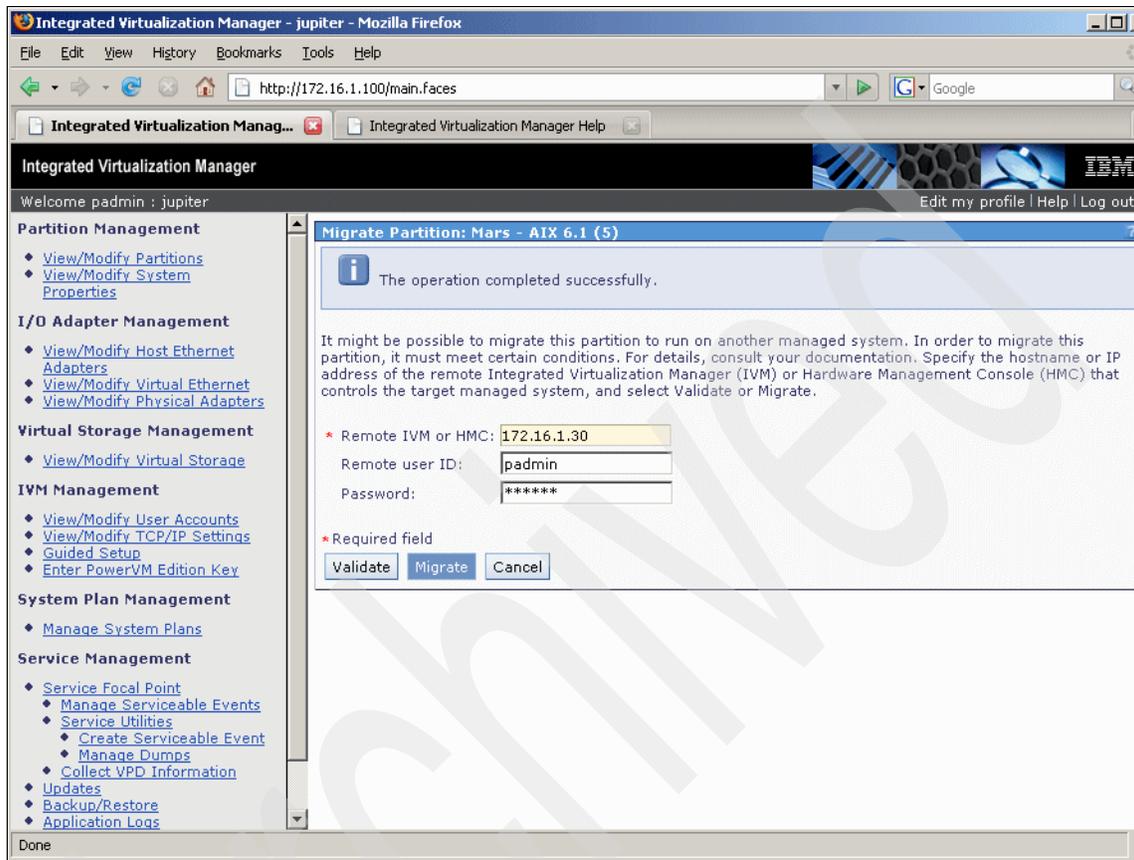


Figure 8-9 Partition Migration validation success

Figure 8-10 shows the results of the validation process that discovered a problem that would prevent a migration. This error message was generated because of a virtual SCSI assignment that could not be migrated. In this example the problem was due to a virtual optical device that had an assignment to the mobile partition. Another example is shown in Figure 8-4 on page 337, 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.

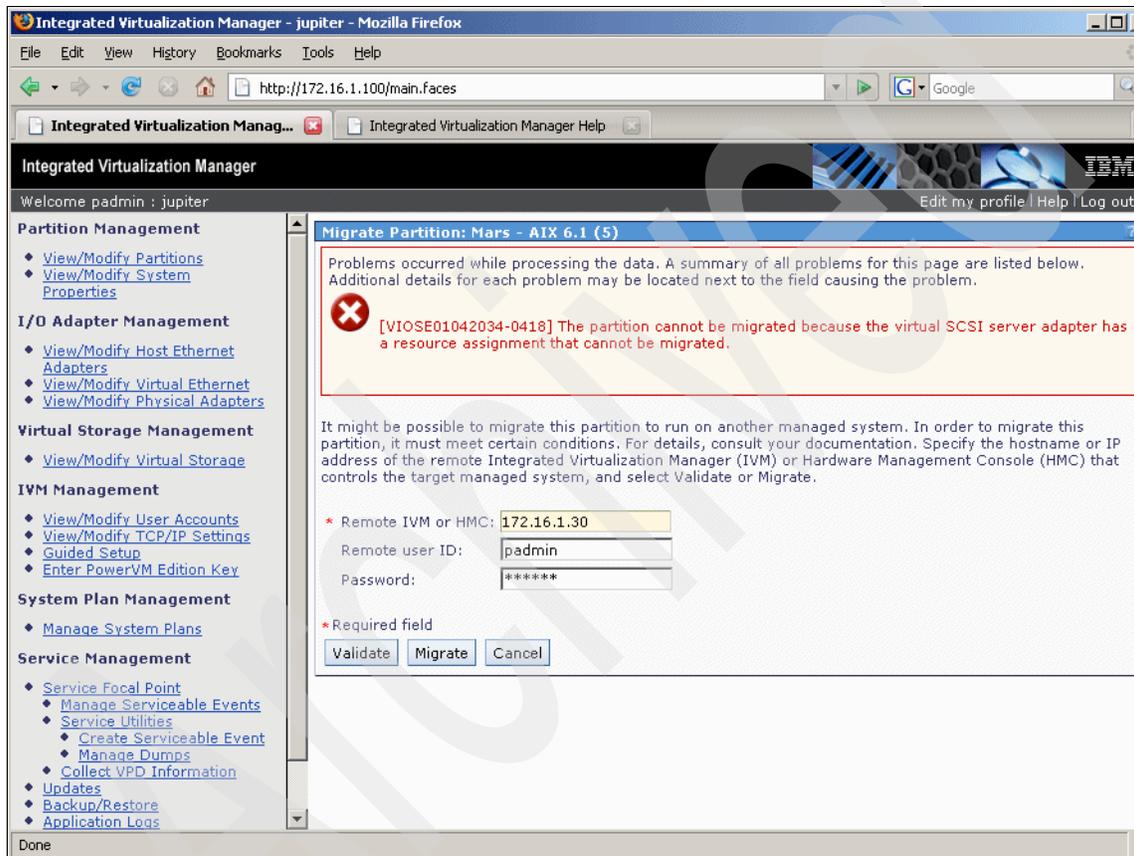


Figure 8-10 Failed validation because of improper virtual SCSI device assignment

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 validate is run again and at the end of this step a Migrate Status view will display, as shown in Figure 8-11.

Integrated Virtualization Manager - jupiter - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://172.16.1.100/main.faces

Integrated Virtualization Manager

Welcome padmin : jupiter Edit my profile | Help | Log out

Partition Management

- View/Modify Partitions
- View/Modify System Properties

I/O Adapter Management

- View/Modify Host Ethernet Adapters
- View/Modify Virtual Ethernet
- View/Modify Physical Adapters

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

- 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 Status

The following partitions are currently migrating. You may stop this operation or continue to monitor it. If the migration status shows an error, you should select Recover, which will attempt to complete the migration, or stop it as appropriate.

Partitions Migrating From This System

Select	Partition ^	Migration Status	Percent Complete	Remote Platform Manager	Remote System
<input type="checkbox"/>	Mars - AIX 6.1 (5)	Migration Starting	0%	172.16.1.30	Server-7998-61X-SN100069A

OK Cancel

Done

Figure 8-11 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 8-12. Also note in this same figure that the state of the mobile partition has changed from *Running* to *Migrating- Running*.

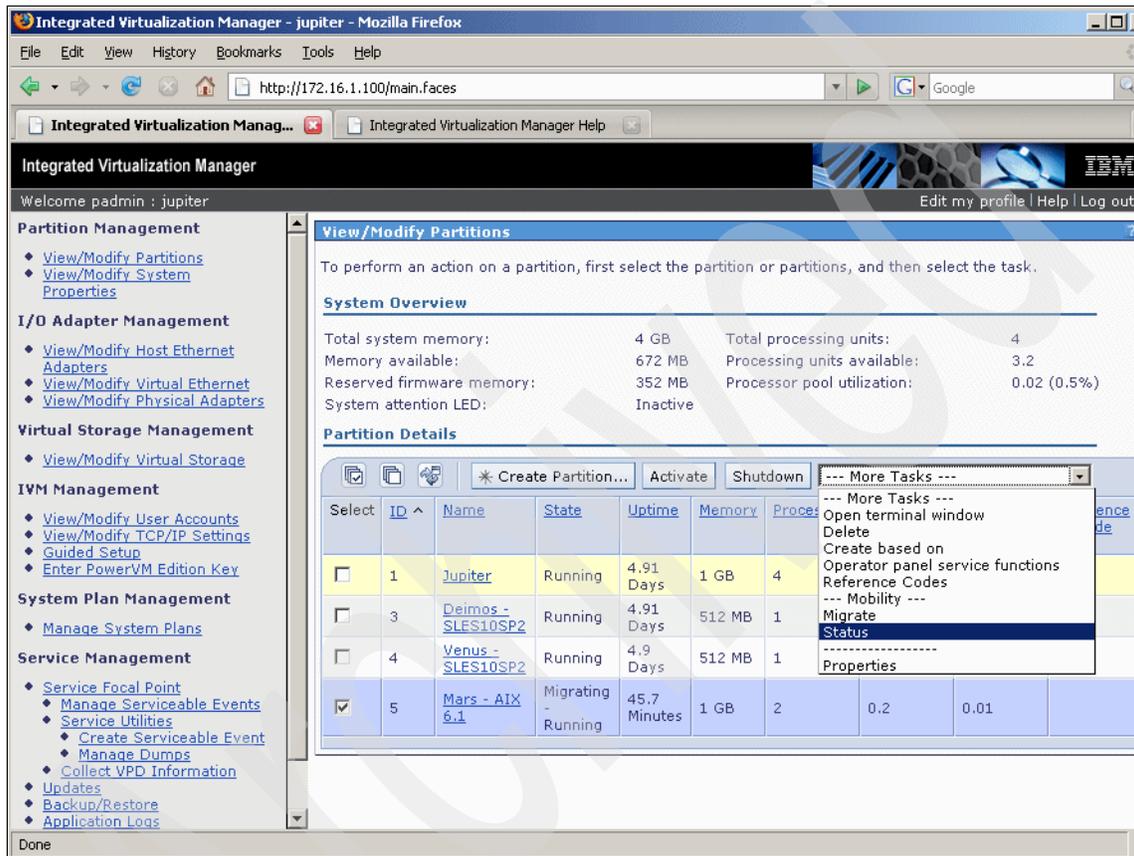


Figure 8-12 Starting the status view for a mobile partition

Figure 8-13 shows the View/Modify Partitions view on the remote IVM, indicating migration has started.

Note: The mobile partition will retain the same LPAR ID number if available on the remote system, otherwise it will be assigned the first available ID number.

The screenshot displays the Integrated Virtualization Manager (IVM) interface in a Mozilla Firefox browser window. The browser address bar shows the URL `http://172.16.1.30/main.faces`. The IVM interface includes a navigation menu on the left with sections for Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, and Service Management. The main content area is titled "View/Modify Partitions" and contains a "System Overview" section with the following data:

Total system memory:	4 GB	Total processing units:	4
Memory available:	1.14 GB	Processing units available:	2.4
Reserved firmware memory:	368 MB	Processor pool utilization:	1.06 (26.5%)
System attention LED:	Inactive		

Below the system overview is the "Partition Details" section, which includes a table of partitions:

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	VIOS-Neptune	Running	1.83 Days	1 GB	4	0.4	0.06	
<input type="checkbox"/>	2	Phobos - RHEL5-U2	Running	6.34 Hours	512 MB	1	1.0	1.00	CA00E105
<input type="checkbox"/>	5	Mars - AIX 6.1	Migrating - Running		1 GB	2	0.2		C20025FE

Figure 8-13 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 8-14 on the formerly remote system. On the original local system the mobile partition is removed from the View/Modify Partition view.

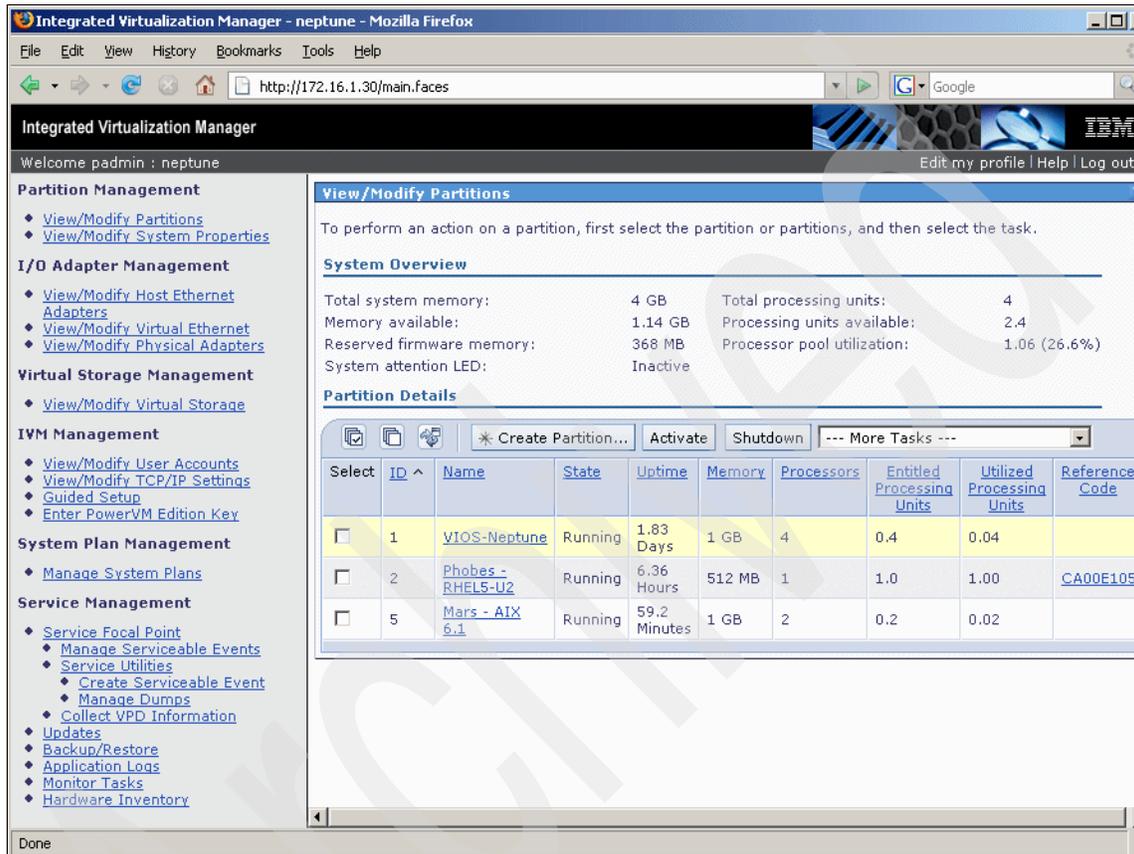


Figure 8-14 Partition migration complete to remote system

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 8-5 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 8-5 `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.

The -o flag or operation has the following options:

- ▶ s - stop
- ▶ m - validate and migrate
- ▶ r - recover
- ▶ v - validate

The -t flag in Example 8-5 on page 348 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 UI 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 8-6 shows the `migr1par` command with the operation flag with the m option for validate and migrate.

Example 8-6 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 use of the `lssyscfg` command as shown in Example 8-7 or the `ls1parmigr` command used in Example 8-8 on page 350 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 8-7 lssyscfg command used to check migrating partition status

```
$ migr1par -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 8-8 *lslparmigr* 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
```

```
$ lslparmigr -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_name=Jupiter,dest_msp_id=1,bytes_transmitted=1117792957,bytes_remaining=17162240,remote_manager=172.16.1.100,remote_user=admin
```

Once the migration is complete, the instance of the mobile LPAR will no longer appear on the original source VIOS.

JS12 and JS22 power management using EnergyScale technology

The EnergyScale technology described in 2.4, “IBM EnergyScale technology” on page 36 can be used by the BladeCenter Advanced Management Module and Active Energy Manager to monitor and control power usage of the IBM BladeCenter JS12 and JS22 blades. This chapter describes how to use the BladeCenter AMM and Active Energy Manager module of IBM Director to utilize these features.

An in-depth discussion of installing, configuring, and using Active Energy Manager can be found in *Going Green with IBM Active Energy Manager*, REDP-4361.

9.1 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 UIs. The information can also be collected by the Active Energy Manager standalone product or the Active Energy Manager module for IBM Director. This section describes how to use both interfaces of the AMM to make changes in the Power Saving and Power Capping modes.

9.1.1 Using the AMM Web UI for blade power management

A detailed description of how to access the AMM will not be covered in this section but detailed access and use information can be found in *IBM eServer BladeCenter Systems Management*, REDP-3582. The Web UI 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 9-1.

System Status Summary

● System is operating normally. All monitored parameters are OK.

The following links can be used to view the status of different components.

- [Blades](#)
- [I/O Modules](#)
- [Storage Modules](#)
- [Management Module](#)
- [Power Modules](#)
- [Power Module Cooling Devices](#)
- [Chassis Cooling Devices](#)
- [Media Tray](#)

Blades

Click the icon in the Status column to view detailed information about each blade.

Bay	Status	Name	Pwr	Owner**		cKVM*	I/O Compatibility	WOL*	Local Control		
				KVM	MT*				Pwr	KVM	MT*
1		No blade present									
2	●	JS22 NFS server	On	X			OK	N/A	X	X	X
3		No blade present									
4	●	JS12-1	On		X		OK	N/A	X	X	X
5		No blade present									
6		No blade present									

Figure 9-1 Selecting the Power Management option on AMM UI

The AMM UI now displays the Power Management view that is divided into the following sections, as shown in Figure 9-2 on page 355 and Figure 9-3 on page 356:

- ▶ BladeCenter Power Domain Summary
- ▶ BladeCenter Power Domain Planning
- ▶ BladeCenter Chassis Power Summary

- ▶ BladeCenter Chassis Configuration Setting
- ▶ BladeCenter Chassis Power Consumption

Note: In this example an IBM BladeCenter S is being used and has a single power domain. IBM BladeCenter H chassis have two power domains.

In each section a complete description of the fields can be obtained by clicking the blue question mark.

The Power Domain Summary section provides information on the status of the power domain, power modules, power management policy, maximum power limit, power in use and the type of power service, total chassis power usage, power allocation, and maximum power available. The link **AC Power Source Redundancy** can be clicked to review the current power domain management policy or change to a different option. This section also provides a link to a power domain detail view that can be accessed by clicking the **Power Domain** link. The detail view is described in “AMM Power domain details view” on page 357.

The Power Domain Planning section shows the maximum power limit, the maximum allocated power, and the remaining power available.

The Chassis Power Summary section shows total DC power available, total AC power in use and total thermal output.

The Chassis Configuration Setting section allows modification of the acoustic mode setting and data sampling interval from the drop-down boxes.

The Chassis Power Consumption section displays a graph of the power consumption history in watts. The horizontal axis or time interval can be changed from the drop-down box to report the last hour up to the last 24 hours.

172.16.0.125 BladeCenter Advanced Management Module - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://172.16.0.125/private/main.php?

IBM BladeCenter[®] S Advanced Management Module

Bay 1: SN#YK168082D1ZG
User: USERID

- 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
 - Admin/Power/Restart
 - Configuration
 - Firmware Update
- Storage Tasks
 - Configuration
- 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
 - AMM Service Data
 - Blade Service Data
 - AMM Status

BladeCenter Power Domain Summary

	Power Domain
Status	● Power domain status is good.
Power Modules	Bay 1: 1450W Bay 2: 1450W Bay 3: 1450W Bay 4: 1450W
Power Management Policy	AC Power Source Redundancy Intended for dual AC power sources into the chassis. Total allowed power limited to the capacity of two Power Modules. This is the most conservative approach and is recommended when all four Power Modules are installed. the chassis is correctly wired with dual AC power sources, one AC power can fail without affecting blade operation. Note that some blades may not be allowed to power on if doing so would exceed the policy power limit.
Maximum Power Limit †	2900W
Power in Use ††	738W
Power Service	220 VAC

BladeCenter Power Domain Planning

	Power Domain
Maximum Power Limit †	2900W
- Allocated Power (Max) †††	1346W
= Remaining Power	1554W

† Maximum power available based on the number of power modules and the Power Management Policy setting.
 †† Represents the maximum worst case and measured power based on the capability of all components.
 ††† Reserved power for all components in this domain.

BladeCenter Chassis Power Summary

Total DC Power Available	2900W
Total AC Power In Use	773W
Total Thermal Output	2,637.5 BTU/Hour

Figure 9-2 AMM Power Management view part 1

172.16.0.125 BladeCenter Advanced Management Module - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://172.16.0.125/private/main.php? Google

IBM BladeCenter[®] S Advanced Management Module

Bay 1: SN#YK168082D1ZG
User: USERID

- 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
 - Admin/Power/Restart
 - Configuration
 - Firmware Update
- Storage Tasks
 - Configuration
- 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
 - AMM Service Data
 - Blade Service Data
 - AMM Status

BladeCenter Chassis Power Summary

Total DC Power Available	2900W
Total AC Power In Use	776W
Total Thermal Output	2,647.7 BTU/Hour

BladeCenter Chassis Configuration Setting

These settings apply to the entire chassis (including the empty bays).

Acoustic mode:

Data Sampling Interval:

BladeCenter Chassis Power Consumption

Power Consumption History

Trend Period:

AC Power (Watts)

Time

Power consumption

Current usage:

Average: 774

Figure 9-3 AMM Power Management view part 2

AMM Power domain details view

The BladeCenter Power Domain Details view in Figure 9-4 on page 358 and Figure 9-5 on page 359 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 shown for each module. Modules that have specific capabilities or collect power trend data appear as a link to a module-specific detail view.

Archived

172.16.0.125 BladeCenter Advanced Management Module - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://172.16.0.125/private/main.php? Google

IBM BladeCenter[®] S Advanced Management Module

Bay 1: SN#YK168082D1ZG
User: USERID

- 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
- IO Module Tasks
 - Admin/Power/Restart
 - Configuration
 - Firmware Update
- Storage Tasks
 - Configuration
- 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
 - AMM Service Data
 - Blade Service Data

BladeCenter Power Domain Details

Links: [Power Summary](#)

Bay(s)	Status	Module	State	Power In Use	Allocated Power		CPU Duty Cycles
					Maximum	Minimum	
<i>Chassis Components</i>							
		Midplane	On	10W	10W	10W	n/a
1		Media Module	On	10W	10W	10W	n/a
<i>Power Module Cooling Devices</i>							
1		Power Module	On	10W	10W	10W	n/a
2		Power Module	On	10W	10W	10W	n/a
3		Power Module	On	10W	10W	10W	n/a
4		Power Module	On	10W	10W	10W	n/a
<i>Chassis Cooling Devices</i>							
1		Cooling Module	On	9W	68W	68W	n/a
2		Cooling Module	On	9W	68W	68W	n/a
3		Cooling Module	On	10W	68W	68W	n/a
4		Cooling Module	On	9W	68W	68W	n/a
<i>Storage</i>							
1		Storage Module	On	100W	100W	100W	n/a
2		Storage Module	On	100W	100W	100W	n/a
<i>Management Modules</i>							
1		SN#YK168082D1ZG	On	25W	25W	25W	n/a
<i>I/O Modules</i>							
1		Ethernet SM	On	45W	45W	45W	n/a
3		SAS Conn Mod	On	60W	60W	60W	n/a
<i>Blades</i>							
2		JS22 NFS server	On	202W	402W	368W	n/a ††
4		JS12-1	On	149W	282W	256W	n/a ††
DOMAIN TOTALS				Power In Use	Maximum	Minimum	
Power Allocation				778W	1346W	1286W	

Done

Figure 9-4 AMM Power Domain detail view part 1

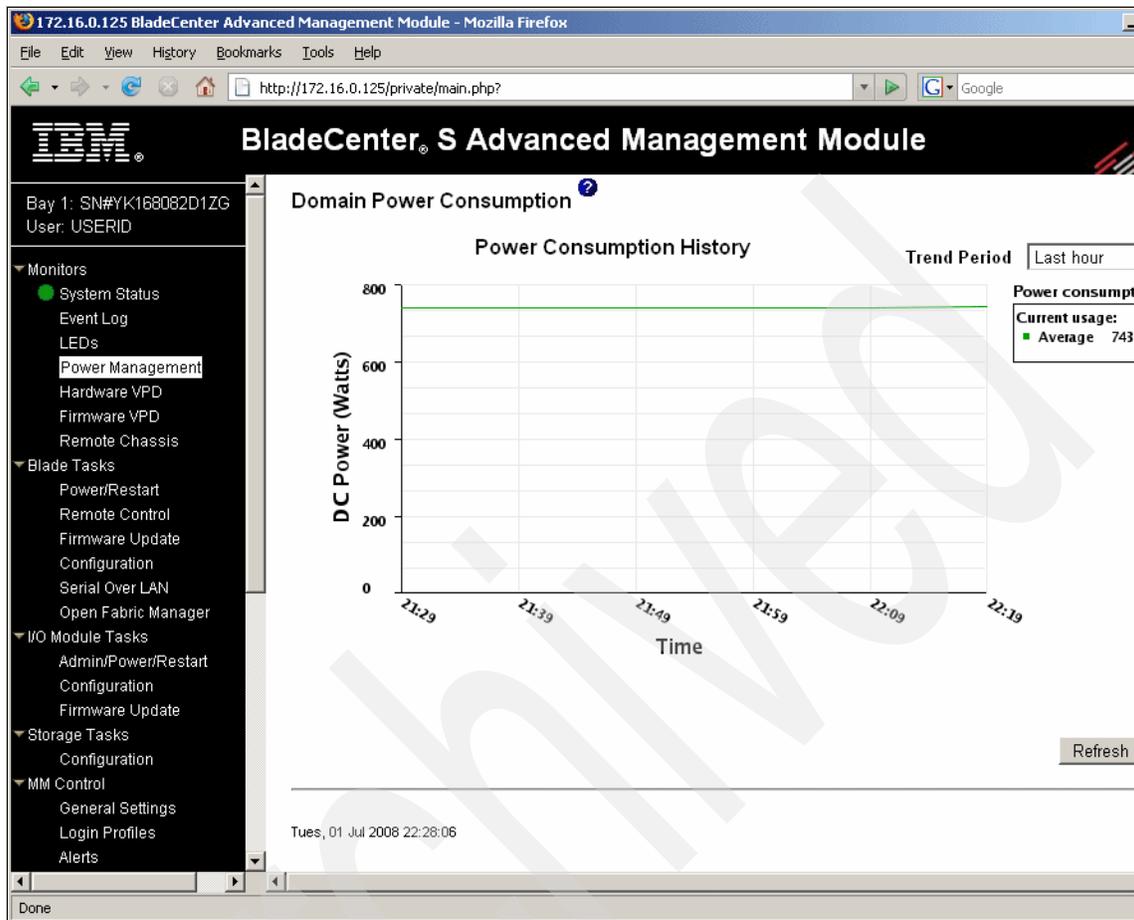


Figure 9-5 AMM Power Domain detail view part 2

AMM Blade bay power summary view

Clicking the link that represents an IBM BladeCenter JS12 or 22 blade from the power domain details view will display current information about the EnergyScale capabilities of the blade. The number of processors, effective and maximum CPU speed, drop-down boxes to enable Power Capping and Power Saver Mode, and power consumption trend data are also shown in this view.

Figure 9-6 on page 360 and Figure 9-7 on page 361 show an example of an IBM BladeCenter JS12. The drop-down boxes in the Configuration Setting section of this view allow the Power Capping and Power Saver Mode to be enabled/disabled. The Power Consumption chart can be modified from the drop-down box to display the last hour up to the last 24 hours of trend data. The chart can also be modified to display average, maximum, or minimum values.

The screenshot displays the IBM BladeCenter S Advanced Management Module web interface in a Mozilla Firefox browser. The browser's address bar shows the URL `http://172.16.0.125/private/main.php?`. The page title is "BladeCenter S Advanced Management Module".

On the left side, there is a navigation menu with the following items:

- Bay 1: SN#YK168082D1ZG
- User: USERID
- 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
 - Admin/Power/Restart
 - Configuration
 - Firmware Update
- Storage Tasks
 - Configuration
- MM Control
 - General Settings

The main content area is titled "BladeCenter JS12-1 (Bay 4) Power Summary" and includes a help icon. Below the title are two links: [Power Summary](#) and [Power Domain](#).

A table displays the power summary data:

Capability	Power metering is supported. Power can be dynamically measured and capped. Power
Processors	2
Effective CPU Speed	3800 MHz
Maximum CPU Speed	3800 MHz

Below the table is the "BladeCenter JS12-1 (Bay 4) Configuration Setting" section, also with a help icon. It contains three configuration items:

- Power Capping:
- Maximum Power (range 256-282):
- Power Saver Mode:

A "Save" button is located at the bottom right of the configuration section.

The browser's status bar at the bottom shows "Done".

Figure 9-6 Blade bay power summary part 1

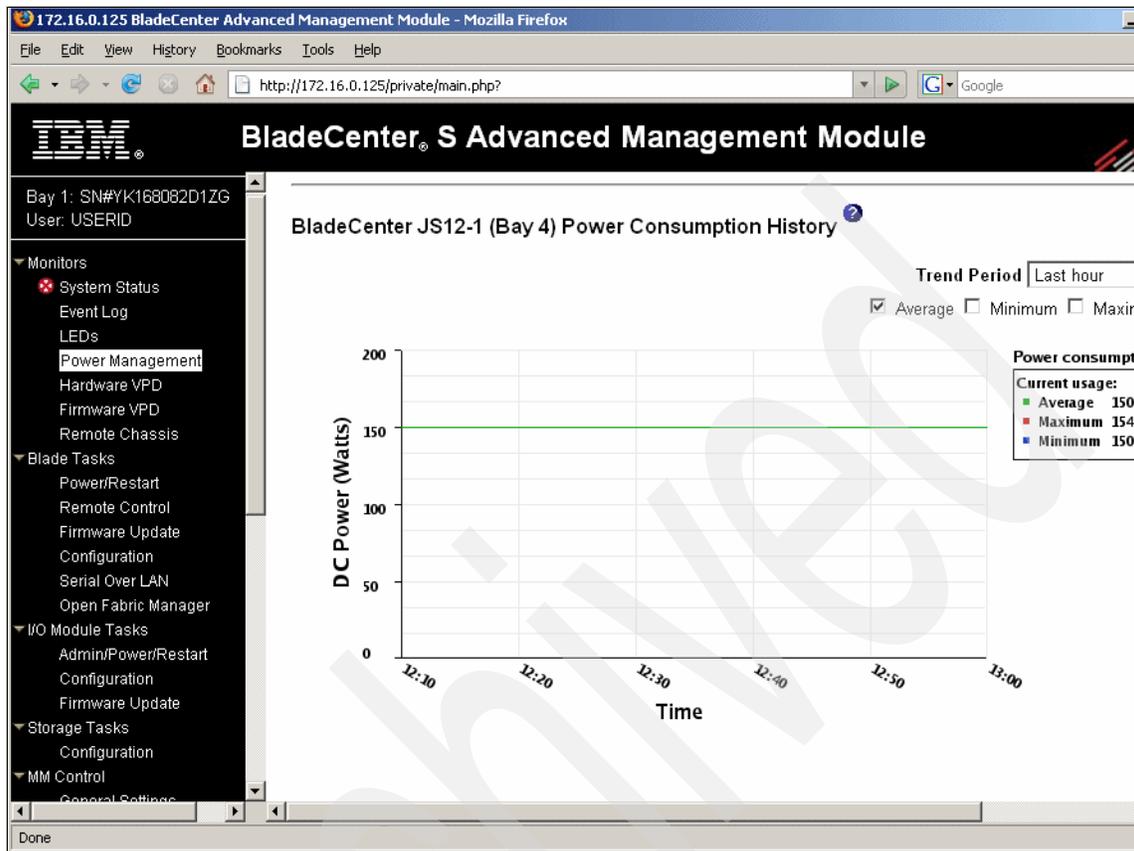


Figure 9-7 Blade bay power summary part 2

9.1.2 Using the AMM CLI UI for blade power management

Similar to the Web UI, the 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. The login is completed by using the same user ID and password that is used for the Web UI.

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 JS12 or 22. This command with all options is shown in Example 9-1.

Example 9-1 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 9-2 shows the **fuelg** command used from the `system>` prompt with no flags to display the BladeCenter Power Domain information.

Example 9-2 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 9-3 shows the environment being set to blade slot 4. Note the prompt change.

Example 9-3 The env command used to set a persistent target

```
system> env -T blade[4]
OK
system:blade[4]>
```

Example 9-4 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 9-4 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 **fuelg** command flags for changing blade specific EnergyScale parameters are:

- ▶ **-int *interval*** - polling interval between 10 and 60 minutes in increments of 5.
- ▶ **-pcap *value*** - power cap for blades, the value between the min and max amounts, setting pcap also enables pme.
- ▶ **-pme *on/off*** - power management and capping enabling/disabled uses min value unless a different pcap value has been used.
- ▶ **-pt *interval*** - power trending data 1,6,12, or 24 hours.

9.2 Power management through AEM

The Active Energy Manager module for IBM Director provides expanded control and monitoring capabilities. In addition, AEM provides scheduling options to control the execution of power saver mode and power capping options automatically at predetermined times.

The following information and examples assume that IBM Director and the Active Energy Manager module have been installed and configured. Complete planning, installation, configuring, and usage information of IBM Director can be found in *Implementing IBM Director 5.20*, SG24-6188. IBM Active Energy Manager installation and use can be found in *Going Green with IBM Active Energy Manager*, REDP-4361.

9.2.1 Active Energy Manager Console

Figure 9-8 on page 365 shows the selection of a target blade from the IBM Director-managed object context menu. Right-click the managed object and select **Active Energy Manager** from the menu. Alternatively, the Active Energy Manager task from the Director task menu can be dragged and dropped on the target blade.

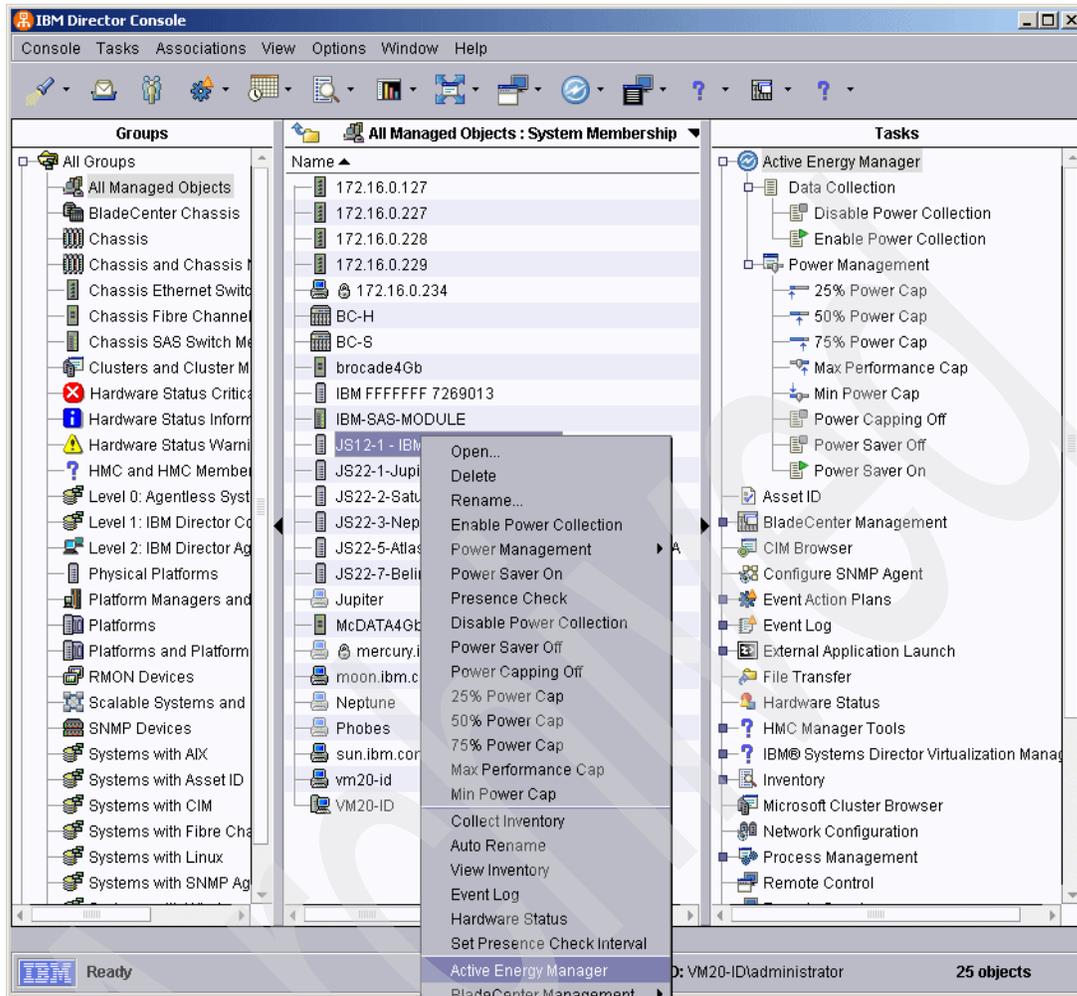


Figure 9-8 Starting the AEM task for a blade object in IBM Director

When the selection is made, the Active Energy Manager console opens with the Current Data view displaying physical target identification and EnergyScale information, as shown in Figure 9-9.

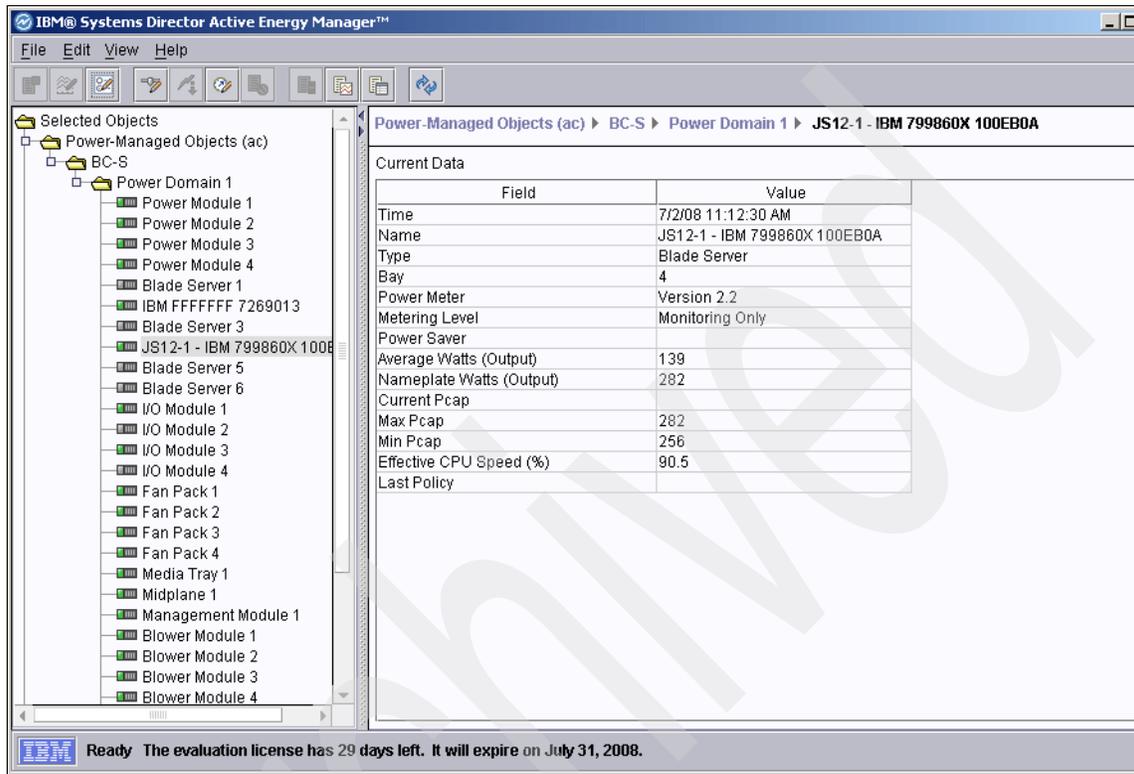


Figure 9-9 AEM console showing Current Data view for target object

Trend data can be viewed in chart or table form by selecting the appropriate icon from the AEM console tool bar, as shown in Figure 9-10 and Figure 9-11.



Figure 9-10 Show Trend Data as Chart Icon



Figure 9-11 Show Trend Data as Table Icon

When selecting the chart icon, the console view changes, as shown in Figure 9-12.

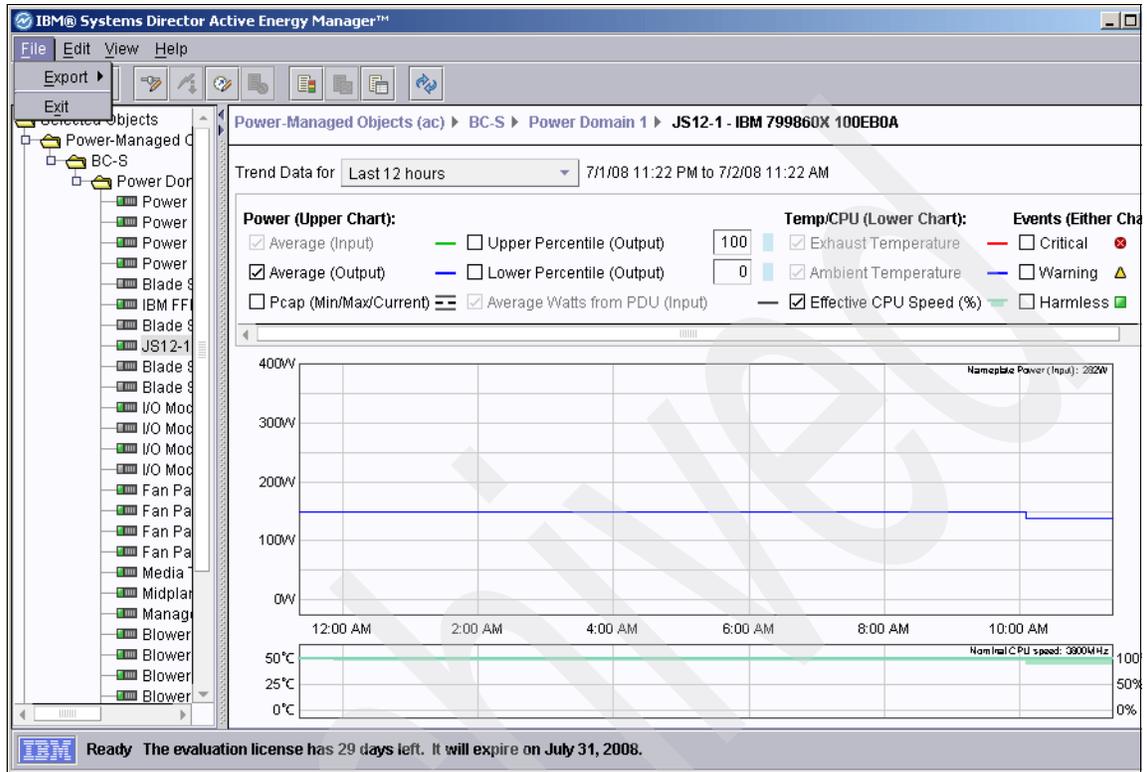


Figure 9-12 AEM console trend data as a chart

Figure 9-13 shows the trend data view in a table format after selecting the table icon from the AEM console window.

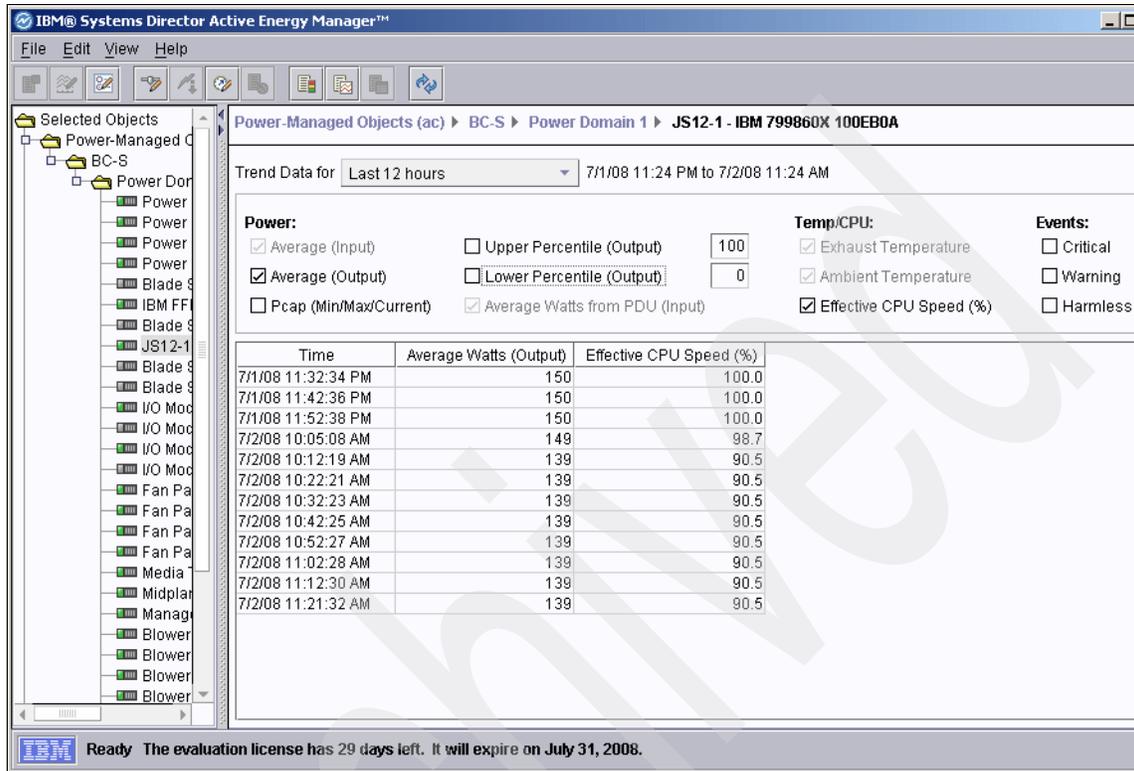


Figure 9-13 AEM console trend data as a table

Changing the EnergyScale features that directly alter the energy usage of the blade can be accessed from the Manage Power dialog box. This box is opened by clicking the **Manage Power** icon on the AEM console tool bar shown in Figure 9-14.



Figure 9-14 Manage Power Icon

When the Manage Power icon is selected, the dialog box shown in Figure 9-15 opens. The Manage Power dialog box provides the ability to create new power policies for a target and enable Power Save Mode and Power Cap modes.

Note: If a new policy is created, it will be available by right-clicking as a selection under a managed object when viewed from the IBM Director Console Managed Objects menu list. In addition, the new policy will be available in the IBM Director Tasks list menu under Active Energy Manager/Power Management.

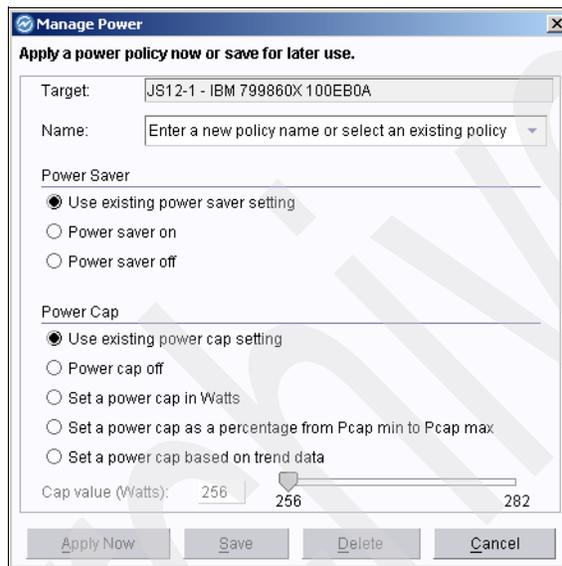


Figure 9-15 Manage Power dialog

9.3 Power Saver Mode example

This example enables Power Saver Mode and shows power trend data using three different methods:

- ▶ Advanced Management Module Web UI
- ▶ Advanced Management Module CLI
- ▶ IBM Director and Active Energy Management

9.3.1 Enabling Power Saver Mode from the AMM

The following example of enabling Power Saver Mode from the AMM Web UI and AMM CLI uses an IBM BladeCenter JS12 installed in an IBM BladeCenter S chassis. The JS12 has 16 GB of memory, two 73-GB 10 K SAS internal drives, and a SAS expansion card. IBM Virtual I/O Server was installed in LPAR 1 and active, IBM i in LPAR 2 but not active, and AIX6.1 in LPAR and active.

The configuration was changed when working with IBM Director and Active Energy Manager from the previous examples. This configuration uses an IBM BladeCenter JS22 installed in an IBM BladeCenter H. The JS22 has 4 GB of memory, one 73-GB 10 K SAS internal drive, one Emulex CFFv Fibre Channel expansion card, and one Qlogic CFFh “combo” expansion card with two Ethernet and two Fibre Channel ports. IBM Virtual I/O Server was installed in LPAR 1 and active, Red Hat Enterprise Linux 5.2 in LPAR 2 but not active, and AIX6.1 in LPAR and active.

Power Saver Mode enabled from the AMM Web UI

To enable Power Saver Mode from the AMM Web UI, use the blade bay power summary view shown in Figure 9-16 on page 371. Use the drop-down box for Power Saver Mode and select **Enabled**. Next click **Save** to complete the selection. Clicking **Refresh** at the bottom of this window will update the Effective CPU Speed. It can take several minutes to stabilize at the final effective CPU speed.

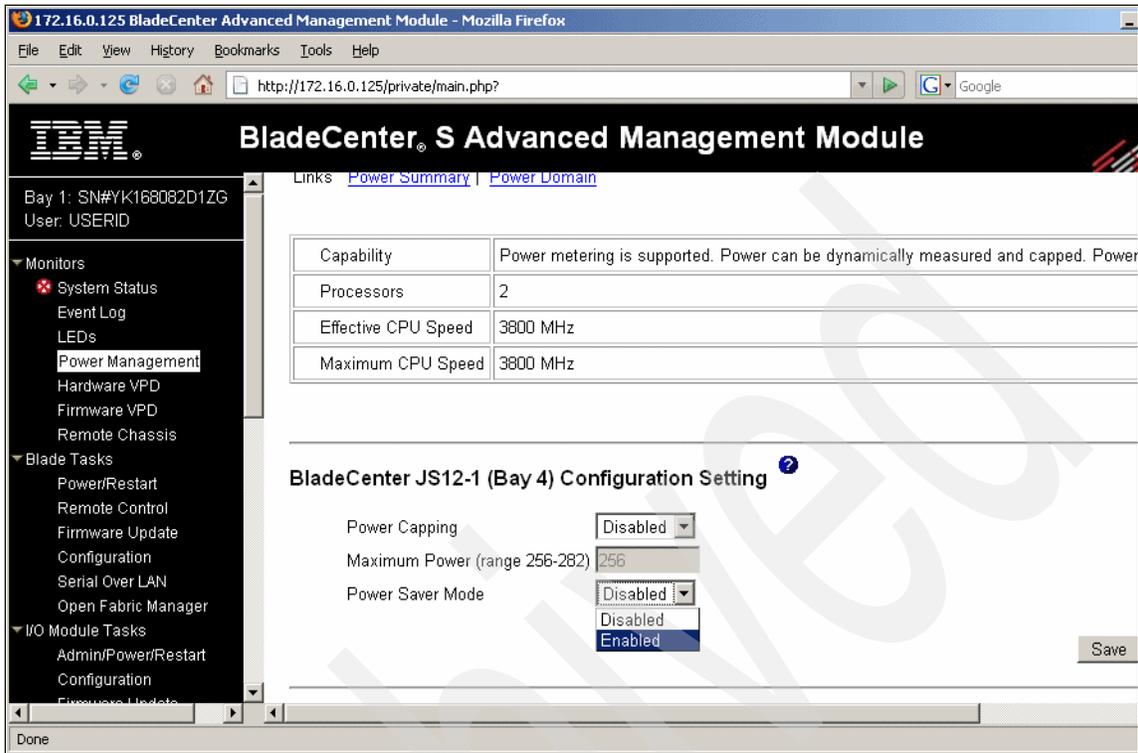


Figure 9-16 AMM Web UI used to enable Power Saver mode

Power trend information can be reviewed as a graphical representation from the lower half of the same view. In this example the previous six hours are shown in Figure 9-17.

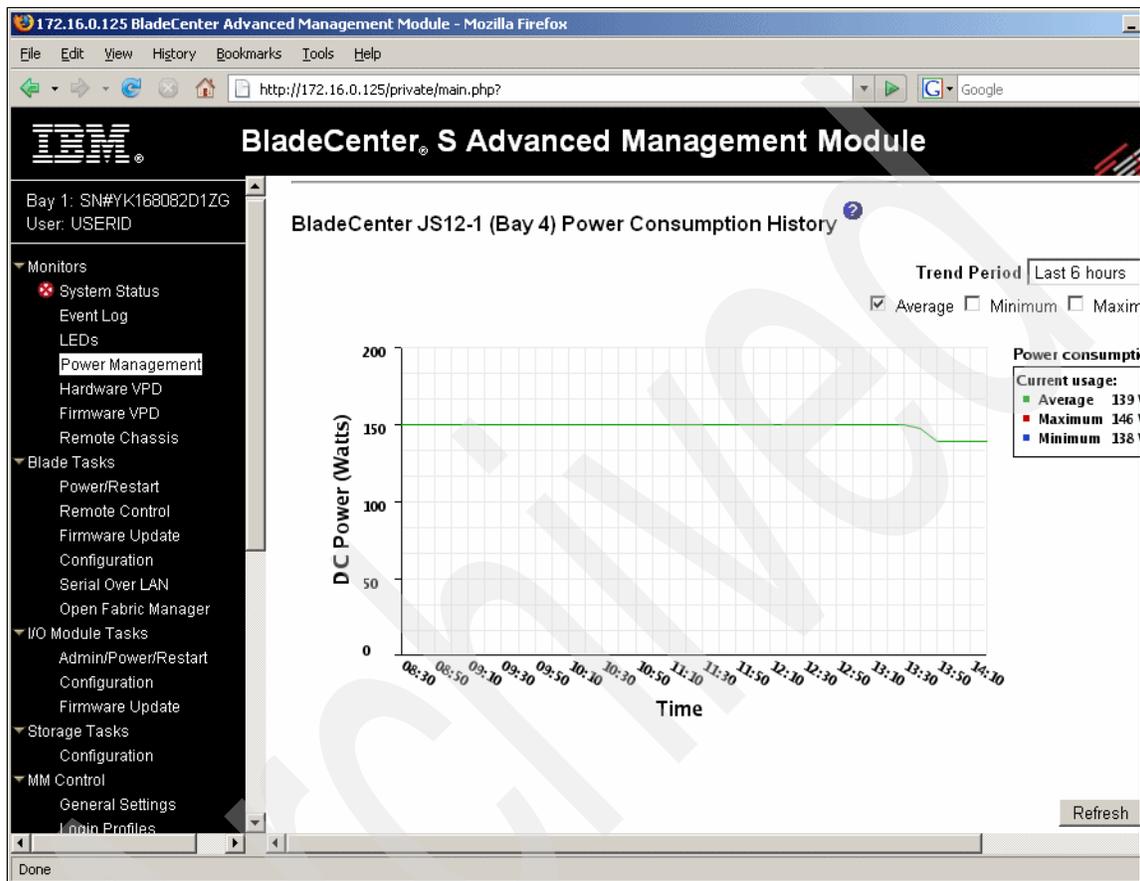


Figure 9-17 Power trend data in Power Saver Mode

Power Saver Mode enabled from the AMM CLI

The baseline values and power consumption are shown in Example 9-4 on page 363. Power Saver Mode was enabled using the `fuelg -ps` command and then the changes reviewed using the `fuelg` command shown in Example 9-5.

Example 9-5 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 9-6.

Example 9-6 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
```

9.3.2 Power Saver Mode enabled from IBM Director/AEM

Enabling Power Saver Mode from IBM Director and AEM is a standard policy that is available from the Director tasks list under Active Energy Manager/Power Management /Power Saver On. This policy can be enabled on a JS12 or JS22 target blade. From the Managed Objects menu list, right-click the desired target as shown in Figure 9-18 on page 374 and click **Power Saver On**. An alternative way to enable this policy is to drag and drop the Power Saver On task from the Tasks menu under Active Energy Manager/Power Management to the managed object target.

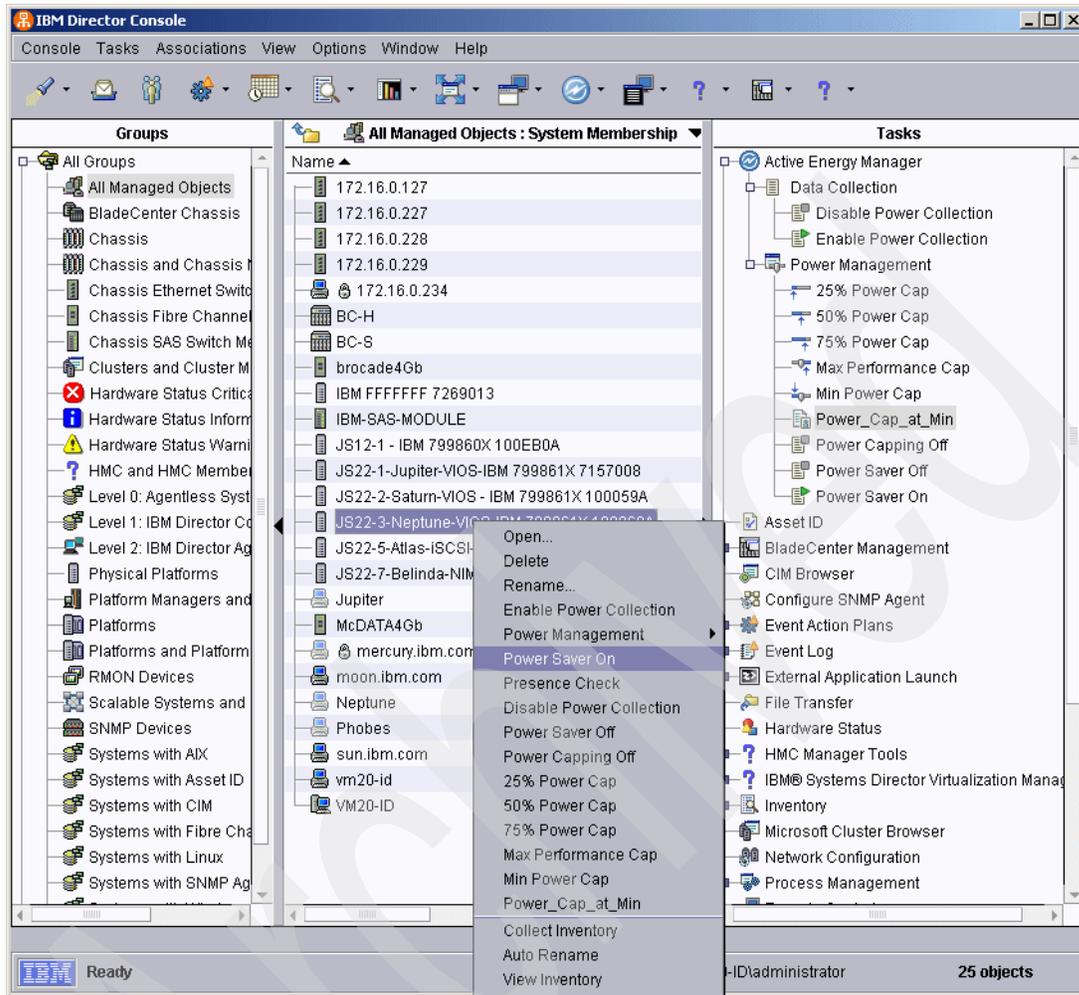


Figure 9-18 Enable Power Saver Mode from the IBM Director console

The actual execution of the task can be immediate or scheduled, as shown in Figure 9-19. For this example we chose the Execute Now option.

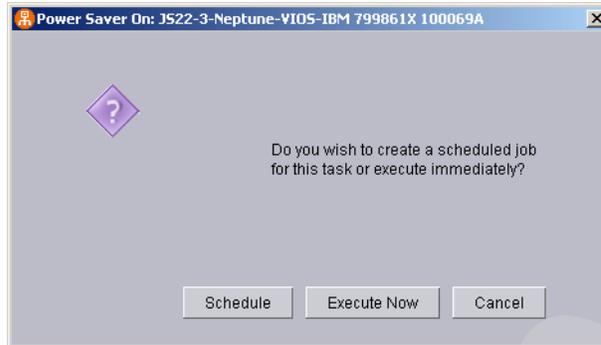


Figure 9-19 Power Saver On execution dialog box

When **Execute Now** is selected, an Execution History window as shown in Figure 9-20 is opened displaying the status of the task. After reviewing the history, this window may be closed.

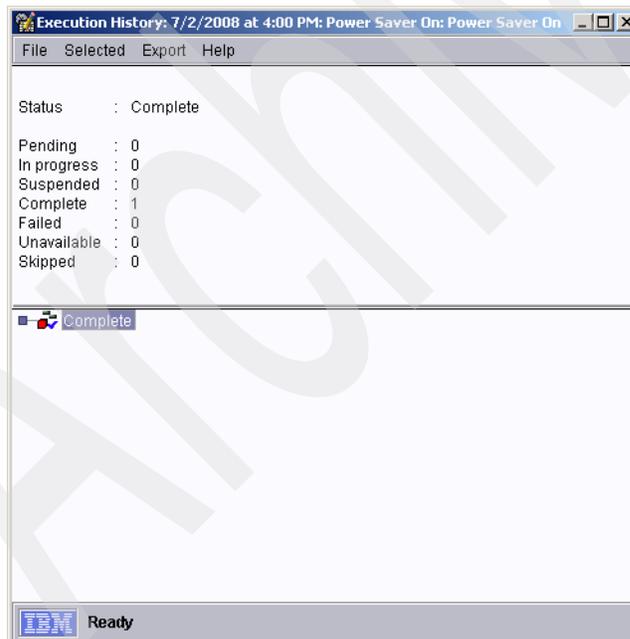


Figure 9-20 Task Execution History window

Power trend data in chart form is shown in Figure 9-21 and table format in Figure 9-22 on page 377. Both of these formats had the Events check boxes marked to display the Power Saver On task execution.

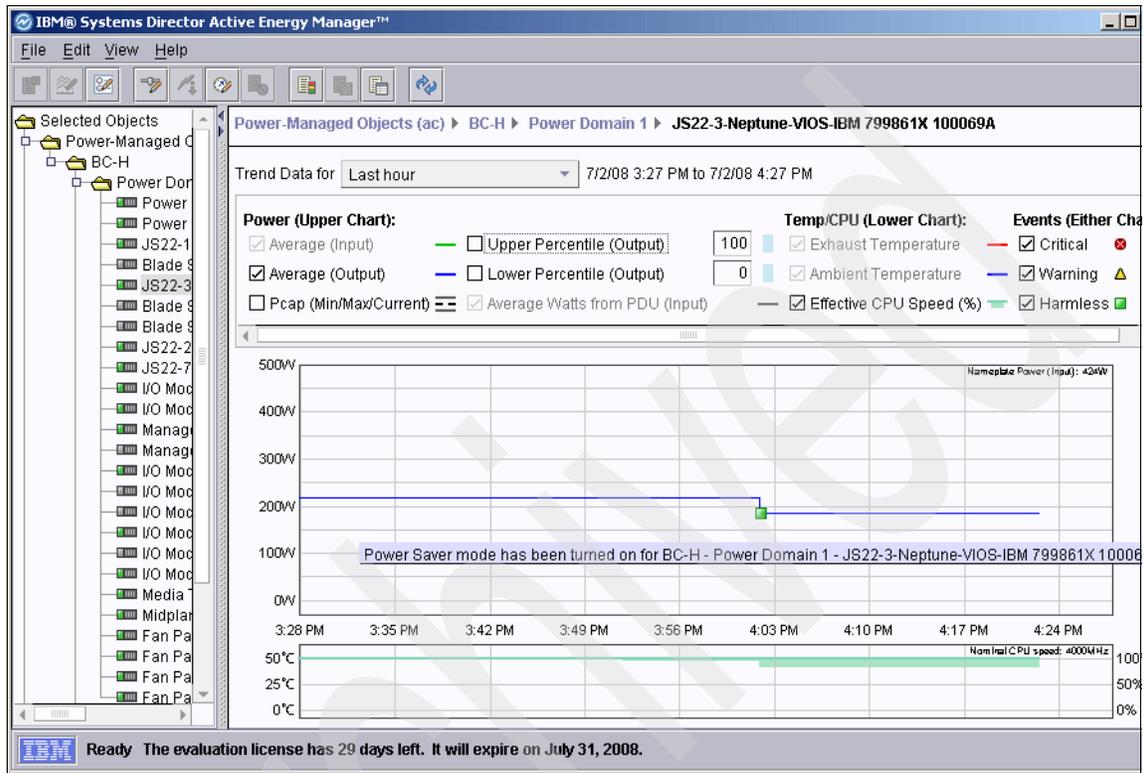


Figure 9-21 Power trend chart showing a Power Saver On event

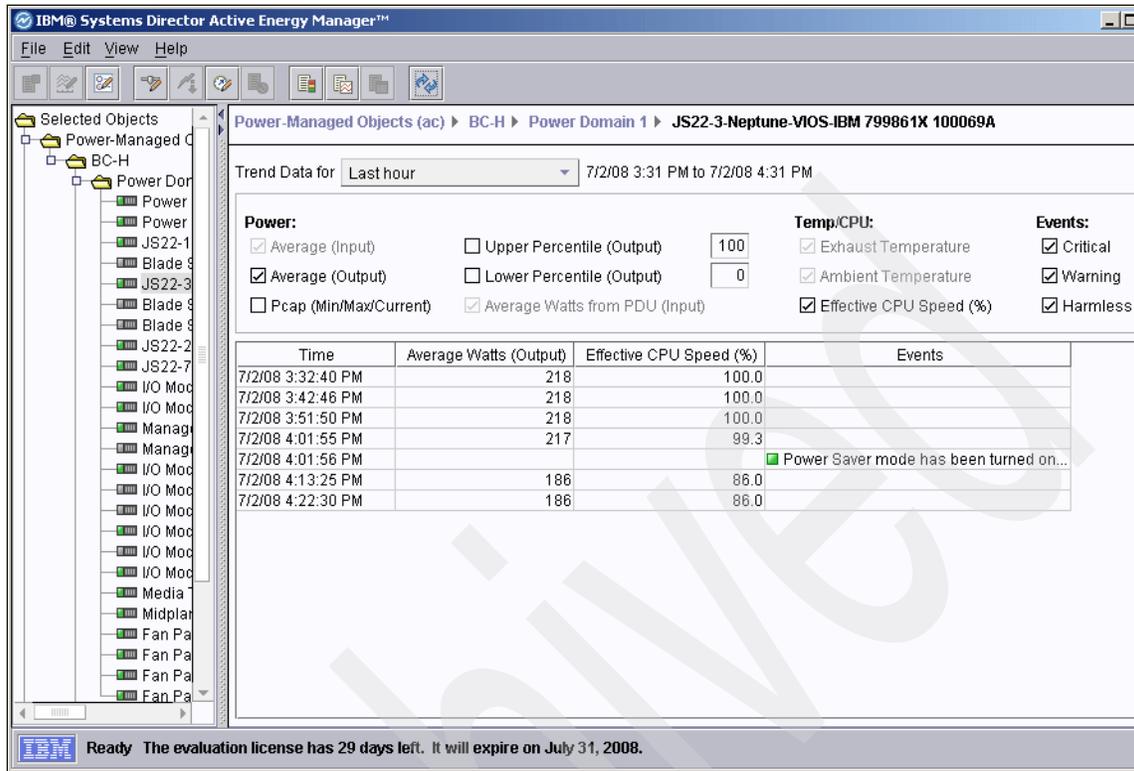


Figure 9-22 Power trend table showing a Power Saver On event

9.3.3 Scheduling Power Saver Mode with IBM Director and AEM

This example uses the Active Energy Module in IBM Director to create and schedule a policy enabling and disabling Power Saver Mode.

The same IBM BladeCenter JS22 configuration used previously in demonstrating manual activation of Power Saver Mode will be used.

Creating a policy with Active Energy Manager

In this simple example, the Power Saver Mode is enabled for a JS22 on Friday evenings at 6:30 p.m. and disabled Monday mornings at 5:00 a.m.

From the Active Energy Manager console that is opened on the target blade click the **Manager Power** icon as shown in Figure 9-14 on page 368 to open the Manage Power dialog box. In the Manage Power dialog box enter a policy name, and select a Power Saver and Power Cap setting.

Note: Note that the selection in the Power Cap section must be changed from the Use existing power cap setting before the new policy can be changed.

The values chosen were (Figure 9-23):

- ▶ A policy name of Weekend_Throttle_Back (no spaces allowed)
- ▶ Power Saver mode to Power saver on
- ▶ Power Cap mode to Power cap off

Manage Power

Apply a power policy now or save for later use.

Target: JS22-3-Neptune-VIOS-IBM 799861X 100069A

Name: Weekend_Throttle_Back

Power Saver

Use existing power saver setting

Power saver on

Power saver off

Power Cap

Use existing power cap setting

Power cap off

Set a power cap in Watts

Set a power cap as a percentage from Pcap min to Pcap max

Set a power cap based on trend data

Cap value (Watts): 424 372 424

Apply Now Save Delete Cancel

Figure 9-23 Creating a new policy in the Manage Power dialog box

The policy can now be applied, either immediately or scheduled or saved. For this example it is saved. Click **Save** and the dialog box closes. The new policy appears in the Director tasks lists under the Active Energy Manager/Power Management structure and as a section when right-clicking any JS12 or JS22 blade, as shown in Figure 9-24.

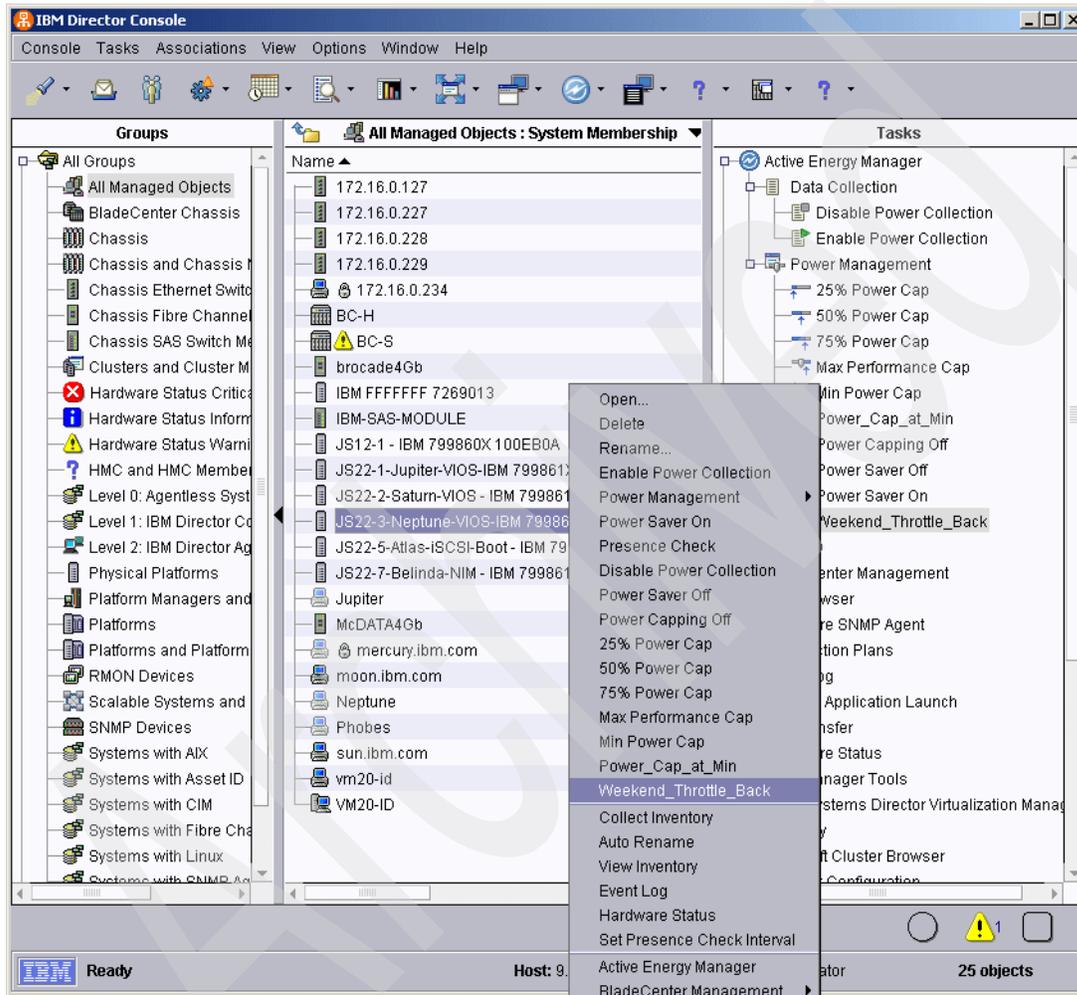


Figure 9-24 Director console showing a new policy addition

Creating a scheduled job with IBM Director

With the new Weekend_Throttle_Back policy created, it can be scheduled with IBM Director. This example only shows one method of creating a job schedule.

From the Director-managed objects list right-click the target blade and select the desired policy. A policy dialog box opens as shown in Figure 9-25, offering the choice of Schedule or Execute Now; for this example we chose **Schedule**.

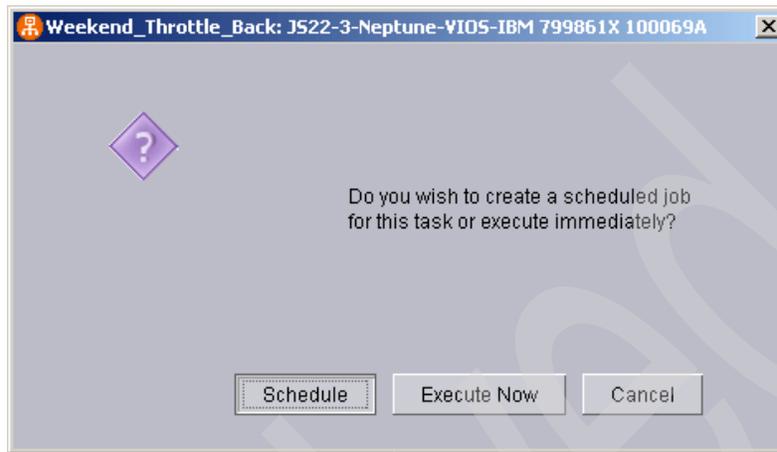


Figure 9-25 Policy execution dialog box

When **Schedule** is clicked, a New Scheduled Job dialog box is opened. Enter a Scheduled Job name and click Advanced as shown in Figure 9-26.

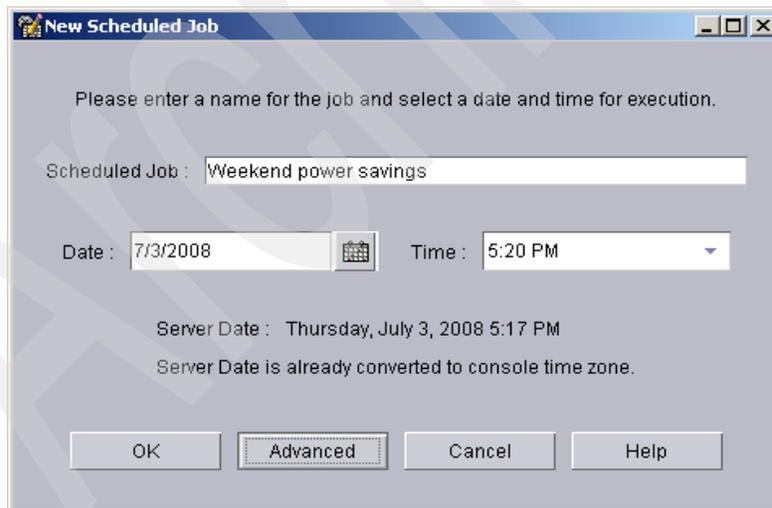


Figure 9-26 New Scheduled Job dialog box

A job name of "Weekend power savings" was chosen and then **Advanced** was clicked.

The Advanced button will alter the New Scheduled Job window to display four tabs to further define the Job. The tabs are:

- ▶ Date/Time
- ▶ Task
- ▶ Targets
- ▶ Options

Figure 9-27 shows the new view of the new dialog.

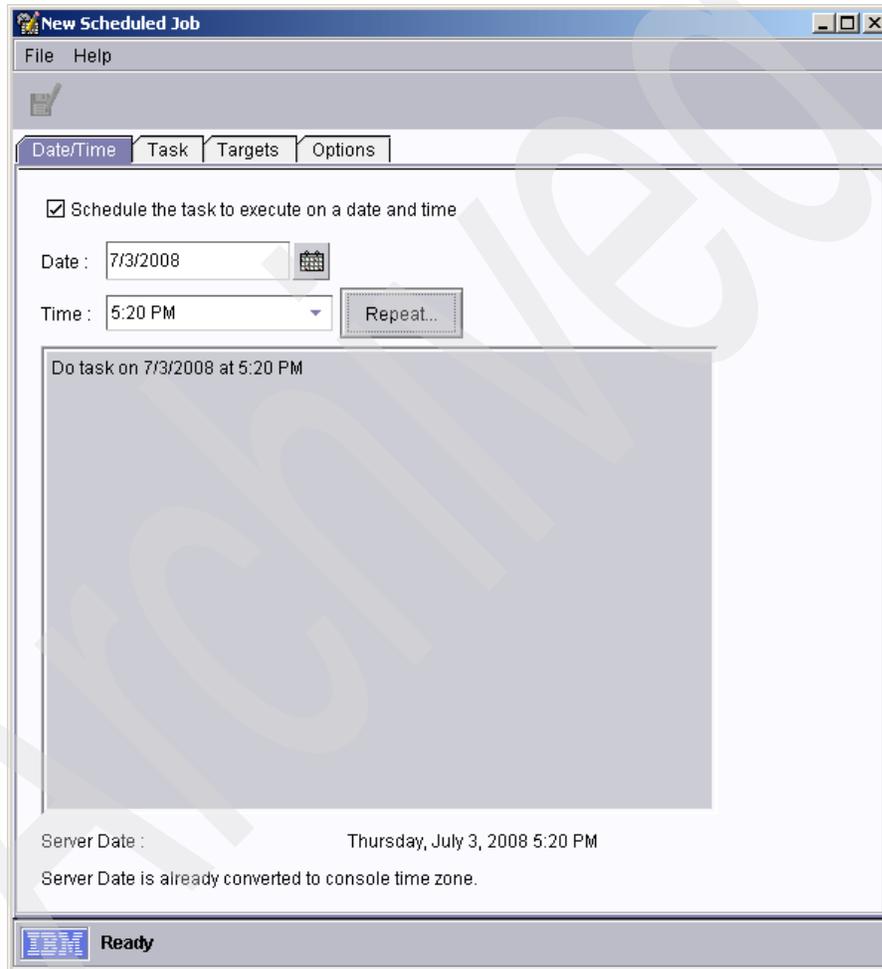


Figure 9-27 New Scheduled Job dialog showing advanced options

Click the **Date/Time** tab to open the scheduling window, as shown in Figure 9-28. The Repeat dialog box opens offering additional scheduling options. For this example the job is configured to execute weekly on Fridays at 06:30 p.m. for two months. With the selections made, click **OK** to continue to the next step.

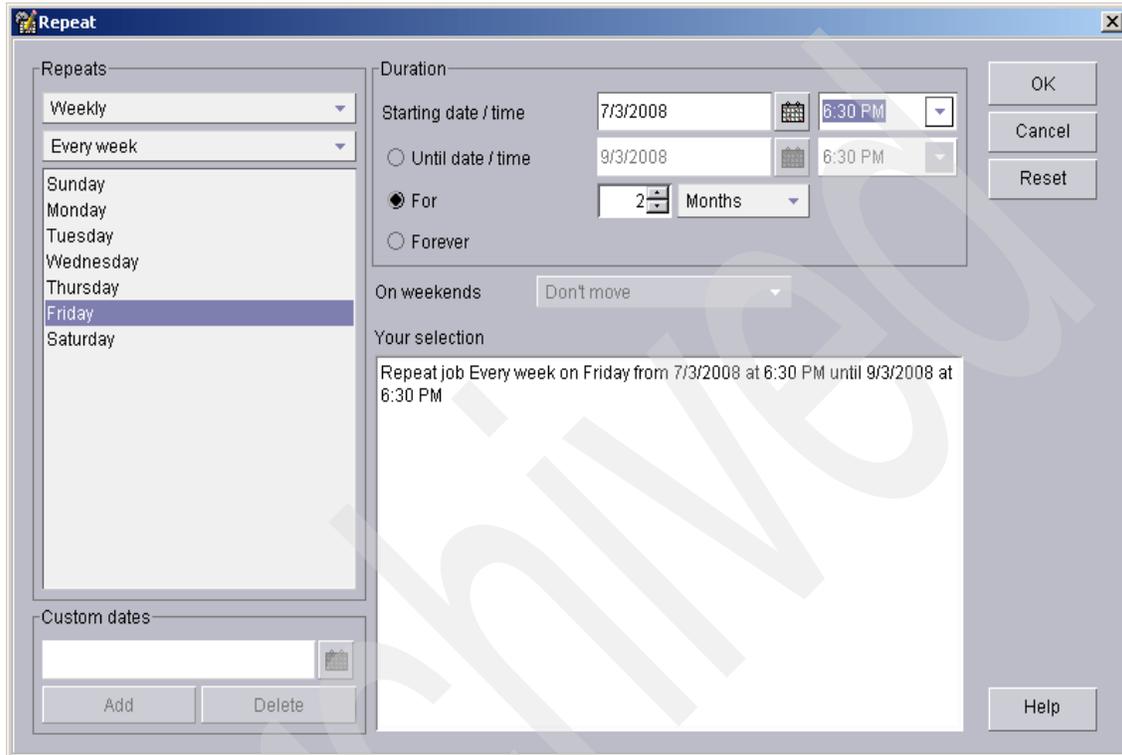


Figure 9-28 Repeat options dialog box for scheduling a policy job

Select the **Task** tab to display available tasks. Figure 9-29 shows this view with the Active Energy Manager tasks expanded. Since the task scheduling process was started with the Weekend_throttle_back policy, it is already selected and displayed in the Selected Task area.

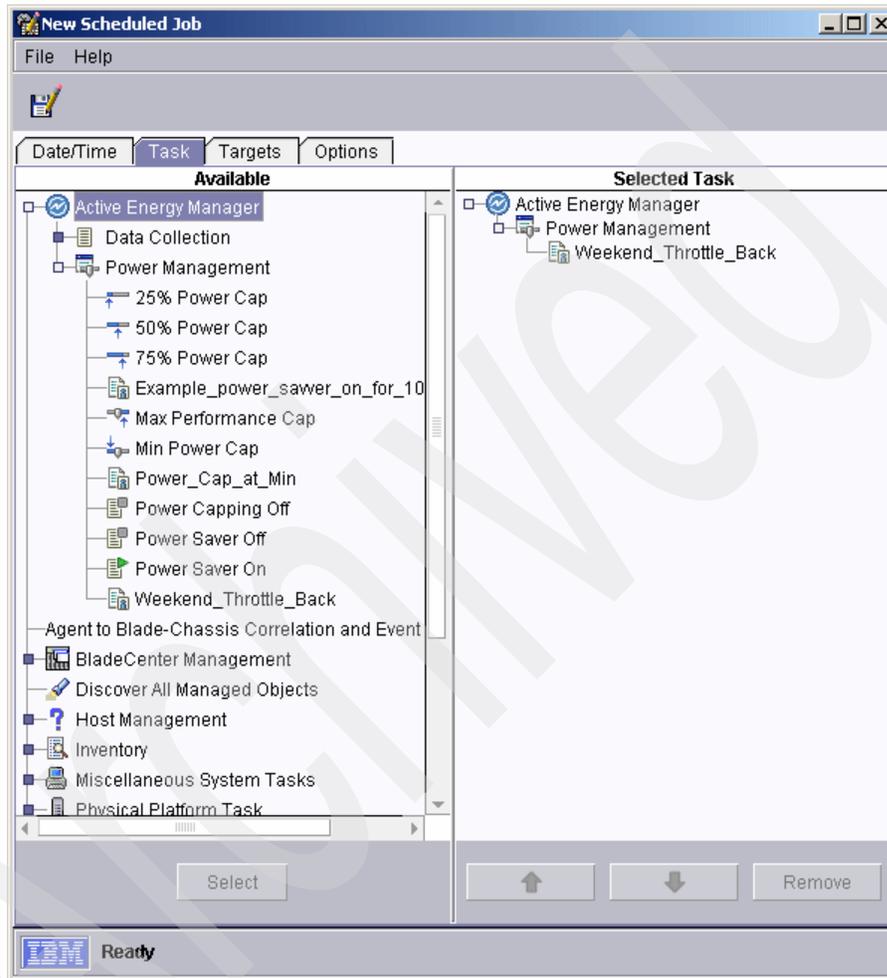


Figure 9-29 scheduled job task selection

Selecting the **Targets** tab displays a list of available targets. Similar to the task selection—since this process started with the server Neptune—it is automatically added to the selected list, as shown in Figure 9-30.

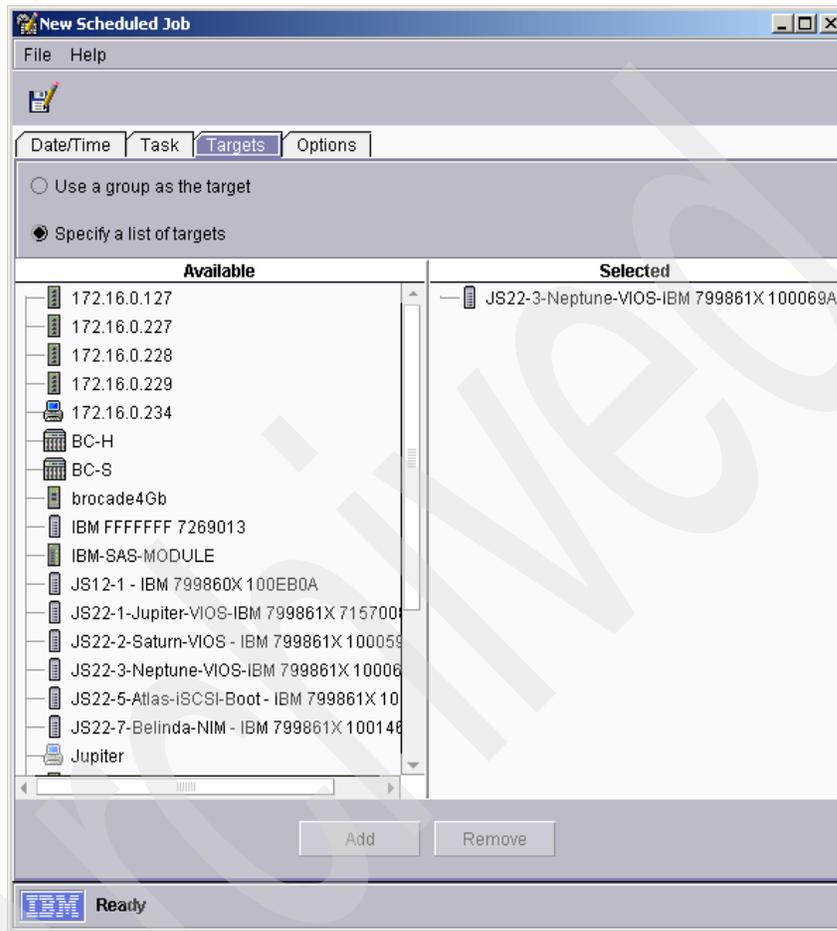


Figure 9-30 Scheduled job target selections

The final tab selection is the **Options** tab. The view for this tab, as shown in Figure 9-31, has sections for Special Execution Options, Execution History record retention, and Events to be generated. This example limits the execution history to the last 10 events and only when the event was successful.

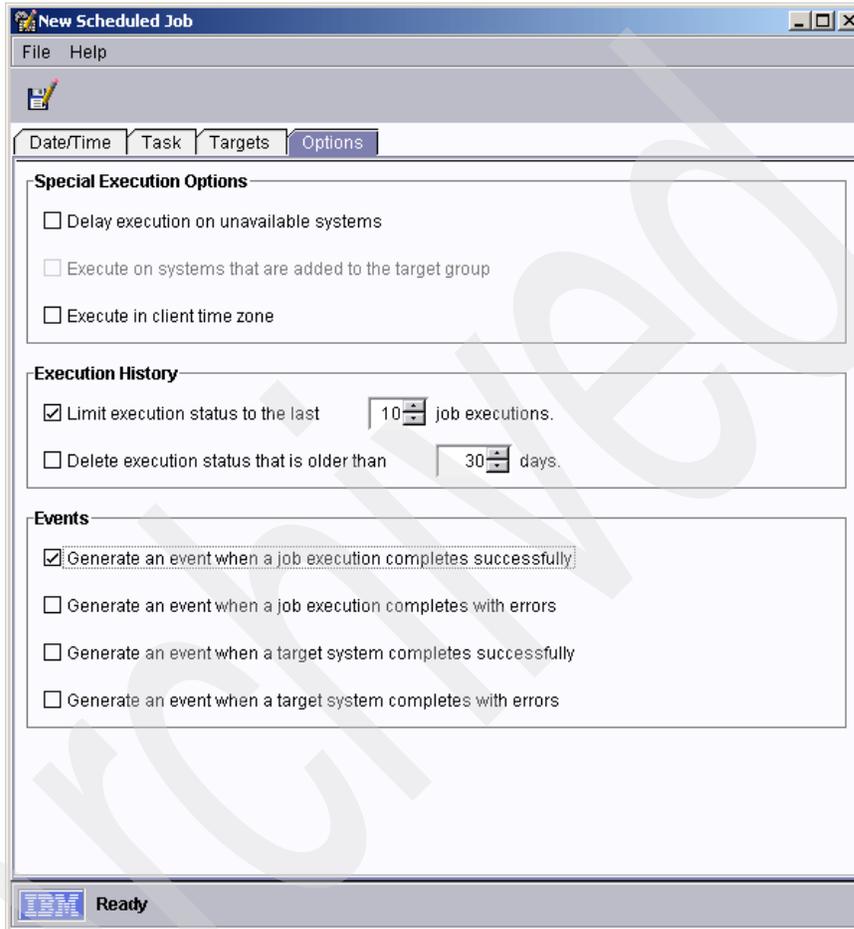


Figure 9-31 Scheduled job options selection

With the scheduled job configuration complete, select **File** from the menu line and **Save As** to save the changes.

From the IBM Director Tasks menu, double-click the **Scheduler** option to view all currently scheduled tasks. The tabs in the view allow the scheduled jobs to be viewed by month, week, or day. The last tab, **Jobs**, when selected, shows a list of scheduled jobs and their execution history, as shown in Figure 9-32.

Selecting the job name from the Jobs list gives the status of the current execution and the time of the next scheduled occurrence. Figure 9-32 shows the scheduled job “Weekend power savings” that executes the policy Weekend_Throttle_Back.

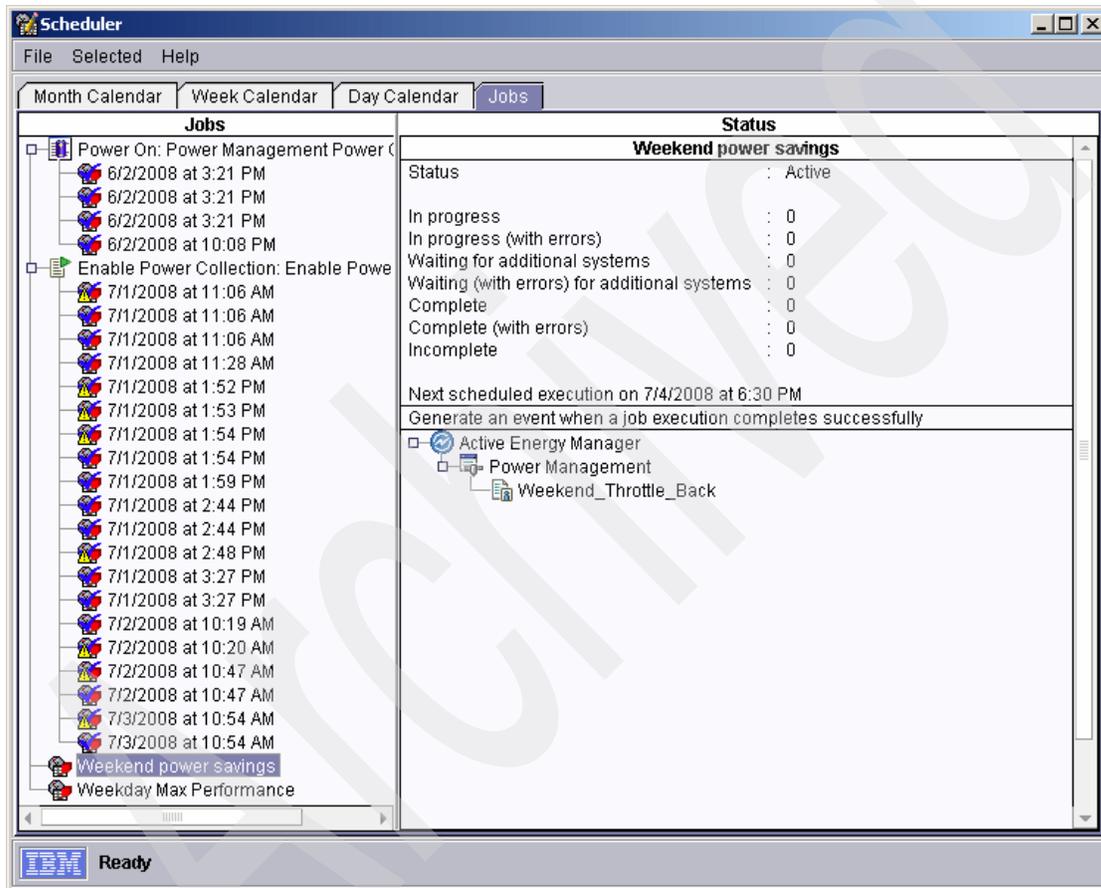


Figure 9-32 Scheduler Jobs view of Weekend power savings job

The corresponding job Weekday Max Performance shown in Figure 9-33 is also listed in the Jobs lists and executes the policy “Weekday_throttle_up”. This policy returns the Power Saver Mode to disabled to allow full CPU speed during the work week for this example.

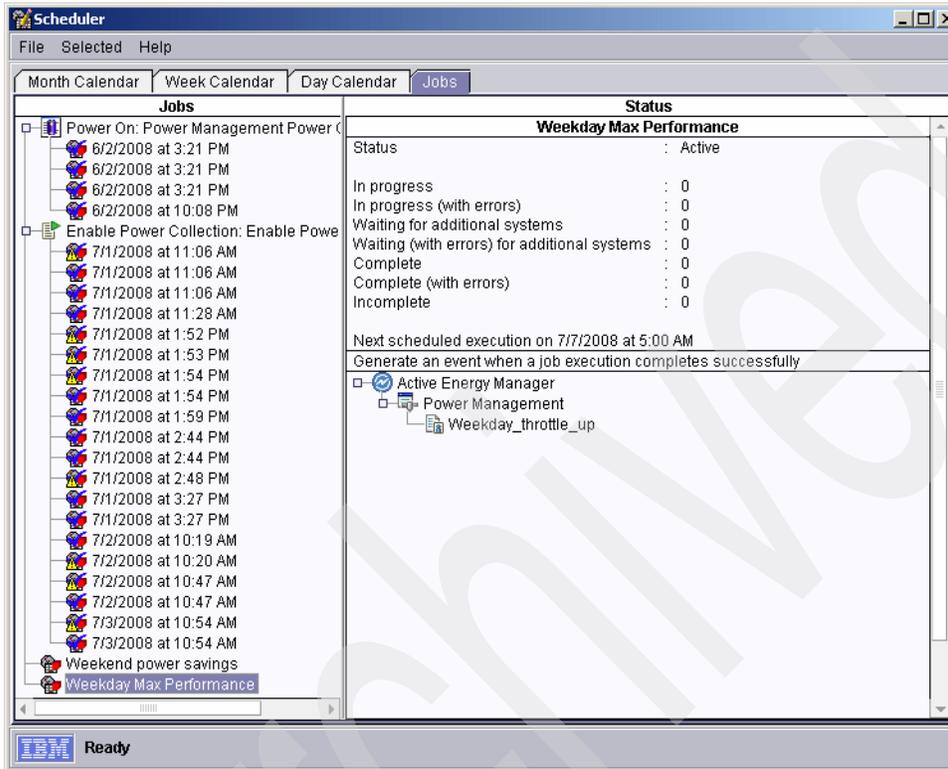


Figure 9-33 scheduler Jobs view of Weekend Max Performance job

Archived

Appendixes

In this part of the book we provide additional technical support information:

Appendix A, “Consoles, SMS, and Open Firmware” on page 391

Appendix B, “Red Hat Enterprise Linux Kickstart” on page 419

Appendix C, “SUSE Linux Enterprise Server AutoYaST” on page 429

Appendix D, “Additional Linux installation configuration options” on page 449

Appendix F, “Service and productivity tools for Linux” on page 465

Archived



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.

Consoles of the IBM BladeCenter JS12 and JS22

IBM BladeCenter JS12 and JS22 are the first blades of a series of blades with Power CPUs that are delivered with a graphics adapter. This graphics adapter makes it possible to use 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 the usage of Serial Over LAN, or SOL.

You can use either the graphical console or the SOL console during POST. A JS22 blade will start the first time as default with the graphical console as long as you do not change this. The blade will ask, after first power on, which console should be used as active console. If you miss this point during POST you need to enter the System Maintenance Services menu from a graphical console and change the console. See “System Management Services menu” on page 399 for information about how to change the active console.

Graphical console

The graphical console is available by either using the connected keyboard, display, and mouse on a management module installed in the BladeCenter chassis, or the remote control function of the Advanced Management Module.

Using keyboard, video, and mouse connected to a management module to gain access to the graphical console

The physical console connected to a management module will most of the time not be used. If you are working locally at the BladeCenter, you may use this console to access blades. There is one analog console available that is switched via the KVM switch that is built into the management module between the blades in the chassis. To switch the console you may use the keyboard/video select button that is placed on the front panel of each blade. Figure A-1 on page 393 shows the location of the front panel itself as well as the control elements on the panel of an IBM BladeCenter JS12 and JS22.

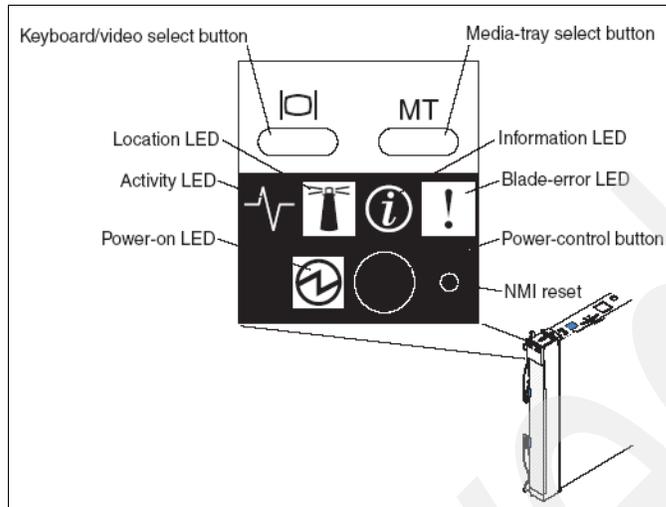


Figure A-1 Control elements of a JS12 and JS22

Pressing the keyboard/video select button switches the console to the blade on which the button was pressed. There is only one blade in a chassis that has lit the keyboard/video select button.

Note: Be sure that you are using the keyboard, video, and mouse connected to the active Advanced Management Module. There is only one management module active at one time. You will recognize this by looking at the management modules' LEDs. Only one management module has the blinking heartbeat 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.

An alternative method to switch the physical console between the blades is the usage of the key combination Shift+Num lock+Num lock+Bay Number+Enter on an IBM Space Saver Keyboard that is typically used in an IBM Rack. On standard keyboards the shift key is not required.

Use the key combination as follows:

1. Click and hold the Shift key.
2. Click Num Lock twice.
3. Release the Shift key.

4. Click the bay number - one of 1-14 depending on the chassis you are using.
5. Click Enter.

Using remote control to access the graphical console

Remote control is a feature of the management module installed in a BladeCenter chassis. It allows to 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 JS12 and JS22.

To open a remote control session, 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 395. The default account with all administrative rights is USERID with password PASSWORD.

Note: It is recommended that each administrator has its own account. Authentication against an LDAP server can be used when more accounts are required than are possible to configure on the management module.

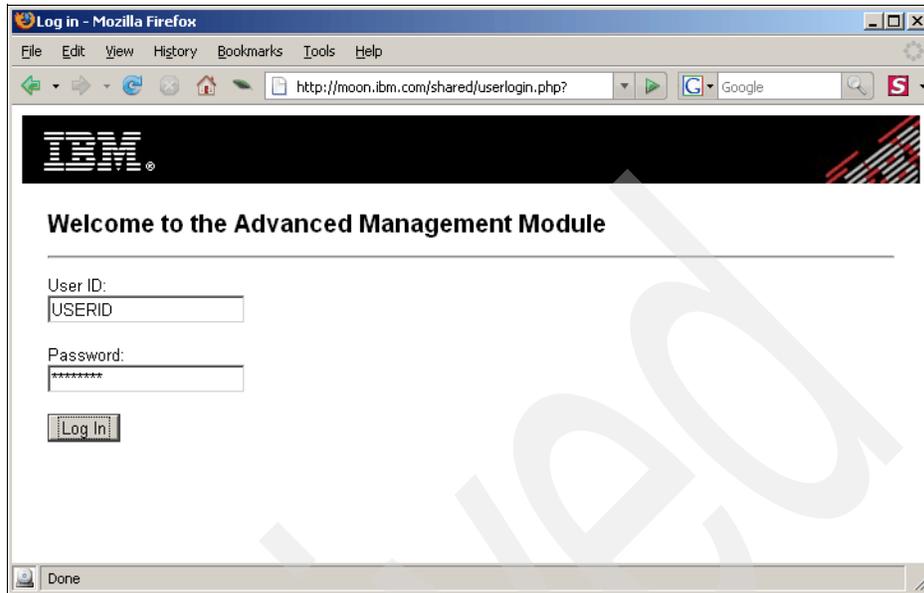


Figure A-2 AMM login panel

Select the time-out parameter that defines after how much idle time the session will be closed. Click **Continue**; Figure A-3.

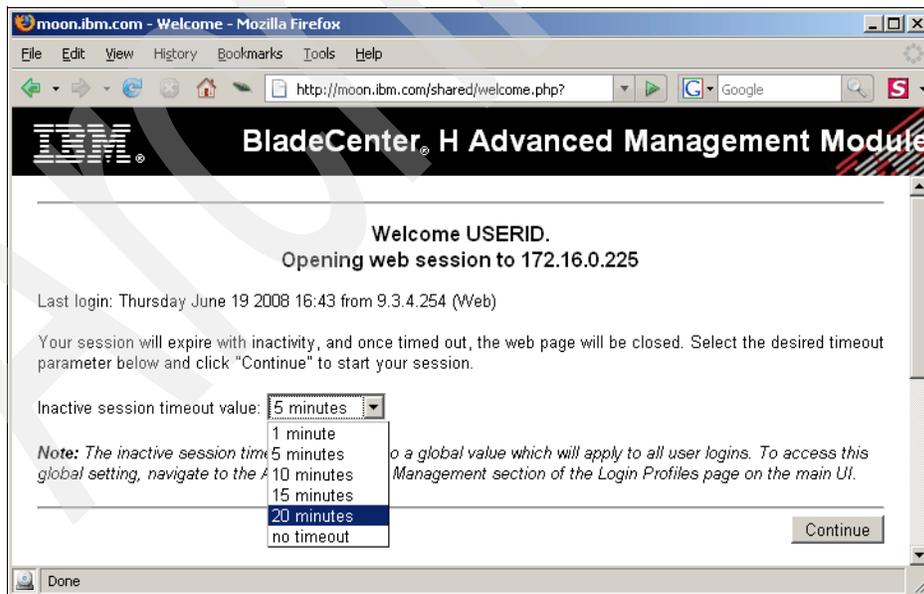


Figure A-3 Select time-out parameter

After successful login you will 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. Verify that there is no remote control session in progress by observing the remote control status. The Refresh button allows to refresh the status. Then scroll down to Start Remote Control.

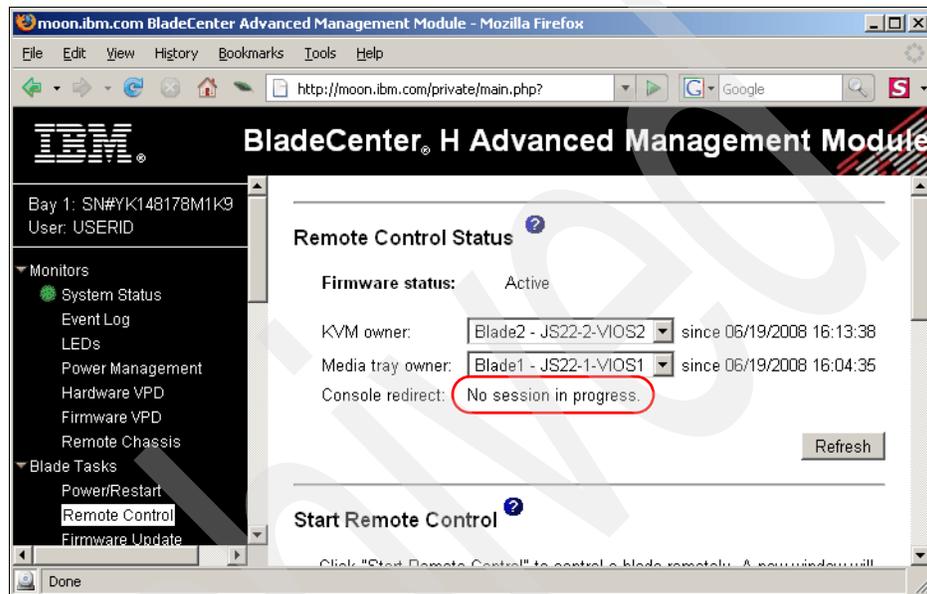


Figure A-4 Remote control status

Click **Start Remote Control** as shown in Figure A-5. A new window will open with the remote control Java applet. Be sure that there are no popup blockers running or configure them to allow the popup windows from the AMM.

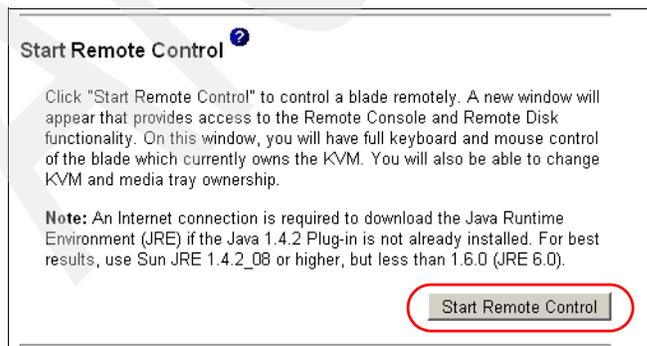


Figure A-5 Start Remote Control

The remote control Java applet will start in a new window. Figure A-6 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, there can only be one analog session active. The physical console shows the same view as the remote console. Concurrent KVM will be shown only via remote control and is represented by a red D. JS22 does not support cKVM.

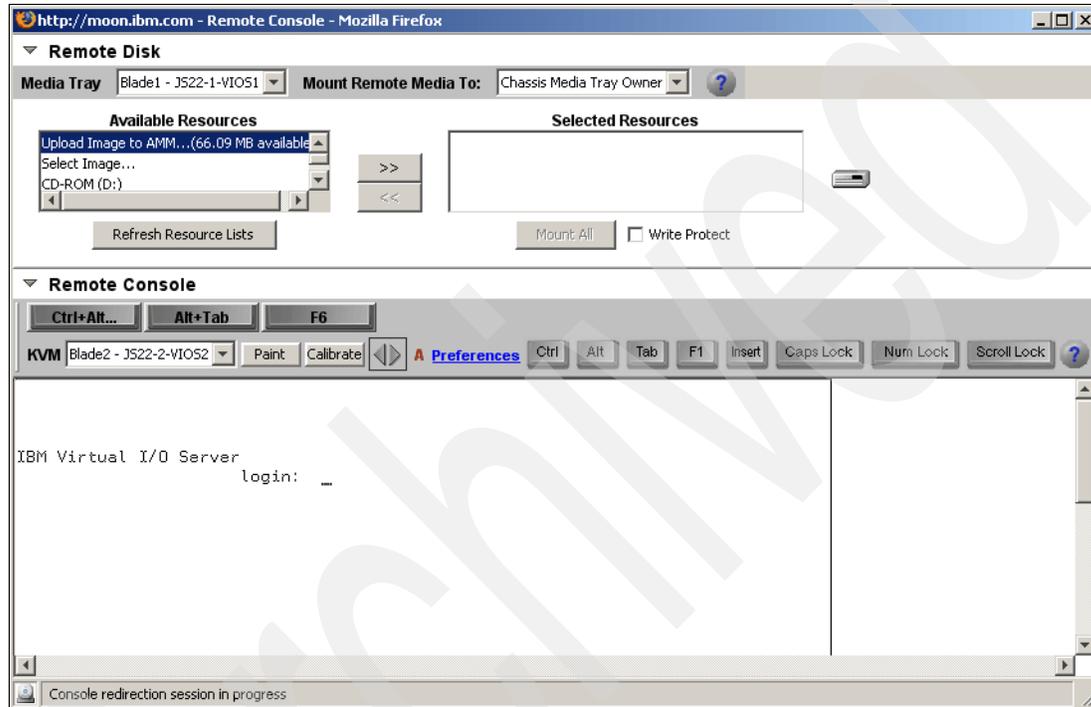


Figure A-6 Remote control - remote console and remote disk

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, such as the BladeCenter JS20 and JS21 blade server. 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 the need for 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 no need for a serial concentrator
- ▶ Standard Telnet interface that eliminates the need for special client software

The IBM BladeCenter management module command-line interfaces 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 command-line interface session with the management module. When this Telnet or SSH command-line interface 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 enables you to have 14 simultaneous SOL sessions active (one for each of up to 14 blade servers) with 6 additional command-line interface sessions available for BladeCenter unit management.

For a BladeCenter T unit, this enables you to have 8 simultaneous SOL sessions active (one for each of up to 8 blade servers) with 12 additional command-line interface sessions available for BladeCenter unit management.

With a BladeCenter S unit you have 6 simultaneous SOL sessions active (one for each of up to 6 blade servers) with 14 additional command-line interface 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 command-line interface 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 JS22.
- ▶ An Ethernet switch module or Intelligent Pass Through module is installed in Bay 1 of a BladeCenter.
- ▶ SOL is enabled for those blades that you wish to connect to via SOL.
- ▶ The Ethernet switch module must be set up correctly.

For details about how to set up SOL, see the Serial over LAN setup Guide, which can be found at:

<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-54666&brandind=5000020>.

The System Management Service menu section contains an example of how to establish a Telnet or SSH connection to the management module and then an SOL console. See “Power on the IBM BladeCenter JS12 and JS22 using a Telnet or SSH session into the Advanced Management Module” on page 402.

System Management Services menu

The System Management Services Menu (SMS menu) is considered the main configuration interface of the IBM BladeCenter JS22. It is described in *Installation and User's Guide for IBM BladeCenter JS22 (Type 7998)*. This guide is delivered on CD with each JS22 and can also be downloaded from the IBM support Web page at:

<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5073042&brandind=5000020>.

One of the first settings you may wish to change is the console that you want to use. When the blade starts the first time it comes up using 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 392 about 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 will only be available on the active console. The nonactive console can only be used to access the operating system.

Power on the IBM BladeCenter JS12 or JS22 using the Advanced Management Module 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 your AMM.

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.

In the left menu, click **Power/Restart** under Blade Tasks as shown in Figure A-7.

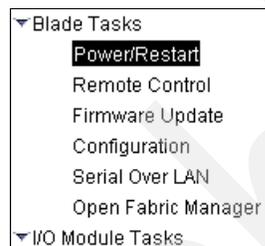


Figure A-7 Power/Restart in the left menu of the AMM Web interface

The power status of the blade is visible. Click the check box in front of the blade. Then choose the power action below the table with the blades. Use Power On Blade to start the blade or Restart Blade to restart the blade. See Figure A-8 on page 401.

Blade Power / Restart ²

Click the checkboxes in the first column to select one or more blades; then, click one of the links below the table to perform the desired action.

<input type="checkbox"/>	Bay	Name	Pwr	Local Pwr Control	Wake on LAN	Console Redirect
<input type="checkbox"/>	1	JS22-1-VIOS1	On	Enabled	N/A	
<input type="checkbox"/>	2	JS22-2-VIOS2	On	Enabled	N/A	
<input checked="" type="checkbox"/>	3	JS22-3-VIOS3	On	Enabled	N/A	1
	4	No blade present				
<input type="checkbox"/>	5	JS22-4-native	On	Enabled	N/A	
	6	No blade present				
<input type="checkbox"/>	7	JS22 native	On	Enabled	N/A	
	8	No blade present				
	9	No blade present				
	10	No blade present				
	11	No blade present				
	12	No blade present				
	13	No blade present				
	14	No blade present				

[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](#)

The following operations can be executed only on some POWER-based blades.

[Restart Blade and clear NVRAM](#)

[Restart Blade with Diagnostic Boot](#)

[Restart Blade with Diagnostic Boot and Default Bootlist](#)

Figure A-8 Power/Restart blade options

Note: The Restart Blade option will perform a power off and a power on of your selected blade. The operating system will not shut down properly. Use this option only when there is no operating system running or the blade is in POST, SMS, or Open Firmware prompt.

The blade will perform the requested action. Refresh this Web page to see a status change. Now use the console of your choice to work with the blade.

Consoles are described in “Consoles of the IBM BladeCenter JS12 and JS22” on page 392.

Power on the IBM BladeCenter JS12 and JS22 using a Telnet or SSH session into the Advanced Management Module

To start or power cycle the blade via 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 default account is USERID with password PASSWORD. See Example A-1 on page 403.

Note: Remember that the 0 in PASSWORD is a zero.

Help is available via the command **help** or **help {command}**. Every command may be executed with one of these options to show the online help for the command:

```
env -h
```

```
env -help
```

```
env ?
```

This example uses the command **env** to show available options to get help. The *Management Module Command-line Interface Reference Guide* you can find online at:

<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-54667&brandind=5000020>.

After logon it is advisable to change the time-out of the Telnet or SSH session with the command **telnetcfg**. The command will be issued to the current primary (active) management module.

To get a list of available targets in a BladeCenter chassis, use the command **list -1 {number of levels}**. The output in Example A-1 on page 403 shows, at the beginning of the **list -1 2** command, that the first management module is the active one. The **telnetcfg** command uses this active AMM as target to extend the time-out of a Telnet session or switch it off. Use 0 to switch the time-out off or use any value between 1 and 4,294,967,295 seconds. During installation it is a good idea to switch the time-out off so that the console will not disconnect.

Every command that is executed has a target. This target is specified by the **-T** option. To make the commands that you work with shorter, 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 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]
      switch[1]
      switch[2]
      switch[3]
      switch[4]
      switch[8]
      switch[9]
      switch[10]
      blade[1]   JS22-1-VIOS1
      blade[2]   JS22-2-VIOS2
      blade[3]   JS22-3-VIOS3
      blade[5]   JS22-4-native
      blade[7]   JS22-5-Install-Srv
      mt[1]

system>
system> telnetcfg -t 0 -T mm[1]
OK
```

```
system> env -T blade[3]
OK
system:blade[3]> power -state
Off
system:blade[3]> 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+[. 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 “Power on the IBM BladeCenter JS12 or JS22 using the Advanced Management Module Web interface” on page 400 or in “Power on the IBM BladeCenter JS12 and JS22 using a Telnet or SSH session into the Advanced Management Module” on page 402 you will see—depending on the console you have chosen—a message to select this console as the active console.

Figure A-9 shows an SOL console and Figure A-10 on page 405 shows a picture from the physical console using Remote Control.

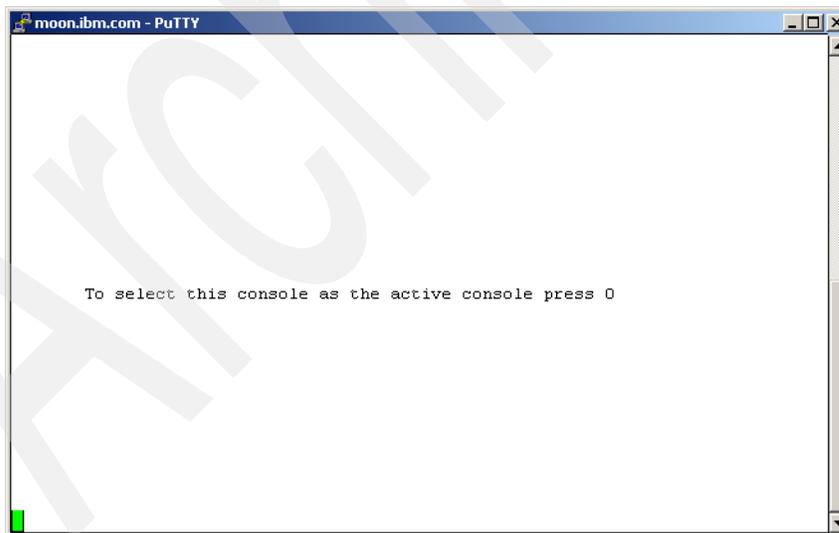


Figure A-9 Serial over a LAN console - select active console

Depending on the console you have open you have to enter a 0 to activate the SOL console or a 1 to activate the physical console. If you do not make your decision, the physical console will be used.

You need to enter the SMS menu over the physical console to change the active console, in this case as described in the next steps.

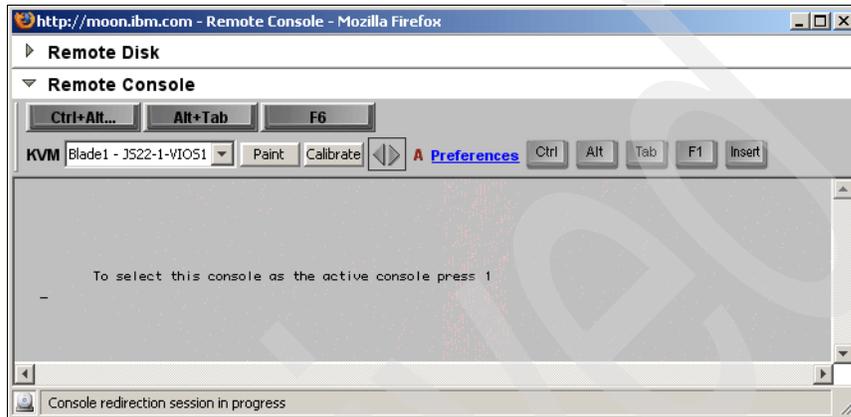


Figure A-10 Physical console shown with remote control - select active console

After a console is chosen as active console, either by the user or automatically, the system will show the Power On Self Test (POST).

IBM BladeCenter JS12 and JS22 Power On Self Test (POST)

As with previous JS2x blades, there are no System Reference Codes (SRC) shown on a console during POST. The System Reference Codes 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 wish to see the System Reference Codes. Click **System Reference Codes**. The table with the System Reference Codes can be refreshed by clicking **Refresh**.

The POST prints the words Memory, Keyboard, Network, SCSI, and Speaker as shown in Figure A-11 on page 406. 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 options from the menu. You may:

- ▶ (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 used to load the operating system will be the default.

Click the number 1 to enter the SMS menu.

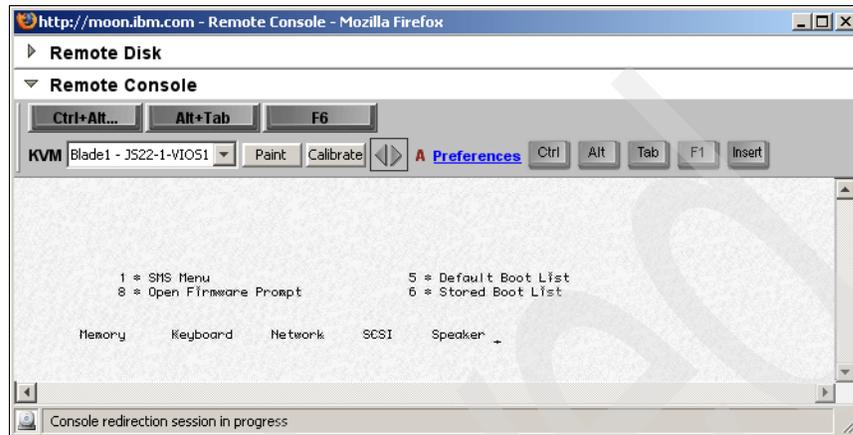


Figure A-11 JS22 POST

Select the active console using the System Maintenance Services menu

When the blade is going through the POST, you can enter the System Maintenance Services menu.

To change the current active console in the SMS menu, click **5** to select the console. See Example A-2.

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 clicked **5** to select a console, you will see on the physical and on the SOL console a message notifying you that you have to enter 0 or 1 depending on the console you use. Figure A-10 on page 405 shows the message on an SOL console and Figure A-9 on page 404 shows the message on the physical console. Switch to the console you want to activate and enter the number as shown on this console. Either use 0 or 1. This console will be the next active console. If you choose to do nothing, the system will revert back 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 JS22. 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 JS22 to build the system firmware may have its own extensions and may also not implement each feature or function. The information on this page may not apply fully to the IBM BladeCenter JS22.

In certain situations it is helpful to use the Open Firmware prompt to change system settings or help troubleshoot problems. Also, the SMS menu is considered the main configuration interface. A few commands will be explained that help getting SAN booted 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 the World Wide Port Name and/or World Wide Node Name.
3. Set the connection type.
4. Set the transfer rate.
5. Query available targets.

This appendix is split into a section about the QLogic host bus adapter and the Emulex host bus adapter. We start with a description of how to get access to the Open Firmware prompt.

Note: You may leave the System Maintenance Services Menu from the main menu with 0 to the Open Firmware prompt.

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 a iSCSI boot device entry.

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

Example: A-5 printenv boot-device output with iSCSI boot device

```
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 will return 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 JS22 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 will be 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 409. The device tree in your system may differ from the example shown here. With this information you can build the command to select the device. Enter the command:

```
“ /pci@800000020000204/fibre-channel@0” select-dev
```

to select the first host adapter port. The second HBA port is selected with the command:

```
“ /pci@800000020000204/fibre-channel@0,1" select-dev
```

Note that there is a space between the first “ and /. Example A-7 shows the output of this command. The link of the adapter port will come up and the adapter will log 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-wnn** 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 shows the output of this command.

Example: A-8 Display World Wide Port Name of a QLogic host bus adapter port

```
0 > my-wnn ok
1 > . 2100001b32005216 ok
0 >
```

Firmware version and FCode level of the HBA can be shown with the command **version**. 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 >
```

In case of the usage of an Optical Pass Through Module, it is necessary to 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 Through Module can only handle transfer rates up to 2 GB. Auto Negotiation will 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 on page 412.

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 as well as node and port names. Data rate is shown as well. 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.

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 may wish 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 click 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 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 JS22 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 WVN 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 need 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 409. The device tree in your system may differ from the example shown here. With this information you can build the command to select the device. Enter the command

```
" /pci@800000020000203/fibre-channel@0" select-dev
```

to select the first host adapter port. The second HBA port is selected with the command:

```
" /pci@800000020000203/fibre-channel@0,1" select-dev
```

Note that there is a space between the leading " and /. Example A-15 shows the output of this command. The link of the adapter port will come up and the adapter will log in to the switch. You are able now 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 via 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. 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**. The command will show the current link speed and the current setting, as shown in Example A-19.

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 click Enter as shown in Example A-20. 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 command **.topology**. 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-fc-1**, **set-auto-ptp**, **set-fc-1** or **set-ptp**. **set-auto-fc-1** sets loop preferred, otherwise point-to-point. **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. In Example A-21 you can see 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 that the described commands require that you have an HBA port selected and that they only have effect on the selected HBA port. You need to perform the necessary actions on both HBA ports.

To leave the Open Firmware prompt and restart the blade, use the command **reset-a11**.

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
```

Archived



Red Hat Enterprise Linux Kickstart

This appendix describes the Red Hat Kickstart tool to perform automated installations of Red Hat Enterprise Linux 5.2 (Red Hat Enterprise Linux 5.2) for the POWER architecture.

We discuss the following topics:

- ▶ Kickstart introduction
- ▶ Kickstart profile creation methods
- ▶ Create a Kickstart file using Kickstart Configurator

Kickstart introduction

The Kickstart configuration tool enables a system administrator to install Red Hat Enterprise Linux on a large number of systems in parallel using an automated process. The Kickstart file is a text file that contains responses to all the system configuration questions typically asked during a manual installation. The file can be stored on a single server system and read by multiple POWER systems concurrently.

This appendix covers Kickstart file creation for Red Hat Enterprise Linux 5.2.

Kickstart profile creation methods

There are three methods to create a Kickstart profile for systems:

1. Copy the sample kickstart file from a manually installed system as shown in 6.5, “Red Hat Enterprise Linux 5.2 automated installation” on page 300.
2. Use the Kickstart Configurator to create the Kickstart file.
3. Use a text editor to create a Kickstart file from scratch.

We cover the steps for option 2 here.

Create a Kickstart file using Kickstart Configurator

In this section, we use the Kickstart Configurator tool with a graphical interface to demonstrate how to create a basic Kickstart text file.

There are a lot of optional settings in this file, but some are mandatory settings or dependencies. It is impossible to cover every configuration option, so we try to provide a general overview so you are 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).

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 may have to make some manual changes to it.

Note: Kickstart is not included in the default software installation. Install the `pykickstart` and `system-config-kickstart` rpm packages from the installation media to get Kickstart working.

1. After you install the required Kickstart packages, 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 B-1. It is also important to define a root password to enable SSH login after installation. This password is encrypted in the configuration file.

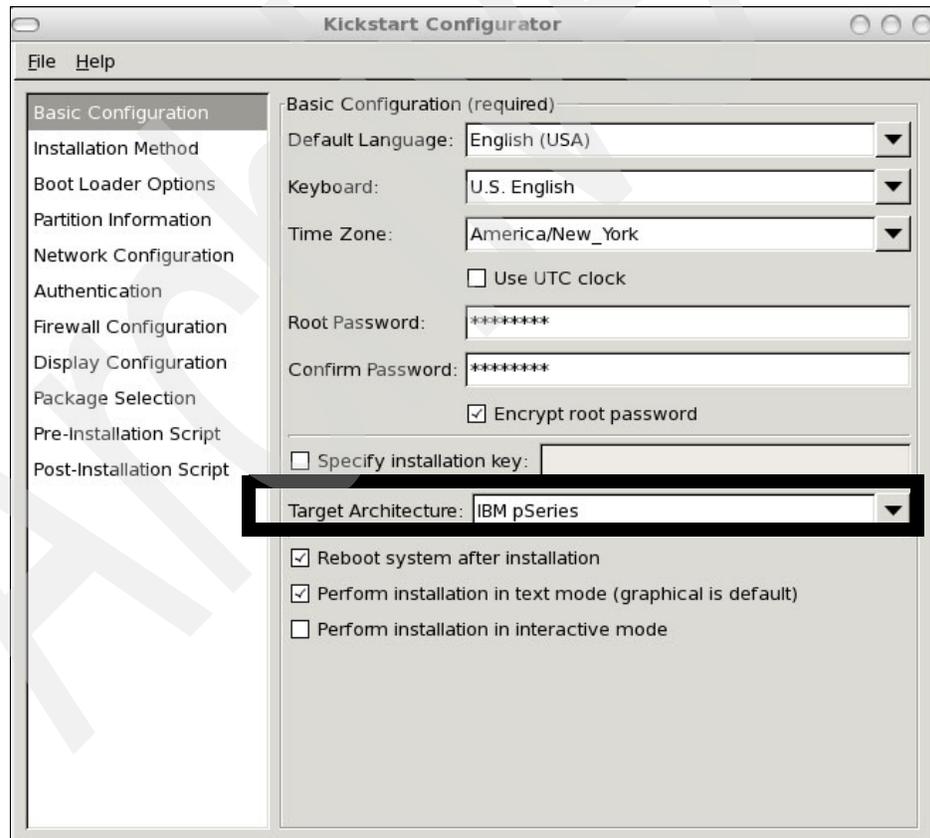


Figure B-1 Kickstart main window with Basic Configuration panel (©2008 Red Hat, Inc.)

3. In the Installation Method panel (shown in Figure B-2), all the basic parameters for a network installation using NFS are shown.

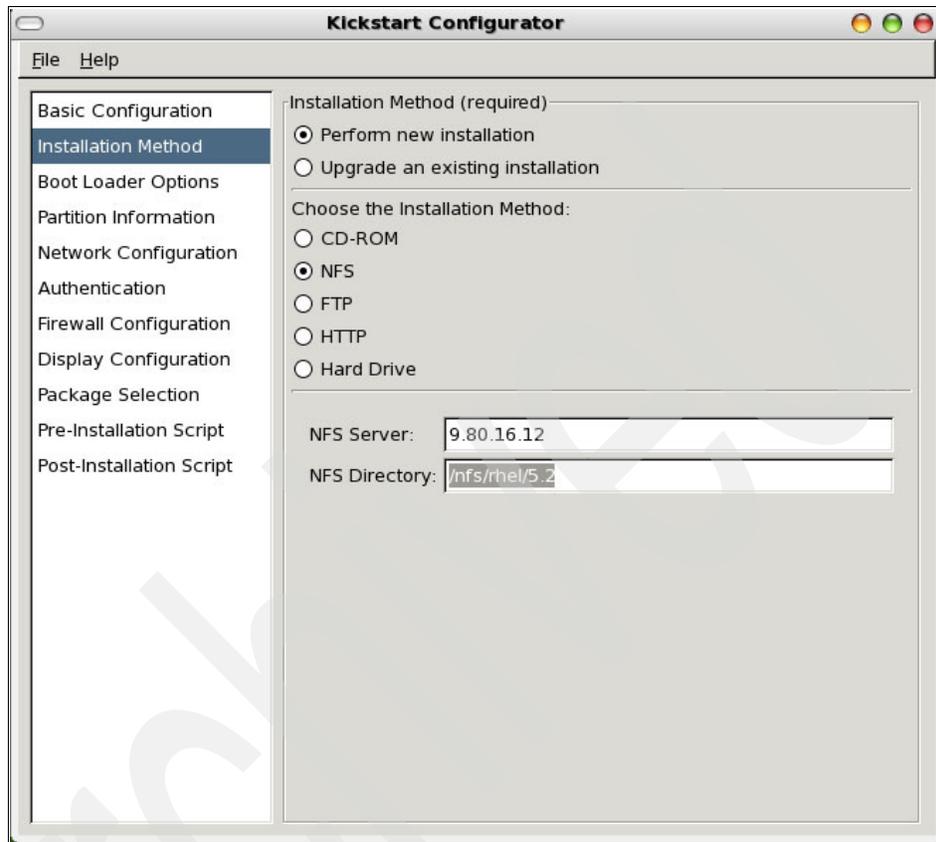


Figure B-2 Installation Method panel (©2008 Red Hat, Inc.)

- The next editable panel is the Partition Information panel, shown in Figure B-3. Press **Add** to create a partition. The tool will help you select the mount point, file system type, and partition size.

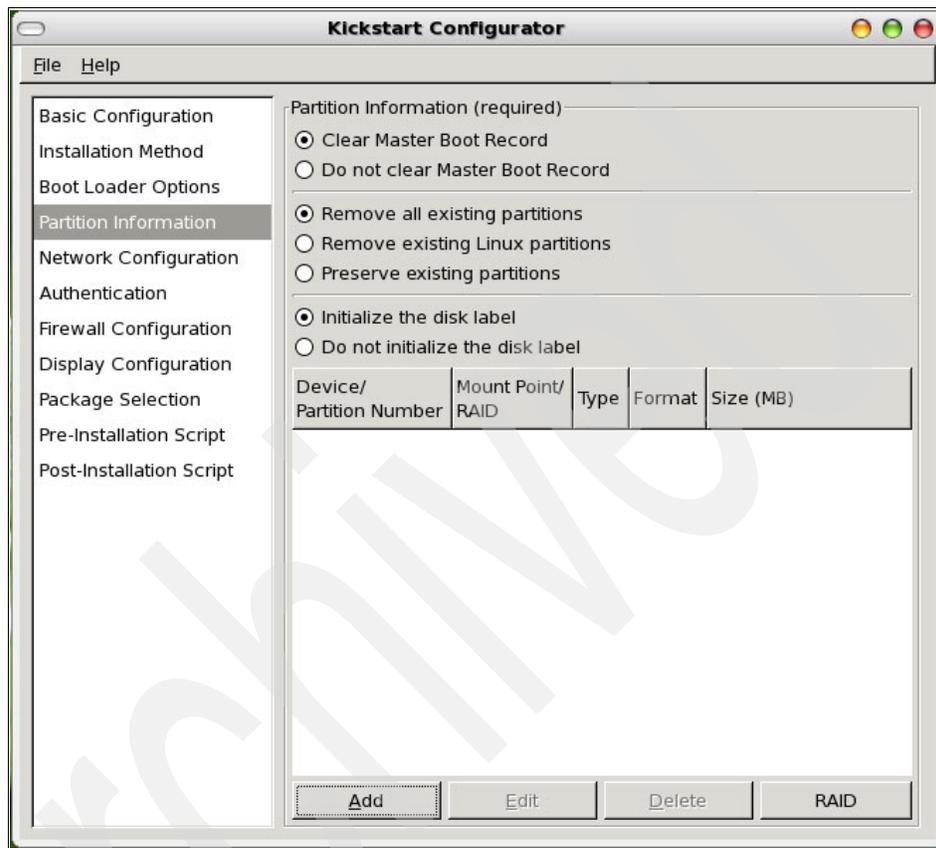


Figure B-3 Partition Information panel (©2008 Red Hat, Inc.)

5. In the Network Configuration panel, click **Add Network Device** to add the devices you are installing from. If you need to go back and make changes to this setup, click **Edit Network Device** (see Figure B-4).



Figure B-4 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 B-5 shows the Firewall Configuration panel. As an example, it is good to enable SSH and to trust interface eth1 at the very minimum to access the system later using the network.

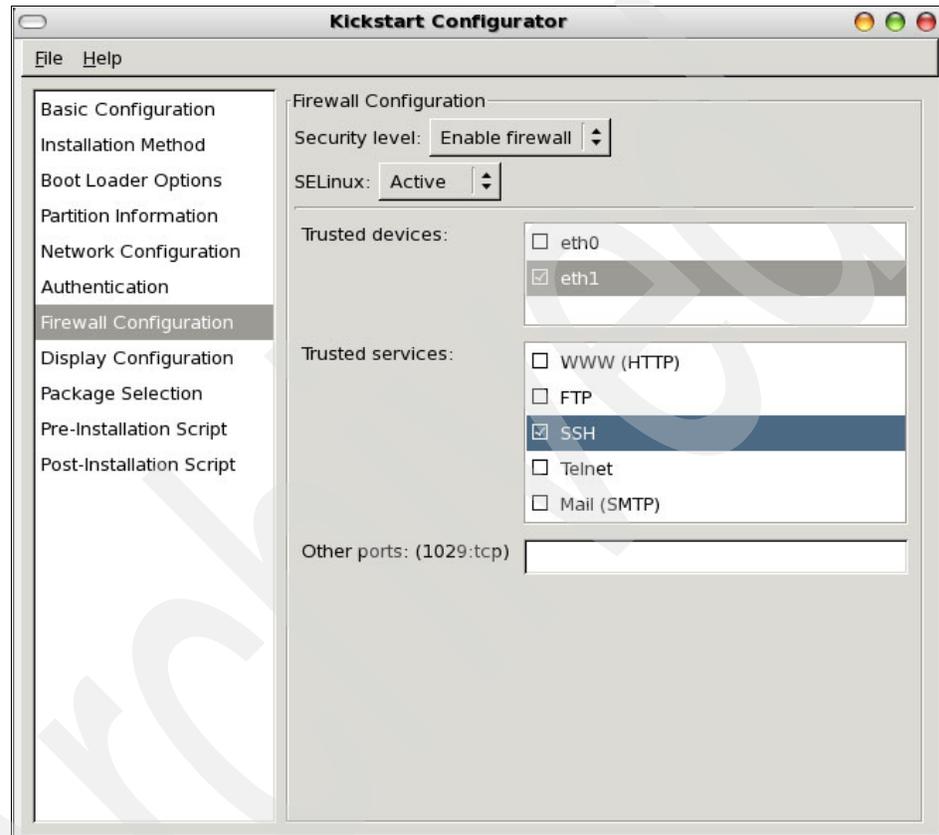


Figure B-5 Firewall Configuration panel (©2008 Red Hat, Inc.)

8. Figure B-6 shows the Package Selection panel. It is not possible to select individual packages from this panel. However, you can add individual packages to the %packages section of the Kickstart file after saving it.

Note: If you see the message

Package selection is disabled due to problems downloading package information.

in the Package Selection panel, it means you have no repositories defined. 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 5.2 installation discs to a directory and use them as a reference as given in the Red Hat Knowledge Base article available at:

http://kbase.redhat.com/faq/FAQ_103_12748.shtm

9. After the package selection, save the configuration using the **File** → **Save** option in the menu.

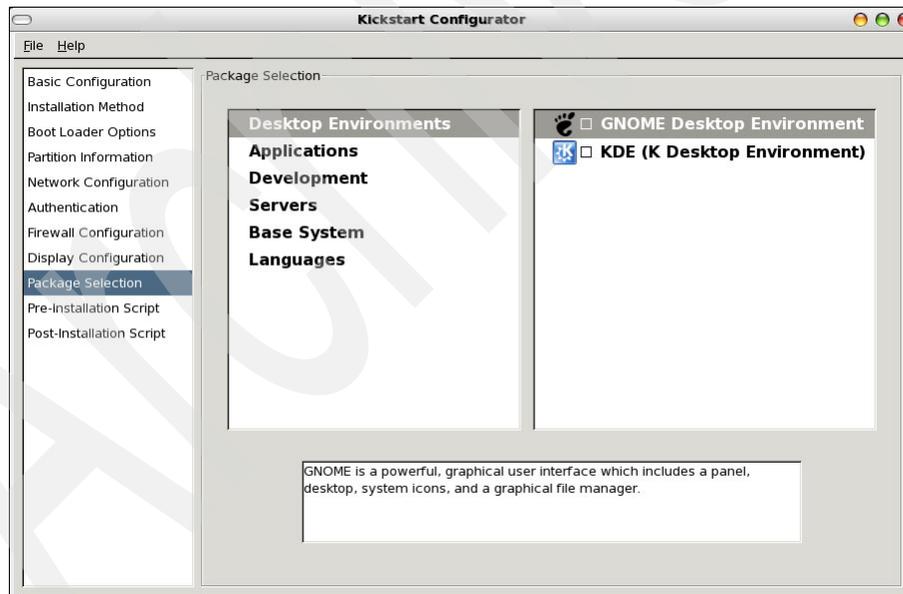


Figure B-6 Package Selection panel (©2008 Red Hat, Inc.)

10. The basic Kickstart configuration file created with the Kickstart Configurator is shown in Example B-1.

Important: The order of the main sections in the Kickstart configuration file is important for the functionality.

Example: B-1 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=9.34.21.18 --dir=/rhel/5.2
# 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
```

11. Manually adjust the Kickstart configuration file that you have created with a text editor if desired.

Note: If you have not defined any disk partition options or you were unsure of your disk partition layout, we recommend that you manually edit the Kickstart file to include the following information after the #Partition clearing information section:

```
#Disk partitioning information
autopart
```

This option will automatically create disk partitions.

Red Hat Enterprise Linux 5.2 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/5.2/html/Installation_Guide/ch-kickstart2.html

SUSE Linux Enterprise Server AutoYaST

This appendix describes the SUSE AutoYaST tool to perform automated installations of SUSE Linux Enterprise Server 10 Service Pack 2 (SLES10 SP2).

We discuss the following topics:

- ▶ AutoYaST introduction
- ▶ AutoYaST profile creation methods
- ▶ Create an AutoYaST profile using YaST Control Center

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 using an automated process. The AutoYaST profile is a file written using the Extensible Markup Language (XML). It contains responses to all the 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 10 SP2.

Note: The procedures covered here are not applicable for previous SLES releases because the XML layouts are different.

AutoYaST profile creation methods

There are three methods to create an AutoYaST profile for systems using *identical* hardware:

1. Clone the install configuration information from a reference machine that was installed manually. This is covered in “Linux network installation (detailed)” on page 311 of Chapter 7, “SUSE Linux Enterprise Server installation” on page 305.
2. Use the YaST Control Center to create and modify the AutoYaST profile.
3. Use an XML editor to create a AutoYaST profile from scratch.

We cover the steps for option 2 here.

Create an AutoYaST profile using 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 while the text mode version requires Tab, Enter, Up/Down Arrow, and Space bar keys to navigate. Otherwise, there is no difference between the two modes and the same configuration options in both will result in the same XML file.

There are a lot of optional settings, but some are mandatory settings or dependencies. It is impossible to cover every possible configuration option, so we try to provide a general overview to make you familiar enough with the tool to navigate on your own.

Starting the YaST graphical interface

1. SSH as root with X11 forwarding enabled into a system running SLES10 SP2. For example:

```
ssh -X root@9.3.20.18
```

Restriction: It is necessary to issue this command on a system with a running X Server to use the graphical interface.

2. From the command line type:

```
yast2
```

and a new window will appear with YaST Control Center at the very top as shown in Figure C-1 on page 432.

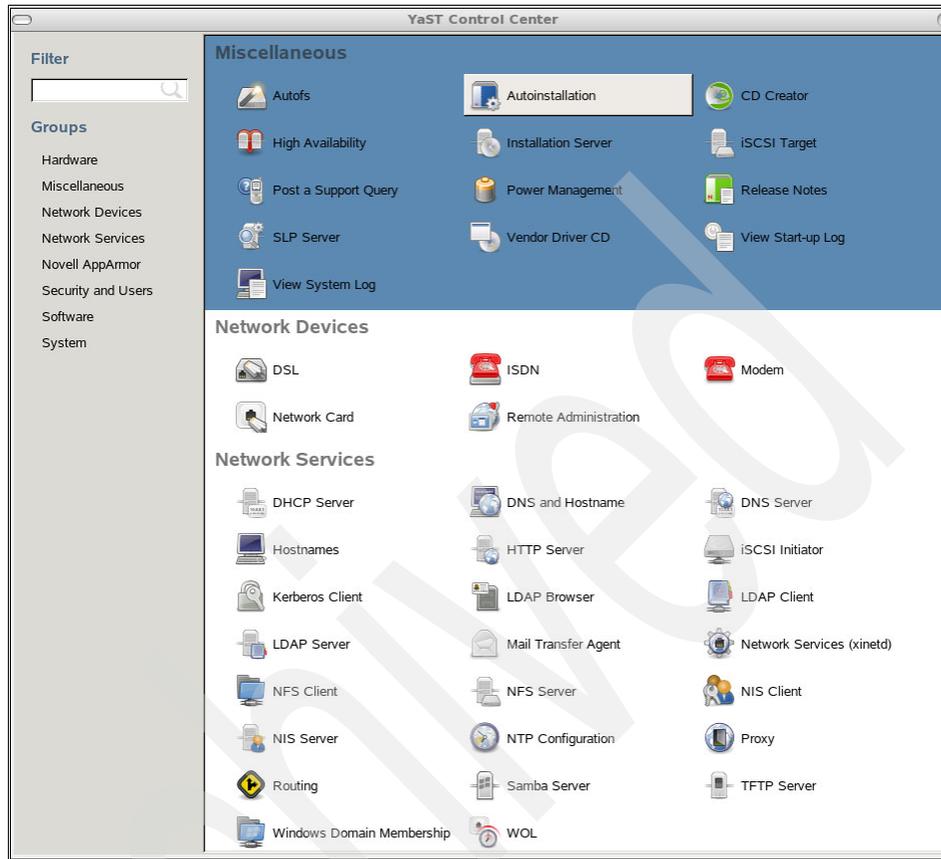


Figure C-1 YaST Control Center in graphics mode

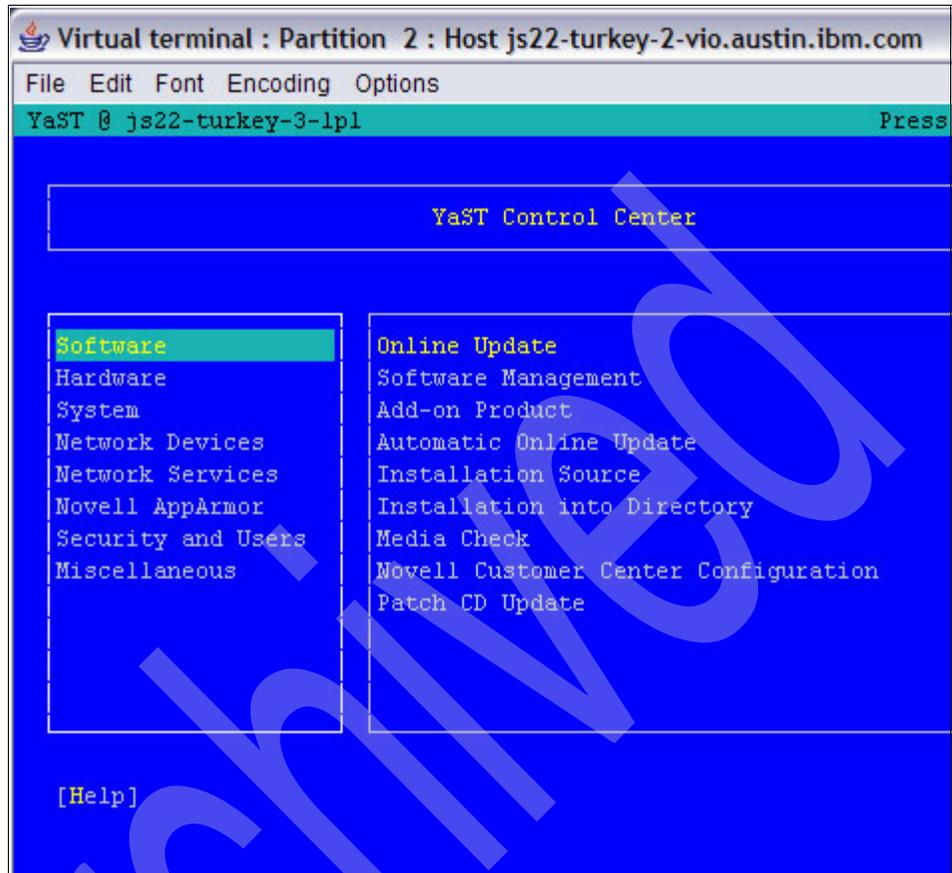


Figure C-2 YaST Control Center in text mode

Navigating the YaST graphical interface

1. Start the YaST application, which opens a window as shown in Figure C-3.
Launch the **Autoinstallation** applet from the Miscellaneous section of YaST.

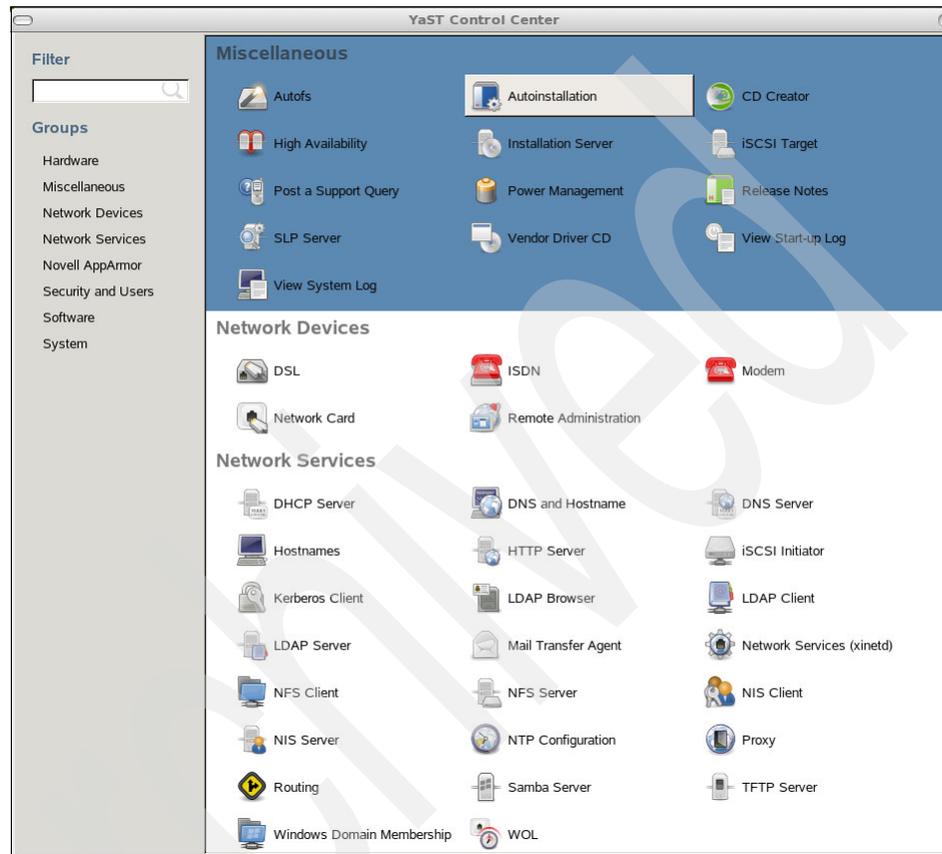


Figure C-3 Selecting the Autoinstallation option

2. After the selection, the main AutoYaST configuration window opens as shown in Figure C-4.

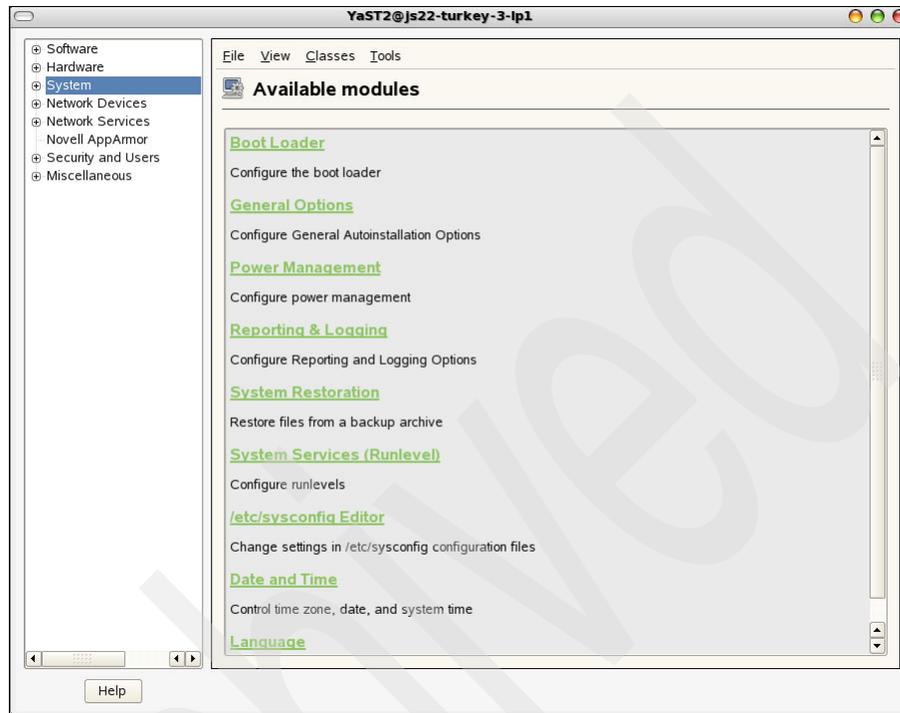


Figure C-4 Main AutoYaST menu (SLES10 SP2)

3. Clone the configuration of the installation server by selecting **Tools** → **Create Reference Profile**, as shown in Figure C-5.

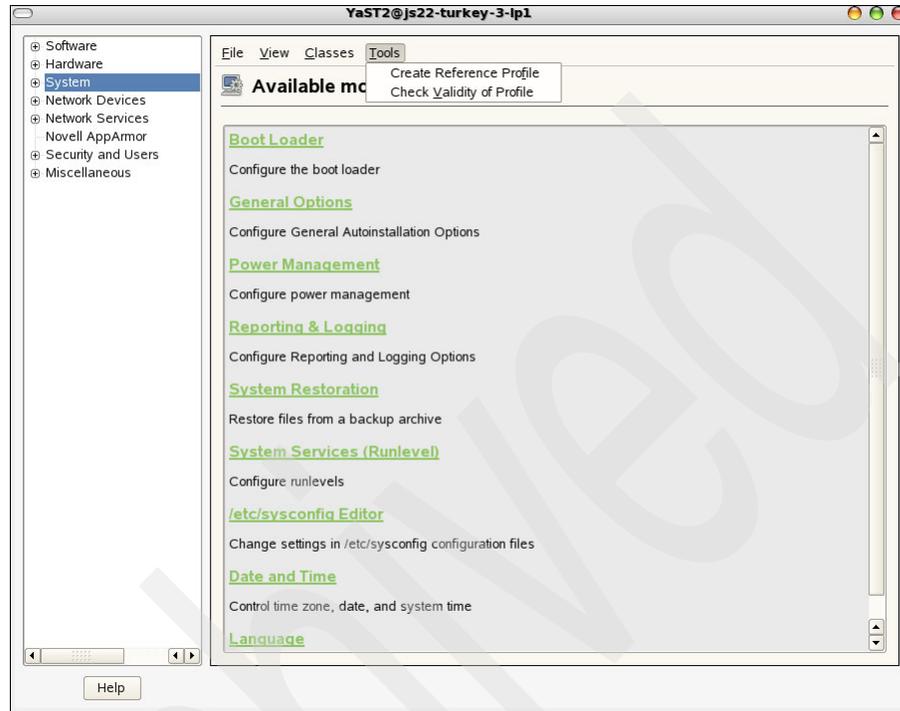


Figure C-5 Create a reference profile

4. A second window opens, as shown in Figure C-6. In addition to the default resources such as boot loader, partitioning, and software selection, it is possible to add other aspects of your system to the profile by checking items in the Select Additional Resources section. When ready, click **Create** so YaST can collect the system information and create the AutoYaST profile.

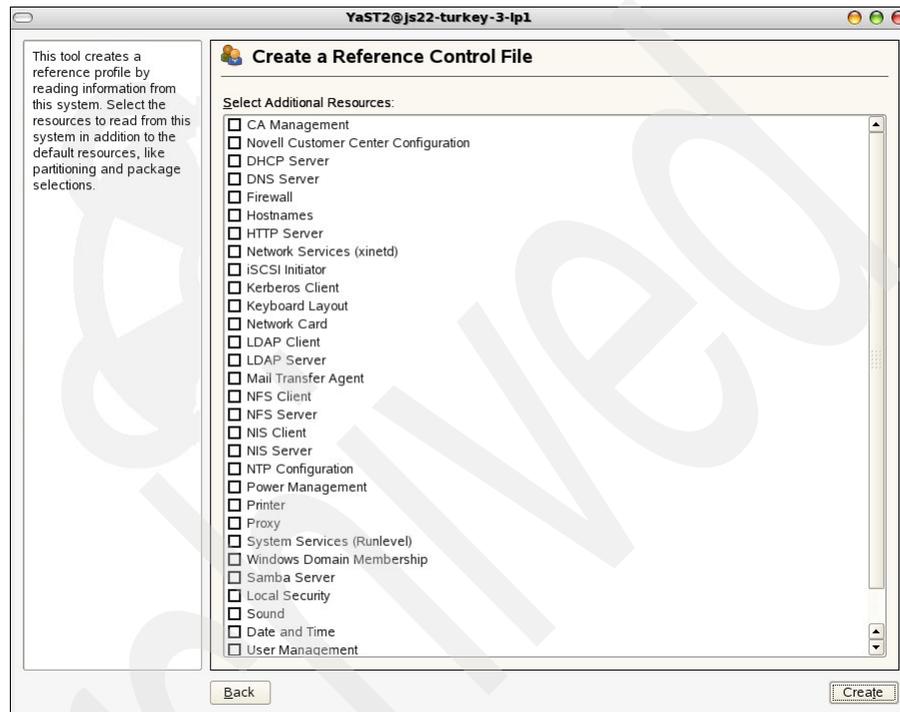


Figure C-6 Selecting additional resources

5. If the profile is complete and meets your requirements, select **File** → **Save** and enter a file name such as

sles10_sp2_autoinst.xml

Or adjust some of the options provided on the left side of the AutoYaST main menu option. Each option is given here with some highlights:

- a. **Software** - Options to select and configure the Online Update and Package Selection sections. Figure C-7 on page 438 shows the **Software** → **Package Selection** window. We chose the minimum software configuration to save time during installation.

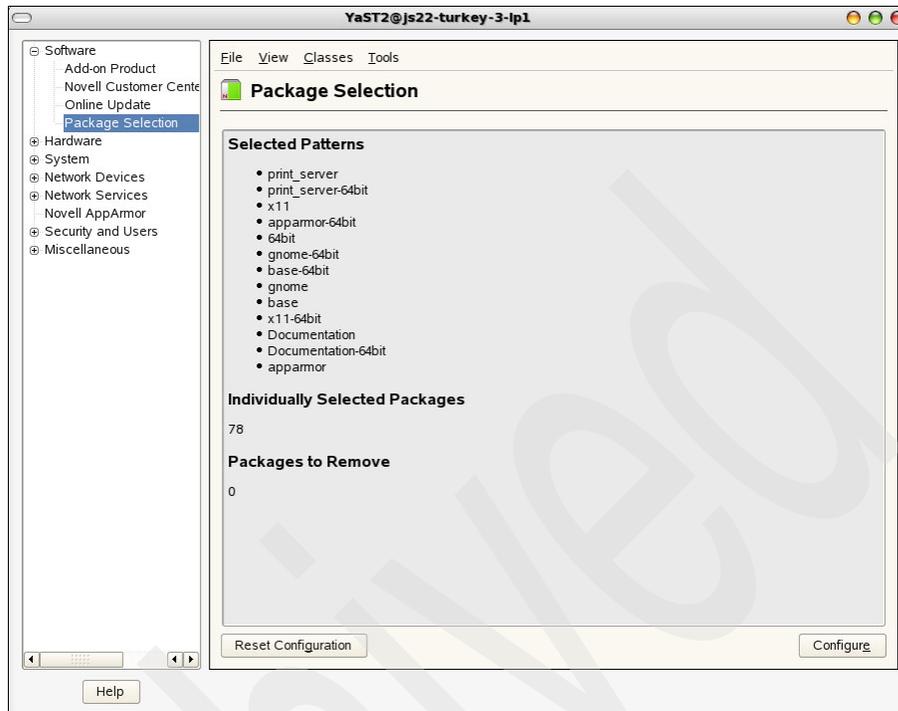


Figure C-7 AutoYaST software selection

- b. Hardware - Configure 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.
- c. System - Set 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 the Boot Loader configuration, which is cloned from the `/etc/lilo.conf` file of the system running the YaST tool so it may require some editing. Edit it by pressing **Configure** → **Edit**. Click **Finish** when you are done.
- d. Network Devices - Set the network adapter information. You can set network module information and IP details here.
 - i. Click **Configure** to open the menu.
 - ii. The Network cards configuration main menu opens. Select **Configure** to add interfaces.
 - iii. Select **Traditional Method with ifup** → **Next**.

- iv. Remove any static IP configurations on the next panel and press **Add**. Some selections are already configured, such as Device Type: Ethernet. Type, for example, ehea, as module name for the adapter and click **Next**.
- v. 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.
- vi. 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.

- e. Network Services - Configures network clients and daemons using this option. There are more than 15 daemons to choose from and all are optional.
- f. Security and Users - Creates users and configures security policies. To allow SSH through the enabled firewall, follow these steps:
 - 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.

It is mandatory to define the root user password to log in a root via SSH and the steps to set the password are:

- i. Select **User Management** from the left menu.
- ii. Click **Configure**.
- iii. Select **Users**, high-light **root** and its row, and change the Filter to **System Users** by clicking the **Set Filter** drop-down button as shown in Figure C-8 on page 440.

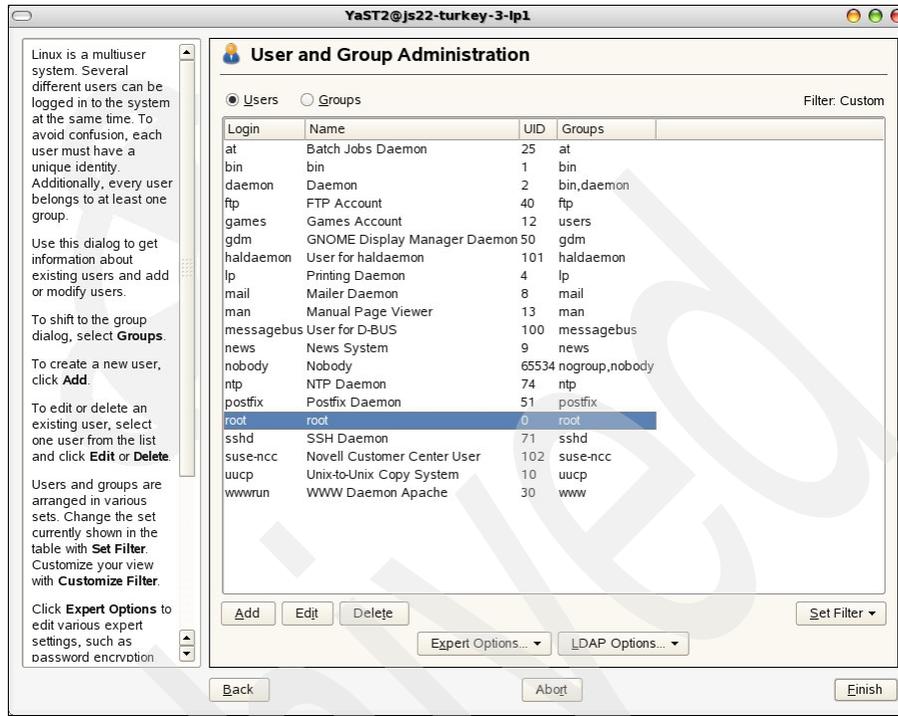


Figure C-8 Configure the root user

- iv. High-light **root** and its row again and press **Edit**.
 - v. Add the root user password. This password is saved encrypted in the XML file. Press **Accept** when finished.
 - vi. Click **Finish** to return to the AutoYaST main menu.
 - g. Misc - Allows 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**.

Example: C-1 Newly created XML file

```
<?xml version="1.0" ?>
<!DOCTYPE profile (View Source for full doctype...)>
- <profile xmlns="http://www.suse.com/1.0/yast2ns"
xmlns:config="http://www.suse.com/1.0/configns">
- <bootloader>
- <global>
<activate>>true</activate>
<boot_chrp_custom>/dev/sda1</boot_chrp_custom>
<default>linux</default>
```

```

<lines_cache_id>1</lines_cache_id>
<timeout config:type="integer">80</timeout>
</global>
- <initrd_modules config:type="list">
- <initrd_module>
<module>ibmvscsic</module>
</initrd_module>
</initrd_modules>
<loader_type>ppc</loader_type>
- <sections config:type="list">
- <section>
<append />
<image>/boot/vmlinux-2.6.16.60-0.21-ppc64</image>
<initial>1</initial>
<initrd>/boot/initrd-2.6.16.60-0.21-ppc64</initrd>
<kernel>/boot/vmlinux</kernel>
<lines_cache_id>0</lines_cache_id>
<name>linux</name>
<original_name>linux</original_name>
<root>/dev/sda3</root>
<type>image</type>
</section>
</sections>
</bootloader>
- <firewall>
<FW_ALLOW_FW_BROADCAST_DMZ />
<FW_ALLOW_FW_BROADCAST_EXT />
<FW_ALLOW_FW_BROADCAST_INT />
<FW_DEV_DMZ />
<FW_DEV_EXT>any eth-id-00:1a:64:44:31:f2 eth0</FW_DEV_EXT>
<FW_DEV_INT>eth1</FW_DEV_INT>
<FW_FORWARD_ALWAYS_INOUT_DEV />
<FW_FORWARD_MASQ />
<FW_IGNORE_FW_BROADCAST_DMZ>no</FW_IGNORE_FW_BROADCAST_DMZ>
<FW_IGNORE_FW_BROADCAST_EXT>yes</FW_IGNORE_FW_BROADCAST_EXT>
<FW_IGNORE_FW_BROADCAST_INT>no</FW_IGNORE_FW_BROADCAST_INT>
<FW_IPSEC_TRUST>no</FW_IPSEC_TRUST>
<FW_LOG_ACCEPT_ALL>no</FW_LOG_ACCEPT_ALL>
<FW_LOG_ACCEPT_CRIT>yes</FW_LOG_ACCEPT_CRIT>
<FW_LOG_DROP_ALL>no</FW_LOG_DROP_ALL>
<FW_LOG_DROP_CRIT>yes</FW_LOG_DROP_CRIT>
<FW_MASQUERADE>no</FW_MASQUERADE>
<FW_PROTECT_FROM_INT>no</FW_PROTECT_FROM_INT>
<FW_ROUTE>no</FW_ROUTE>
<FW_SERVICES_ACCEPT_DMZ />
<FW_SERVICES_ACCEPT_EXT />
<FW_SERVICES_ACCEPT_INT />
<FW_SERVICES_DMZ_IP />
<FW_SERVICES_DMZ_RPC />

```

```

<FW_SERVICES_DMZ_TCP />
<FW_SERVICES_DMZ_UDP />
<FW_SERVICES_EXT_IP />
<FW_SERVICES_EXT_RPC />
<FW_SERVICES_EXT_TCP>ssh</FW_SERVICES_EXT_TCP>
<FW_SERVICES_EXT_UDP />
<FW_SERVICES_INT_IP />
<FW_SERVICES_INT_RPC />
<FW_SERVICES_INT_TCP />
<FW_SERVICES_INT_UDP />
<enable_firewall config:type="boolean">true</enable_firewall>
<start_firewall config:type="boolean">true</start_firewall>
</firewall>
- <general>
- <mode>
<confirm config:type="boolean">false</confirm>
</mode>
- <mouse>
<id>none</id>
</mouse>
- <signature-handling>
<accept_file_without_checksum
config:type="boolean">true</accept_file_without_checksum>
<accept_non_trusted_gpg_key
config:type="boolean">true</accept_non_trusted_gpg_key>
<accept_unknown_gpg_key config:type="boolean">true</accept_unknown_gpg_key>
<accept_unsigned_file config:type="boolean">true</accept_unsigned_file>
<accept_verification_failed
config:type="boolean">false</accept_verification_failed>
<import_gpg_key config:type="boolean">true</import_gpg_key>
</signature-handling>
</general>
<groups config:type="list" />
- <networking>
- <dns>
<dhcp_hostname config:type="boolean">false</dhcp_hostname>
<dhcp_resolv config:type="boolean">false</dhcp_resolv>
</dns>
- <interfaces config:type="list">
- <interface>
<bootproto>dhcp</bootproto>
<device>eth0</device>
<startmode>auto</startmode>
<usercontrol>no</usercontrol>
</interface>
- <interface>
<bootproto>dhcp</bootproto>
<device>eth1</device>
<startmode>auto</startmode>

```

```

<usercontrol>no</usercontrol>
</interface>
</interfaces>
<managed config:type="boolean">>false</managed>
- <modules config:type="list">
- <module_entry>
<device>static-0</device>
<module>ehea</module>
<options />
</module_entry>
- <module_entry>
<device>static-1</device>
<module>e1000</module>
<options />
</module_entry>
</modules>
- <routing>
<ip_forward config:type="boolean">>false</ip_forward>
</routing>
</networking>
- <partitioning config:type="list">
- <drive>
<device>/dev/sda</device>
- <partitions config:type="list">
- <partition>
<create config:type="boolean">>true</create>
<format config:type="boolean">>false</format>
<partition_id config:type="integer">65</partition_id>
<partition_nr config:type="integer">1</partition_nr>
<partition_type>primary</partition_type>
<size>-7848448</size>
</partition>
- <partition>
<create config:type="boolean">>true</create>
<filesystem config:type="symbol">swap</filesystem>
<format config:type="boolean">>true</format>
<mount>swap</mount>
<mountby config:type="symbol">device</mountby>
<partition_id config:type="integer">130</partition_id>
<partition_nr config:type="integer">2</partition_nr>
<partition_type>primary</partition_type>
<size>2146798080</size>
</partition>
- <partition>
<create config:type="boolean">>true</create>
<filesystem config:type="symbol">reiser</filesystem>
<format config:type="boolean">>true</format>
<mount>/</mount>
<mountby config:type="symbol">device</mountby>

```

```

<partition_id config:type="integer">131</partition_id>
<partition_nr config:type="integer">3</partition_nr>
<partition_type>primary</partition_type>
<size>13917173248</size>
</partition>
</partitions>
<use>all</use>
</drive>
</partitioning>
- <report>
- <errors>
<log config:type="boolean">>true</log>
<show config:type="boolean">>true</show>
<timeout config:type="integer">0</timeout>
</errors>
- <messages>
<log config:type="boolean">>true</log>
<show config:type="boolean">>true</show>
<timeout config:type="integer">0</timeout>
</messages>
- <warnings>
<log config:type="boolean">>true</log>
<show config:type="boolean">>true</show>
<timeout config:type="integer">0</timeout>
</warnings>
- <yesno_messages>
<log config:type="boolean">>true</log>
<show config:type="boolean">>true</show>
<timeout config:type="integer">0</timeout>
</yesno_messages>
</report>
- <security>
<console_shutdown>reboot</console_shutdown>
<cracklib_dict_path>/usr/lib/cracklib_dict</cracklib_dict_path>
<cwd_in_root_path>yes</cwd_in_root_path>
<cwd_in_user_path>yes</cwd_in_user_path>
<displaymanager_remote_access>no</displaymanager_remote_access>
<displaymanager_shutdown>all</displaymanager_shutdown>
<enable_sysrq>no</enable_sysrq>
<fail_delay>3</fail_delay>
<gid_max>60000</gid_max>
<gid_min>1000</gid_min>
<group_encryption>des</group_encryption>
<lastlog_enab>yes</lastlog_enab>
<obscure_checks_enab>yes</obscure_checks_enab>
<pass_max_days>99999</pass_max_days>
<pass_min_days>0</pass_min_days>
<pass_min_len>5</pass_min_len>
<pass_warn_age>7</pass_warn_age>

```

```

<passwd_encryption>blowfish</passwd_encryption>
<passwd_remember_history>0</passwd_remember_history>
<passwd_use_cracklib>yes</passwd_use_cracklib>
<permission_security>secure</permission_security>
<run_updatedb_as>nobody</run_updatedb_as>
<system_gid_max>499</system_gid_max>
<system_gid_min>100</system_gid_min>
<system_uid_max>499</system_uid_max>
<system_uid_min>100</system_uid_min>
<uid_max>60000</uid_max>
<uid_min>500</uid_min>
<useradd_cmd>/usr/sbin/useradd.local</useradd_cmd>
<userdel_postcmd>/usr/sbin/userdel-post.local</userdel_postcmd>
<userdel_precmd>/usr/sbin/userdel-pre.local</userdel_precmd>
</security>
- <software>
- <packages config:type="list">
<package>xorg-x11-Xnest</package>
<package>yast2-trans-en_US</package>
<package>zypper</package>
<package>libgdiplus</package>
<package>libgimpprint</package>
<package>libgsm</package>
<package>libicu</package>
<package>libiniparser</package>
<package>libnl</package>
<package>libogg</package>
<package>libogg-64bit</package>
<package>librtas</package>
<package>libssui</package>
<package>libsvg</package>
<package>libsvg-cairo</package>
<package>libvorbis</package>
<package>libvorbis-64bit</package>
<package>libzyp-zmd-backend</package>
<package>limal-nfs-server</package>
<package>limal-nfs-server-perl</package>
<package>mDNSResponder-lib</package>
<package>mDNSResponder-lib-64bit</package>
<package>makedumpfile</package>
<package>mono-winforms</package>
<package>net-snmp</package>
<package>pciutils-ids</package>
<package>perl-Bit-Vector</package>
<package>perl-Bootloader</package>
<package>perl-Carp-Clan</package>
<package>perl-Compress-Zlib</package>
<package>perl-Config-Crontab</package>
<package>perl-Config-IniFiles</package>

```

```

<package>perl-Crypt-SmbHash</package>
<package>perl-DBD-SQLite</package>
<package>perl-DBI</package>
<package>perl-Date-Calc</package>
<package>perl-File-Tail</package>
<package>perl-Digest-MD4</package>
<package>perl-Digest-SHA1</package>
<package>perl-Net-Daemon</package>
<package>perl-Parse-RecDescent</package>
<package>perl-PIRPC</package>
<package>perl-TermReadKey</package>
<package>perl-TimeDate</package>
<package>perl-X500-DN</package>
<package>perl-URI</package>
<package>perl-XML-Parser</package>
<package>perl-XML-Writer</package>
<package>perl-gettext</package>
<package>perl-libapparmor</package>
<package>powerpc-utils</package>
<package>python-cairo</package>
<package>qt-qt3support</package>
<package>samba-client</package>
<package>susehelp_en</package>
<package>taglib</package>
<package>unrar</package>
<package>update-alternatives</package>
<package>wdiff</package>
<package>flac-64bit</package>
<package>flac</package>
<package>aspell-en</package>
<package>iprutils</package>
<package>cifs-mount</package>
<package>audiofile</package>
<package>audiofile-64bit</package>
<package>boost</package>
<package>iso-codes</package>
<package>esound-64bit</package>
<package>esound</package>
<package>dmapi</package>
<package>OpenEXR-64bit</package>
<package>OpenEXR</package>
<package>dbus-1-qt</package>
<package>a2ps</package>
<package>aalib</package>
<package>aalib-64bit</package>
<package>dbus-1-mono</package>
</packages>
- <patterns config:type="list">
<pattern>print_server</pattern>

```

```

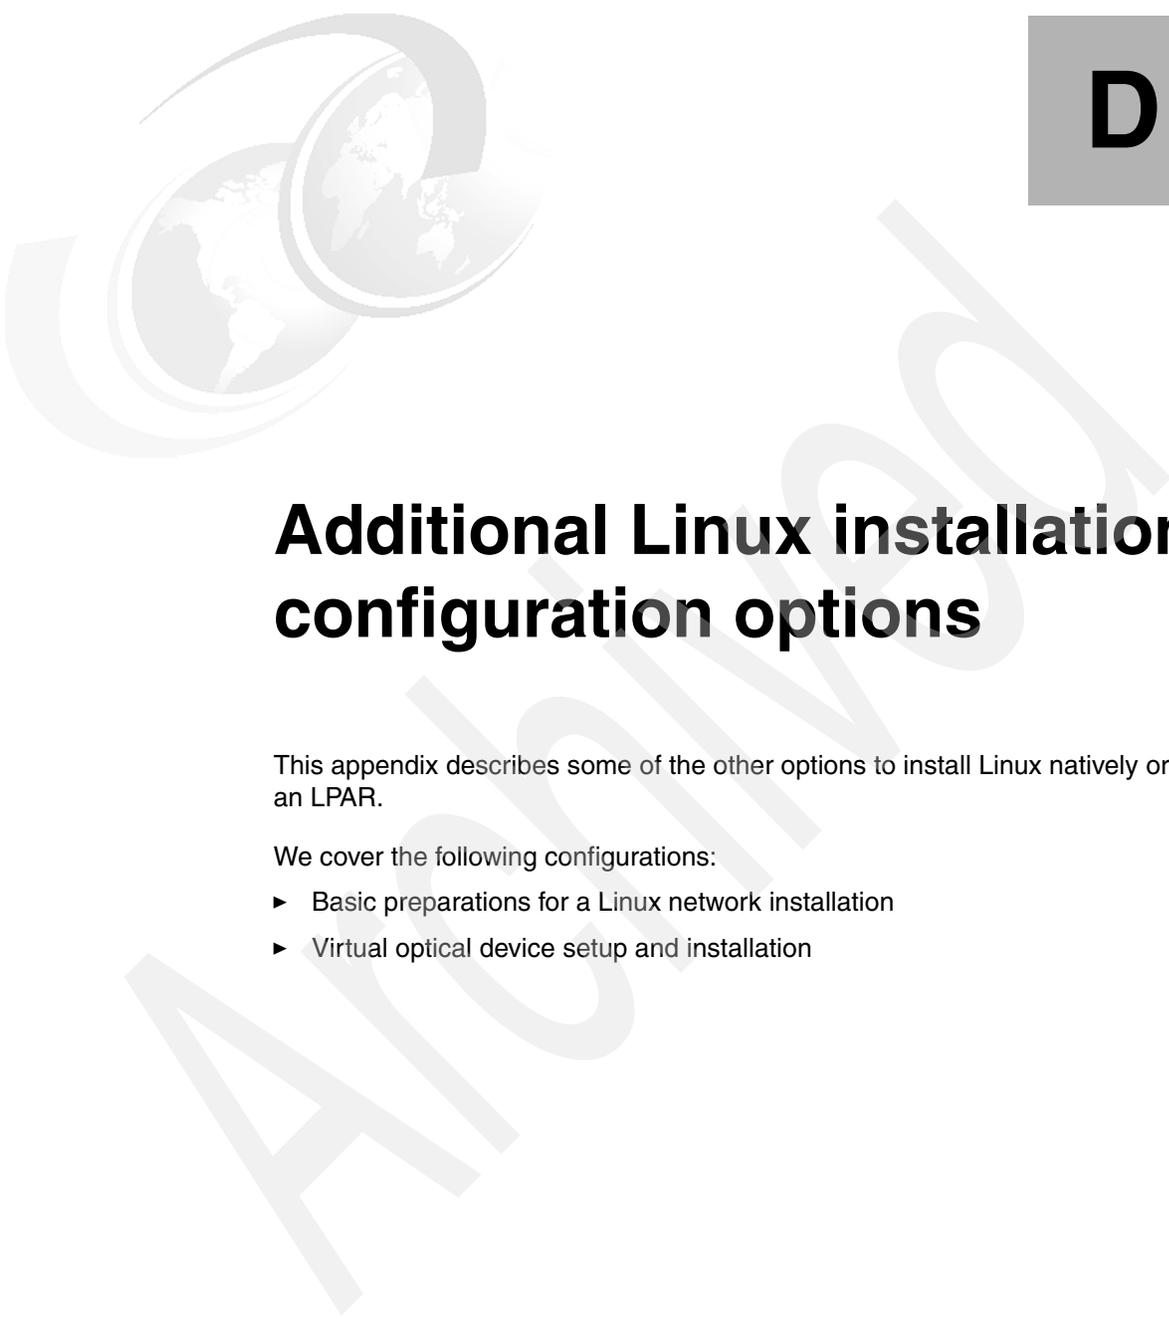
<pattern>print_server-64bit</pattern>
<pattern>x11</pattern>
<pattern>apparmor-64bit</pattern>
<pattern>64bit</pattern>
<pattern>gnome-64bit</pattern>
<pattern>base-64bit</pattern>
<pattern>gnome</pattern>
<pattern>base</pattern>
<pattern>x11-64bit</pattern>
<pattern>Documentation</pattern>
<pattern>Documentation-64bit</pattern>
<pattern>apparmor</pattern>
</patterns>
</software>
- <timezone>
<hwclock>UTC</hwclock>
<timezone>US/Eastern</timezone>
</timezone>
- <user_defaults>
<expire />
<group>500</group>
<groups>video</groups>
<home>/home</home>
<inactive />
<shell />
<skel />
</user_defaults>
- <users config:type="list">
- <user>
<encrypted config:type="boolean">>true</encrypted>
<user_password>$2a$05$8ASD5r59p019t9v1SsRGXerICQaZb3UL5f1y7cKbau069dzIa2y9C</us
er_password>
<username>root</username>
</user>
</users>
</profile>

```

Tip: A more detailed description of the automated installation setup is available in the SUSE Linux Enterprise Server 10 SP2 documentation and at this Web site:

http://www.novell.com/documentation/sles10/sles_admin/index.html?page=/documentation/sles10/sles_admin/data/cha_deployment_autoinst.html

Archived



Additional Linux installation configuration options

This appendix describes some of the other options to install Linux natively or on an LPAR.

We cover the following configurations:

- ▶ Basic preparations for a Linux network installation
- ▶ Virtual optical device setup and installation

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.

Installing Linux using the network - General remarks

You always require the following services to perform a network installation:

- ▶ 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 JS22.
- ▶ A running Trivial File Transfer Protocol (TFTP) service to serve the boot image to a BladeCenter JS22.
- ▶ It is necessary to set up *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)

This section shows how to configure 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 using the network. The standard DHCP daemon is called *dhcpcd*, but there are other DHCP daemons.

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 called *dhcpcd3-server*, you can find the configuration in */etc/dhcpcd3-server.conf* and */etc/sysconfig/dhcpcd3-server*, and the start/stop script is in */etc/init.d/dhcp3-server*.

The standard DHCP daemon is configured through two files: The */etc/sysconfig/dhcpd* file, which stores the basic configuration, and the */etc/dhcpd.conf* file, which contains the 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 D-1 for a simple client configuration stored in *dhcpd.conf*.

Note: Keep in mind that Example D-1 contains environment-specific Internet Protocol (IP) and Media Access Control (MAC) address information. One way to learn the MAC address of a JS22 BladeCenter is to use the BladeCenter management module. Select **Monitors** → **Hardware VPD** from the left-side options and then click your blade bay and then 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 **man dhcpd** on the command prompt.

Example: D-1 dhcpd.conf example

```
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 JS22 {
    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 to double-check or troubleshoot a configuration in general.

- ▶ To trace messages of running services, type `tail -f -n 10 /var/log/messages` to get the last 10 messages and auto update if there are new messages.
- ▶ Connect to a running service with a local client, remote client, or both these clients and try to receive the data that you want.
- ▶ 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. There are several implementations of TFTP daemons available. The standard TFTP daemon is called *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 SLES10 SP2:

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 D-1.

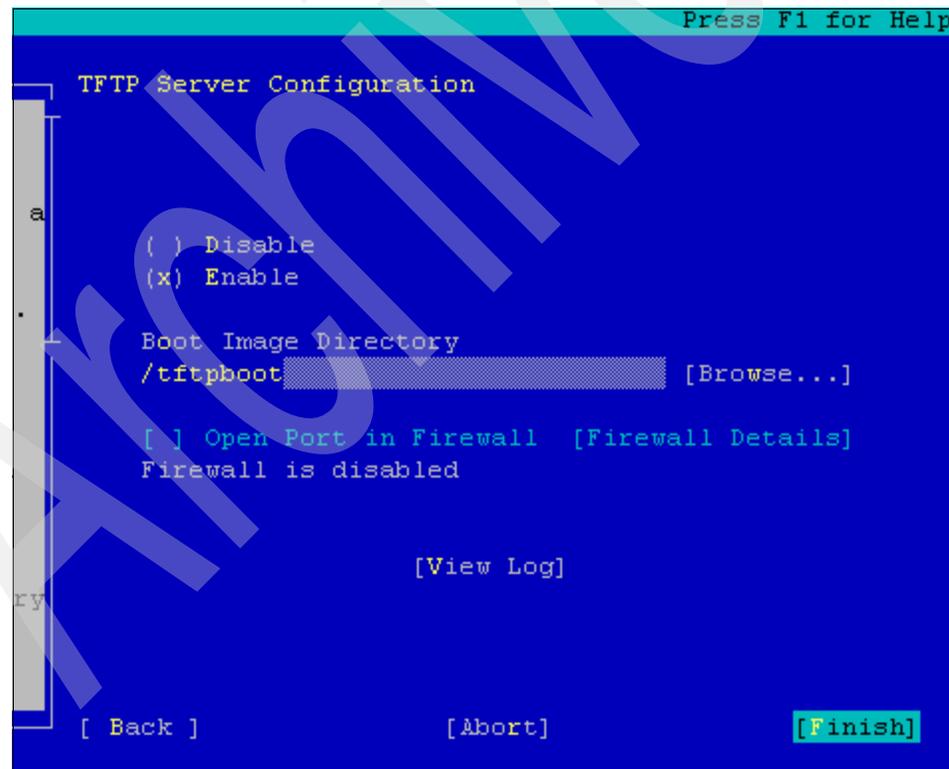


Figure D-1 Configure a TFTP server in SLES10 SP2

5. Finally, scroll down to **[Finish]** and press the Enter key.

Example D-2 shows a TFTP daemon configuration for xinetd stored in `/etc/xinet.d/tftpd`.

Example: D-2 Configuring a TFTP daemon in the `/etc/xinet.d/tftp` file on SLES10 SP2

```
# 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 using a text editor as well.

Red Hat Enterprise Linux and SLES create TFTP from the xinetd daemon. Restart the xinetd service after the tftp configuration is complete by:

```
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 D-1 on page 451. The process is slightly different for Red Hat Enterprise Linux and SLES, so we document both separately here.

Copying the SLES10 install kernel

To copy the SLES10 install kernel, use the following procedure:

1. Mount the SLES10 SP2 DVD1 on the system running the tftp server. For example, on a system running SLES, type:

```
mount /dev/sr0 /mnt

```

2. Then enter

```
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, use the following procedure:

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. Then enter

```
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 need 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 10 SP2. The configuration steps are:

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 location source location directory. This example uses /install as shown in Figure D-2 on page 456.

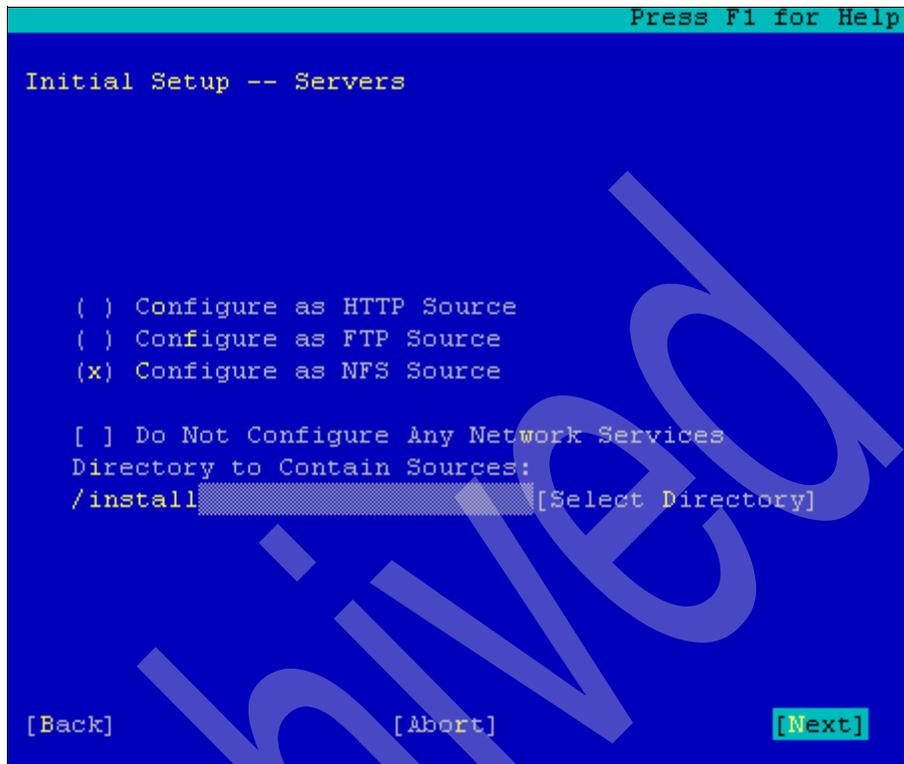


Figure D-2 Initial setup of SLES NFS installation server

4. Then click **[Next]**.
5. Leave the defaults for **Host Wild Card** and **Options**.
6. Click **[Next]**. With this, an NFS server serving /install is set up automatically.
7. Click **Add** to configure an installation source.
8. As **Source Name**, enter the desired name for this installation source, for example, sles10sp2. This creates a subdirectory sles10sp2 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 via **Select Directory** to the directory that contains all ISO images of all CDs (see Figure D-3 on page 457).
10. Click **[Next]** when finished.

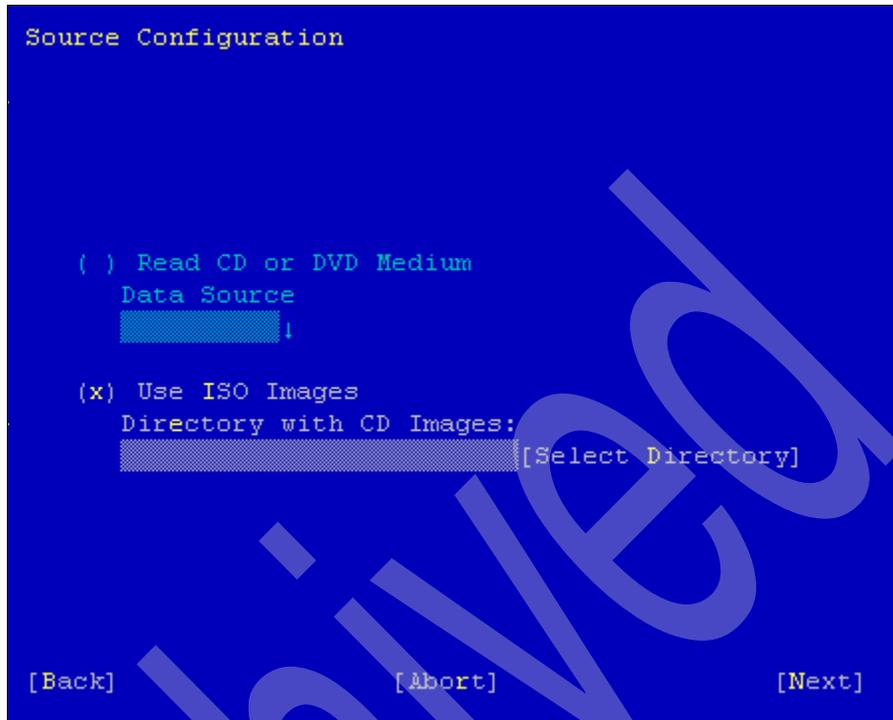


Figure D-3 Source configuration window

11. If you chose the Read CD or DVD Medium option given in Figure D-3, you will be prompted to insert the first DVD.
12. Insert SLES10 SP2 DVD1 into the BladeCenter media tray and press **[Continue]**. The data from DVD1 is copied to the `/install/sles10sp2/CD1` directory.

Note: If you used the CD option instead of a DVD, you will be 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

This subsection provides the steps 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 via NFS entry in `/etc/exports`. For example:

```
/install/RHEL5.2 *(ro, async, no_root squash)
```

3. Then 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 3.6.4, "Optical devices" on page 108 to copy the CD/DVD image onto the VIOS virtual media library. Once the CD/DVD image is copied onto the VIOS hard disk, assign the virtual optical device to the LPAR. The install process is the same as in 6.2, "Linux LPAR installation using DVD" on page 285 for Red Hat Enterprise Linux and 7.3, "Linux LPAR installation using DVD" on page 307 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 JS22 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.

VIOS/IVM Command Help

This appendix shows available methods for obtaining command line help from the restricted shell in VIOS.

The help command

Enter `help` to show an overview of the available commands:

```
$ help
```

Install Commands

`ioslevel`

`license`

`lssw`

`oem_platform_level`

`oem_setup_env`

`remote_management`

`updateios`

LAN Commands

`cfglnagg`

`cfgnamesrv`

`entstat`

`fcstat`

`hostmap`

Security Commands

`lsfailedlogin`

`lsgcl`

`viosecure`

`mkldap`

`ldapadd`

`ldapsearch`

`snmpv3_ssw`

`mkkrb5c1nt`

UserID Commands

`chuser`

`lsuser`

`mkuser`

`passwd`

hostname
lsnetsvc
lstcpip
mktcpip
chtcpip
netstat
optimizenet
ping
rmtcpip
seastat
startnetsvc
stopnetsvc
traceroute
vasistat

Device Commands

chdev
chpath
cfgdev
lsdev
lsmap
lspath
mkpath
mkvdev
mkvt
rmdev
rmpath
rmvdev
rmvt

Physical Volume Commands

lspv
migratepv

Logical Volume Commands

chlv
cplv
extendlv
lslv
mklv
mklvcopy
rmlv
rmlvcopy

Volume Group Commands

rmuser

Maintenance Commands

alt_root_vg
backup
backupios
bootlist
cattracerpt
chdate
chlang
cfgassist
cl_snmp
cpvdi
dsmc
diagmenu
errlog
fsck
invscout
ldfware
loginmsg
lsfware
lslparinfo
motd
mount
pdump
replphyvol
restore
restorevgstruct
save_base
savevgstruct
showmount
shutdown
snap
snmp_info
snmp_trap
startsysdump
starttrace
stoptrace
svmon
sysstat
topas
uname
unmount
viostat
vmstat

```
activatevg  
chvg  
deactivatevg  
exportvg  
extendvg  
importvg  
lsvg  
mirrorios  
mkgv  
redefvg  
reducevg  
syncvg  
unmirrorios
```

Storage Pool Commands

```
chbdsp  
chsp  
lssp  
mkbdsp  
mksp  
rmbdsp  
rmsp
```

Virtual Media Commands

```
chrep  
chvopt  
loadopt  
lsrep  
lsvopt  
mkrep  
mkvopt  
rmrep  
rmvopt  
unloadopt
```

```
wkldmgr  
wkldagent  
wkldout
```

Monitoring Commands

```
cfgsvc  
lssvc  
startsvc  
stopsvc
```

Shell Commands

```
awk  
cat  
chmod  
clear  
cp  
crontab  
date  
ftp  
grep  
head  
ls  
man  
mkdir  
more  
mv  
rm  
sed  
stty  
tail  
tee  
vi  
wall  
wc  
who
```

Using the -h flag

Help relating to an individual command is available with the -h flag as shown in the following example using the **mkvt** command.

```
$ mkvt -h  
Usage: mkvt {-id PartitionID}
```

Opens a virtual terminal session to an operating system or virtual I/O server partition in a managed system.

`-id` The ID of the partition for which to open the virtual terminal session.

Using the man command

Detailed help relating to an individual command is available using the **man** command as shown in the following example using the **man mkvt** command.

```
$ man mkvt
```

```
mkvt Command
```

```
Purpose
```

```
Create a virtual terminal connection to a partition.
```

```
Syntax
```

```
mkvt { -id lparID }
```

```
Description
```

```
The mkvt command opens a virtual terminal connection to the target partition. You can terminate the virtual terminal connection in one of the following ways:
```

- * The virtual terminal contains an escape sequence that allows you to break out of the command. The escape sequence is `<cr>~.`, or more explicitly: the Enter key, the tilde (~), and the period (.
- * You can use the `rmvt` command to force the session to be closed.

```
A partition can only have one open virtual terminal session.
```

This command requires additional HMC configuration if used in an HMC environment.

Flags

`-id lparID`
The ID of the partition for which to open the virtual terminal session.

Exit Status

27
Unexpected error

28
Virtual terminal is already connected

29
Virtual terminal device is not found

30
Permission denied

31
Specified device does not exist

Security

This command is not accessible by users with the `ViewOnly` role.

Examples

1 Create a virtual terminal connection to the partition with ID 3:

```
mkvt -id 3
```

Related Information

The `rmvt` command.

Archived

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 JS22 BladeCenter.

This appendix contains the following:

- ▶ Overview of IBM service diagnostic aids and productivity tools
- ▶ Installing tools on Red Hat Enterprise Linux or SUSE Linux Enterprise Server running on BladeCenter servers
- ▶ Installing tools on Red Hat Enterprise Linux or SUSE Linux Enterprise Server running on IVM-managed servers

IBM service diagnostic aids and productivity tools 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) functionality as well as the ability to modify logical partition (LPAR) profiles with hotplug, Dynamic Logical Partitioning (DLPAR), and Live Partition Migration 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 F-1 on page 467 shows how to determine the appropriate packages for your environment.

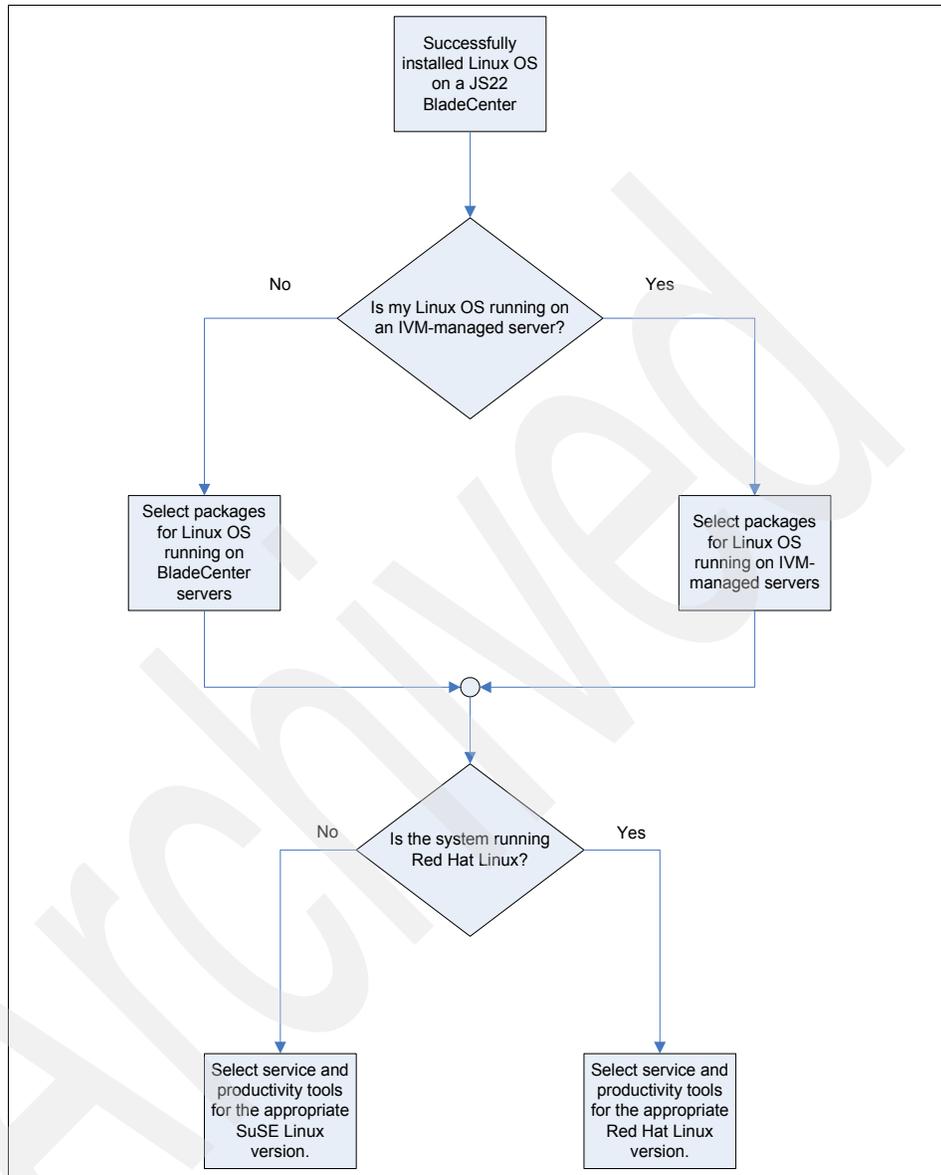


Figure F-1 Service diagnostic and productivity packages decision tree

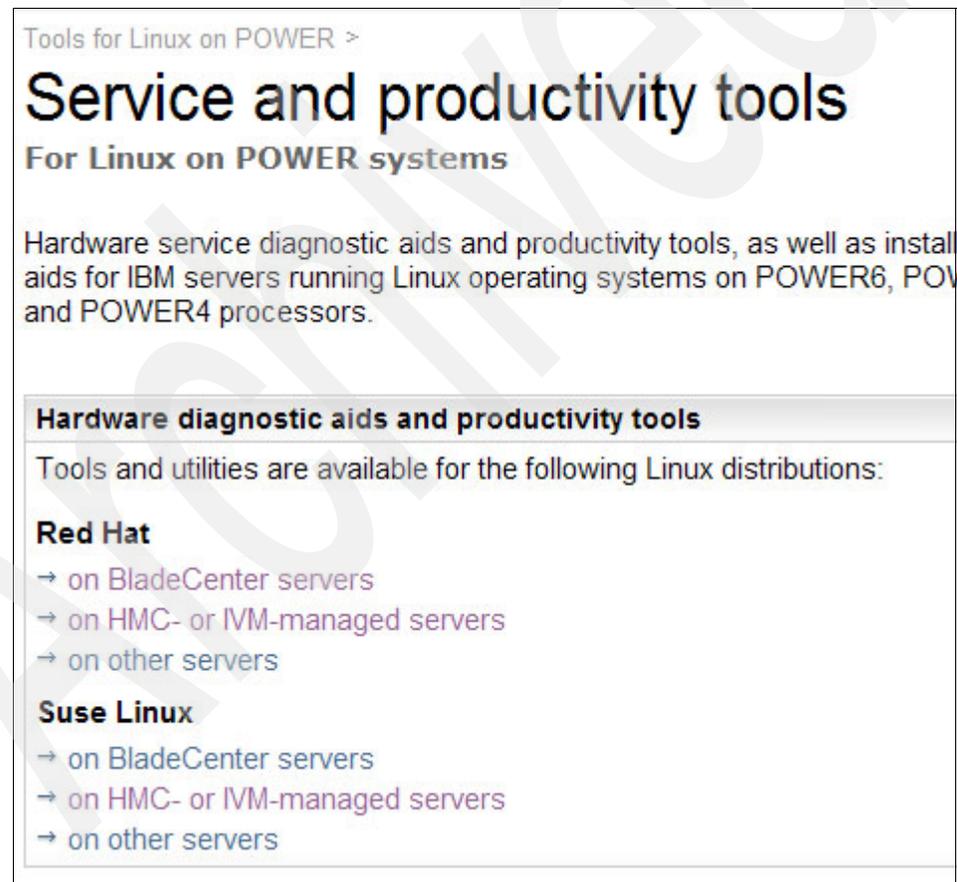
Install tools on Red Hat Enterprise Linux 5/SLES 10 running on BladeCenter servers

This section describes the steps to configure a JS22 BladeCenter running on a BladeCenter server with the service aids and productivity tools. These steps are applicable for systems running a native Red Hat Enterprise Linux 5/SLES 10 (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 **BladeCenter servers** link under the Red Hat or SuSE Linux distribution headings, as shown in Figure F-2.



Tools for Linux on POWER >

Service and productivity tools

For Linux on POWER systems

Hardware service diagnostic aids and productivity tools, as well as install aids for IBM servers running Linux operating systems on POWER6, POWER7, and POWER4 processors.

Hardware diagnostic aids and productivity tools

Tools and utilities are available for the following Linux distributions:

Red Hat

- [on BladeCenter servers](#)
- [on HMC- or IVM-managed servers](#)
- [on other servers](#)

Suse Linux

- [on BladeCenter servers](#)
- [on HMC- or IVM-managed servers](#)
- [on other servers](#)

Figure F-2 Select “on BladeCenter servers”

- Click the tab (as shown in Figure F-3) that matches your Linux operating system (OS) level. The packages under each tab are unique to that Linux OS level.

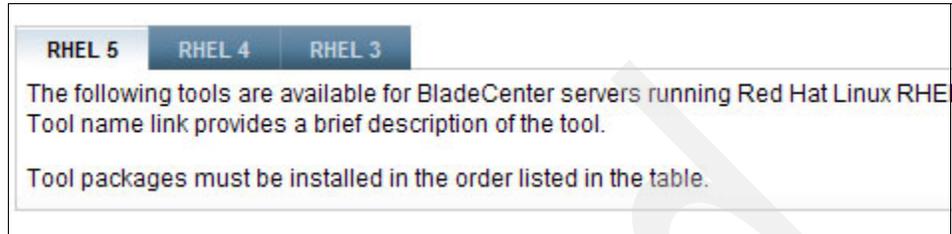


Figure F-3 OS level selection tabs example

- Click and save each of the packages under the Package downloads column. At the time of this publication the packages were:

RHEL 5 tools		
Tool name	Package downloads	Last u
Platform Enablement Library	librtas-1.3.3-0.ppc64.rpm	2008.0
Hardware Inventory	lsvpd-0.15.1-1.ppc.rpm	2006.0
Service log	servicelog-0.2.9-0.ppc64.rpm	2008.0
Error Log Analysis	diagela-2.2.2-0.ppc64.rpm	2008.0

Figure F-4 Available packages for Red Hat on BladeCenter servers

Tool	Basic Information
Platform Enablement Library	A library that allows 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 may require service.

Tool	Basic Information
Error Log Analysis	Provides automatic analysis and notification of errors reported by the platform firmware.

Tip: Click the links under the Tool name column for the latest detailed description of each tool.

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 the CD/DVD tip below).
6. Install each rpm package with
rpm -Uvh <packagename>.rpm

Important: These packages *must* be installed in the order listed in the table. Otherwise there will be dependency failures.

7. We recommend a system shutdown and restart after installing all the packages.
8. The service aids and productivity tools section is complete.

Tip1: If you are placing the rpms on a CD and DVD, here are some steps to access the files:

1. Assign the JS22 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/
or 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: We recommend placing these rpms in a yum repository to quickly update or install these tools on a large number of machines.

Install tools on Red Hat Enterprise Linux 5/SLES 10 running on IVM-managed servers

This section describes the steps to configure a JS22 BladeCenter LPAR running on a 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/1opdiags/home.html>

2. On this Web site, click **HMC- or IVM-managed servers** in the Red Hat/SuSE Linux section, depending on what OS is running on the LPAR.
3. Click the tab (as shown in Figure F-3 on page 469) 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 shown in Figure F-5 on page 472.

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	lsvpd-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 F-5 Available packages for Red Hat/SuSE Linux on IVM-managed server

Tool	Basic Information
Platform Enablement Library	A library that 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 needed to monitor and manage one or more Linux systems
RSCT core	See description above.
CSM core	CSM packages provide for the exchange of host-based authentication security keys.
CSM-client	See description above.
ServiceRM	Service Resource Manager is a Reliable, Scalable, Cluster Technology (RSCT) resource manager that creates the Serviceable Events from the output of the Error Log Analysis Tool (diagela).
DynamicRM	Dynamic Resource Manager is a Reliable, Scalable, Cluster Technology (RSCT) resource manager that allows a 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 may require service.

Tool	Basic Information
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.

Tip: Click the links under the Tool name column for the latest detailed description of each tool.

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 are placing the rpms on a CD and DVD, here are some steps to access the files.” on page 470).
6. Install each rpm package with

```
rpm -Uvh <packagename>.rpm
```

Important: These packages *must* be installed in the order listed in the table. Otherwise there will be dependency failures.

7. We recommend a complete LPAR shutdown and reactivation after installing all the 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 F-6 on page 475 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 F-6 DLPAR and Live Partition mobility services are enabled

See Chapter 3, “System configuration using VIOS and IVM” on page 55 for more information on IVM options and functions.

9. Installation of the service aids and productivity tools is complete.

Tip: We recommend placing these rpms in a yum repository to quickly update or install these tools on a large number of machines.

Archived

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 484. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM BladeCenter Products and Technology*, SG24-7523
- ▶ *IBM System i and System p*, SG24-7487
- ▶ *IBM System Storage DS4000 and Storage Manager V10.10*, SG24-7010
- ▶ *PowerVM Live Partition Mobility on IBM System p*, 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
- ▶ *PowerVM Virtualization on IBM System p Managing and Monitoring*, SG24-7590
- ▶ *Implementing IBM Director 5.20*, SG24-6188
- ▶ *Going Green with IBM Active Energy Manager*, REDP-4361
- ▶ *Integrated Virtualization Manager on IBM System p5*, REDP-4061
- ▶ *Implementing the IBM BladeCenter S Chassis*, REDP-4357
- ▶ *IBM System i Overview: Models 515, 525, 550, 570, 595, and More*, REDP-5052

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. For more information about the IBM Active Energy Manager see:
<http://www.ibm.com/systems/management/director/extensions/actengmrg.html>
- ▶ BM periodically releases maintenance packages for the AIX 5L operating system. These packages are available on CD-ROM, or you can download them from the following Web site:
<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. For more information about the suma command functionality, visit the following Web site:
<http://www.ibm.com/systems/p/os/aix/whitepapers/suma.html>
- ▶ For information about the features and external devices supported by Linux on the IBM BladeCenter JS22 and IBM BladeCenter JS12 Express server visit the following Web site:
<http://www.ibm.com/systems/power/software/linux/index.html/>
- ▶ For information about SUSE Linux Enterprise Server 10, visit the following Web site:
<http://developer.novell.com/yesssearch/Search.jsp>
- ▶ For information about Red Hat Enterprise Linux, visit the following Web site:
<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 the following Web site:
http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html
- ▶ For more information of IBM i 6.1 operating systems running on IBM BladeCenter JS22 and IBM BladeCenter JS12 Express see:
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>
- ▶ Visit the following Web site to download IBM Director 5.2:
<http://www.ibm.com/systems/management/director/download.html>
- ▶ For more information about CSM visit the following Web site:
<http://www.ibm.com/systems/cluster/software/csm/>

- ▶ Complete VIOS configuration and maintenance information can be found in *System i and System p Using the Virtual I/O Server*, found at this link:
<http://publib.boulder.ibm.com/infocenter/systems/topic/iphb1/iphb1pdf.pdf>
- ▶ The AIX process can be reviewed in the White Paper *Installing AIX6 on the POWER6 BladeCenter JS22 using Keyboard Video Mouse or Serial Over LAN* available at:
ftp://w3.ibm.com/support/americas/pseries/js22AIX61_v2_122607.pdf
- ▶ The BladeCenter Interoperability Guide can be found at:
<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ndocid=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. The data sheet can be found at:
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>
- ▶ All supported hardware and operating systems are listed on IBM ServerProven. ServerProven can be found at:
<http://www-03.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 at:
<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/eserver.html>
- ▶ To verify which of those 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>
- ▶ The DS8000 interoperability matrix can be found at:
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>
- ▶ The DS6000 interoperability matrix can be found at:
<http://www.ibm.com/servers/storage/disk/ds6000/interop.html>
- ▶ The DS4000 interoperability matrix can be found at:
<http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html>
- ▶ The DS3000 interoperability matrix can be found at:
<http://www-03.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf>

- ▶ The ESS interoperability matrix can be found at:
<http://www.ibm.com/servers/storage/disk/ess/interop-matrix.html>
 The N series interoperability matrix can be found at:
<http://www-03.ibm.com/systems/storage/nas/interophome.html>
- ▶ The SAN volume Controller support matrix can be found at:
<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskind=3&brandind=5000033&familyind=5329743>
- ▶ The SAN switch interoperability matrix can be found at:
<http://www-03.ibm.com/systems/storage/san/index.html>
- ▶ When you plan to implement IBM i on a JS12 or JS22 verify in addition to the previous steps the support matrix of IBM i on blades. Some restrictions apply when IBM i is used.
<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>
http://www-03.ibm.com/systems/power/hardware/blades/supported_environments.pdf
- ▶ The System Storage Interoperation Center (SSIC) helps to identify a supported storage environment. You find this web based tool at:
<http://www-03.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. They can be found at:
<http://www.ibm.com/systems/support/storage/config/hba/index.wss>
- ▶ 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>
- ▶ The PowerVM Editions Operations Guide contains useful information:
<http://publib.boulder.ibm.com/infocenter/systems/topic/iphdx/sa76-0100.pdf>
- ▶ The AIX process can be reviewed in the White Paper *Installing AIX6 on the POWER6 BladeCenter JS22 using Keyboard Video Mouse or Serial Over LAN* available at:
ftp://w3.ibm.com/support/americas/pseries/js22AIX61_v2_122607.pdf
- ▶ For more information on supported devices on a BladeCenter JS12 server refer to:
<http://www.ibm.com/systems/i/os/i5os/v6r1/blades/config.html>

- ▶ For more information on supported components and procedure to install updates to IBM BladeCenter JS22 or IBM BladeCenter JS12 Express server, BladeCenter S or BladeCenter H chassis refer to:

<http://www.ibm.com/systems/i/os/i5os/v6r1/blades/config.html>

- ▶ For firmware updates for IBM BladeCenter JS22 or IBM BladeCenter JS12 Express, BladeCenter S, and BladeCenter H chassis and integrated components refer to:

<http://www.ibm.com/system/support/supportsite.wss/brandmain?brandind=5000020>

- ▶ The latest fix pack are part of the BladeCenter JS12 order. Consult the supported environments page to verify you have the minimum supported release of VIOS:

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

- ▶ To get the latest recommended fixes and information refer to:

http://www-912.ibm.com/s_dir/s1kbase.nsf/recommendedfiles

- ▶ To receive detailed information on how to define a SAN environment for an IBM BladeCenter JS22 or JS12 blade installed in a BladeCenter H chassis using LUN attached disks to install IBM i 6.1 refer to:

http://www.ibm.com/systems/power/hardware/blades/i_on_blade_readme.pdf

- ▶ The Storage Configuration Manager (SCM) is proposed 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.ibm.com/systems.support/supportsite.wss/docdisplay?Indocid=MIGR-5502070?brandid=5000008>

- ▶ To meet all the System i Access for Windows hardware requirements follow the instructions described in:

<http://publib.boulder.ibm.com/infocenter/systems/scopt/i5os/topic/rzajr/rzajrhardwarereq.htm#hardwarereq>

- ▶ To meet all the System i Access for Windows software requirements follow the instructions described in:

<http://publib.boulder.ibm.com/infocenter/systems/scopt/i5os/topic/rzajr/rzajrsoftwarereq.htm#softwarereq>

- ▶ For more information to 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:

<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 6.1 user profiles. Refer to Service tools user IDs at:
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzanh/rzanhwhatuserids.htm>
- ▶ The IBM Workload Estimator is available at:
<http://www.ibm.com/systems/support/tools/estimator/index.html>
- ▶ Performance actions related to disk formatting are described in InfoCenter at the following:
<https://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzajy/rzajoverview.html>
- ▶ Language feature codes are available at:
<http://publib.boulder.ibm.com/infocenter/scope/i5os/topic/rzahc/rzahcnlvfeaturecodes.htm#rzhcnlvfeaturecodes>
- ▶ Software installation information for IBM i:
<http://publib.boulder.ibm.com/infocenter/series/v5r3/topic/rzahc/rzahcswinstallprocess.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/slkbase.nsf/recommendedfixes
- ▶ The primary Web site for downloading fixes for all operating systems and applications refer to:
<http://www-912.ibm.com/eserver/support/fixes>
- ▶ Refer to the Ethernet on System i 6.1 topic in the IBM Information Center to configure the IBM i 6.1 communication:
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzajy/rzajoverview.html>
- ▶ More detailed information to the IBM Systems Director Navigator for i functionality can be found at:
<http://www.ibm.com/systems/i/software/navigator/index.html>
- ▶ A detailed procedure using the IBM SAS LTO tape library attached to a BladeCenter S chassis refer to:
http://www.ibm.com/systems/power/hardware/blades/i_on_blade_readme.html
- ▶ 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>

- ▶ The Red Hat Recommended Partitioning Scheme is available at:
http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/ch11s03.html
- ▶ A 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>
- ▶ The Red Hat Recommended Partitioning Scheme is available at:
http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/ch11s03.html
- ▶ RedHat installation instructions are available here:
http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/pt-install-info-ppc.html
- ▶ 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>
- ▶ The Novell Web site has additional installation preparation information for SLES10 Sp2 available at:
<http://www.novell.com/documentation/sles10/index.html>
- ▶ Your awesome Setup Tool (YaST) will assist with the completion of a SLES install. More detailed installation instructions are available here:
http://www.novell.com/documentation/sles10/sles_admin/index.html?page=/documentation/sles10/sles_admin/data/sles_admin.html
- ▶ The link to the Virtual I/O Server download site is also available here:
<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 available at:
<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>
- ▶ For details about how to setup SOL see the Serial over LAN setup Guide. This guide can be found at:
<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-54666&brandind=5000020>

- ▶ The *Management Module command-line interface reference guide* can you find online at:
<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-54667&brandind=5000020>
- ▶ For a technical overview and full details, as well as latest updates on IBM i on Power blades, see the Read-me First:
<http://www-03.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

(IVM)

View/Modify Host Ethernet Adapters

Properties

General tab 87

/etc/dhcpd.conf 451

/etc/init.d/dhcpd restart command 452

/etc/init.d/dhcpd start command 452

/etc/xinet.d/tftpd 454

/sbin/service nfs reload command 458

/sbin/service nfs start command 458

A

ABAP 34

AC Power Source Redundancy 354

active console 404

Active Energy Manager 38, 377

Active Energy Manager (AEM) 3, 364

Active Energy Manager Console 364

current data 366

Manage Power 368

policies 369

policies, creating 377

trend data 366

Advanced Management Module 352

Advanced Management Web interface 49

Advanced Peer-to-Peer Networking 270

Advanced Program-to-Program Communication
270

AEM 3, 38, 364, 373

AIX 47

Installation 149

AIX 6.1

LPAR Installation on VIOS 169

AIX install kernel load 172

Alternate processor retry 32

Altivec 35

Altivec (SIMD) 35

Altivec SIMD accelerator 5

AMM 216

APPC 270

APPN 270

automated monitoring 50

AutoYaST 429–430

Autoyast 326

B

Backup/Restore

IBM i 6.1 271

Basic Zone Permission Table 220

BladeCenter H 13

BladeCenter Hardware preparation 214

BladeCenter Power Domain 357

boot

from iSCSI LUN 200

boot device 66

boot net autoyast=nfs

//193.200.1.80/home/autoinst.xml install=nfs

//192.168.1.254/SuSE/SLES10_SP2 com-
mand 327

boot net ks=nfs

//192.168.1.254/ks.cfg ksdevice=eth1 ip=dhcp

command 303

Bootstrap Protocol (BOOTP) 450

bridge 162

BRMS 272

C

cache availability 32

cards 44

CDROM

virtual 166

CFFh 57

CFFv 57

chassis 11

Chassis Power 354

Chipkill technology 3

chsyscfg 56

chsysstate 142, 147

Cisco 44

client partition

create using IVM 156

Cluster Systems Management 49–50

Combined Form Factor horizontal 57

Combined Form Factor vertical 57

command

- install_assist 177, 200
- Commands
 - smitty
 - smitty install_latest 178
- commands
 - AMM
 - env 363
 - fuelg 361
 - chsyscfg 56
 - IBM i
 - SHTDWN SYS 144
 - IVM
 - chsyscfg 85
 - chsysstate 142
 - LPAR shutdown 147
 - lpcfgop 74
 - lshwres 85
 - lslparmigr 349
 - lsrefcode 143
 - lssyscfg, determine RMC status 340
 - lssyscfg, determine system name 349
 - lssyscfg, LPAR state 143
 - lssyscfg, migrate status 349
 - lsvet 332
 - migrpar 348
 - power 403
 - printenv 409
 - printenv boot-device 409
 - show-devs 409
 - smitty
 - smitty tmiscsi 183
 - suma 48
 - telnetcfg 402
 - VIOS
 - cfgassist 74–75
 - chdev, change hdisk reserve policy 335
 - entstat 56
 - help 73
 - ioslevel 331
 - license 74
 - lsdev 100
 - lsdev, check hdisk reserve policy 335
 - lspv 100
 - lstcpip 74, 77
 - mktcpip 74–75
 - mkvt 140
 - rmtcpip 80
- configuration file
 - Kickstart 427
 - configuration mask 192
 - Connect the System i LAN console 236
 - Configuration wizard 237
 - dedicated service tools (DST) 243
 - hosts file 241
 - Service tool device ID 243
 - System serial number 242
 - Target partition 242
 - Console Identification 248
 - Create virtual media library for backup
 - IBM i V6R1 Backup/Restore
 - Create virtual media library for backup 272
 - Creating IBM i partition 247
 - Load Source and Console Identification 248
 - CSM 49–50

D

- Decimal floating point 33
- Dedicated Device Tools 248
- default ID 73
- devices.tmiscsw 181
- DHCP
 - AutoYaST configuration 439
 - configuration 451
 - restarting 452
- diagnostics 51
- digital KVM 394
- Director 50
- Disk configuration in BladeCenter S
 - Basic Zone Permission Table 220
 - Individual BladeCenter S disk configuration 224
 - Predefined Config 223
 - Storage Configuration Manager (SCM) 481
- disk configuration in BladeCenter S 219
 - Zoning 220
- Disk consideration in BladeCenter S 217
- disk storage types 61
- DS4800
 - live partition mobility 66
- DST 248
- DVD Installation
 - AIX V6.1 on iSCSI target 189
- Dynamic Host Configuration Protocol (DHCP) 450
- Dynamic Logical Partitioning (DLPAR) 304, 328, 466
- Dynamic LPAR (DLPAR) 129
 - capabilities, retrieve 130
 - ethernet 133

- memory 131
 - optical devices 135
 - physical adapters 137
 - processing 132
 - storage 135
- E**
- eConfigurator 4
 - Emulex 44
 - EnergyScale 37, 351
 - Energyscale technology 36
 - entstat 56
 - Ethernet bridge 162
 - Ethernet I/O module 216
 - ETSI 24
 - Expansion cards 44, 57
 - CFFh
 - Qlogic combo card 60
 - CFFv
 - QLogic 4 Gbs fibre channel HBA 57
 - HSFF
 - Cisco 4x Infiniband 59
 - OS and adapter suport 45
 - partnumber conversion table 46
 - StFF
 - Qlogic iSCSI HBA 62
 - Expansion Cards form factor 44
 - external storage 62
- F**
- Fibre Channel 44
 - fibre channel storage 64
 - File Transfer Protocol (FTP)
 - installation server 450
 - Firmware 391
 - firmware management 50
 - firmware prompt
 - accessing 408
 - firmware updates download for JS12/22 330
 - FPSCR 33
 - fuelg 361, 363
- G**
- Gigabit Ethernet 12
- H**
- hdisk reserve policy 335
 - HEA 43
 - help 73
 - High Speed Form Factor 57
 - Host Ethernet Adapter 43
 - Host Ethernet Adapter (HEA) 56, 86, 339
 - configuring as a SEA 89
 - host group 68
 - hotplug 466
 - HSFF 57
 - Hypertext Transfer Protocol (HTTP) 450
- I**
- I/O hotplug 304, 328
 - IBM BladeCenter chassis 11
 - BladeCenter E 20
 - BladeCenter H 12
 - BladeCenter HT 27
 - BladeCenter S 16, 62
 - BladeCenter T 24
 - IBM Director 49–50, 364, 373
 - Creating a scheduled job 379
 - Scheduler 386
 - IBM i 49
 - IBM i 6.1 49
 - IBM i V6R1 Backup/Restore
 - IBM SAS LTO 271
 - IBM i V6R1 install media preparation 249
 - IBM i V6R1 installation 209
 - BladeCenter Hardware preparation 214
 - Disk configuration in BladeCenter S 219
 - Disk consideration in BladeCenter S 217
 - Hardware environments 212
 - Preparation 210
 - Software installation process 210
 - Storage consideration BladeCenter H 217
 - IBM i V6R1 installing PTFs 269
 - IBM ID 270
 - IBM i V6R1 IPL types 250
 - IBM i V6R1 Navigator for i 270
 - IBM i V6R1 Shutdown and restart 279
 - Shutdown IBM i V6R1 partition 279
 - Start an IBM i V6R1 partition 280
 - IBM i V6R1 TCP/IP setup 270
 - IBM i, storage considerations 65, 480
 - IBM ID 270
 - IBM SAS LTO 272
 - IBM ServerProven 64, 479
 - IBM service and productivity tools 327

- IBM service and productivity tools are packages 304
- IBM System Access for Windows V6R1 225
 - Installation System i Access for Windows 226
 - Preparation 225
- IBM Tivoli Storage Manager
 - TSM client check 278
- IBM Total Storage 64
- InfiniBand 12
- Infiniband 44
- Initial Program Load 200
- Install IBM i V6R1 operating system 261
- install_assist 177, 200
- installation
 - IBM i 6.1 operating system 261
 - iSCSI target 177
- installation assistant 177
- installation of Red Hat Enterprise Linux 5.2 over the network 289
- installation of Red Hat Enterprise Linux 5.2 using a DVD 285
- installation of SLES 10 SP2 over the network 311
- installation of SLES 10 SP2 using a DVD 307
- installation summary
 - AIX 174
- Installation System i Access for Windows 226
- Installing License internal code (LIC) 250
- integrated virtual ethernet 43
- Integrated Virtual Ethernet (IVE) 43
- Integrated Virtualization Manager 87
- Integrated Virtualization Manager (IVM) 311
 - command line interface 72
 - Create Partition button 118
 - Guided Setup view 81
 - Partition Properties 83
 - General tab 83
 - Memory tab 83
 - Processing tab 83
 - user interface 71
 - navigation area 71
 - work area 71
- View/Modify Host Ethernet Adapters 86
 - Properties 87
 - Connected Partitions tab 88
- View/Modify Partitions link 83
- View/Modify Physical Adapters link 94
- View/Modify Virtual Ethernet, SEA configuration 90
- View/Modify Virtual Storage 99
 - Physical Volumes tab 99
 - Storage Pools tab 101
- VIOS management 70
- Intelligent Copper Pass-Thru Module 56
- Internal SAS 61
- interoperability matrix 65
- interposer card 59
- IPL 200
- IQN (iSCSI qualified name) 189
- iSCSI 44, 62, 177
 - iSCSI boot 63
 - iSCSI HBA 178
 - iSCSI LUN boot 200
 - iSCSI software targets 178
 - iSCSI target
 - AIX Installation 177
- IVE 43
- IVM 46, 55

K

- Kickstart Configurator 420
- KVM 57, 394

L

- Linux 48
- live partition migration 304, 328, 466
- Live Partition Mobility 47, 64, 329
 - migrate using CLI 349
 - migrate using IVM 345
 - status using CLI 349
 - status using IVM 346
 - validate using CLI 348
 - validate using IVM 340
- logical drives 68
- logical partition (LPAR) 466
 - non-dynamic operations 70
- logical partitioning 46
- LPAR activation
 - from the CLI 142
 - from the UI 141
- LPAR creation
 - Optical Devices 125
 - Partition Ethernet 122
 - Partition Memory 120
 - Partition Name 119
 - Partition Processors 121
 - Partition Storage Types and Assignments 123
 - Partition Summary 127

- Physical Adapters 126
- LPAR shutdown 144
 - from the CLI 147
 - from the UI 144
- lpcfgop 74
- lpcfgop command 74
- lsparmigr 349
- lsrefcode 143
- lssyscfg 143, 340, 349
- LUN 188
 - creating on an iSCSI target 188

M

- man dhcpd command 451
- mask 192
- media 57
- media library 69
 - adding new media 112
 - create 110
 - create blank media 113
 - delete 111
 - extend 111
 - modify media assignment 115
- media tray 57
- Mellanox 44
- memory
 - placement rules 39
- memory buffer 38
- memory controller 8, 10
- memory region size 334
- Memory subsystem 38
 - JS12 memory subsystem 40
 - JS22 memory subsystem 39
- Micro-Partitioning 46
- migrpar 348
- monitoring
 - automated 50
 - resources 50
- MSIM 59
- Multi Switch Interconnect Module (MSIM) 59, 67
- Multi-Switch Interconnect Module 59

N

- native Red Hat Enterprise Linux 5.2 installation 300
- native SLES10 SP2 installation 325
- Navigator for i 270
- NEBS 24
- Network considerations

- AMM 216
- Ethernet I/O module 216
- IBM i V6R1 installation 216
- PC for LAN Console 217
- V6R1 LAN console 217
- V6R1 production interface 217
- VIOS/IVM 217
- Network File System (NFS) 450, 455
- network file system (NFS) 289, 311, 450
- network installation 51
 - configuration file 451
 - installation source 450
 - Linux 450
 - Trivial File Transfer Protocol (TFTP) 453
- NIM server 57
- Novell® SuSE Linux Enterprise Server 47

O

- Open Firmware 391
 - boot parameters 303, 327
 - interface 407
- Operating system support 47
 - AIX 47
 - IBM i 49
 - Linux 48
- Optical devices 108, 151
 - physical 108
- Optical Pass Thru Module 58
- OS and adapter support 45

P

- partition
 - Creating IBM i 6.1 246
 - preparation 156
- Partition Management 159
- Partition Mobility 56
- partnumber conversion table 46
- PCI express 42
- PCI-e 42
- PCI-x 42
- physical optical
 - remote media 108
- physical optical device 153
- physical volumes 99
 - assigning 164
- policy 377
- POST 405
- Power Cap 378

- Power cap 37
- Power Capping 36
- Power Domain 354
- Power Domain Planning 354
- Power On Self Test 405
- Power Saver Mode 36, 369–370, 372–373
- Power Trending 36
- POWER6 processor 32–33
 - Altivec (SIMD) 35
 - Decimal floating point 33
 - Energyscale technology 36
 - Simultaneous Multi Threading 34
- PowerVM 46, 332
 - PowerVM editions for BladeCenter JS22 and JS12 47
- PowerVM editions for BladeCenter JS22 and JS12 47
- PowerVM key 332
- printenv 409
 - printenv boot-device 409
- Processor Core Nap 37
- Processor instruction retry 32
- PTF 49
- PWRDWN SYS 279

Q

- Qlogic 44

R

- Red Hat Enterprise Linux 47–48
 - configuration tool 420
 - Kickstart 428
- Red Hat Enterprise Linux 5.2 automated installation 300
- Red Hat Enterprise Linux for POWER Version 4.6 284
- Red Hat Enterprise Linux for POWER Version 5.1 284
- Red Hat Enterprise Linux Kickstart 419
- Redbooks Web site 484
 - Contact us xiii
- Reference Code 280
- reliability, availability, and serviceability (RAS) 304, 327, 466
- remote command execution 51
- remote control 394
- reset-all 413
- resource balancing 35

- resource monitoring 50
- Resource Monitoring and Control (RMC) 339
- rootvg 69
- rpm -Uvh command 470, 474

S

- SAS 3, 43
- SAS Modules 217
- SAS/SATA 61
- SCM 224
- Serial Attached SCSI (SAS) 43
- Serial over LAN (SoL) 57, 72, 397
- ServerProven 64
- Service and Productivity Tools for Linux 465
- service xinetd restart command 454
- set-connection-mode 412
- Shared Dedicated Capacity 47
- Shared Ethernet Adapter (SEA) 89, 339
- show-devs 409
- Shutdown IBM i V6R1 partition 279
- Simultaneous Multi Threading 34
- single processor checkstopping 32
- smitty
 - smitty install_latest 178
 - smitty tmiscsi 183
- SMS 391
- SMT 34
- SoL 57, 397
- SSH 393
- SSIC 66
- SST 248
- Standard Form Factor 57
- Start an IBM i V6R1 partition 280
 - Reference Code 280
- StFF 57
- storage area network (SAN) 289, 311
- Storage Configuration Manager 224
- Storage Configuration Manager (SCM) 481
- Storage consideration BladeCenter H 217
- storage considerations 64
- Storage Management 150
- storage partitions 68
- Storage Pool 150
- Storage Pools 101
 - delete 102
 - reduce 102
- suma 48
- support matrixes 64

- SUSE Linux Enterprise Server 48
- SuSE Linux Enterprise Server 47
- SUSE Linux Enterprise Server (SLES) 305
- SUSE Linux Enterprise Server 10 SP2
 - AutoYaST 447
- system access for Windows 225
- System Management Service (SMS)
 - options 399
- System Service Tool 248
- System specificatio
 - BladeCenter JS12 memory support 10
- System specification 4
 - BladeCenter JS12 5
 - BladeCenter JS22 5
 - BladeCenter JS22 & JS12 Express physical packaging 6
 - BladeCenter JS22 memory support 8
- System Storage Interoperation Center 66
- system-config-kickstart command 420
- Systems management 49
 - Advanced Management Web interface 49
 - Cluster Systems Management 50
 - IBM Director 50

T

- Target Protocol Device 183
- telnet 393
- telnetcfg 402
- Thermal Power Management Device
 - POWER6 processor
 - Thermal Power Management Device 37
- thread priority 35
- TPMD 37
- Trivial File Transfer Protocol (TFTP) 450
 - configuration 453
 - network installation 450
- TSM client check 278

U

- Universal Serial Bus (USB) subsystem 42
- USB 42

V

- VIOS 47, 55
 - planning 56
- VIOS/IVM 217
- Virtual CDROM 166

- virtual disk
 - create 104
- virtual disks 103
 - assigning 164
 - delete 107
 - extend 105
- Virtual Ethernet
 - Integrated Virtualization Manager (IVM) 71
- Virtual Ethernet Adapters 89
- Virtual I/O Server (VIOS)
 - data sheet 64, 479
 - default user 73, 81
 - fixpack download location 330
- Virtual Input/Output Server 55
- virtual media library 150
- virtual optical device on the Integrated Virtual Manager (IVM) to perform a CD/DVD installation of a Linux operating system image 458
- virtual optical devices 338
- Virtual Storage Management 98
- Virtual Terminal 138
- virtualization 47
- VMLibrary 154

W

- Windows 225
- WRKOPTVOL 277

Y

- yast command 453
- yast2 command 431

Z

- Zone Group 220
- Zoning 220

Archived



Redbooks

IBM BladeCenter JS12 and JS22 Implementation Guide

(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages

IBM BladeCenter JS12 and JS22 Implementation Guide



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

Supporting PowerVM including Live Partition Mobility on POWER6 processor-based blades

Achieving greater energy efficiency through consolidation of server workloads

This IBM Redbooks publication provides a detailed technical guide for configuring and using the IBM BladeCenter JS12 and IBM BladeCenter JS22 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 will be useful for data center system architects; network, storage and facilities engineers; system administrators; and application architects that need a detailed system bring-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 6.1 (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

SG24-7655-00

ISBN 0738431907