IBM

IBM Flex System p260 and p460 Planning and Implementation Guide

- Describes the new POWER7 compute nodes for IBM PureFlex System
- Provides detailed product and planning information
- Set up partitioning and OS installation

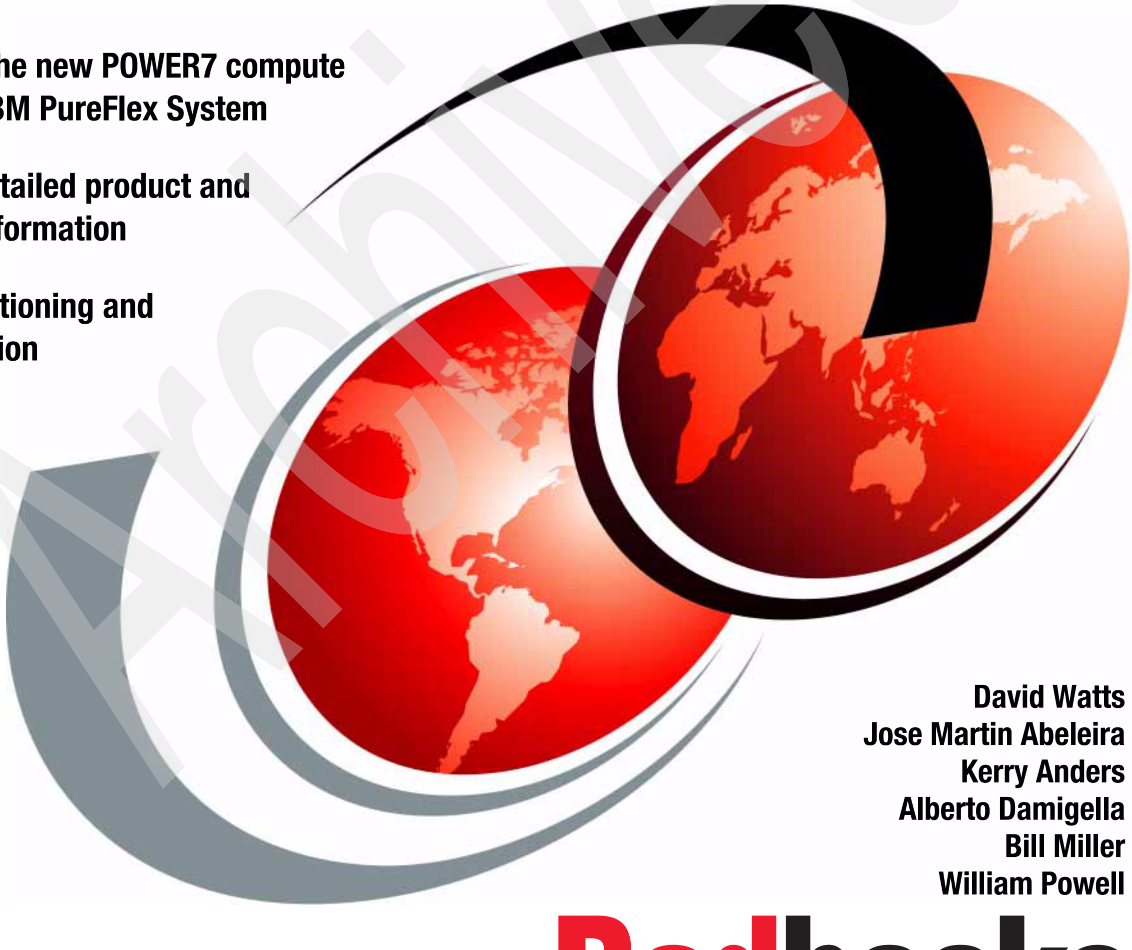

David Watts
Jose Martin Abeleira
Kerry Anders
Alberto Damigella
Bill Miller
William Powell

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM Flex System p260 and p460 Planning and Implementation Guide

June 2012

SG24-7989-00

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

First Edition (June 2012)

This edition applies to:

IBM PureFlex System
IBM Flex System Enterprise Chassis
IBM Flex System Manager
IBM Flex System p260 Compute Node
IBM Flex System p24L Compute Node
IBM Flex System p460 Compute Node

Contents

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

Active Memory™
AIX 5L™
AIX®
BladeCenter®
BNT®
Electronic Service Agent™
EnergyScale™
HACMP™
IBM Flex System™
IBM SmartCloud™
IBM Systems Director Active Energy Manager™
IBM®
iDataPlex®
Language Environment®
Micro-Partitioning®
POWER Hypervisor™
Power Systems™
Power Systems Software™
POWER6+™
POWER6®
POWER7®
PowerHA®
PowerVM®
POWER®
PureApplication™
PureFlex™
PureSystems™
Redbooks®
Redbooks (logo) ®
ServerProven®
Storwize®
System Storage®
System x®
System z®
Tivoli®
VMready®
Workload Partitions Manager™

The following terms are trademarks of other companies:

Intel Xeon, 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 and other countries.

Linux is a trademark of Linus Torvalds 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.

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

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

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

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

Preface

To meet today's complex and ever-changing business demands, you need a solid foundation of compute, storage, networking, and software resources that is simple to deploy and can quickly and automatically adapt to changing conditions. You also need to be able to take advantage of broad expertise and proven preferred practices in systems management, applications, hardware maintenance, and more.

IBM® PureFlex™ System combines no-compromise system designs with built-in expertise and integrates them in complete, optimized solutions. At the heart of PureFlex System is the IBM Flex System™ Enterprise Chassis, a fully integrated infrastructure platform that supports a mix of compute, storage, and networking resources to meet the demands of your applications.

The IBM Flex System p260 and p460 Compute Nodes are IBM Power Systems™ servers optimized for virtualization, performance, and efficiency. The nodes support IBM AIX®, IBM i, or Linux operating environments, and are designed to run various workloads in IBM PureFlex System.

This IBM Redbooks® publication is a comprehensive guide to IBM PureFlex System and the Power Systems compute nodes. We introduce the offerings and describe the compute nodes in detail. We then describe planning and implementation steps and go through some of the key the management features of the IBM Flex System Manager management node.

This book is for customers, IBM Business Partners, and IBM technical specialists that want to understand the new offerings and to plan and implement an IBM Flex System installation that involves the Power Systems compute nodes.

The team who wrote this book

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

David Watts is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks publications on hardware and software topics related to IBM Flex System, IBM System x®, and IBM BladeCenter® servers and associated client platforms. He has authored over 200 books, papers, and Product Guides. He holds a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM both in the US and Australia since 1989. David is an IBM Certified IT Specialist and a member of the IT Specialist Certification Review Board.

Jose Martin Abeleira is a Certified Consulting IT Specialist, and IBM Certified Systems Expert Enterprise Technical Support for AIX and Linux in Montevideo, Uruguay. He has been with IBM for 6 years and has 10 years of AIX experience. During this time, he has worked for Global Delivery as a Senior AIX/UNIX Administrator, Uruguay Service Delivery Center as a Senior AIX/UNIX Administrator, and a UNIX/Storage Team Leader.

Kerry Anders is a Consultant for IBM POWER® systems and IBM PowerVM® in Lab Services for the IBM Systems and Technology Group, based in Austin, Texas. He is part of the Lab Service core team that implements IBM PureFlex System and supports clients in implementing IBM Power Systems blades using Virtual I/O Server, Integrated Virtualization Manager, and AIX. Previously, he was the Systems Integration Test Team Lead for the IBM BladeCenter JS21 blade with IBM SAN storage using AIX and Linux. His prior work includes test experience with the JS20 blade, and 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. Kerry has authored four other IBM Redbooks publications, the most recent being *IBM BladeCenter PS703 and PS704 Technical Overview and Introduction*, REDP-4744.

Alberto Damigella is an IT Specialist working in the IBM Power Systems and IBM Flex System pre-sales team for IBM Italy as a Customer Technical Support (CTS) specialist. He regularly performs Power Systems presentations and demonstrations, and he helps the sales team write technical descriptions and system consolidation studies. He is based in Milan, Italy. Alberto started at IBM as a Computer Aided Three-dimensional Interactive Application (CATIA) pre-sales and post-sales specialist in 1985.

Bill Miller is an IT Specialist in Lab Services Technical Training. He has been with IBM since 1983. He has had an array of responsibilities, starting in development, and then moving to roles as a Systems Engineer and IBM Global Services consultant that focuses on AIX, IBM Tivoli® Storage Manager, and IBM HACMP™ (IBM PowerHA®) planning and implementation. He is currently responsible for course development, maintenance, and delivery for the PowerHA and Flex System curriculums.

Will Powell has been a specialist in hardware and warranty support for System x, BladeCenter, POWER blades, and IBM iDataPlex® since 2004 at the IBM Technical Support Center in Atlanta, Georgia. He has particular expertise and experience with integrated networking, storage, Fibre Channel, InfiniBand, clustering, RAID, and high-availability computing. He is a corporate member of the Technology Association of Georgia (TAG). He has provided technical consulting to Rivers of the World since 2000. Will holds a Bachelor of Science in Computer Science degree from North Georgia College & State University.

Figure 1 The team (l-r) - David, Martin, Kerry, Will, Alberto, Bill

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

From IBM marketing:

- TJ Aspden
- Michael Bacon
- John Biebelhausen
- Bruce Corregan
- Mary Beth Daughtry

- Mike Easterly
- Diana Cunniffe
- Kyle Hampton
- Botond Kiss
- Shekhar Mishra
- Justin Nguyen
- Sander Kim
- Dean Parker
- Hector Sanchez
- David Tareen
- David Walker
- Randi Wood
- Bob Zuber

From IBM Power Systems development:

- Chris Austen
- Kaena Freitas
- Jim Gallagher
- Ned Gamble
- Bill Johnson
- Rick McBride
- Lenny Nichols
- Amartey Pearson
- Dean Price
- Mike Stys
- Richard Vasconi

Others from IBM around the world

- Bill Champion
- Michael L. Nelson
- Mike Schambureck
- Matt Slavin
- Rochelle White

From the International Technical Support Organization:

- Kevin Barnes
- Tamikia Barrow
- Mary Comianos
- Shari Deiana
- Cheryl Gera
- Ilya Krutov
- Karen Lawrence
- Julie O'Shea
- Linda Robinson

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

IBM PureSystems

During the last 100 years, information technology moved from a specialized tool to a pervasive influence on nearly every aspect of life. From tabulating machines that counted with mechanical switches or vacuum tubes to the first programmable computers, IBM has been a part of this growth, while always helping customers solve problems.

Information technology (IT) is a constant part of business and of our lives. IBM expertise in delivering IT solutions has helped the planet become smarter. As organizational leaders seek to extract more real value from their data, business processes, and other key investments, IT is moving to the strategic center of business.

To meet those business demands, IBM introduces a new category of systems: Systems that combine the flexibility of general-purpose systems, the elasticity of cloud computing, and the simplicity of an appliance that is tuned to the workload. Expert integrated systems are the building blocks of this capability. This new category of systems represents the collective knowledge of thousands of deployments, established preferred practices, innovative thinking, IT leadership, and distilled expertise.

The offerings in IBM PureSystems™ are designed to deliver value in the following ways:

- Built-in expertise helps you address complex business and operational tasks automatically.
- Integration by design helps you tune systems for optimal performance and efficiency.
- Simplified experience, from design to purchase to maintenance, creates efficiencies quickly.

The IBM PureSystems offerings are optimized for performance and virtualized for efficiency. These systems offer a no-compromise design with system-level upgradeability. IBM PureSystems is built for cloud computing, containing “built-in” flexibility and simplicity.

At IBM, expert integrated systems come in two types:

- IBM PureFlex System: Infrastructure systems deeply integrate the IT elements and expertise of your system infrastructure.
- IBM PureApplication™ System: Platform systems include middleware and expertise for deploying and managing your application platforms

1.1 IBM PureFlex System

To meet today’s complex and ever-changing business demands, you need a solid foundation of server, storage, networking, and software resources that is simple to deploy and can quickly and automatically adapt to changing conditions. You also need access to, and the ability to take advantage of, broad expertise and proven preferred practices in systems management, applications, hardware maintenance, and more.

IBM PureFlex System is a comprehensive infrastructure system that provides an expert integrated computing system, combining servers, enterprise storage, networking, virtualization, and management into a single structure. Its built-in expertise enables organizations to manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management. These systems are ideally suited for customers interested in a system that delivers the simplicity of an integrated solution, but that also want control over tuning middleware and the runtime environment.

IBM PureFlex System recommends workload placement based on virtual machine compatibility and resource availability. Using built-in virtualization across servers, storage, and networking, the infrastructure system enables automated scaling of resources and true workload mobility.

IBM PureFlex System undergoes significant testing and experimentation, so it can mitigate IT complexity without compromising the flexibility to tune systems to the tasks businesses demand. By providing both flexibility and simplicity, IBM PureFlex System can provide extraordinary levels of IT control, efficiency, and operating agility that enable businesses to rapidly deploy IT services at a reduced cost. Moreover, the system is built on decades of expertise, enabling deep integration and central management of a comprehensive, open-choice infrastructure system, and dramatically cutting down on the skills and training required for managing and deploying the system.

IBM PureFlex System combines advanced IBM hardware and software along with patterns of expertise and integrates them into three optimized configurations that are simple to acquire and deploy so you get fast time to value for your solution.

The three PureFlex System configurations are:

- IBM PureFlex System Express: Designed for small and medium businesses, it is the most affordable entry point for PureFlex System.
- IBM PureFlex System Standard: Optimized for application servers with supporting storage and networking, it is designed to support your key ISV solutions.
- IBM PureFlex System Enterprise: Optimized for transactional and database systems and has built-in redundancy for highly reliable and resilient operation to support your most critical workloads.

These three configurations are summarized in Table 1-1.

Table 1-1 IBM PureFlex System

Component	IBM PureFlex System Express	IBM PureFlex System Standard	IBM PureFlex System Enterprise
IBM PureFlex System 42U Rack	1	1	1
IBM Flex System Enterprise Chassis	1	1	1
IBM Flex System Fabric EN4093 10Gb Scalable Switch	1	1	Two with both port-count upgrades

Component	IBM PureFlex System Express	IBM PureFlex System Standard	IBM PureFlex System Enterprise
IBM Flex System FC3171 8Gb SAN Switch	1	2	2
IBM Flex System Manager Node	1	1	1
IBM Flex System Manager software license	IBM Flex System Manager with 1-year service and support	IBM Flex System Manager Advanced with 3-year service and support	Flex System Manager Advanced with 3-year service and support
Chassis Management Module	2	2	2
Chassis power supplies (std/max)	2 / 6	4 / 6	6 / 6
Chassis 80 mm fan modules (std/max)	4 / 8	6 / 8	8 / 8
IBM Storwize® V7000 Disk System	Yes (redundant controller)	Yes (redundant controller)	Yes (redundant controller)
IBM Storwize V7000 Software	Base with 1-year software maintenance agreement	Base with 3-year software maintenance agreement	Base with 3-year software maintenance agreement

The fundamental building blocks of IBM PureFlex System solutions are the IBM Flex System Enterprise Chassis, complete with compute nodes, networking, and storage.

For more details about IBM PureFlex System, see Chapter 2, “IBM PureFlex System” on page 15.

1.2 IBM PureApplication System

IBM PureApplication System is a platform system that pre-integrates a full application platform set of middleware and expertise with the IBM PureFlex System in a single management console. It is a workload-aware, flexible platform that is easy to deploy, customize, safeguard, and manage in a traditional or private cloud environment, ultimately providing superior IT economics.

With the IBM PureApplication System, you can provision your own patterns of software, middleware, and virtual system resources. You can provision these patterns within a unique framework that is shaped by IT preferred practices and industry standards that are culled from many years of IBM experience with clients and from a deep understanding of smarter computing. These IT preferred practices and standards are infused throughout the system.

With IBM PureApplication System:

- IBM builds expertise into preintegrated deployment patterns, which can speed the development and delivery of new services.
- By automating key processes, such as application deployment, PureApplication System built-in expertise capabilities can reduce the cost and time required to manage an infrastructure.
- Built-in application optimization expertise reduces the number of unplanned outages through preferred practices and automation of the manual processes identified as sources of those outages.
- Administrators can use built-in application elasticity to scale up or to scale down automatically. Systems can use data replication to increase availability.

You can use patterns of expertise to automatically balance, manage, and optimize the necessary elements, from the underlying hardware resources up through the middleware and software. You can use these patterns of expertise to help deliver and manage business processes, services, and applications by encapsulating preferred practices and expertise into a repeatable and deployable form. This preferred practice knowledge and expertise was gained from decades of optimizing the deployment and management of data centers, software infrastructures, and applications around the world.

You can use these patterns to achieve the following types of value:

- Agility. As you seek to innovate to bring products and services to market faster, you need fast time-to-value. You can use expertise built into a solution to eliminate manual steps, automate delivery, and support innovation.
- Efficiency. To reduce costs and conserve valuable resources, you must get the most out of your systems with energy efficiency, simple management, and fast, automated response to problems. With built-in expertise, you can optimize your critical business applications and get the most out of your investments.

- Increased simplicity. You need a less complex environment. You can use patterns of expertise to help you easily consolidate diverse servers, storage, and applications onto an easier-to-manage, integrated system.
- Control. With optimized patterns of expertise, you can accelerate cloud implementations to lower risk by improving security and reducing human error.

IBM PureApplication System is available in four configurations. You can use these configuration options to choose the size and compute power that meets your needs for application infrastructure. You can upgrade to the next size when your organization requires more capacity, and in most cases, you can do so without taking an application downtime.

Table 1-2 provides a high-level overview of the configurations.

Table 1-2 IBM PureApplication System configurations

	IBM PureApplication System W1500-96	IBM PureApplication System W1500-192	IBM PureApplication System W1500-384	IBM PureApplication System W1500-608
Cores	96 cores	192 cores	384 cores	608 cores
Memory	1.5 TB	3.1 TB	6.1 TB	9.7 TB
SSD Storage	6.4 TB			
HDD Storage	48.0 TB			
Application Services Entitlement	Included			

IBM PureApplication System is outside the scope of this book. For more details about it, see the following website:

http://ibm.com/expert

1.3 IBM Flex System: The building blocks for IBM PureSystems

IBM PureSystems is built from no-comprise building blocks based on reliable IBM technology that support open standards and offer confident roadmaps: IBM Flex System. IBM Flex System is designed for multiple generations of technology, supporting your workload today, and ready for the future demands of your business.

1.3.1 Management

IBM Flex System Manager is designed to optimize the physical and virtual resources of the IBM Flex System infrastructure while simplifying and automating repetitive tasks. From easy system set-up procedures with wizards and built-in expertise, to consolidated monitoring for all of your resources (compute, storage, networking, virtualization, and energy), IBM Flex System Manager provides core management functionality along with automation. It is an ideal solution that you can use to reduce administrative expense and focus your efforts on business innovation.

From a single user interface, you get:

- Intelligent automation
- Resource pooling
- Improved resource utilization
- Complete management integration
- Simplified setup

1.3.2 Compute nodes

Taking advantage of the full capabilities of IBM POWER7® processors or Intel Xeon processors, the compute nodes are designed to offer the performance you need for your critical applications.

With support for a range of hypervisors, operating systems, and virtualization environments, the compute nodes provide the foundation for:

- Virtualization solutions
- Database applications
- Infrastructure support
- Line of business applications

1.3.3 Storage

You can use the storage capabilities of IBM Flex System to gain advanced functionality with storage nodes in your system while taking advantage of your existing storage infrastructure through advanced virtualization.

IBM Flex System simplifies storage administration by using a single user interface for all your storage through a management console that is integrated with the comprehensive management system. You can use these management and storage capabilities to virtualize third-party storage with nondisruptive migration of the current storage infrastructure. You can also take advantage of intelligent tiering so you can balance performance and cost for your storage needs. The solution also supports local and remote replication and snapshots for flexible business continuity and disaster recovery capabilities.

1.3.4 Networking

With a range of available adapters and switches to support key network protocols, you can configure IBM Flex System to fit in your infrastructure while still being ready for the future. The networking resources in IBM Flex System are standards-based, flexible, and fully integrated into the system, so you get no-compromise networking for your solution. Network resources are virtualized and managed by workload. These capabilities are automated and optimized to make your network more reliable and simpler to manage.

Key capabilities include:

- Supports the networking infrastructure you have today, including Ethernet, Fibre Channel, and InfiniBand
- Offers industry-leading performance with 1 Gb, 10 Gb, and 40 Gb Ethernet, and 8 Gb and 16 Gb Fibre Channel and FDR InfiniBand
- Provides pay-as-you-grow scalability so you can add ports and bandwidth when needed

1.3.5 Infrastructure

The IBM Flex System Enterprise Chassis is the foundation of the offering, supporting intelligent workload deployment and management for maximum business agility. The 14-node, 10 U chassis delivers high-performance connectivity for your integrated compute, storage, networking, and management resources. The chassis is designed to support multiple generations of technology and offers independently scalable resource pools for higher utilization and lower cost per workload.

1.4 IBM Flex System overview

The expert integrated system of IBM PureSystems is based on a new hardware and software platform called IBM Flex System.

1.4.1 IBM Flex System Manager

The IBM Flex System Manager (FSM) is a high performance scalable systems management appliance with a preloaded software stack. As an appliance, the hardware is closed (on a dedicated compute node platform) and designed to provide a specific purpose: configure, monitor, and manage IBM Flex System resources in multiple IBM Flex System Enterprise Chassis (Enterprise Chassis), optimizing time-to-value.

The FSM provides a world-class user experience with a truly “single pane of glass” approach for all chassis components. Featuring an instant resource-oriented view of the Enterprise Chassis and its components, the FSM provides vital information for real-time monitoring.

An increased focus on optimizing time-to-value is evident in such features as:

- Setup wizards, including initial setup wizards, which provide intuitive and quick setup of the FSM.
- A Chassis Map, which provides multiple view overlays to track health, firmware inventory, and environmental metrics.
- Configuration management for a repeatable setup of compute, network, and storage devices.
- Remote presence application for remote access to compute nodes with single sign-on.
- Quick search provides results as you type.

Beyond the physical world of inventory, configuration, and monitoring, IBM Flex System Manager enables virtualization and workload optimization for a new class of computing:

- Resource utilization: Within the network fabric, FSM detects congestions, notification policies, and relocation of physical and virtual machines, including storage and network configurations.
- Resource pooling: FSM pools network switching, with placement advisors that consider VM compatibility, processor, availability, and energy.
- Intelligent automation: FSM has automated and dynamic VM placement based on utilization, energy, hardware predictive failure alerts, or host failures.

Figure 1-1 shows the IBM Flex System Manager.

Figure 1-1 IBM Flex System Manager

1.4.2 IBM Flex System Enterprise Chassis

The IBM Flex System Enterprise Chassis (Enterprise Chassis) offers compute, networking, and storage capabilities far exceeding products that are currently available in the market. With the ability to handle up 14 compute nodes, intermixing POWER7 and Intel x86, the Enterprise Chassis provides flexibility and tremendous compute capacity in a 10 U package. Additionally, the rear of the chassis accommodates four high speed networking switches. Interconnecting compute, networking, and storage through a high performance and scalable mid-plane, Enterprise Chassis can support 40 Gb speeds.

The ground-up design of the Enterprise Chassis reaches new levels of energy efficiency through innovations in power, cooling, and air flow. Smarter controls and futuristic designs allow the Enterprise Chassis to break free of “one size fits all” energy schemes.

The ability to support the workload demands of tomorrow’s workloads is built in to a new I/O architecture, providing choice and flexibility in fabric and speed. With the ability to use Ethernet, InfiniBand, FC, FCoE, and iSCSI, the Enterprise Chassis is uniquely positioned to meet the growing I/O needs of the IT industry.

Figure 1-2 shows the IBM Flex System Enterprise Chassis.

Figure 1-2 The IBM Flex System Enterprise Chassis

1.4.3 Compute nodes

IBM Flex System offers compute nodes that vary in architecture, dimension, and capabilities. The new, no-compromise nodes feature leadership designs for current and future workloads. Optimized for efficiency, density, performance, reliability, and security, the portfolio includes the following IBM POWER7 based and Intel Xeon based nodes:

- IBM Flex System x240 Compute Node, a two socket Intel Xeon based compute node
- IBM Flex System p260 Compute Node, a two socket IBM POWER7 based compute node
- IBM Flex System p24L Compute Node, a two socket IBM POWER7 based compute node optimized for Linux installations
- IBM Flex System p460 Compute Node, a four socket IBM POWER7 based compute node

Figure 1-3 shows the IBM Flex System p460 Compute Node.

Figure 1-3 IBM Flex System p460 Compute Node

The nodes have complementary leadership I/O capabilities of up to 16 x 10 Gb lanes per node.

Here are the I/O adapters offered:

- IBM Flex System EN2024 4-port 1Gb Ethernet Adapter
- IBM Flex System EN4054 4-port 10Gb Ethernet Adapter
- IBM Flex System EN4132 2-port 10Gb Ethernet Adapter
- IBM Flex System CN4054 10Gb Virtual Fabric Adapter
- IBM Flex System FC3052 2-port 8Gb FC Adapter
- IBM Flex System FC3172 2-port 8Gb FC Adapter
- IBM Flex System FC5022 2-port 16Gb FC Adapter
- IBM Flex System IB6132 2-port FDR InfiniBand Adapter
- IBM Flex System IB6132 2-port QDR InfiniBand Adapter

You have the opportunity to gain future expansion capabilities to existing and new compute nodes.

1.4.4 I/O modules

Networking in data centers is undergoing a transition from a discrete, traditional model to a more flexible, optimized model. The network architecture in IBM Flex System is designed to address the key challenges customers are facing today in their data centers. The key focus areas of the network architecture on this platform are unified network management, optimized and automated network virtualization, and a simplified network infrastructure.

Providing innovation, leadership, and choice in the I/O module portfolio uniquely positions IBM Flex System to provide meaningful solutions to address customer needs.

Here are the I/O Modules offered with IBM Flex System:

- IBM Flex System Fabric EN4093 10Gb Scalable Switch
- IBM Flex System EN2092 1Gb Ethernet Scalable Switch
- IBM Flex System EN4091 10Gb Ethernet Pass-thru
- IBM Flex System FC3171 8Gb SAN Switch
- IBM Flex System FC3171 8Gb SAN Pass-thru
- IBM Flex System FC5022 16Gb SAN Scalable Switch
- IBM Flex System FC5022 24-port 16Gb ESB SAN Scalable Switch
- IBM Flex System IB6131 InfiniBand Switch
- IBM Flex System IB6132 2-port QDR InfiniBand Adapter

Figure 1-4 shows the IBM Flex System Fabric EN4093 10Gb Scalable Switch.

Figure 1-4 IBM Flex System Fabric EN4093 10Gb Scalable Switch

1.5 This book

This book describes the IBM Flex System components in detail. It describes the technology and features of the chassis, compute nodes, management features, and connectivity and storage options. We start with a description of the systems management features of the product portfolio.

This book is a comprehensive guide to IBM PureFlex System and the Power Systems compute nodes. We introduce the offerings and describe the compute nodes in detail. We then describe the management features of IBM PureFlex System and describe partitioning and installing an operating system.

2

IBM PureFlex System

IBM PureFlex System provides an integrated computing system that combines servers, enterprise storage, networking, virtualization, and management into a single structure. You can use its built-in expertise to manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management.

The themes of PureFlex System are:

- Configurations that ease acquisition experience and match your needs
- Optimized to align with targeted workloads and environments
- Designed for cloud with SmartCloud Entry included on Standard and Enterprise
- Choice of architecture, operating system, and virtualization engine
- Designed for simplicity with integrated, single-system management across physical and virtual resources
- Simplified ordering that accelerates deployment into your environments
- Ships as a single integrated entity directly to you
- Includes factory integration and lab services optimization

IBM PureFlex System has three preintegrated offerings that support compute, storage, and networking requirements. You can select from these offerings, which are designed for key client initiatives and help simplify ordering and configuration. As a result, PureFlex System helps cut the cost, time, and complexity of system deployments.

The IBM PureFlex System offerings are as follows:

- Express: An infrastructure system for small-sized and midsized businesses; the most cost-effective entry point. See 2.1, “IBM PureFlex System Express” on page 17.
- Standard: An infrastructure system for application servers with supporting storage and networking. See 2.2, “IBM PureFlex System Standard” on page 26.
- Enterprise: An infrastructure system optimized for scalable cloud deployments with built-in redundancy for highly reliable and resilient operation to support critical applications and cloud services. See 2.3, “IBM PureFlex System Enterprise” on page 35.

The main components of a PureFlex System configuration are:

- A preinstalled and configured IBM Flex System Enterprise Chassis
- Compute nodes with either IBM POWER or Intel Xeon processors
- IBM Flex System Manager, preinstalled with management software and licenses for software activation
- IBM Storwize V7000 external storage unit
- All hardware components preinstalled in an IBM PureFlex System 42U rack
- Choice of:
  - Operating system: AIX, IBM i, Microsoft Windows, Red Hat Enterprise Linux, or SUSE Linux Enterprise Server
  - Virtualization software: PowerVM, KVM, VMware ESX, or Microsoft Hyper V
  - SmartCloud Entry (see 2.4, “IBM SmartCloud Entry” on page 44).
- Complete pre-integrated software and hardware
- On-site services included to get you up and running quickly

IBM Flex System configurations with Power Systems compute node: Orders for Power Systems compute node must be through one of the three IBM PureFlex System configurations. Build-to-order configurations are not available.

2.1 IBM PureFlex System Express

The tables in this section represent the hardware, software, and services that make up IBM PureFlex System Express. We describe the following items:

- Chassis
- Top-of-Rack Ethernet switch
- Top-of-Rack SAN switch
- Compute nodes
- IBM Flex System Manager
- IBM Storwize V7000
- Rack cabinet
- Software
- Services

To specify IBM PureFlex System Express in the IBM ordering system, specify the indicator feature code listed in Table 2-1 for each machine type.

Table 2-1 Express indicator feature code

AAS feature code	XCC feature code	Description
EFD1	A2VS	IBM PureFlex System Express Indicator Feature Code - First of each MTM (for example, first compute node)

2.1.1 Chassis

Table 2-2 lists the major components of the IBM Flex System Enterprise Chassis including the switches and options.

Feature codes: The tables in this section do not list all feature codes. Some features are not listed here for brevity.

Table 2-2 Components of the chassis and switches

AAS feature code	XCC feature code	Description	Minimum quantity
7893-92X	8721-HC1	IBM Flex System Enterprise Chassis	1
3593	A0TB	IBM Flex System Fabric EN4093 10Gb Scalable Switch	1
3282	5053	10GbE 850 nm Fiber SFP+ Transceiver (SR)	2
EB29	3268	IBM BNT® SFP RJ45 Transceiver	5

AAS feature code	XCC feature code	Description	Minimum quantity
3595	A0TD	IBM Flex System FC3171 8Gb SAN Switch	1
3286	5075	IBM 8Gb SFP+ Short-Wave Optical Transceiver	2
3590	A0UD	Additional PSU 2500 W	0
4558	6252	2.5 m, 16A/100-240V, C19 to IEC 320-C20 power cord	2
9039	A0TM	Base Chassis Management Module	1
3592	A0UE	Additional Chassis Management Module	1
9038	None	Base Fan Modules (four)	1
7805	A0UA	Additional Fan Modules (two)	0

2.1.2 Top-of-Rack Ethernet switch

If more than one chassis is configured, then a Top-of-Rack (TOR) Ethernet switch is added to the configuration. If only one chassis is configured, then the TOR switch is optional. Table 2-3 lists the switch components.

Table 2-3 Components of the Top-of-Rack Ethernet switch

AAS feature code	XCC feature code	Description	Minimum quantity
7309-HC3	1455-64C	IBM System Networking RackSwitch G8264	0[a]
	1455-48E	IBM System Networking RackSwitch G8052	0[a]
ECB5	A1PJ	3m IBM Passive DAC SFP+ Cable	1 per EN4093 switch
EB25	A1PJ	3m IBM QSFP+ DAC Break Out Cable	0

a. One is required when two or more Enterprise Chassis are configured.

2.1.3 Top-of-Rack SAN switch

If more than one chassis is configured, then a Top-of-Rack SAN switch is added to the configuration. If only one chassis is configured, then the SAN switch is optional. Table 2-4 lists the switch components.

Table 2-4 Components of the Top-of-Rack SAN switch

AAS feature code	XCC feature code	Description	Minimum quantity
2498-B24	2498-B24	24-port SAN Switch	0
5605	5605	5m optic cable	1
2808	2808	8 Gb SFP transceivers (8 pack)	1

2.1.4 Compute nodes

The PureFlex System Express requires either of the following compute nodes:

- IBM Flex System p260 Compute Node (IBM POWER7 based) (Table 2-5)
- IBM Flex System x240 Compute Node (Intel Xeon based) (Table 2-7 on page 21)

Table 2-5 lists the major components of the IBM Flex System p260 Compute Node.

Table 2-5 Components of IBM Flex System p260 Compute Node

AAS feature code	Description	Minimum quantity
IBM Flex System p260 Compute Node		
7895-22x	IBM Flex System p260 Compute Node	1
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter	1
1762	IBM Flex System EN4054 4-port 10Gb Ethernet Adapter	1
Base Processor 1 Required, select only one, Min 1, Max 1		
EPR1	8 Cores, (2x 4 core), 3.3 GHz + 2-socket system board	1
EPR3	16 Cores, (2x 8 core), 3.2 GHz + 2-socket system board	
EPR5	16 Cores, (2x 8 core), 3.55 GHz + 2-socket system board	

AAS feature code	Description	Minimum quantity
Memory - 8 GB per core minimum with all DIMM slots filled with same memory type		
8145	32GB (2x 16GB), 1066MHz, LP RDIMMs (1.35V)	
8199	16GB (2x 8GB), 1066MHz, VLP RDIMMs (1.35V)	

Table 2-6 lists the major components of the IBM Flex System p24L Compute Node.

Table 2-6 Components of IBM Flex System p24L Compute Node

AAS feature code	Description	Minimum quantity
IBM Flex System p24L Compute Node		
1457-7FL	IBM Flex System p24L Compute Node	1
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter	1
1762	IBM Flex System EN4054 4-port 10Gb Ethernet Adapter	1
Base Processor 1 Required, select only one, Min 1, Max 1		
EPR7	12 cores, (2x 6core), 3.7 GHz + 2-socket system board	1
EPR8	16 cores, (2x 8 core), 3.2 GHz + 2-socket system board	
EPR9	16 cores, (2x 8 core), 3.55 GHz + 2-socket system board	
Memory - 2 GB per core minimum with all DIMM slots filled with same memory type		
8145	32GB (2x 16GB), 1066MHz, LP RDIMMs (1.35V)	
8199	16GB (2x 8GB), 1066MHz, VLP RDIMMs (1.35V)	
8196	8GB(2X4GB), 1066MHz, DDR3, VLP RDIMMS(1.35V)	
EM04	4GB (2x2GB), 1066MHz, DDR3 DRAM, (RDIMM, 1Rx8)	

Table 2-7 lists the major components of the IBM Flex System x240 Compute Node.

Table 2-7 Components of IBM Flex System x240 Compute Node

AAS feature code	XCC feature code	Description	Minimum quantity
IBM Flex System x240 Compute Node			
7863-10X	8737AC1	IBM Flex System x240 Compute Node	
EN20 EN21	A1BC A1BD	x240 with embedded 10Gb Virtual Fabric x240 without embedded 10Gb Virtual Fabric (select one of these base features)	1
1764	A2N5	IBM Flex System FC3052 2-port 8Gb FC Adapter	1
1759	A1R1	IBM Flex System CN4054 10Gb Virtual Fabric Adapter (select if x240 without embedded 10Gb Virtual Fabric is selected - EN21/A1BD)	1
EBK2	49Y8119	IBM Flex System x240 USB Enablement Kit	
EBK3	41Y8300	2GB USB Hypervisor Key (VMware 5.0)	

2.1.5 IBM Flex System Manager

Table 2-8 lists the major components of the IBM Flex System Manager.

Table 2-8 Components of the IBM Flex System Manager

AAS feature code	XCC feature code	Description	Minimum quantity
7955-01M	8731AC1	IBM Flex System Manager	1
EB31	9220	Platform Bundle preload indicator	1
EM09 None	None 8941	8GB (2x 4GB) 1333 MHz RDIMMs (1.35V) 4GB (1x 4GB) 1333 MHz RDIMMs (1.35V)	4[a] 8
None	A1CW	Intel Xeon E5-2650 8C 2.0GHz 20MB 1600MHz 95W	1
1771	5420	200GB, 1.8", SATA MLC SSD	2
3767	A1AV	1TB 2.5” SATA 7.2K RPM hot-swap 6 Gbps HDD	1

a. In the AAS system, FC EM09 has pairs of DIMMs. In the XCC system, FC 8941 has single DIMMs. The DIMMs are otherwise identical.

2.1.6 IBM Storwize V7000

Table 2-9 lists the major components of the IBM Storwize V7000 storage server.

Table 2-9 Components of the IBM Storwize V7000 storage server

AAS feature code	XCC feature code	Description	Minimum quantity
2076-124	2076-124	IBM Storwize V7000 Controller	1
5305	5305	5m Fiber Optic Cable	2
3512 3514	3512 3514	200GB 2.5 INCH SSD or 400GB 2.5 INCH SSD	2[a]
0010	0010	Storwize V7000 Software Preinstall	1
6008	6008	8 GB Cache	2
9730	9730	power cord to PDU (includes 2 power cord)	1
9801	9801	Power supplies	2

a. If a Power Systems compute node is selected, then at least eight drives must be installed in the Storwize V7000. If an Intel Xeon based compute node is selected with SmartCloud Entry, then four drives must be installed in the Storwize V7000.

2.1.7 Rack cabinet

Table 2-10 lists the major components of the rack and options.

Table 2-10 Components of the rack

AAS feature code	XCC feature code	Description	Minimum quantity
7953-94X	93634AX	IBM 42U 1100mm Flex System Dynamic Rack	1
EC06	None	Gray Door	1
EC03	None	Side Cover Kit (Black)	1
EC02	None	Rear Door (Black/flat)	1
7196 7189+6492 7189+6491 7189+6489 7189+6667 7189+6653	5897 5902 5904 5903 5906 None	Combo PDU C19/C13 3 Phase 60A Combo PDU C19/C13 1 Phase 60A Combo PDU C19/C13 1 Phase 63A International Combo PDU C19/C13 3 Phase 32A International Combo PDU C19/C13 1 Phase 32A Australia and NZ Combo PDU C19/C13 3 Phase 16A International	2[a] 2 2 2 2 4

a. Select one PDU line item from this list. These items are mutually exclusive. Most of them have a quantity of 2, except for the 16A PDU, which has a quantity of 4. the selection depends on the customer’s country and utility power requirements.

2.1.8 Software

This section lists the software features of IBM PureFlex System Express.

AIX and IBM i

Table 2-11 lists the software features included with the Express configuration on POWER7 processor-based compute nodes for AIX and IBM i.

Table 2-11 Software features for IBM PureFlex System Express with AIX and IBM i on Power

	AIX V6	AIX V7	IBM i V6.1	IBM i V7.1
Standard components - Express				
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM1 1 year software maintenance (SWMA)			
IBM Flex System Manager	► 5765-FMX FSM Standard ► 5660-FMX 1 year software maintenance			
Operating system	► 5765-G62 AIX Standard V6 ► 5771-SWM 1 yr SWMA	► 5765-G98 AIX Standard V7 ► 5771-SWM 1 yr SWMA	► 5761-SS1 IBM i V6.1 ► 5733-SSP 1 yr SWMA	► 5770-SS1 IBM i V7.1 ► 5733-SSP 1 yr SWMA
Virtualization	► 5765-PVS PowerVM Standard ► 5771-PVS 1 yr SWMA			
Security (PowerSC)	► 5765-PSE PowerSC Standard ► 5660-PSE 1 yr SWMA		Not applicable	Not applicable
Cloud Software (optional)	► None standard in Express configurations. Optional.			
Optional components - Express Expansion				
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring			
IBM Flex System Manager	► 5765-FMS FSM Advanced			
Operating system	► 5765-AEZ AIX V6 Enterprise	► 5765-G99 AIX V7 Enterprise		
Virtualization	► 5765-PVE PowerVM Enterprise			

	AIX V6	AIX V7	IBM i V6.1	IBM i V7.1
Security (PowerSC)	Not applicable	Not applicable	Not applicable	Not applicable
Cloud Software (optional)	► 5765-SCP SmartCloud Entry ► 5660-SCP 1 yr SWMA ► Requires upgrade to 5765-FMS IBM Flex System Manager Advanced	► 5765-SCP SmartCloud Entry ► 5660-SCP 1 yr SWMA ► Requires upgrade to 5765-FMS IBM Flex System Manager Advanced	Not applicable	Not applicable

RHEL and SUSE Linux on Power

Table 2-12 lists the software features included with the Express configuration on POWER7 processor-based compute nodes for Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) on Power.

Table 2-12 Software features for IBM PureFlex System Express with RHEL and SLES on Power

	Red Hat Enterprise Linux (RHEL)	SUSE Linux Enterprise Server (SLES)
Standard components - Express		
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM1 1 year software maintenance (SWMA)	
IBM Flex System Manager	► 5765-FMX FSM Standard ► 5660-FMX 1 year software maintenance	
Operating system	► 5639-RHP RHEL 5 and 6	► 5639-S11 SLES 11
Virtualization	► 5765-PVS PowerVM Standard ► 5771-PVS 1 yr SWMA	
Cloud Software (optional)	► 5765-SCP SmartCloud Entry ► 5660-SCP 1 yr SWMA ► Requires upgrade to 5765-FMS IBM Flex System Manager Advanced	
Optional components - Express Expansion		
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring	
IBM Flex System Manager	► 5765-FMS FSM Advanced	

	Red Hat Enterprise Linux (RHEL)	SUSE Linux Enterprise Server (SLES)
Virtualization	► 5765-PVE PowerVM Enterprise	

Intel Xeon based compute nodes

Table 2-13 lists the software features included with the Express configuration on Intel Xeon based compute nodes.

Table 2-13 Software features for IBM PureFlex System Express on Intel Xeon based compute nodes

	Intel Xeon based compute nodes (AAS)	Intel Xeon based compute nodes (HVEC)
Standard components - Express		
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM1 1 year software maintenance (SWMA)	
IBM Flex System Manager	► 5765-FMX FSM Standard ► 5660-FMX 1 year software maintenance	► 94Y9782 FSM Standard 1 year SWMA
Operating system	► Varies	► Varies
Virtualization	Not applicable	
Cloud Software (optional)	Not applicable	
Optional components - Express Expansion		
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring	
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced	► 94Y9783 IBM Flex System Manager Advanced
Operating system	► 5639-OSX RHEL for x86 ► 5639-W28 Windows 2008 R2 ► 5639-CAL Windows 2008 Client Access	► 5731RSI RHEL for x86 - L3 support only ► 5731RSR RHEL for x86 - L1-L3 support ► 5731W28 Windows 2008 R2 ► 5731CAL Windows 2008 Client Access
Virtualization	VMware ESXi selectable in the hardware configuration	
Cloud Software	► 5765-SCP SmartCloud Entry ► 5660-SCP 1 yr SWMA	► 5641-SC1 SmartCloud Entry with 1 year software maintenance

2.1.9 Services

IBM PureFlex System Express includes the following services:

- Service & Support offerings:
  - Software maintenance: 1 year 9x5 (9 hours per day, 5 days per week).
  - Hardware maintenance: 3 years 9x5 Next Business Day service.
- Maintenance and Technical Support (MTS) offerings:
  - Three years with one microcode analysis per year.
- Lab Services:
  - Three days of on-site lab services
  - If the first compute node is a p260 or p460, 6911-300 is specified.
  - If the first compute node is a x240, 6911-100 is specified.

2.2 IBM PureFlex System Standard

The tables in this section show the hardware, software, and services that make up IBM PureFlex System Standard. We describe the following items:

- Chassis
- Top-of-Rack Ethernet switch
- Top-of-Rack SAN switch
- Compute nodes
- IBM Flex System Manager
- IBM Storwize V7000
- Rack cabinet
- Software
- Services

To specify IBM PureFlex System Standard in the IBM ordering system, specify the indicator feature code listed in Table 2-14 for each machine type.

Table 2-14 Standard indicator feature code

AAS feature code	XCC feature code	Description
EFD2	A2VT	IBM PureFlex System Standard Indicator Feature Code - First of each MTM (for example, first compute node)

2.2.1 Chassis

Table 2-15 lists the major components of the IBM Flex System Enterprise Chassis, including the switches and options.

> **Feature codes:** The tables in this section do not list all feature codes. Some features are not listed here for brevity.

Table 2-15 Components of the chassis and switches

AAS feature code	XCC feature code	Description	Minimum quantity
7893-92X	8721-HC1	IBM Flex System Enterprise Chassis	1
3593	A0TB	IBM Flex System Fabric EN4093 10Gb Scalable Switch	1
3282	5053	10GbE 850 nm Fiber SFP+ Transceiver (SR)	4
EB29	3268	IBM BNT SFP RJ45 Transceiver	5
3595	A0TD	IBM Flex System FC3171 8Gb SAN Switch	2
3286	5075	IBM 8GB SFP+ Short-Wave Optical Transceiver	4
3590	A0UD	Additional PSU 2500W	2
4558	6252	2.5m, 16A/100-240V, C19 to IEC 320-C20 power cord	4
9039	A0TM	Base Chassis Management Module	1
3592	A0UE	Additional Chassis Management Module	1
9038	None	Base Fan Modules (four)	1
7805	A0UA	Additional Fan Modules (two)	1

2.2.2 Top-of-Rack Ethernet switch

If more than one chassis is configured, then a Top-of-Rack (TOR) Ethernet switch is added to the configuration. If only one chassis is configured, then the TOR switch is optional.

Table 2-16 lists the switch components.

Table 2-16 Components of the Top-of-Rack Ethernet switch

AAS feature code	XCC feature code	Description	Minimum quantity
7309-HC3	1455-64C	IBM System Networking RackSwitch G8264	0[a]
7309-G52	1455-48E	IBM System Networking RackSwitch G8052	0[a]
ECB5	A1PJ	3m IBM Passive DAC SFP+ Cable	1 per EN4093 switch
EB25	A1PJ	3m IBM QSFP+ DAC Break Out Cable	0

a. One required when a two or more Enterprise Chassis are configured

2.2.3 Top-of-Rack SAN switch

If more than one chassis is configured, then a Top-of-Rack SAN switch is added to the configuration. If only one chassis is configured, then the SAN switch is optional. Table 2-17 lists the switch components.

Table 2-17 Components of the Top-of-Rack SAN switch

AAS feature code	XCC feature code	Description	Minimum quantity
2498-B24	2498-B24	24-port SAN Switch	0
5605	5605	5m optic cable	1
2808	2808	8 Gb SFP transceivers (8 pack)	1

2.2.4 Compute nodes

The PureFlex System Standard requires either of the following compute nodes:

- IBM Flex System p460 Compute Node (IBM POWER7 based) (Table 2-18 on page 29)
- IBM Flex System x240 Compute Node (Intel Xeon based) (Table 2-19 on page 29)

Table 2-18 lists the major components of the IBM Flex System p460 Compute Node.

Table 2-18 Components of IBM Flex System p460 Compute Node

AAS feature code	Description	Minimum quantity
IBM Flex System p460 Compute Node		
7895-42x	IBM Flex System p460 Compute Node	1
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter	2
1762	IBM Flex System EN4054 4-port 10Gb Ethernet Adapter	2
Base Processor 1 Required, select only one, Min 1, Max 1		
EPR2	16 Cores, (4x 4 core), 3.3 GHz + 4-socket system board	1
EPR4	32 Cores, (4x 8 core), 3.2 GHz + 4-socket system board	
EPR6	32 Cores, (4x 8 core), 3.55 GHz + 4-socket system board	
Memory - 8 GB per core minimum with all DIMM slots filled with same memory type		
8145	32GB (2x 16GB), 1066MHz, LP RDIMMs (1.35V)	
8199	16GB (2x 8GB), 1066MHz, VLP RDIMMs (1.35V)	

Table 2-19 lists the major components of the IBM Flex System x240 Compute Node.

Table 2-19 Components of IBM Flex System x240 Compute Node

AAS feature code	XCC feature code	Description	Minimum quantity
IBM Flex System x240 Compute Node			
7863-10X	8737AC1	IBM Flex System x240 Compute Node	
EN20 EN21	A1BC A1BD	x240 with embedded 10Gb Virtual Fabric x240 without embedded 10Gb Virtual Fabric (select one of these base features)	1
1764	A2N5	IBM Flex System FC3052 2-port 8Gb FC Adapter	1
1759	A1R1	IBM Flex System CN4054 10Gb Virtual Fabric Adapter (select if x240 without embedded 10Gb Virtual Fabric is selected - EN21/A1BD)	1

AAS feature code	XCC feature code	Description	Minimum quantity
EBK2	49Y8119	IBM Flex System x240 USB Enablement Kit	
EBK3	41Y8300	2GB USB Hypervisor Key (VMware 5.0)	

2.2.5 IBM Flex System Manager

Table 2-20 lists the major components of the IBM Flex System Manager.

Table 2-20 Components of the IBM Flex System Manager

AAS feature code	XCC feature code	Description	Minimum quantity
7955-01M	8731AC1	IBM Flex System Manager	1
EB31	9220	Platform Bundle preload indicator	1
EM09 None	None 8941	8GB (2x 4GB) 1333 MHz RDIMMs (1.35V) 4GB (1x 4GB) 1333 MHz RDIMMs (1.35V)	4[a] 8
None	A1CW	Intel Xeon E5-2650 8C 2.0GHz 20MB 1600MHz 95W	1
1771	5420	200GB, 1.8", SATA MLC SSD	2
3767	A1AV	1TB 2.5” SATA 7.2K RPM hot-swap 6 Gbps HDD	1

a. In the AAS system, FC EM09 are pairs of DIMMs. In the XCC system, FC 8941 are single DIMMs. The DIMMs are otherwise identical.

2.2.6 IBM Storwize V7000

Table 2-21 lists the major components of the IBM Storwize V7000 storage server.

Table 2-21 Components of the IBM Storwize V7000 storage server

AAS feature code	XCC feature code	Description	Minimum quantity
2076-124	2076-124	IBM Storwize V7000 Controller	1
5305	5305	5m Fiber Optic Cable	2
3512 3514	3512 3514	200GB 2.5 INCH SSD or 400GB 2.5 INCH SSD	2[a]
0010	0010	Storwize V7000 Software Preinstall	1

AAS feature code	XCC feature code	Description	Minimum quantity
6008	6008	8 GB Cache	2
9730	9730	power cord to PDU (includes 2 power cord)	1
9801	9801	Power supplies	2

a. If a Power Systems compute node is selected, then at least eight drives must be installed in the Storwize V7000. If an Intel Xeon based compute node is selected with SmartCloud Entry, then four drives must be installed in the Storwize V7000.

2.2.7 Rack cabinet

Table 2-22 lists the major components of the rack and options.

Table 2-22 Components of the rack

AAS feature code	XCC feature code	Description	Minimum quantity
7953-94X	93634AX	IBM 42U 1100mm Flex System Dynamic Rack	1
EC06	None	Gray Door	1
EC03	None	Side Cover Kit (Black)	1
EC02	None	Rear Door (Black/flat)	1
7196 7189+6492 7189+6491 7189+6489 7189+6667 7189+6653	5897 5902 5904 5903 5906 None	Combo PDU C19/C13 3 Phase 60A Combo PDU C19/C13 1 Phase 60A Combo PDU C19/C13 1 Phase 63A International Combo PDU C19/C13 3 Phase 32A International Combo PDU C19/C13 1 Phase 32A Australia and NZ Combo PDU C19/C13 3 Phase 16A International	2[a] 2 2 2 2 4

a. Select one PDU line item from this list. These items are mutually exclusive. Most of them have a quantity of 2, except for the 16A PDU, which has a quantity of 4. The selection depends on the customer's country and utility power requirements.

2.2.8 Software

This section lists the software features of IBM PureFlex System Standard.

AIX and IBM i

Table 2-23 lists the software features included with the Standard configuration on POWER7 processor-based compute nodes for AIX and IBM i.

Table 2-23 Software features for IBM PureFlex System Standard with AIX and IBM i on Power

	AIX V6	AIX V7	IBM i V6.1	IBM i V7.1
Standard components - Standard				
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM3 3 year software maintenance (SWMA)			
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced ► 5662-FMS 3 year software maintenance			
Operating system	► 5765-G62 AIX Standard V6 ► 5773-SWM 3 year SWMA	► 5765-G98 AIX Standard V7 ► 5773-SWM 3 year SWMA	► 5761-SS1 IBM i V6.1 ► 5773-SWM 3 year SWMA	► 5770-SS1 IBM i V7.1 ► 5773-SWM 3 year SWMA
Virtualization	► 5765-PVE PowerVM Enterprise ► 5773-PVE 3 year SWMA			
Security (PowerSC)	► 5765-PSE PowerSC Standard ► 5662-PSE 3 year SWMA		Not applicable	Not applicable
Cloud Software (default but optional)	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 year SWMA	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 year SWMA	Not applicable	Not applicable
Optional components - Standard Expansion				
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring			
IBM Flex System Manager	Not applicable			
Operating system	► 5765-AEZ AIX V6 Enterprise	► 5765-G99 AIX V7 Enterprise		
Virtualization	► 5765-PVE PowerVM Enterprise			

	AIX V6	AIX V7	IBM i V6.1	IBM i V7.1
Security (PowerSC)	Not applicable	Not applicable	Not applicable	Not applicable
Cloud Software (optional)	Not applicable	Not applicable	Not applicable	Not applicable

RHEL and SUSE Linux on Power

Table 2-24 lists the software features included with the Standard configuration on POWER7 processor-based compute nodes for Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) on Power.

Table 2-24 Software features for IBM PureFlex System Standard with RHEL and SLES on Power

	Red Hat Enterprise Linux (RHEL)	SUSE Linux Enterprise Server (SLES)
Standard components - Standard		
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM3 3 year software maintenance agreement (SWMA)	
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced ► 5662-FMS 3 year SWMA	
Operating system	► 5639-RHP RHEL 5 & 6	► 5639-S11 SLES 11
Virtualization	► 5765-PVE PowerVM Enterprise ► 5773-PVE 3 year SWMA	
Cloud Software (optional)	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 year SWMA	
Optional components - Standard Expansion		
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring	
IBM Flex System Manager	Not applicable	
Virtualization	Not applicable	

Intel Xeon based compute nodes

Table 2-25 lists the software features included with the Standard configuration on Intel Xeon based compute nodes.

Table 2-25 Software features for IBM PureFlex System Standard on Intel Xeon based compute nodes

	Intel Xeon based compute nodes (AAS)	Intel Xeon based compute nodes (HVEC)
Standard components - Standard		
IBM Storwize V7000 Software	► 5639-VM1 - V7000 Base PID ► 5639-SM3 - three years of software maintenance agreement (SWMA)	
IBM Flex System Manager	► 5765-FMX FSM Standard ► 5662-FMX 3 year software maintenance	► 94Y9787 FSM Standard, 3 year SWMA
Operating system	► Varies	► Varies
Virtualization	► VMware ESXi selectable in the hardware configuration	
Cloud Software (optional) (Windows and RHEL only)	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 yr SWMA	► 5641-SC3 SmartCloud Entry, 3 yr SWMA
Optional components - Standard Expansion		
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring	
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced	► 94Y9783 IBM Flex System Manager Advanced
Operating system	► 5639-OSX RHEL for x86 ► 5639-W28 Windows 2008 R2 ► 5639-CAL Windows 2008 Client Access	► 5731RSI RHEL for x86 - L3 support only ► 5731RSR RHEL for x86 - L1-L3 support ► 5731W28 Windows 2008 R2 ► 5731CAL Windows 2008 Client Access
Virtualization	VMware ESXi selectable in the hardware configuration	
Cloud Software	Not applicable	Not applicable

2.2.9 Services

IBM PureFlex System Standard includes the following services:

- Service & Support offerings:
  - Software maintenance: 1 year 9x5 (9 hours per day, 5 days per week).
  - Hardware maintenance: 3 years 9x5 Next Business Day service.
- Maintenance and Technical Support (MTS) offerings:
  - 3 years with one microcode analysis per year.
- Lab Services:
  - 5 days of on-site Lab services
  - If the first compute node is a p260 or p460, 6911-300 is specified.
  - If the first compute node is a x240, 6911-100 is specified.

2.3 IBM PureFlex System Enterprise

The tables in this section represent the hardware, software. and services that make up IBM PureFlex System Enterprise. We describe the following items:

- Chassis
- Top-of-Rack Ethernet switch
- Top-of-Rack SAN switch
- Compute nodes
- IBM Flex System Manager
- IBM Storwize V7000
- Rack cabinet
- Software
- Services

To specify IBM PureFlex System Enterprise in the IBM ordering system, specify the indicator feature code listed in Table 2-26 for each machine type.

Table 2-26 Enterprise indicator feature code

AAS feature code	XCC feature code	Description
EFD3	A2VU	IBM PureFlex System Enterprise Indicator Feature Code: first of each MTM (for example, first compute node)

2.3.1 Chassis

Table 2-27 lists the major components of the IBM Flex System Enterprise Chassis, including the switches and options.

Feature codes: The tables in this section do not list all feature codes. Some features are not listed here for brevity.

Table 2-27 Components of the chassis and switches

AAS feature code	XCC feature code	Description	Minimum quantity
7893-92X	8721-HC1	IBM Flex System Enterprise Chassis	1
3593	A0TB	IBM Flex System Fabric EN4093 10Gb Scalable Switch	2
3596	A1EL	IBM Flex System Fabric EN4093 10Gb Scalable Switch Upgrade 1	2
3597	A1EM	IBM Flex System Fabric EN4093 10Gb Scalable Switch Upgrade 2	2
3282	5053	10GbE 850 nm Fiber SFP+ Transceiver (SR)	4
EB29	3268	IBM BNT SFP RJ45 Transceiver	6
3595	A0TD	IBM Flex System FC3171 8Gb SAN Switch	2
3286	5075	IBM 8GB SFP+ Short-Wave Optical Transceiver	8
3590	A0UD	Additional PSU 2500W	4
4558	6252	2.5m, 16A/100-240V, C19 to IEC 320-C20 power cord	6
9039	A0TM	Base Chassis Management Module	1
3592	A0UE	Additional Chassis Management Module	1
9038	None	Base Fan Modules (four)	1
7805	A0UA	Additional Fan Modules (two)	2

2.3.2 Top-of-Rack Ethernet switch

A minimum of two Top-of-Rack (TOR) Ethernet switches are required in the Enterprise configuration. Table 2-28 lists the switch components.

Table 2-28 Components of the Top-of-Rack Ethernet switch

AAS feature code	XCC feature code	Description	Minimum quantity
7309-HC3	1455-64C	IBM System Networking RackSwitch G8264	2[a]
	1455-48E	IBM System Networking RackSwitch G8052	2[a]
ECB5	A1PJ	3m IBM Passive DAC SFP+ Cable	1 per EN4093 switch
EB25	A1PJ	3m IBM QSFP+ DAC Break Out Cable	1

a. For Power Systems configurations, two are required. For System x configurations, two are required when two or more Enterprise Chassis are configured.

2.3.3 Top-of-Rack SAN switch

A minimum of two Top-of-Rack (TOR) SAN switches are required in the Enterprise configuration. Table 2-29 lists the switch components.

Table 2-29 Components of the Top-of-Rack SAN switch

AAS feature code	XCC feature code	Description	Minimum quantity
2498-B24	2498-B24	24-port SAN Switch	0
5605	5605	5m optic cable	1
2808	2808	8 Gb SFP transceivers (8 pack)	1

2.3.4 Compute nodes

The PureFlex System Enterprise requires either of the following compute nodes:

- IBM Flex System p460 Compute Node (IBM POWER7 based) (Table 2-30 on page 38)
- IBM Flex System x240 Compute Node (Intel Xeon based) (Table 2-31 on page 38)

Table 2-30 lists the major components of the IBM Flex System p260 Compute Node.

Table 2-30 Components of IBM Flex System p460 Compute Node

AAS feature code	Description	Minimum quantity
IBM Flex System p460 Compute Node		
7895-42x	IBM Flex System p460 Compute Node	2
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter	2
1762	IBM Flex System EN4054 4-port 10Gb Ethernet Adapter	2
Base Processor 1 Required, select only one, Min 1, Max 1		
EPR2	16 Cores, (4x 4 core), 3.3 GHz + 4-socket system board	1
EPR4	32 Cores, (4x 8 core), 3.2 GHz + 4-socket system board	
EPR6	32 Cores, (4x 8 core), 3.55 GHz + 4-socket system board	
Memory - 8 GB per core minimum with all DIMM slots filled with same memory type		
8145	32GB (2x 16GB), 1066MHz, LP RDIMMs (1.35V)	
8199	16GB (2x 8GB), 1066MHz, VLP RDIMMs (1.35V)	

Table 2-31 lists the major components of the IBM Flex System x240 Compute Node.

Table 2-31 Components of IBM Flex System x240 Compute Node

AAS feature code	XCC feature code	Description	Minimum quantity
IBM Flex System x240 Compute Node			
7863-10X	8737AC1	IBM Flex System x240 Compute Node	2
EN20 EN21	A1BC A1BD	x240 with embedded 10Gb Virtual Fabric x240 without embedded 10Gb Virtual Fabric (select one of these base features)	1 per
1764	A2N5	IBM Flex System FC3052 2-port 8Gb FC Adapter	1 per
1759	A1R1	IBM Flex System CN4054 10Gb Virtual Fabric Adapter (select if x240 without embedded 10Gb Virtual Fabric is selected - EN21/A1BD)	1 per

AAS feature code	XCC feature code	Description	Minimum quantity
EBK2	49Y8119	IBM Flex System x240 USB Enablement Kit	
EBK3	41Y8300	2GB USB Hypervisor Key (VMware 5.0)	

2.3.5 IBM Flex System Manager

Table 2-32 lists the major components of the IBM Flex System Manager.

Table 2-32 Components of the IBM Flex System Manager

AAS feature code	XCC feature code	Description	Minimum quantity
7955-01M	8731AC1	IBM Flex System Manager	1
EB31	9220	Platform Bundle preload indicator	1
EM09 None	None 8941	8GB (2x 4GB) 1333 MHz RDIMMs (1.35V) 4GB (1x 4GB) 1333 MHz RDIMMs (1.35V)	4[a] 8
None	A1CW	Intel Xeon E5-2650 8C 2.0GHz 20MB 1600MHz 95W	1
1771	5420	200GB, 1.8", SATA MLC SSD	2
3767	A1AV	1TB 2.5" SATA 7.2K RPM hot-swap 6 Gbps HDD	1

a. In the AAS system, FC EM09 has pairs of DIMMs. In the XCC system, FC 8941 has single DIMMs. The DIMMs are otherwise identical.

2.3.6 IBM Storwize V7000

Table 2-33 lists the major components of the IBM Storwize V7000 storage server.

Table 2-33 Components of the IBM Storwize V7000 storage server

AAS feature code	XCC feature code	Description	Minimum quantity
2076-124	2076-124	IBM Storwize V7000 Controller	1
5305	5305	5m Fiber Optic Cable	4
3512 3514	3512 3514	200GB 2.5 INCH SSD or 400GB 2.5 INCH SSD	2[a]
0010	0010	Storwize V7000 Software Preinstall	1

AAS feature code	XCC feature code	Description	Minimum quantity
6008	6008	8 GB Cache	2
9730	9730	power cord to PDU (includes 2 power cord)	1
9801	9801	Power supplies	2

a. If Power Systems compute node is selected, then at least eight drives must be installed in the Storwize V7000. If an Intel Xeon based compute node is selected with SmartCloud Entry, then four drives must be installed in the Storwize V7000.

2.3.7 Rack cabinet

Table 2-34 lists the major components of the rack and options.

Table 2-34 Components of the rack

AAS feature code	XCC feature code	Description	Minimum quantity
7953-94X	93634AX	IBM 42U 1100mm Flex System Dynamic Rack	1
EC06	None	Gray Door	1
EC03	None	Side Cover Kit (Black)	1
EC02	None	Rear Door (Black/flat)	1
7196 7189+6492 7189+6491 7189+6489 7189+6667 7189+6653	5897 5902 5904 5903 5906 None	Combo PDU C19/C13 3 Phase 60A Combo PDU C19/C13 1 Phase 60A Combo PDU C19/C13 1 Phase 63A International Combo PDU C19/C13 3 Phase 32A International Combo PDU C19/C13 1 Phase 32A Australia and NZ Combo PDU C19/C13 3 Phase 16A International	2[a] 2 2 2 2 4

a. Select one PDU line item from this list. These items are mutually exclusive. Most of these items have a quantity of 2, except for the 16A PDU, which has a quantity of 4. The selection depends on the customer's country and utility power requirements.

2.3.8 Software

This section lists the software features of IBM PureFlex System Enterprise.

AIX and IBM i

Table 2-35 lists the software features included with the Enterprise configuration on POWER7 processor-based compute nodes for AIX and IBM i.

Table 2-35 Software features for IBM PureFlex System Enterprise with AIX and IBM i on Power

	AIX 6	AIX 7	IBM i 6.1	IBM i 7.1
Standard components - Standard				
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM3 3 year software maintenance (SWMA)			
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced ► 5662-FMS 3 year software maintenance			
Operating system	► 5765-G62 AIX Standard V6 ► 5773-SWM 3 year SWMA	► 5765-G98 AIX Standard V7 ► 5773-SWM 3 year SWMA	► 5761-SS1 IBM i V6.1 ► 5773-SWM 3 year SWMA	► 5770-SS1 IBM i V7.1 ► 5773-SWM 3 year SWMA
Virtualization	► 5765-PVE PowerVM Enterprise ► 5773-PVE 3 year SWMA			
Security (PowerSC)	► 5765-PSE PowerSC Standard ► 5662-PSE 3 year SWMA		Not applicable	Not applicable
Cloud Software (default but optional)	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 year SWMA	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 year SWMA	Not applicable	Not applicable
Optional components - Standard Expansion				
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring			
IBM Flex System Manager	Not applicable			
Operating system	► 5765-AEZ AIX V6 Enterprise	► 5765-G99 AIX V7 Enterprise		
Virtualization	5765-PVE PowerVM Enterprise			
Security (PowerSC)	Not applicable	Not applicable	Not applicable	Not applicable
Cloud Software (optional)	Not applicable	Not applicable	Not applicable	Not applicable

RHEL and SUSE Linux on Power

Table 2-36 lists the software features included with the Enterprise configuration on POWER7 processor-based compute nodes for Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) on Power.

Table 2-36 Software features for IBM PureFlex System Enterprise with RHEL and SLES on Power

	Red Hat Enterprise Linux (RHEL)	SUSE Linux Enterprise Server (SLES)
Standard components - Standard		
IBM Storwize V7000 Software	► 5639-VM1 V7000 Base PID ► 5639-SM3 3 year software maintenance agreement (SWMA)	
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced ► 5662-FMS 3 year SWMA	
Operating system	► 5639-RHP RHEL 5 and 6	► 5639-S11 SLES 11
Virtualization	► 5765-PVE PowerVM Enterprise ► 5773-PVE 3 year SWMA	
Cloud Software (optional)	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 year SWMA	
Optional components - Standard Expansion		
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring	
IBM Flex System Manager	Not applicable	
Virtualization	Not applicable	

Intel Xeon based compute nodes

Table 2-37 lists the software features included with the Enterprise configuration on Intel Xeon based compute nodes.

Table 2-37 Software features for IBM PureFlex System Enterprise on Intel Xeon based compute nodes

	Intel Xeon based compute nodes (AAS)	Intel Xeon based compute nodes (HVEC)
Standard components - Enterprise		
IBM Storwize V7000 Software	► 5639-VM1 - V7000 Base PID ► 5639-SM3 - 3 years of software maintenance agreement (SWMA)	
IBM Flex System Manager	► 5765-FMX FSM Standard ► 5662-FMX 3 year software maintenance	► 94Y9787 FSM Standard, 3 year SWMA

	Intel Xeon based compute nodes (AAS)	Intel Xeon based compute nodes (HVEC)
Operating system	► Varies	► Varies
Virtualization	► VMware ESXi selectable in the hardware configuration	
Cloud Software (optional)	► 5765-SCP SmartCloud Entry ► 5662-SCP 3 yr SWMA	► 5641-SC3 SmartCloud Entry, 3 yr SWMA
Optional components - Enterprise Expansion		
IBM Storwize V7000 Software	► 5639-EV1 V7000 External virtualization software ► 5639-RM1 V7000 Remote Mirroring	
IBM Flex System Manager	► 5765-FMS IBM Flex System Manager Advanced	► 94Y9783 IBM Flex System Manager Advanced
Operating system	► 5639-OSX RHEL for x86 ► 5639-W28 Windows 2008 R2 ► 5639-CAL Windows 2008 Client Access	► 5731RSI RHEL for x86 - L3 support only ► 5731RSR RHEL for x86 - L1-L3 support ► 5731W28 Windows 2008 R2 ► 5731CAL Windows 2008 Client Access
Virtualization	VMware ESXi selectable in the hardware configuration	
Cloud Software	Not applicable	Not applicable

2.3.9 Services

IBM PureFlex System Enterprise includes the following services:

- Service & Support offerings:
  - Software maintenance: 1 year 9x5 (9 hours per day, 5 days per week).
  - Hardware maintenance: 3 years 9x5 Next Business Day service.
- Maintenance and Technical Support (MTS) offerings:
  - 3 years with one microcode analysis per year.
- Lab Services:
  - 7 days of on-site lab services
  - If the first compute node is a p260 or p460, 6911-300 is specified.
  - If the first compute node is a x240, 6911-100 is specified.

2.4 IBM SmartCloud Entry

In IT environments, you face the challenges of delivering new capabilities with growth of data, and the increase of applications and the amount of physical hardware, such as servers, storages, and networks. The traditional means of deploying, provisioning, managing, and maintaining physical and virtual resources can no longer meet the demands of increasingly complex IT infrastructure. Virtualization simplifies and improves efficiency and utilization, and helps you manage growth beyond physical resource boundaries.

With SmartCloud Entry, you can build on your current virtualization strategies to continue to gain IT efficiency, flexibility, and control.

There are many advantages to adapting a cloud environment:

- You can reduce a data center’s footprint and management cost.
- You can use an automated server request/provisioning solution.
- You can improve utilization and workload management, and deliver new services.
- You have rapid service deployment, where you reduce the time to deployment form several weeks to just days or hours.
- You have a built-in metering system.
- You can improve IT governance and risk management.

IBM simplifies your journey from server consolidation to cloud management by providing complete cloud solutions, hardware, software technologies, and services for implementing a private cloud with added value on top of a virtualized infrastructure. IBM SmartCloud™ Entry for Cloud offerings provides these cloud solutions, and have features that set them apart from other products. This single product provides a comprehensive cloud software stack with capabilities that can be otherwise provided only by using multiple products from other providers, such as VMware. It enables customers to quickly deploy their cloud environment; in addition, IBM offers advanced cloud features as you need them.

You can use existing IBM server investments and virtualized environments to deploy IBM SmartCloud Entry with the essential cloud infrastructure capabilities:

When you create images using this product, you:

- Simplify the storage of thousands of images.
- Easily create new ‘golden master’ images and software appliances using corporate standard operating systems.
- Convert images from physical systems or between various x86 hypervisors.

- Reliably track images to ensure compliance and minimize security risks.
- Optimize resources, reducing the number of virtualized images and the storage required for them.

When you deploy VMs using this product, you:

- Slash time to value for new workloads from months to a few days.
- Deploy application images across compute and storage resources.
- Provide user self-service for improved responsiveness.
- Ensure security through VM isolation and project-level user access controls.
- Have an easy to use solution; there no need to know all the details of the infrastructure.
- Have investment protection with the full support of existing virtualized environments.
- Optimize performance on IBM systems with dynamic scaling, expansive capacity, and continuous operation.

When you operate a private cloud using this product, you:

- Cut costs with efficient operations.
- Delegate provisioning to authorized users to improve productivity.
- Maintain full oversight to ensure an optimally run and safe cloud through automated approval or rejection.
- Standardize deployment and configuration to improve compliance and reduce errors by setting policies, defaults, and templates.
- Simplify administration with an intuitive interface for managing projects, users, workloads, resources, billing, approvals, and metering.

IBM Cloud and virtualization solutions offer flexible approaches to cloud computing. Where you start your journey depends on your business needs.

For more information about IBM SmartCloud Entry, go to the following website:

http://ibm.com/systems/cloud

3

Introduction to IBM Flex System

IBM Flex System is a solution composed of hardware, software, and expertise. The IBM Flex System Enterprise Chassis, the major hardware component, is the next generation platform that provides new capabilities in many areas:

- Scalability
- Current and future processors
- Memory
- Storage
- Bandwidth and I/O speeds
- Power
- Energy efficiency and cooling
- Systems management

Figure 3-1 shows the front and rear views of the Enterprise Chassis.

Figure 3-1 IBM Flex System Enterprise Chassis - front and rear

The chassis provides locations for 14 half-wide nodes, four scalable I/O switch modules, and two Chassis Management Modules. Current node configurations include half-wide and full-wide options. The chassis supports other configurations, such as full-wide by double-high. Power and cooling can be scaled up in a modular fashion as additional nodes are added.

Table 3-1 shows the specifications of the Enterprise Chassis:

Table 3-1 Enterprise Chassis specifications

Feature	Specifications
Machine type-model	System x ordering sales channel: 8721-A1x. Power Systems sales channel: 7893-92X.
Form factor	10 U rack mounted unit.
Maximum number of compute nodes supported	14 half-wide (single bay), or 7 full-wide (two bays) or 3 double-height full-wide (four bays). Mixing is supported.
Chassis per 42U rack	4.
Nodes per 42U rack	56 half-wide, or 28 full-wide.

Feature	Specifications
Management	One or two Chassis Management Modules, for basic chassis management. Two CMMs form a redundant pair; one CMM is standard in 8721-A1x. The CMM interfaces with the integrated management module (IMM) or flexible service processor (FSP) integrated in each compute node in the chassis. There is an optional IBM Flex System Manager management appliance for comprehensive management, including virtualization, networking, and storage management.
I/O architecture	Up to 8 lanes of I/O to an I/O adapter, with each lane capable of up to 16 Gbps bandwidth. Up to 16 lanes of I/O to a half wide-node with two adapters. There are a wide variety of networking solutions, including Ethernet, Fibre Channel, FCoE, and InfiniBand
Power supplies	Six 2500-watt power modules that provide N+N or N+1 redundant power; two power modules are standard in model 8721-A1x. The power supplies are 80 PLUS Platinum certified and provides 95% efficiency at 50% load and 92% efficiency at 100% load. There is a power capacity of 2500 W output rated at 200 VAC. Each power supply contains two independently powered 40 mm cooling fan modules.
Fan modules	Ten fan modules (eight 80 mm fan modules and two 40 mm fan modules); Four 80 mm and two 40 mm fan modules are standard in model 8721-A1x.
Dimensions	▶ Height: 440 mm (17.3 in.). ▶ Width: 447 mm (17.6 in.). ▶ Depth, measured from front bezel to rear of chassis: 800 mm (31.5 in.). ▶ Depth, measured from node latch handle to the power supply handle: 840 mm (33.1 in.).
Weight	▶ Minimum configuration: 96.62 kg (213 lb). ▶ Maximum configuration: 220.45 kg (486 lb).
Declared sound level	6.3 to 6.8 bels.
Temperature	Operating air temperature 5 - 40 °C.
Electrical power	Input power: 200 - 240 V AC (nominal), 50 or 60 Hz. Minimum configuration: 0.51 kVA (two power supplies). Maximum configuration: 13 kVA (six power supplies).
Power consumption	12,900 W maximum,

3.1 Compute nodes

The IBM Flex System portfolio of servers, or *compute nodes*, includes IBM POWER7 and Intel Xeon processors. Depending on the compute node design, there are two form factors:

- Half-wide node: This node occupies one chassis bay, or half of the chassis width. An example is the IBM Flex System p260 Compute Node.
- Full-wide node: This node occupies two chassis bays side-by-side, or the full width of the chassis. An example is the IBM Flex System p460 Compute Node.

Figure 3-2 shows a front view of the chassis, with the bay locations identified and several half-wide nodes installed.

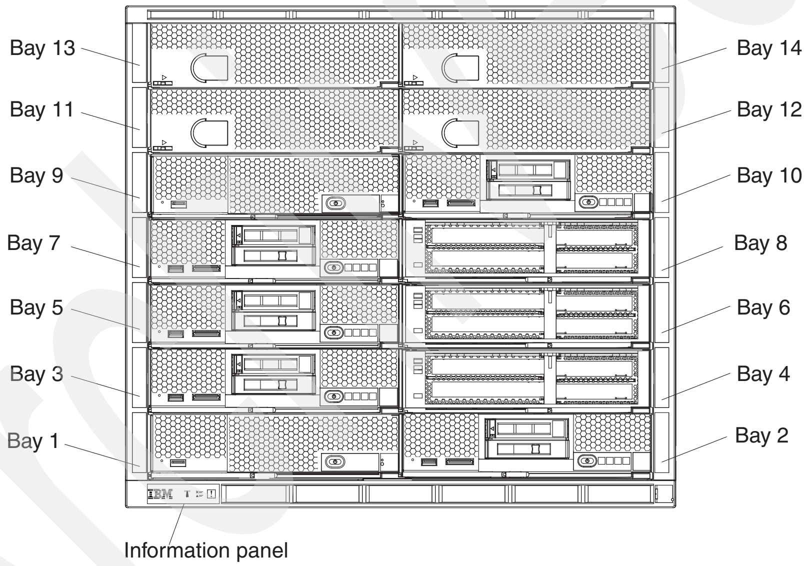

Figure 3-2 Enterprise Chassis - front view

Compute nodes in the Intel and IBM processor types have options for processor speeds, quantities of memory, expansion cards, and internal disk types and sizes.

Virtualization technologies supported are PowerVM on Power Systems compute node and VMware ESX, KVM, and Hyper V on Intel based compute nodes.

3.2 I/O modules

The I/O modules or switches provide external connectivity to nodes outside the chassis and internal connectivity to the nodes in the chassis. These switches are scalable in terms of the number of internal and external ports that can be enabled, and how these ports can be used to aggregate bandwidth and create virtual switches within a physical switch. The number of internal and external physical ports available exceeds previous generations of products. These additional ports can be scaled or enabled as requirements grow, and additional capability can be introduced.

The Enterprise Chassis can accommodate a total of four I/O modules, which are installed in a vertical orientation into the rear of the chassis, as shown in Figure 3-3.

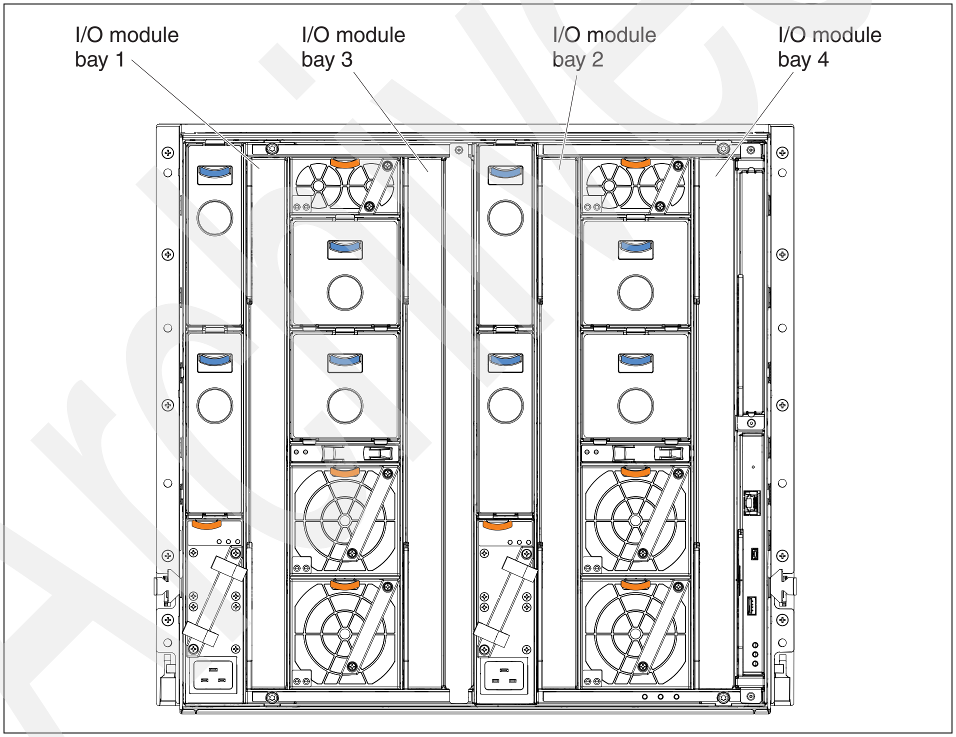

Figure 3-3 Enterprise Chassis I/O module locations

The internal connections between the node ports and the I/O module internal ports are defined by:

- I/O modules 1 and 2

  These two modules connect to the ports on an I/O expansion card in slot position 1 for a half-wide compute node (such as the p260) or slot positions 1 and 3 for a full-wide compute node (such as the p460).

  > **Intel computer nodes:** Certain Intel compute nodes offer an integrated local area network (LAN) on the system board (LOM). POWER based compute nodes do not have the LOM option.

- I/O modules 3 and 4

  These two modules are connected to the ports on an I/O expansion card in slot position 2 for a half-wide compute node or slots positions 2 and 4 for a full-wide compute node.

This connectivity is shown in Figure 3-4.

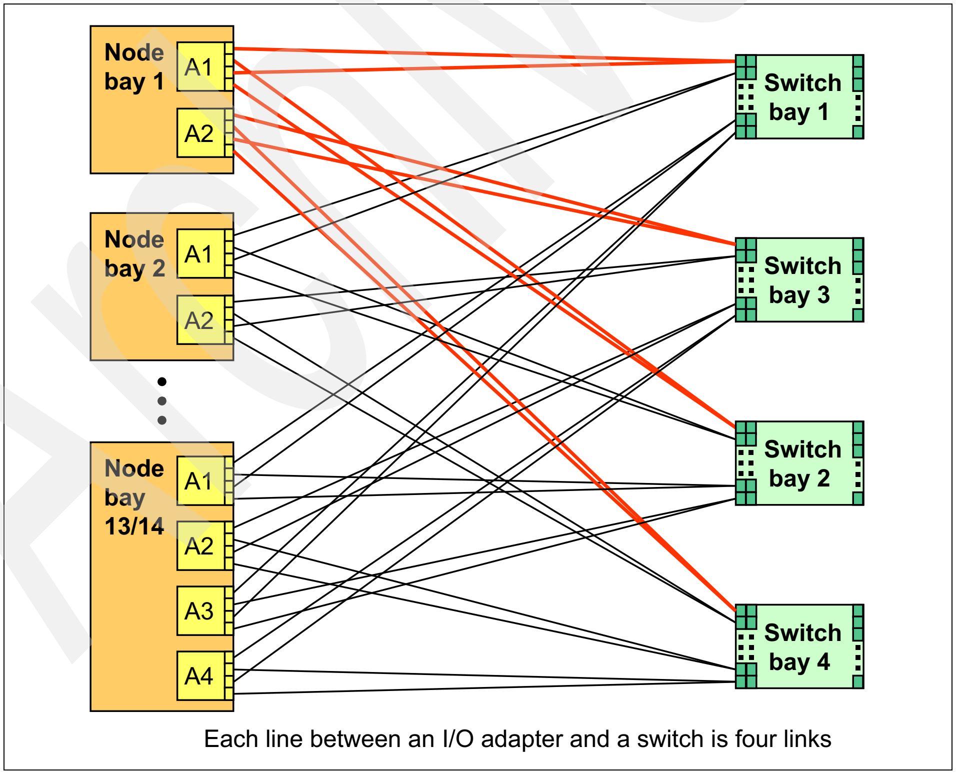

Figure 3-4 Connectivity between I/O adapter slots and switch bays

The following Ethernet switches were announced at the time of writing:

- IBM Flex System Fabric EN4093 10Gb Scalable Switch
  - 42x internal ports, 14x 10 Gb and 2x 40 Gb (convertible to 8x 10 Gb) uplinks
  - Base switch: 10x external 10 Gb uplinks, 14x 10 Gb internal 10 Gb ports
  - Upgrade 1: Adds 2x external 40 Gb uplinks and 14x internal 10 Gb ports
  - Upgrade 2: Adds 4x external 10 Gb uplinks, 14x internal 10 Gb ports
- IBM Flex System EN2092 1Gb Ethernet Scalable Switch
  - 28 Internal ports, 20 x 1 Gb and 4 x 10 Gb uplinks
  - Base: 14 internal 1 Gb ports, 10 external 1 Gb ports
  - Upgrade 1: Adds 14 internal 1 Gb ports, 10 external 1 Gb ports
  - Uplinks upgrade: Adds four external 10 Gb uplinks
- IBM Flex System EN4091 10Gb Ethernet Pass-thru
  - 14x 10 Gb internal server ports
  - 14x 10 Gb external SFP+ ports

The following Fibre Channel switches were announced at the time of writing:

- IBM Flex System FC3171 8Gb SAN Pass-thru
  - 28 internal and six external ports: 2, 4, and 8 Gb capable
- IBM Flex System FC3171 8Gb SAN Switch
  - 28 internal and six external ports: 2, 4, and 8 Gb capable
- IBM Flex System FC5022 16Gb SAN Scalable Switch and IBM Flex System FC5022 24-port 16Gb ESB SAN Scalable Switch
  - 28 internal and 20 external ports; 4, 8, and 16 Gb capable
  - FC5022 16Gb SAN Scalable Switch: Any 12 ports
  - FC5022 16Gb ESB Switch: Any 24 ports

The following InfiniBand switches were announced at the time of writing:

- IBM Flex System IB6131 InfiniBand Switch
  - 14 internal QDR ports (up to 40 Gbps)
  - 18 external QDR ports
  - Upgradeable to FDR speeds (56 Gbps)

For details about the available switches, see *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984.

3.3 Systems management

IBM Flex System uses a tiered approach to overall system management. The tiers are:

- Private management network within each chassis
- Firmware and management controllers for nodes and scalable switches
- Chassis Management Module for basic chassis management
- IBM Flex System Manager for advanced chassis management
- IBM Systems Director
- Upward integration with IBM Tivoli products

3.3.1 Private management network

At a physical level, the private management network is a dedicated 1 Gb Ethernet network within the chassis. This network is only accessible by the management controllers in the compute nodes or switch elements, the Chassis Management Modules, and the IBM Flex System Manager management appliance. This private network ensures a separation of the chassis management network from the data network.

The private management network is the connection for all traffic related to the remote presence of the nodes, delivery of firmware packages, and a direct connection to the management controller on each component.

3.3.2 Management controllers

At the next level, chassis components have their own core firmware and management controllers. Depending on the processor type of the compute nodes, either an Integrated Management Module 2 (IMMv2) or Flexible Service Processor (FSP) serves as the management controller. Additionally, each switch has a controller. In each case, the management controller provides an access point for the next level of system managers and a direct user interface.

3.3.3 Chassis Management Module

The Chassis Management Module (CMM) is a hot-swap module that is central to the management of the chassis and is required in each chassis. The CMM automatically detects any installed modules in the chassis and stores vital product data (VPD) from the modules. The CMM also acts as an aggregation point for the chassis nodes and switches, including enabling all of the management communications by Ethernet connection.

The CMM is also the key component that enables the internal management network. The CMM has a multiport, L2, 1 Gb Ethernet switch with dedicated links to all 14 node bays, the four switch bays, and the optional second CMM.

The second optional CMM provides redundancy in active and standby modes, has the same internal connections as the primary CMM, and is aware of all activity of the primary CMM through the trunk link between the two CMMs. This situation ensures that the backup CMM is ready to take over in a failover situation.

3.3.4 IBM Flex System Manager

The next tier in the management stack is the IBM Flex System Manager (FSM) management appliance. The FSM a dedicated, special purpose, half-wide compute node that can be installed in any chassis node slot and provides full management capabilities for up to four chassis. All functions and software are preinstalled and are initially configured with Quick Start wizards, integrating all components of the chassis, nodes, and switch modules.

The highlighted features of the FSM are:

- A single pane of glass to manage multiple nodes and multiple chassis
- Automatic discovery of nodes and local chassis
- Integrated x86 and POWER servers and storage and network management
- Virtualization management (VMControl)
- Energy management (Active Energy Manager)
- Upward integration to existing an IBM Systems Director or Tivoli environment

The FSM is considered a hardware appliance with a specific hardware configuration and preinstalled software stack. The appliance concept is similar to the Hardware Management Console and the Systems Director Management Console in POWER systems environments. However, FSM expands the capability of these products.

The hardware platform for FSM, although based on a Intel compute node, is not interchangeable with any other compute node. A unique expansion card, not available on other compute nodes, allows the software stack to communicate on the private management network.

The FSM is available in two editions: IBM Flex System Manager and IBM Flex System Manager Advanced.

The IBM Flex System Manager base feature set offers the following functionality:

- Support up to four managed chassis
- Support up to 5,000 managed elements
- Auto-discovery of managed elements
- Overall health status
- Monitoring and availability
- Hardware management
- Security management
- Administration
- Network management (Network Control)
- Storage management (Storage Control)
- Virtual machine lifecycle management (VMControl Express)

The IBM Flex System Manager advanced feature set offers all the capabilities of the base feature set plus:

- Image management (VMControl Standard)
- Pool management (VMControl Enterprise)

FSM management: FSM can discover and manage POWER based resources outside of the IBM Flex System environment, but this feature is not supported at this time.

3.4 Power supplies

A minimum of two and a maximum of six power supplies can be installed in the Enterprise Chassis (Figure 3-5). All power supply modules are combined into a single power domain in the chassis, which distributes power to each of the compute nodes and I/O modules through the Enterprise Chassis midplane.

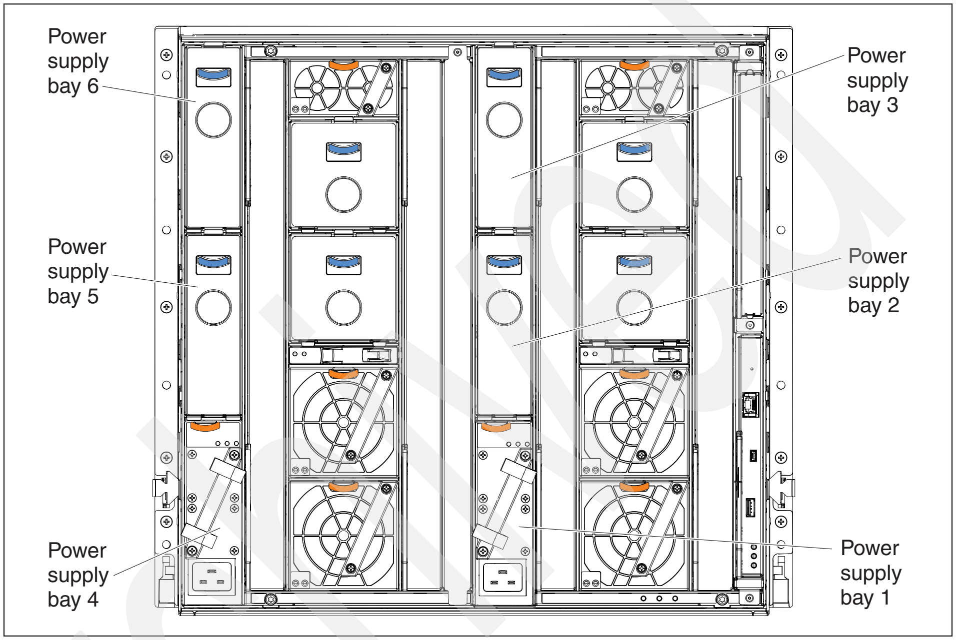

Figure 3-5 Enterprise Chassis power supply locations

The power supplies are 80 PLUS Platinum certified and are rated at 2500 W output rated at 200 VAC, with oversubscription to 3538 W output at 200 VAC. A C20 socket is provided for connection to a power cable, such as a C19-C20.

The 80 PLUS Platinum standard is a performance specification for power supplies used in servers and computers. To meet this standard, the power supply must have an energy efficiency rating of 90% or greater at 20% of rated load, 94% or greater at 50% of rated load, and 91% or greater at 100% of rated load, with a power factor of 0.9 or greater. Further information about the 80 PLUS Platinum standard is at the following website:

https://www.80PLUS.org

The Enterprise Chassis allows configurations of power policies to give N+N or N+1 redundancy.

Tip: N+1 in this context means a single backup device for N number of devices. Any component can replace any other component, but only one time.

N+N means that there are N backup devices for N devices, where N number of devices can fail and each has a backup.

The redundancy options are configured from the Chassis Management Module and can be changed nondisruptively. The five policies are shown in Table 3-2.

Table 3-2 Chassis power management policies

Power management policy	Function
Basic	Allows chassis to fully use power supplies (no N+N or N+1 redundancy).
Redundancy	Single power supply redundancy, and no blade throttling (N+1 setting).
Redundancy with throttling	Single power supply redundancy. Blades can be throttled to stay within the available power. This setting provides higher power availability over simple redundancy (N+1 setting).
AC power source redundancy	Maximum power available, limited to one-half of the installed number of power supplies (N+N setting).
AC power source redundancy with throttling	Maximum power available, limited to one-half of the installed number of power supplies. Blades can be throttled to stay within available power. This setting provides higher power availability compared with simple AC power source redundancy (N+N setting).

Throttling: Blade throttling in the simplest terms is an IBM EnergyScale™ feature of POWER architecture that allows the processor frequency to be varied to reduce power requirements.

In addition to the redundancy settings, a power limiting and capping policy can be enabled by the Chassis Management Module to limit the total amount of power that a chassis requires.

3.5 Cooling

On the topic of Enterprise Chassis cooling, the flow of air in the Enterprise Chassis follows a front to back cooling path, where cool air is drawn in at the front of the chassis and warm air is exhausted to the rear. Air movement is controlled by hot-swappable fan modules in the rear of the chassis and a series of internal dampers.

The cooling is scaled up as required, based upon the number of nodes installed. The number of cooling fan modules required for a number of nodes is described in Table 3-3 on page 61.

Chassis cooling is adaptive and is node-based rather than chassis-based. Inputs into the cooling algorithm are determined from:

- Node configurations
- Power monitor circuits
- Component temperatures
- Ambient temperature

With these inputs, each fan module has greater independent granularity in fan speed control. This results in lower airflow volume (CFM) and lower cooling energy spent at the chassis level for any configuration and workload.

Figure 3-6 shows the fan modules' locations.

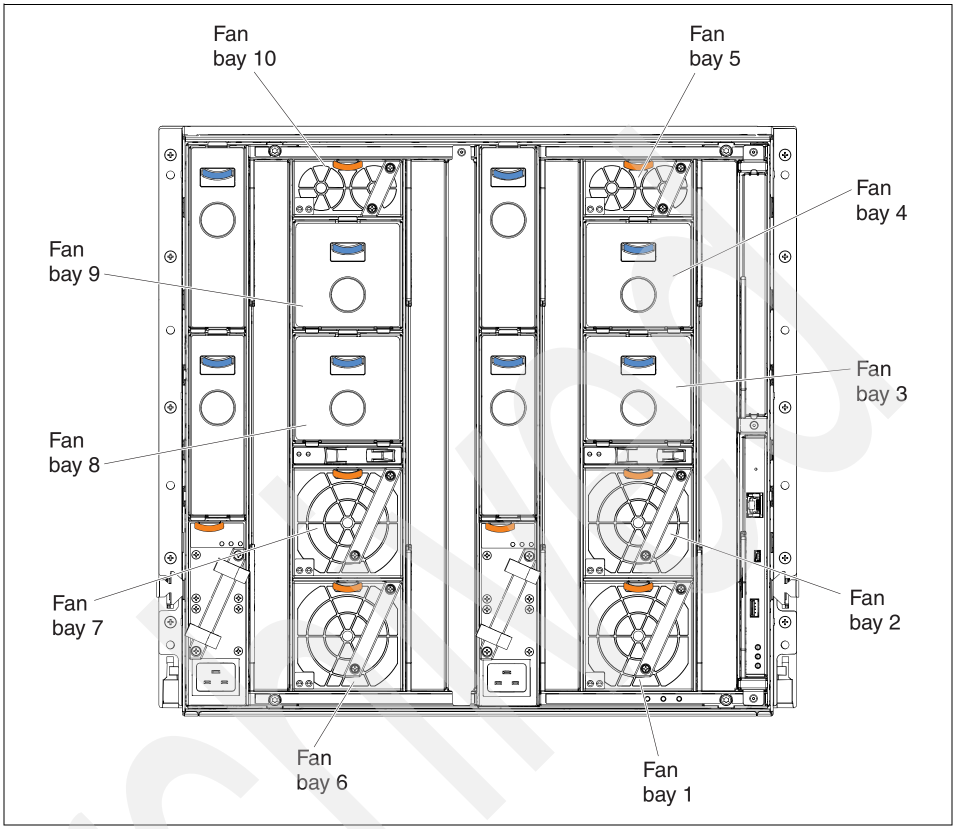

Figure 3-6 Enterprise Chassis fan module locations

3.5.1 Node cooling

There are two compute node cooling zones: zone 1 on the right side of the chassis, and zone 2 on the left side of the chassis (both viewed from the rear). The chassis can contain up to eight 80 mm fan modules across the two zones. Four 80 mm fan modules are included in the base configuration for node cooling. Additional fan modules are added in pairs across the two zones.

Figure 3-7 shows the node cooling zones and fan module locations.

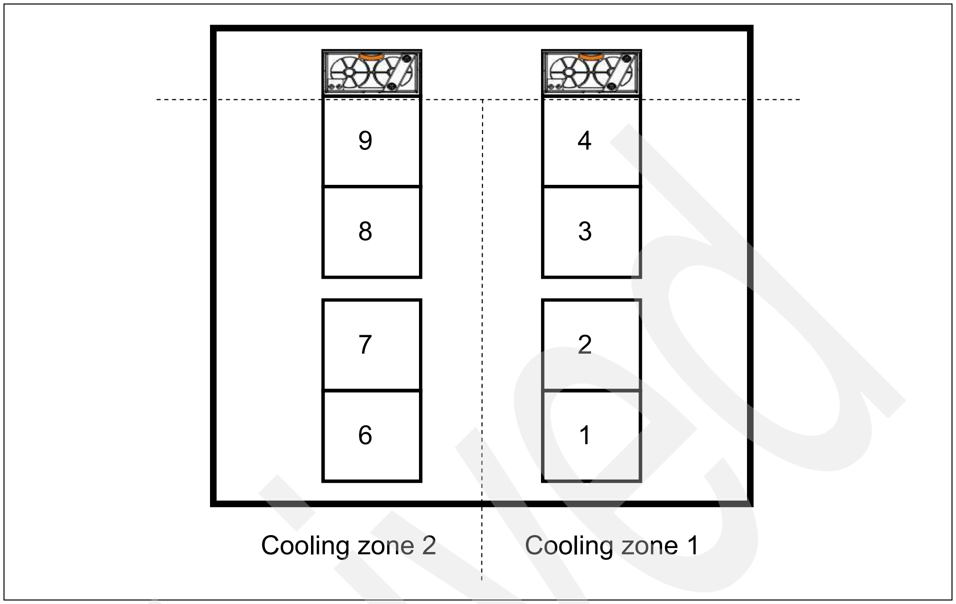

Figure 3-7 Enterprise Chassis node cooling zones and fan module locations

When a node is not inserted in a bay, an airflow damper closes in the midplane to prevent air from being drawn through the unpopulated bay. By inserting a node into a bay, the damper is opened, thus allowing cooling of the node in that bay.

Table 3-3 shows the relationship between the number of fan modules and the number of nodes supported.

Table 3-3 Fan module options and numbers of supported nodes

Fan module option	Total number of fan modules	Total number of nodes supported
Base	4	4
First option	6	8
Second option	8	14

Chassis area: The chassis area for the node is effectively one large chamber. The nodes can be placed in any slot; however, preferred practices indicate that the nodes need to be placed as close together as possible to be inline with the fan modules.

3.5.2 Switch and Chassis Management Module cooling

There are two additional cooling zones for the I/O switch bays. These zones, zones 3 and 4, are on the right and left side of the bays, as viewed from the rear of the chassis. Cooling zones 3 and 4 are serviced by 40 mm fan modules that are included in the base configuration and cool the four available I/O switch bays.

Upon hot-swap removal of a 40 mm fan module, an opening is exposed to the 80 mm fan modules located below, and a back flow damper in the fan bay closes. The backflow damper prevents hot air from entering the system from the rear of the chassis. During the time when the fan module is being replaced, the 80 mm fan modules cool the switch modules and the Chassis Management Module. Figure 3-8 shows cooling zones 3 and 4 that service the I/O switch modules.

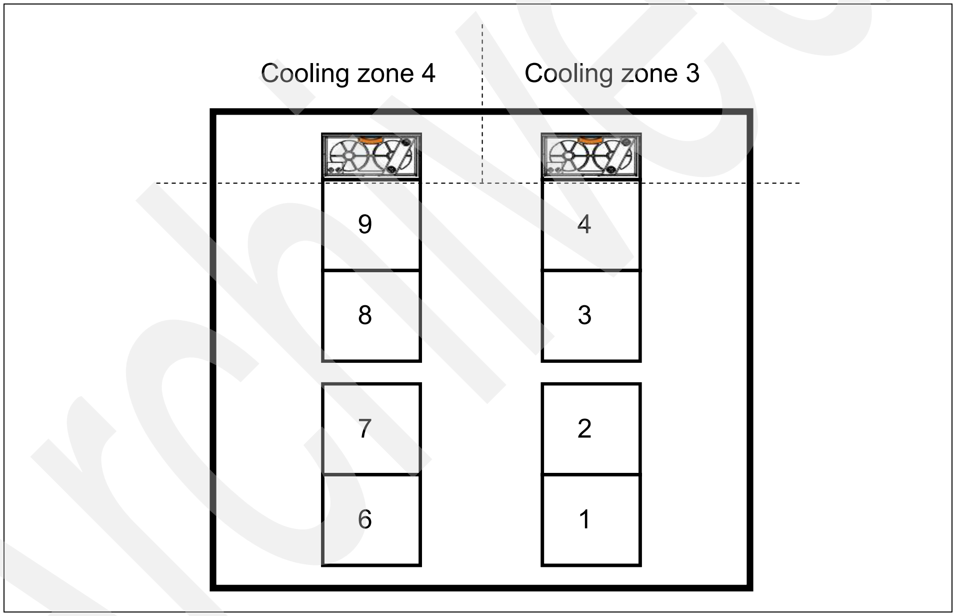

Figure 3-8 Enterprise Chassis switch module and Chassis Management Module cooling zones 3 and 4

3.5.3 Power supply cooling

The power supply modules have two integrated 40 mm fans. Installation or replacement of a power supply and fans is done as a single unit.

The integral power supply fans are not dependent upon the power supply being functional. Rather, they are powered independently from the midplane.

4

Product information and technology

The IBM Flex System p260, p460, and p24L Compute Nodes are based on IBM POWER architecture technologies. These compute nodes run in IBM Flex System Enterprise Chassis units to provide a high-density, high-performance compute node environment, using advanced processing technology.

In this chapter, we describe the server offerings and the technology used in their implementation. We cover the following topics in this chapter:

- Overview
- Front panel
- Chassis support
- System architecture
- IBM POWER7 processor
- Memory subsystem
- Active Memory Expansion
- Storage
- I/O adapters
- System management
- Integrated features
- IBM EnergyScale
- Operating system support

- Warranty and maintenance agreements
- Software support and remote technical support

4.1 Overview

The Power Systems compute nodes for IBM Flex System have three variations tailored to your business needs. They are shown in Figure 4-1.

- IBM Flex System p24L Compute Node: A half-wide compute node
- IBM Flex System p260 Compute Node: A half-wide compute node
- IBM Flex System p460 Compute Node: A full-wide compute node

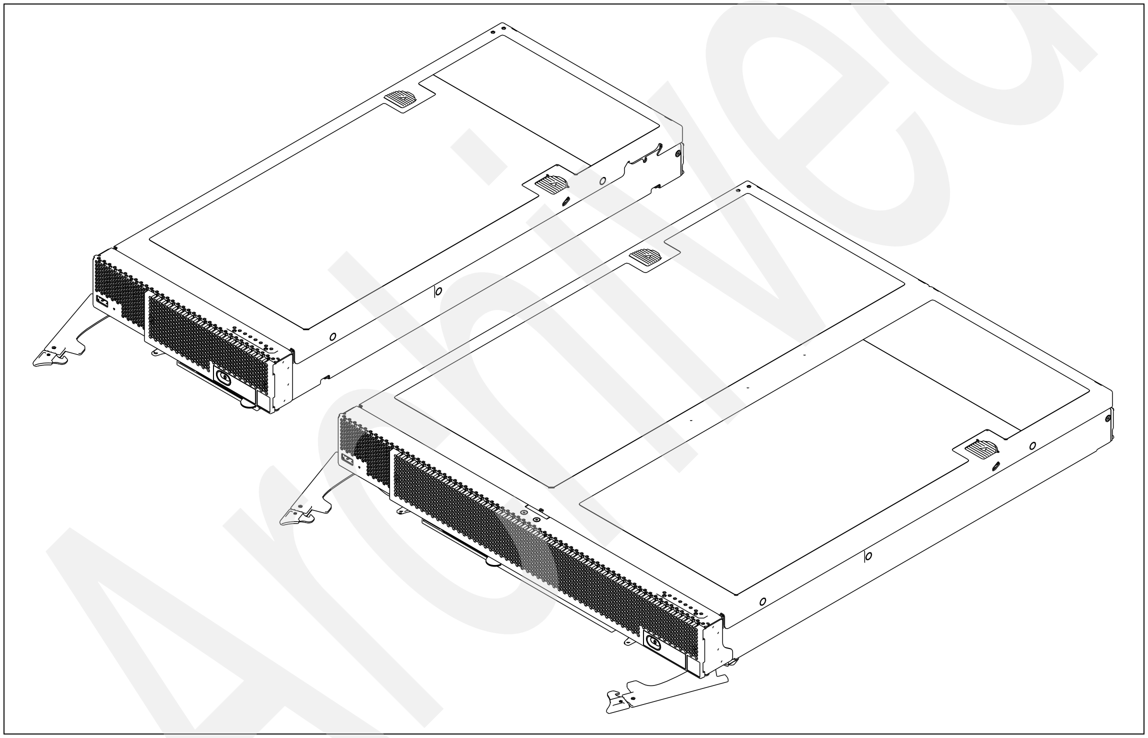

Figure 4-1 POWER7 based compute nodes - The IBM Flex System p260 Compute Node and IBM Flex System p24L Compute Node (left) and the IBM Flex System p460 Compute Node (right)

4.1.1 IBM Flex System p260 Compute Node

IBM Flex System p260 Compute Node, 7895-22X, is a half-wide, Power Systems compute node with two POWER7 processor sockets, 16 memory slots, two I/O adapter slots, and an option for up to two internal drives for local storage.

The IBM Flex System p260 Compute Node has the following features:

- Two processors with up to 16 POWER7 processing cores
- Sixteen DDR3 memory DIMM slots that support IBM Active Memory Expansion
- Supports Very Low Profile (VLP) and Low Profile (LP) DIMMs
- Two P7IOC I/O hubs
- A RAID-capable SAS controller that supports up to two solid-state drives (SSDs) or hard disk drives (HDDs)
- Two I/O adapter slots
- Flexible Support Processor (FSP)
- System management alerts
- IBM Light Path Diagnostics
- USB 2.0 port
- IBM EnergyScale technology

Figure 4-2 shows the system board layout of the IBM Flex System p260 Compute Node.

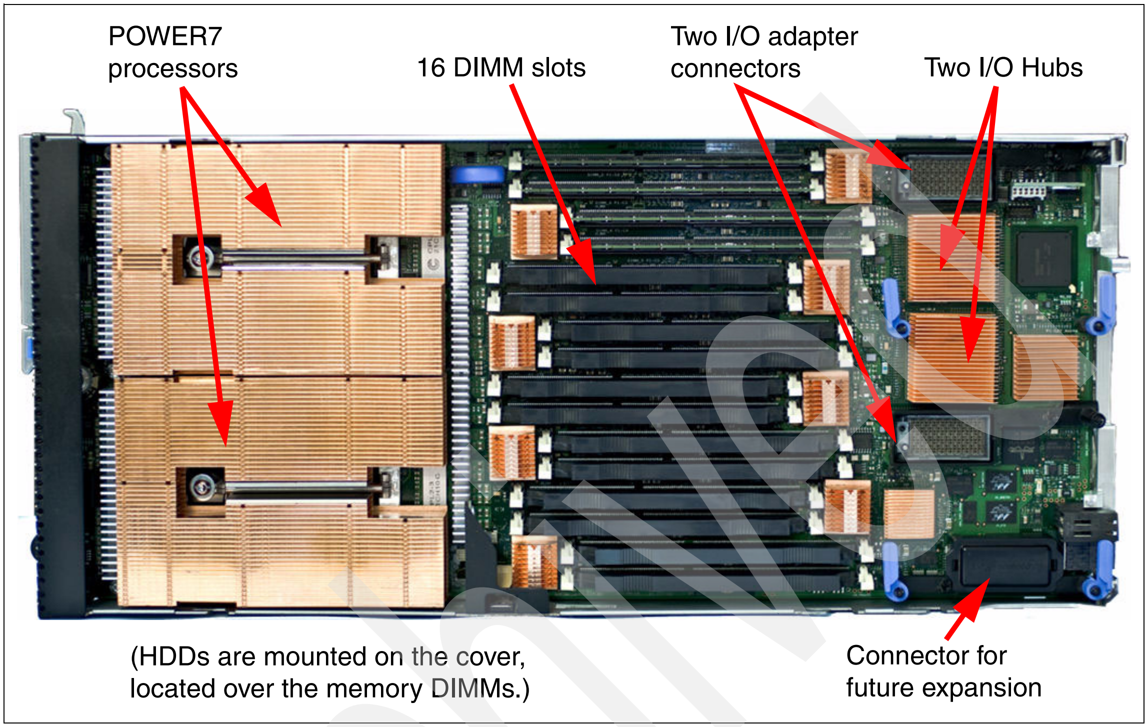

Figure 4-2 System board layout of the IBM Flex System p260 Compute Node

4.1.2 IBM Flex System p460 Compute Node

The IBM Flex System p460 Compute Node, 7895-42X, is a full-wide, Power Systems compute node with four POWER7 processor sockets, 32 memory slots, four I/O adapter slots, and an option for up to two internal drives for local storage.

The IBM Flex System p460 Compute Node has the following features:

- Four processors with up to 32 POWER7 processing cores
- Thirty-two DDR3 memory DIMM slots that support IBM Active Memory™ Expansion
- Supports Very Low Profile (VLP) and Low Profile (LP) DIMMs
- Four P7IOC I/O hubs
- RAID-capable SAS controller that support up to two SSD or HDD drives
- Four I/O adapter slots
- Flexible Support Processor (FSP)
- System management alerts

- IBM Light Path Diagnostics
- USB 2.0 port
- IBM EnergyScale technology

Figure 4-3 shows the system board layout of the IBM Flex System p460 Compute Node.

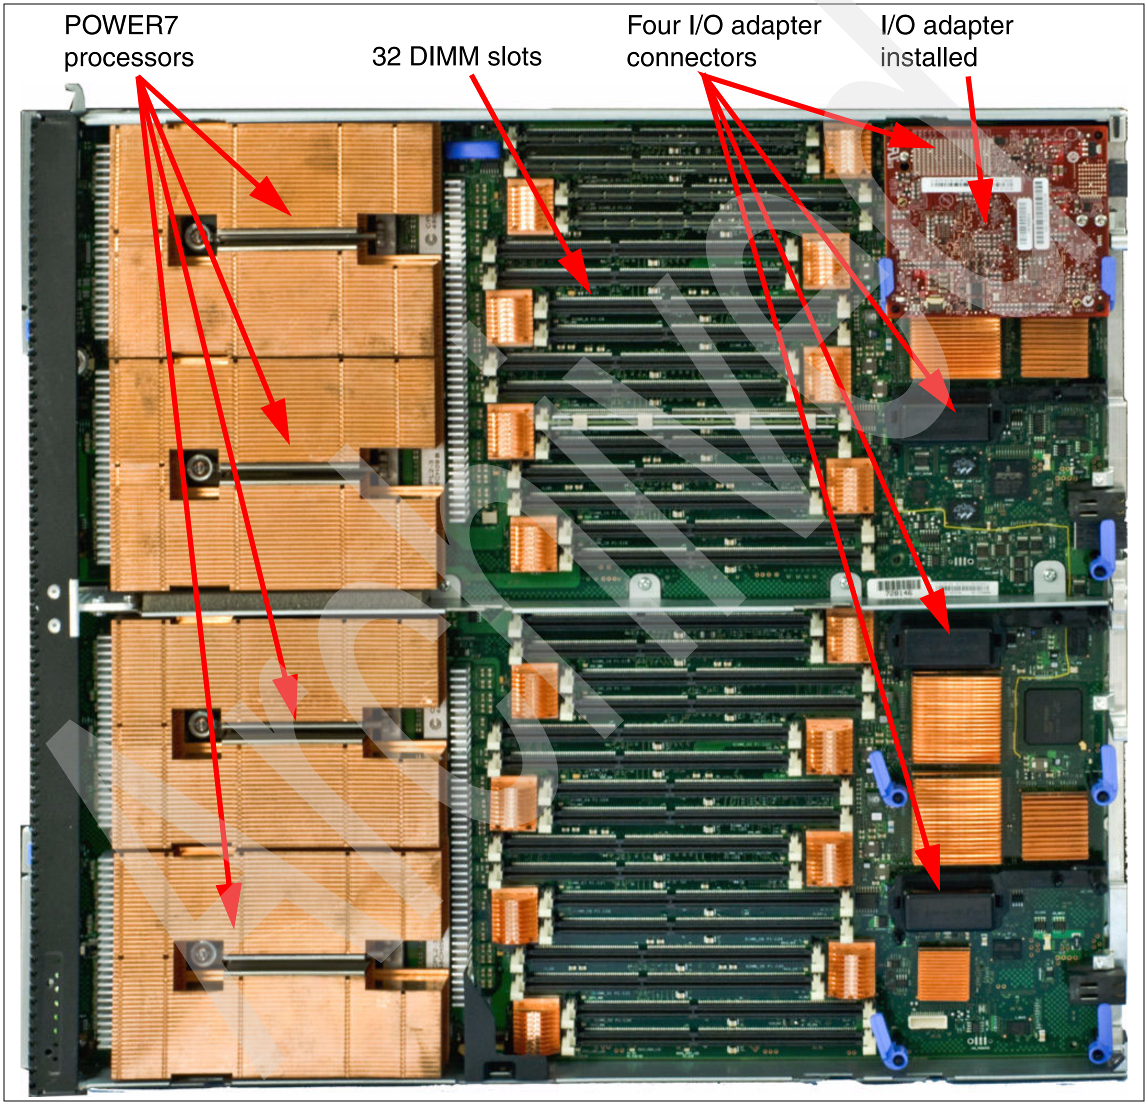

Figure 4-3 System board layout of the IBM Flex System p460 Compute Node

4.1.3 IBM Flex System p24L Compute Node

The IBM Flex System p24L Compute Node shares several similarities to the IBM Flex System p260 Compute Node in that it is a half-wide, Power Systems compute node with two POWER7 processor sockets, 16 memory slots, two I/O adapter slots, and an option for up to two internal drives for local storage. The IBM Flex System p24L Compute Node is optimized for low-cost Linux on Power Systems Servers installations.

The IBM Flex System p24L Compute Node has the following features:

- Up to 16 POWER7 processing cores
- Sixteen DDR3 memory DIMM slots that support Active Memory Expansion
- Supports Very Low Profile (VLP) and Low Profile (LP) DIMMs
- Two P7IOC I/O hubs
- RAID-compatible SAS controller that support up to two SSD or HDD drives
- Two I/O adapter slots
- Flexible Support Processor (FSP)
- System management alerts
- IBM Light Path Diagnostics
- USB 2.0 port
- IBM EnergyScale technology

The system board layout for the IBM Flex System p24L Compute Node is identical to the IBM Flex System p260 Compute Node and is shown in Figure 4-2 on page 66.

4.2 Front panel

The front panel of Power Systems compute nodes has the following common elements, as shown in Figure 4-4:

- One USB 2.0 port
- Power button and light path, light-emitting diode (LED) (green)
- Location LED (blue)
- Information LED (amber)
- Fault LED (amber)

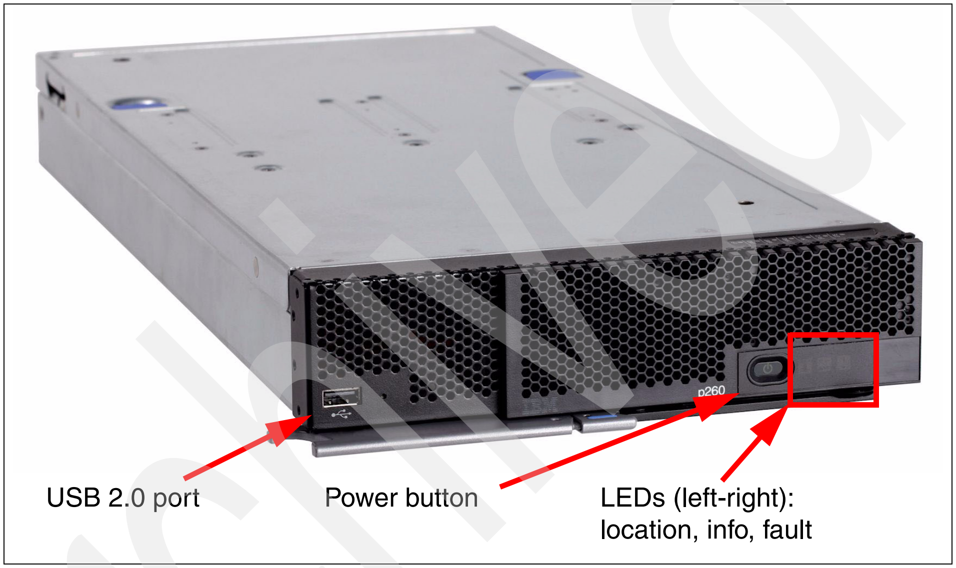

Figure 4-4 Front panel of the IBM Flex System p260 Compute Node

The USB port on the front of the Power Systems compute nodes is useful for various tasks, including out-of-band diagnostic tests, hardware RAID setup, operating system access to data on removable media, and local OS installation. It might be helpful to obtain a USB optical (CD or DVD) drive for these purposes, in case the need arises.

The front panel of the p460 is similar and is shown in Figure 1-3 on page 12.

Tip: There is no optical drive in the IBM Flex System Enterprise Chassis.

4.2.1 Light path diagnostic LED panel

The power button on the front of the server (Figure 4-4 on page 69) has two functions:

- When the system is fully installed in the chassis: Use this button to power the system on and off.
- When the system is removed from the chassis: Use this button to illuminate the light path diagnostic panel on the top of the front bezel, as shown in Figure 4-5.

Figure 4-5 Light path diagnostic panel

The LEDs on the light path panel indicate the following LEDs:

- LP: Light Path panel power indicator
- S BRD: System board LED (might indicate trouble with the processor or memory as well)
- MGMT: Flexible Support Processor (or management card) LED
- D BRD: Drive (or Direct Access Storage Device (DASD)) board LED
- DRV 1: Drive 1 LED (SSD 1 or HDD 1)
- DRV 2: Drive 2 LED (SSD 2 or HDD 2)
- ETE: Sidecar connector LED (not present on the IBM Flex System p460 Compute Node)

If problems occur, you can use the light path diagnostics LEDs to identify the subsystem involved. To illuminate the LEDs with the compute node removed, press the power button on the front panel. This action temporarily illuminates the LEDs of the troubled subsystem to direct troubleshooting efforts towards a resolution.

Typically, an administrator has already obtained this information from the IBM Flex System Manager or Chassis Management Module before removing the node, but having the LEDs helps with repairs and troubleshooting if on-site assistance is needed.

For more information about the front panel and LEDs, see the *IBM Flex System p260 and p460 Compute Node Installation and Service Guide*, available from:

http://www.ibm.com/support

4.2.2 Labeling

IBM Flex System offers several options for labeling your server inventory to track your machines. It is important to not put stickers on the front of the server across the bezel’s grating, as this action inhibits proper airflow to the machine.

We provide the following labeling features:

- Vital Product Data (VPD) sticker

  On the front bezel of the server is a vital product data sticker that lists the following information about the machine, as shown in Figure 4-6:

  – Machine type
  – Model
  – Serial number

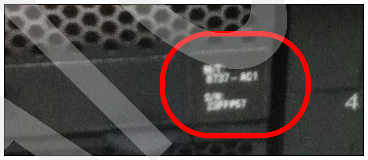

Figure 4-6 Vital Product Data sticker

- Node bay labeling on IBM Flex System Enterprise Chassis

  Each bay of the IBM Flex System Enterprise Chassis has space for a label to be affixed to identify or provide information about each Power Systems compute node, as shown in Figure 4-7.

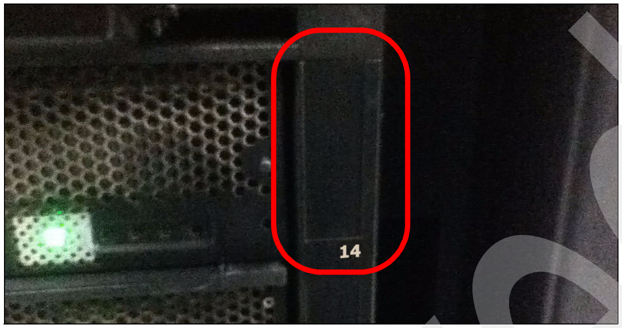

Figure 4-7 Chassis bay labeling

- Pull-out labeling

  Each Power Systems compute node has two pull-out tabs that can also accommodate labeling for the server. The benefit of using these tabs is that they are affixed to the node itself rather than the chassis, as shown in Figure 4-8.

Figure 4-8 Pull-out labeling on the Power Systems compute node

4.3 Chassis support

The Power Systems compute nodes can be used only in the IBM Flex System Enterprise Chassis. They do not fit in the previous IBM modular systems, such as IBM iDataPlex or IBM BladeCenter.

There is no onboard video capability in the Power Systems compute nodes. The machines are designed to use Serial Over LAN (SOL) or the IBM Flex System Manager (FSM).

For more information about the IBM Flex System Enterprise Chassis, see Chapter 3, "Introduction to IBM Flex System" on page 47. For information about FSM, see 6.4, "IBM Flex System Manager" on page 192.

4.4 System architecture

This section covers the system architecture and layout of Power Systems compute nodes.

The overall system architecture for the p260 and p24L is shown in Figure 4-9.

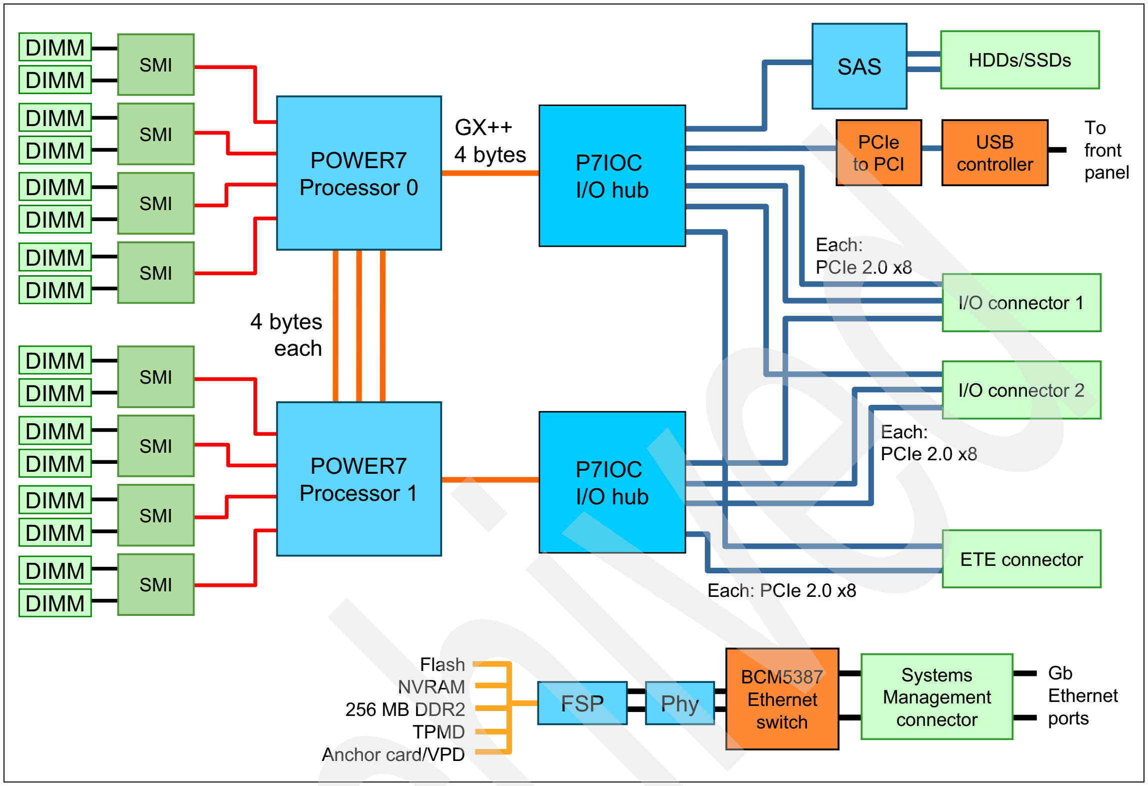

Figure 4-9 IBM Flex System p260 Compute Node block diagram

In this diagram, you can see the two processor slots, with eight memory slots for each processor. Each processor is connected to a P7IOC I/O hub, which connects to the I/O subsystem (I/O adapters and local storage). At the bottom, you can see a representation of the service processor (FSP) architecture.

The IBM Flex System p460 Compute Node shares many of the same components as the IBM Flex System p260 Compute Node. The IBM Flex System p460 Compute Node is a full-wide node, and adds additional processors and memory along with two more adapter slots. It has the same local storage options as the IBM Flex System p260 Compute Node.

The IBM Flex System p460 Compute Node system architecture is shown in Figure 4-10.

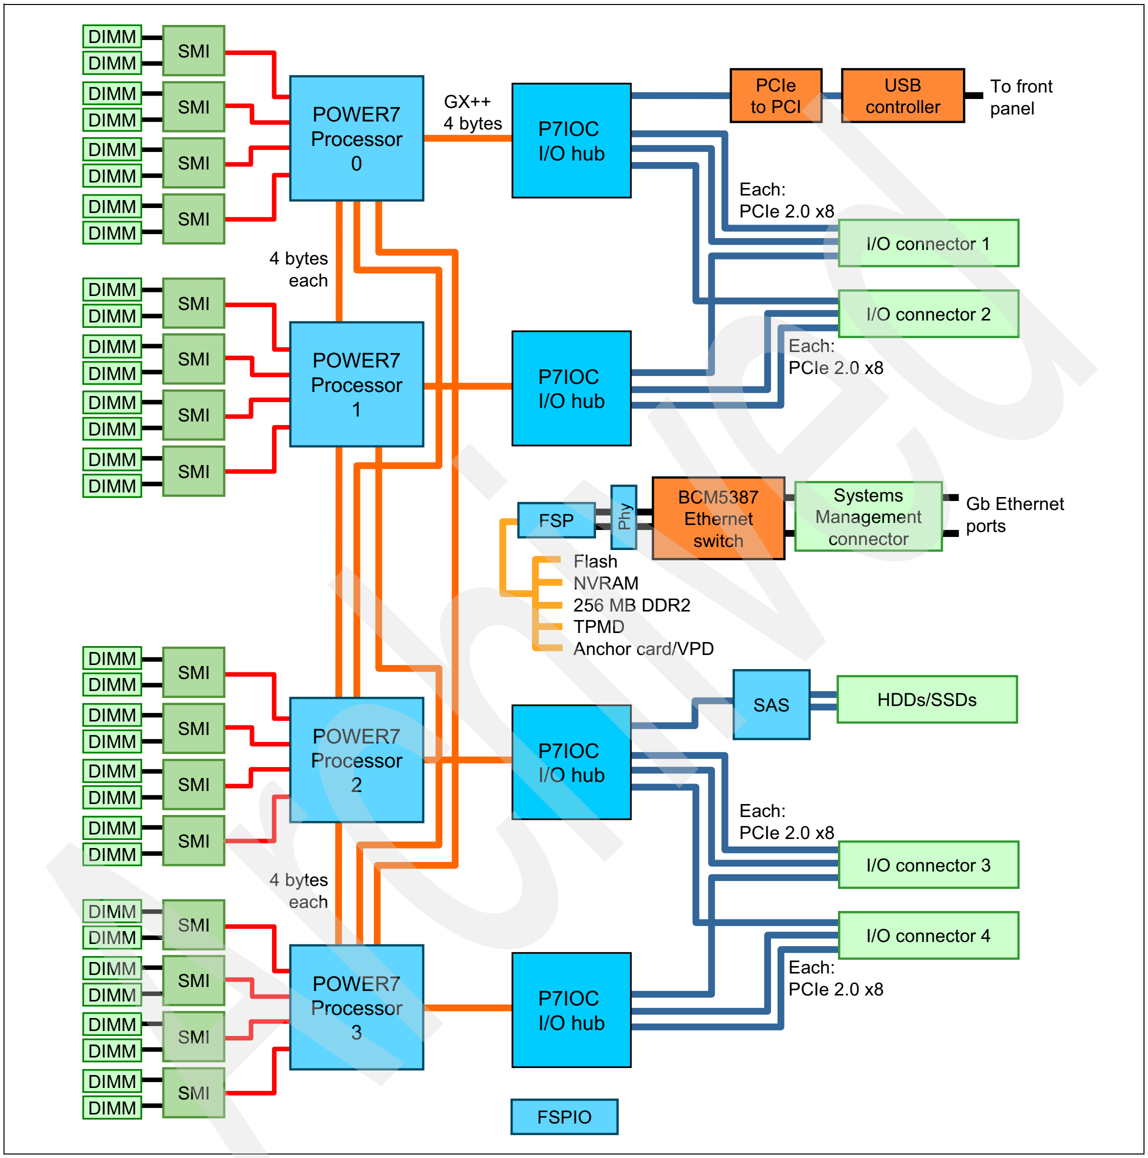

Figure 4-10 IBM Flex System p460 Compute Node block diagram

The four processors in the IBM Flex System p460 Compute Node are connected in a cross-bar formation, as shown in Figure 4-11.

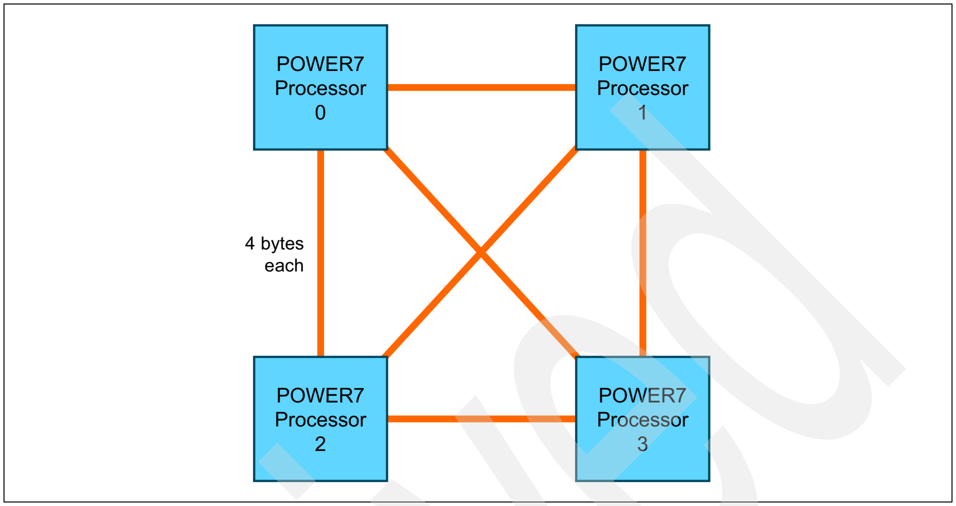

Figure 4-11 IBM Flex System p460 Compute Node processor connectivity

4.5 IBM POWER7 processor

The IBM POWER7 processor represents a leap forward in technology and associated computing capability. The multicore architecture of the POWER7 processor is matched with a wide range of related technologies to deliver leading throughput, efficiency, scalability, and reliability, availability, and serviceability (RAS).

Although the processor is an important component in servers, many elements and facilities must be balanced across a server to deliver maximum throughput. As with previous generations of systems based on POWER processors, the design philosophy for POWER7 processor-based systems is one of system-wide balance in which the POWER7 processor plays an important role.

4.5.1 Processor options for Power Systems compute nodes

Table 4-1 defines the processor options for the Power Systems compute nodes.

Table 4-1 Processor options

Feature code	Cores per POWER7 processor	Number of POWER7 processors	Total cores	Core frequency	L3 cache size per POWE7 processor
IBM Flex System p260 Compute Node					
EPR1	4	2	8	3.3 GHz	16 MB
EPR3	8	2	16	3.2 GHz	32 MB
EPR5	8	2	16	3.55 GHz	32 MB
IBM Flex System p460 Compute Node					
EPR2	4	4	16	3.3 GHz	16 MB
EPR4	8	4	32	3.2 GHz	32 MB
EPR6	8	4	32	3.55 GHz	32 MB
IBM Flex System p24L Compute Node					
EPR8	8	2	16	3.2 GHz	32 MB
EPR9	8	2	16	3.55 GHz	32 MB
EPR7	6	2	12	3.7 GHz	24 MB

4.5.2 Unconfiguring

You can order p260 or p460 with Feature Code #2319, which reduces the number of active processor cores in the compute node to reduce software licensing costs.

Feature Code #2319 is listed in Table 4-2.

Table 4-2 Reconfiguration of cores

Feature code	Description	Minimum	Maximum
2319	Factory Deconfiguration of one core	0	1 less than the total number of cores (for example, for EPR5, the maximum is 7)

This core deconfiguration feature can also be updated after installation by using the field core override option.

The field core override option specifies the number of functional cores that are active in the compute node. The field core override option provides the capability to increase or decrease the number of active processor cores in the compute node. The compute node firmware sets the number of active processor cores to the entered value. The value takes effect when the compute node is rebooted. The field core override value can be changed only when the compute node is powered off.

The advanced system management interface (ASMI) is used to change the number of functional override cores in the compute node. For more information, go to the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.psm.hosts.doc/dpsm_managing_hosts_launch_asm.html

For detailed information about the field core override feature, go to the following website:

http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hby/fieldcore.htm

You can find more related information at the following website:

http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hby/viewprocconfig.htm

System maintenance: The configuration information about this feature is stored in the anchor card (see 4.10.3, “Anchor card” on page 110) and the system board.

If the system board is replaced, transfer the anchor card from the old system board to the new system board. If the anchor card is replaced, the information is transferred from the system board to the new anchor card upon the next boot.

If *both* the system board *and* the anchor card are replaced, then the field core override option must be used to reset the core count back to the previous value.

4.5.3 Architecture

IBM uses innovative methods to achieve the required levels of throughput and bandwidth. Areas of innovation for the POWER7 processor and POWER7 processor-based systems include (but are not limited to) the following elements:

- On-chip L3 cache implemented in embedded dynamic random access memory (eDRAM)
- Cache hierarchy and component innovation
- Advances in memory subsystem
- Advances in off-chip signaling

The superscalar POWER7 processor design also provides various capabilities, including:

- Binary compatibility with the prior generation of POWER processors
- Support for PowerVM virtualization capabilities, including PowerVM Live Partition Mobility to and from IBM POWER6® and IBM POWER6+™ processor-based systems

Figure 4-12 shows the POWER7 processor die layout with major areas identified: eight POWER7 processor cores, L2 cache, L3 cache and chip power bus interconnect, SMP links, GX++ interface, and memory controller.

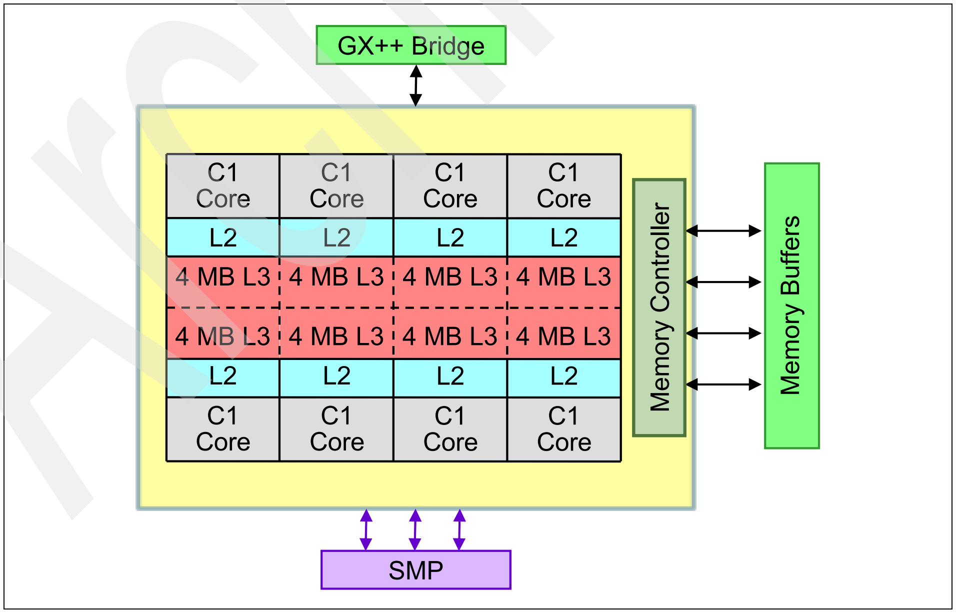

Figure 4-12 POWER7 processor architecture

POWER7 processor overview

The POWER7 processor chip is fabricated with the IBM 45 nm silicon-on-insulator technology, using copper interconnects, and uses an on-chip L3 cache with eDRAM.

The POWER7 processor chip is 567 mm^2 and is built using 1,200,000,000 components (transistors). Eight processor cores are on the chip, each with 12 execution units, 256 KB of L2 cache, and access to up to 32 MB of shared on-chip L3 cache.

For memory access, the POWER7 processor includes an integrated Double Data Rate 3 (DDR3) memory controller, each with four memory channels. To scale effectively, the POWER7 processor uses a combination of local and global high-bandwidth SMP links.

Table 4-3 summarizes the technology characteristics of the POWER7 processor.

Table 4-3 Summary of POWER7 processor technology

Technology	POWER7 processor
Die size	567 mm^2
Fabrication technology	► 45 nm lithography ► Copper interconnect ► Silicon-on-insulator ► eDRAM
Components	1,200,000,000 components (transistors) offering the equivalent function of 2,700,000,000 (For more details, see "On-chip L3 intelligent cache" on page 85)
Processor cores	4
Max execution threads core/chip	4/32
L2 cache per core/per chip	256 KB / 2 MB
On-chip L3 cache per core per chip	4 MB / 32 MB
DDR3 memory controllers	One per processor
Compatibility	Compatible with prior generations of the POWER processor

POWER7 processor core

Each POWER7 processor core implements aggressive out of order (OoO) instruction execution to drive high efficiency in the use of available execution paths. The POWER7 processor has an instruction sequence unit that can dispatch up to six instructions per cycle to a set of queues. Up to eight instructions per cycle can be issued to the instruction execution units. The POWER7 processor has a set of 12 execution units, as follows:

- Two fixed-point units
- Two load store units
- Four double precision floating point units
- One vector unit
- One branch unit
- One condition register unit
- One decimal floating point unit

The caches that are tightly coupled to each POWER7 processor core are as follows:

- Instruction cache: 32 KB
- Data cache: 32 KB
- L2 cache: 256 KB, implemented in fast SRAM
- L3 cache: 4 MB eDRAM

Simultaneous multithreading

An enhancement in the POWER7 processor is the addition of simultaneous multithreading (SMT) mode, known as SMT4 mode, which enables four instruction threads to run simultaneously in each POWER7 processor core. Thus, the instruction thread execution modes of the POWER7 processor are as follows:

- SMT1: Single instruction execution thread per core
- SMT2: Two instruction execution threads per core
- SMT4: Four instruction execution threads per core

SMT4 mode enables the POWER7 processor to maximize the throughput of the processor core by offering an increase in processor-core efficiency. SMT4 mode is the latest step in an evolution of multithreading technologies introduced by IBM.

Figure 4-13 shows the evolution of simultaneous multithreading.

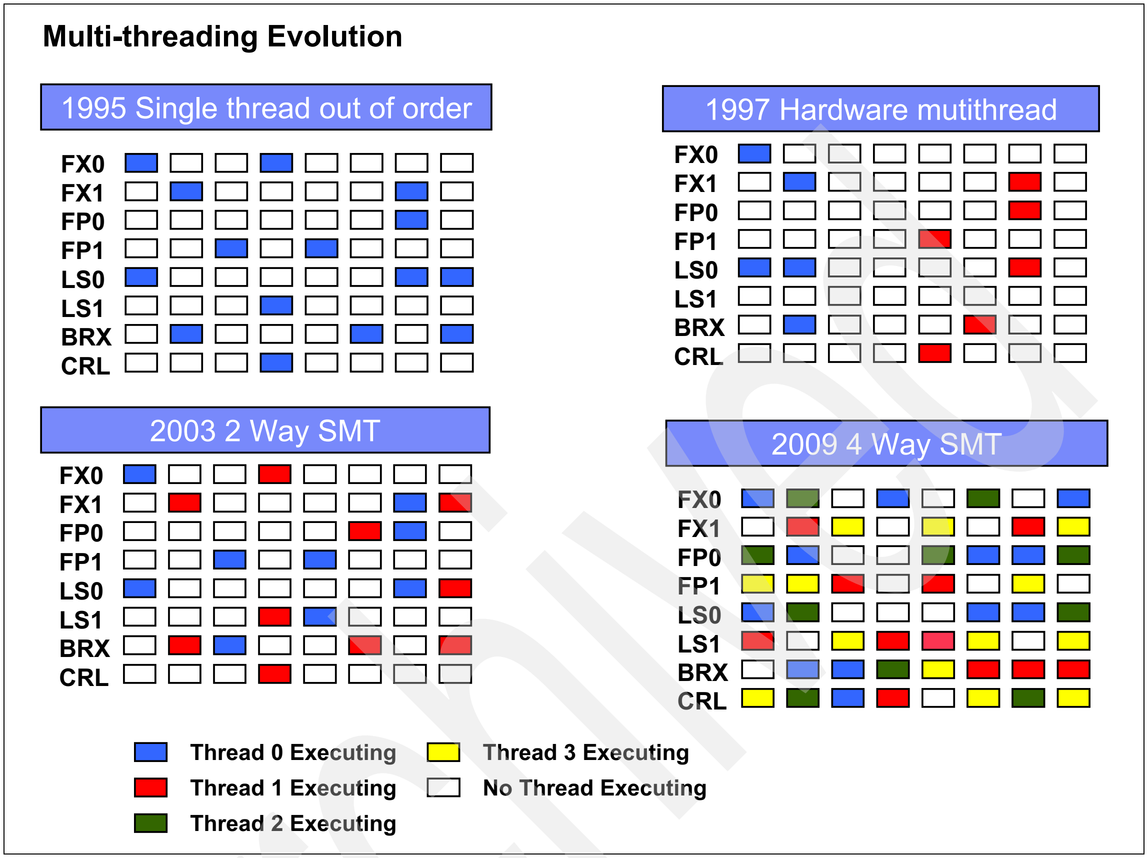

Figure 4-13 Evolution of simultaneous multithreading

The various SMT modes offered by the POWER7 processor provide flexibility, where you can select the threading technology that meets a combination of objectives, such as performance, throughput, energy use, and workload enablement.

Intelligent threads

The POWER7 processor features *intelligent threads*, which can vary based on the workload demand. The system automatically selects (or the system administrator manually selects) whether a workload benefits from dedicating as much capability as possible to a single thread of work, or if the workload benefits more from having this capability spread across two or four threads of work. With more threads, the POWER7 processor delivers more total capacity because more tasks are accomplished in parallel. With fewer threads, workloads that require fast, individual tasks get the performance they need for maximum benefit.

Memory access

The POWER7 processor chip in the compute nodes has one DDR3 memory controller enabled (the second controller is not used, as shown in Figure 4-14), with four memory channels. Each channel operates at 6.4 Gbps and can address up to 32 GB of memory. Thus, the POWER7 processor used in these compute nodes can address up to 128 GB of memory. Figure 4-14 gives a simple overview of the POWER7 processor memory access structure.

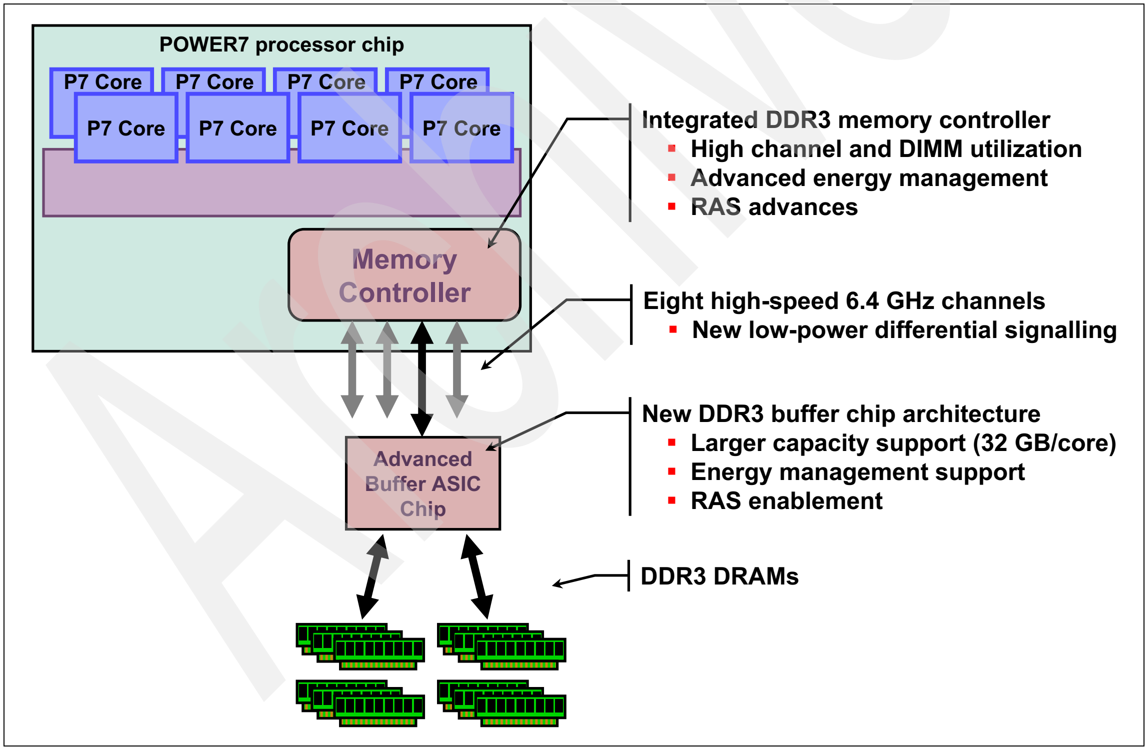

Figure 4-14 Overview of POWER7 memory access structure

Flexible POWER7 processor packaging and offerings

POWER7 processors have the unique ability to optimize to various workload types. For example, database workloads typically benefit from fast processors that handle high transaction rates at high speeds. Web workloads typically benefit more from processors with many threads that allow the breakdown of web requests into many parts and handle them in parallel. POWER7 processors have the unique ability to provide leadership performance in either case.

POWER7 processor cores

The base design for the POWER7 processor is an 8-core processor with 32 MB of on-chip L3 cache (4 MB per core). However, the architecture allows for differing numbers of processor cores to be active: four cores or six cores, in addition to the full 8-core version.

The L3 cache associated with the implementation depends on the number of active cores. For the 8-core version, this situation means that 8 x 4 = 32 MB of L3 cache is available.

Optimized for servers

The POWER7 processor forms the basis for a flexible compute platform and can be offered in a number of configurations to address various system requirements.

The POWER7 processor can be offered with a single active memory controller with four channels for servers for which higher degrees of memory parallelism are not required.

Similarly, the POWER7 processor can be offered with various SMP bus capacities appropriate to the scaling-point of particular server models.

Figure 4-15 shows the physical packaging options that are supported with POWER7 processors.

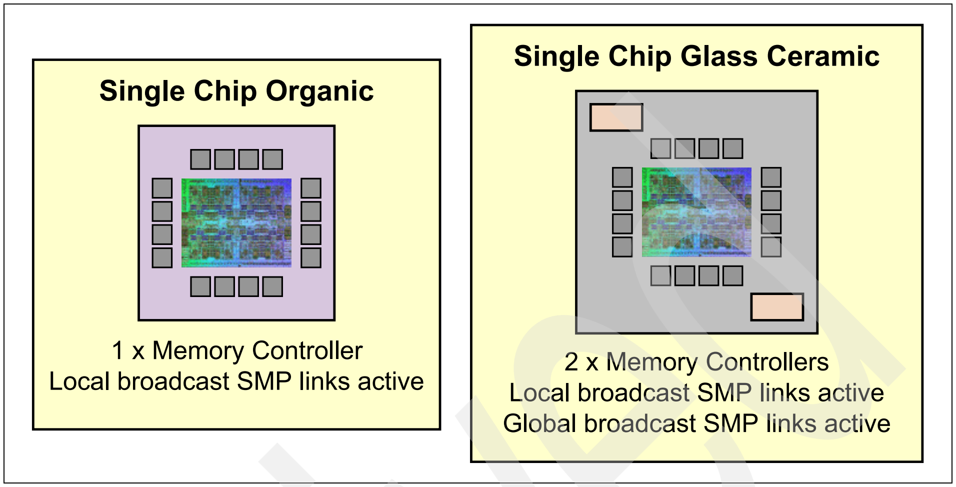

Figure 4-15 Outline of the POWER7 processor physical packaging

On-chip L3 intelligent cache

A breakthrough in material engineering and microprocessor fabrication has enabled IBM to implement the L3 cache in eDRAM and place it on the POWER7 processor die. L3 cache is critical to a balanced design, as is the ability to provide good signaling between the L3 cache and other elements of the hierarchy, such as the L2 cache or SMP interconnect.

The on-chip L3 cache is organized into separate areas with differing latency characteristics. Each processor core is associated with a Fast Local Region of L3 cache (FLR-L3), but also has access to other L3 cache regions as shared L3 cache. Additionally, each core can negotiate to use the FLR-L3 cache associated with another core, depending on reference patterns. Data can also be cloned to be stored in more than one core's FLR-L3 cache, again, depending on reference patterns. This *intelligent cache* management enables the POWER7 processor to optimize the access to L3 cache lines and minimize overall cache latencies.

Figure 4-16 shows the FLR-L3 cache regions for the cores on the POWER7 processor die.

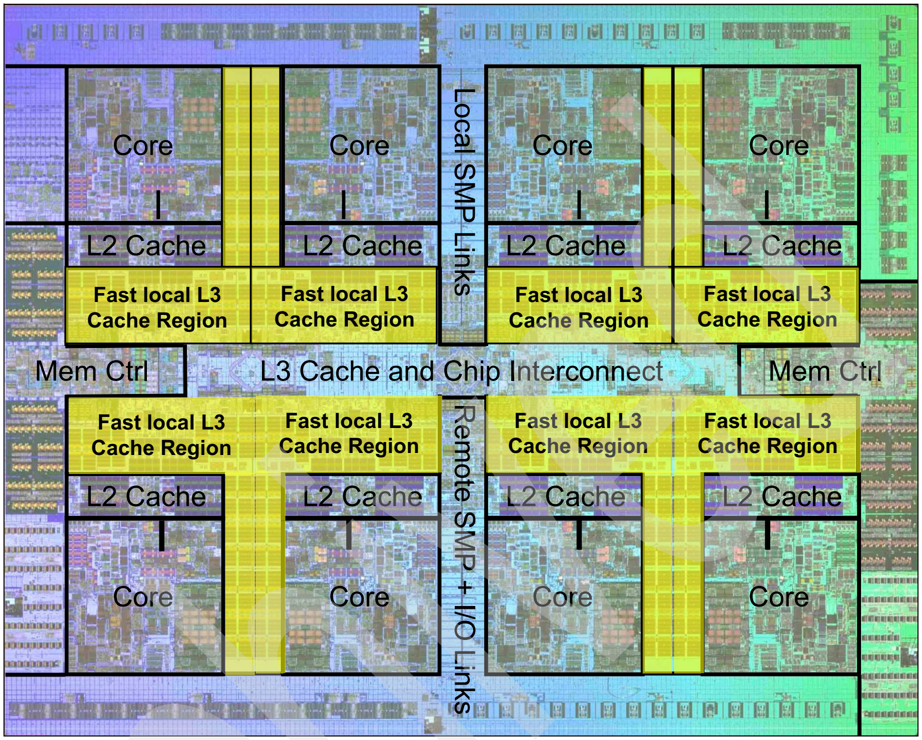

Figure 4-16 FLR-L3 cache regions on the POWER7 processor

The innovation of using eDRAM on the POWER7 processor die is significant for several reasons:

- Latency improvement

  A six-to-one latency improvement occurs by moving the L3 cache on-chip, compared to L3 accesses on an external (on-ceramic) application-specific integrated circuit (ASIC).

- Bandwidth improvement

  A 2x bandwidth improvement occurs with on-chip interconnect. Frequency and bus sizes are increased to and from each core.

- No off-chip drivers or receivers

  Removing drivers and receivers from the L3 access path lowers interface requirements, conserves energy, and lowers latency.

- Small physical footprint

  The performance of eDRAM when implemented on-chip is similar to conventional SRAM but requires far less physical space. IBM on-chip eDRAM uses only one-third of the components used in conventional SRAM, which has a minimum of six transistors to implement a 1-bit memory cell.

- Low energy consumption

  The on-chip eDRAM uses only 20% of the standby power of SRAM.

POWER7 processor and intelligent energy

Energy consumption is an important area of focus for the design of the POWER7 processor, which includes intelligent energy features that help to optimize energy usage and performance dynamically, so that the best possible balance is maintained. Intelligent energy features (such as EnergyScale) work with the IBM Flex System Manager or IBM Systems Director Active Energy Manager™ to optimize processor speed dynamically, based on thermal conditions and system usage.

For more information about the POWER7 energy management features, see *Adaptive Energy Management Features of the POWER7 Processor*, found at the following website:

http://researcher.watson.ibm.com/researcher/files/us-lefurgy/hotchips22_power7.pdf

4.6 Memory subsystem

Each POWER7 processor used in the compute nodes has an integrated memory controller. Industry standard DDR3 Registered DIMM (RDIMM) technology is used to increase reliability, speed, and density of memory subsystems.

4.6.1 Memory placement rules

The recommended memory minimum and maximum for each server is listed in Table 4-4.

Table 4-4 Recommended memory limits

Model	Minimum memory	Maximum memory
p260	8 GB	256 GB (16x 16 GB DIMMs)
p460	32 GB	512 GB (32x 16 GB DIMMs)

Model	Minimum memory	Maximum memory
p24L	24 GB	256 GB (16 x 16 GB DIMMs)

Use a minimum of 2 GB of memory per core. The functional minimum memory configuration for the machine is 4 GB (two 2 GB DIMMs), but that is not sufficient for reasonable production usage of the machine.

Low Profile and Very Low Profile form factors

One benefit of deploying IBM Flex System systems is the ability to use Low Profile (LP) memory DIMMs. This design allows for more choices to configure the machine to match your needs.

Table 4-5 lists the available memory options for the Power Systems compute nodes.

Table 4-5 Memory options

Feature code	Description	Speed	Form factor
EM04	2x 2 GB DDR3 DIMM	1066 MHz	LP
8196	2x 4 GB DDR3 DIMM	1066 MHz	VLP
8199	2x 8 GB DDR3 DIMM	1066 MHz	VLP
8145	2x 16 GB DDR3 DIMM	1066 MHz	LP

DASD/local storage option dependency on memory form factor: Because of the design of the on-cover storage connections, clients that seek to use SAS HDDs *must* use VLP DIMMs (4 GB or 8 GB). The cover is not able to close properly if VLP DIMMs and SAS hard disk drives are configured in the same system. However, solid-state drives (SSDs) and LP DIMMs can be used together.

For more information, see 4.8, “Storage” on page 95.

There are 16 buffered DIMM slots on the p260 and p24L, as shown in Figure 4-17. The p460 adds two more processors and 16 additional DIMM slots, divided evenly (eight memory slots) per processor.

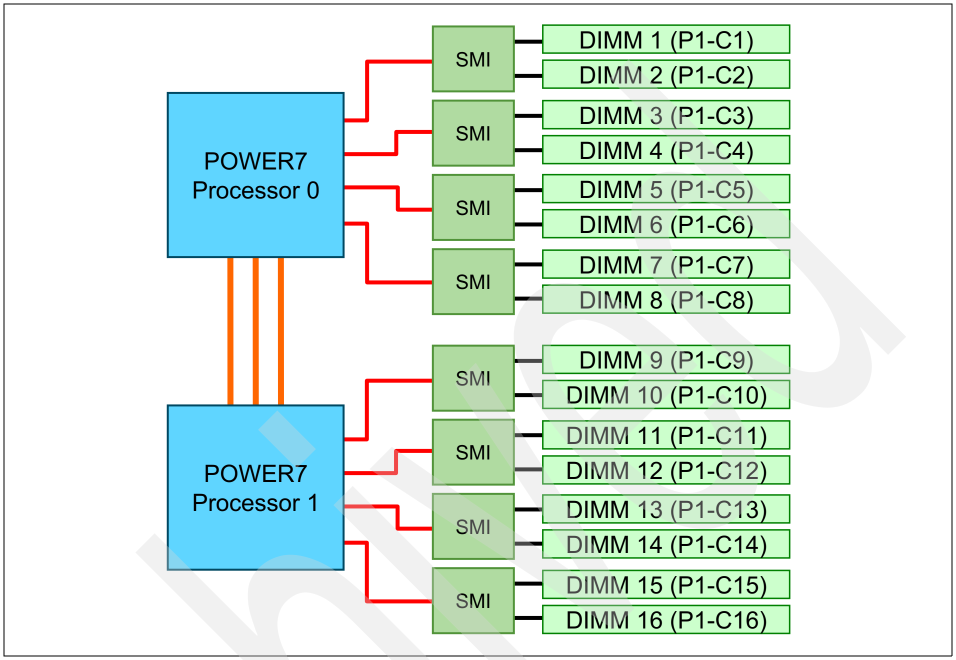

Figure 4-17 Memory DIMM topology (IBM Flex System p260 Compute Node)

The memory-placement rules are as follows:

- Install DIMM fillers in unused DIMM slots to ensure proper cooling.
- Install DIMMs in pairs. Both need to be the same size.
- Both DIMMs in a pair must be the same size, speed, type, and technology. You can mix compatible DIMMs from multiple manufacturers.
- Install only supported DIMMs, as described on the IBM ServerProven® website:

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

For the p260 and p24L, Table 4-6 shows the required placement of memory DIMMs, depending on the number of DIMMs installed.

Table 4-6 DIMM placement - p260 and p24L

Number of DIMMs	Processor 0								Processor 1							
	DIMM 1	**DIMM 2**	**DIMM 3**	**DIMM 4**	**DIMM 5**	**DIMM 6**	**DIMM 7**	**DIMM 8**	**DIMM 9**	**DIMM 10**	**DIMM 11**	**DIMM 12**	**DIMM 13**	**DIMM 14**	**DIMM 15**	**DIMM 16**
2	x			x												
4	x			x					x			x				
6	x			x	x			x	x			x				
8	x			x	x			x	x			x	x			x
10	x	x	x	x	x			x	x			x	x			x
12	x	x	x	x	x			x	x	x	x	x	x			x
14	x	x	x	x	x	x	x	x	x	x	x	x	x			x
16	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

For the IBM Flex System p460 Compute Node, Table 4-7 shows the required placement of memory DIMMs, depending on the number of DIMMs installed.

Table 4-7 DIMM placement - IBM Flex System p460 Compute Node

Number of DIMMs	Processor 0								Processor 1								Processor 2								Processor 3							
	DIMM 1	**DIMM 2**	**DIMM 3**	**DIMM 4**	**DIMM 5**	**DIMM 6**	**DIMM 7**	**DIMM 8**	**DIMM 9**	**DIMM 10**	**DIMM 11**	**DIMM 12**	**DIMM 13**	**DIMM 14**	**DIMM 15**	**DIMM 16**	**DIMM 17**	**DIMM 18**	**DIMM 19**	**DIMM 20**	**DIMM 21**	**DIMM 22**	**DIMM 23**	**DIMM 24**	**DIMM 25**	**DIMM 26**	**DIMM 27**	**DIMM 28**	**DIMM 29**	**DIMM 30**	**DIMM 31**	**DIMM 32**
2	x			x																												
4	x			x													x			x												
6	x			x					x			x					x			x												
8	x			x					x			x					x			x					x			x				
10	x			x	x			x	x			x					x			x					x			x				

Number of DIMMs	Processor 0								Processor 1								Processor 2								Processor 3							
	DIMM 1	DIMM 2	DIMM 3	DIMM 4	DIMM 5	DIMM 6	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 13	DIMM 14	DIMM 15	DIMM 16	DIMM 17	DIMM 18	DIMM 19	DIMM 20	DIMM 21	DIMM 22	DIMM 23	DIMM 24	DIMM 25	DIMM 26	DIMM 27	DIMM 28	DIMM 29	DIMM 30	DIMM 31	DIMM 32
12	x			x	x			x	x			x					x			x	x			x	x			x				
14	x			x	x			x	x			x	x			x	x			x	x			x	x			x				
16	x			x	x			x	x			x	x			x	x			x	x			x	x			x	x			x
18	x	x	x	x	x			x	x			x	x			x	x			x	x			x	x			x	x			x
20	x	x	x	x	x			x	x			x	x			x	x	x	x	x	x			x	x			x	x			x
22	x	x	x	x	x			x	x	x	x	x	x			x	x	x	x	x	x			x	x			x	x			x
24	x	x	x	x	x			x	x	x	x	x	x			x	x	x	x	x	x			x	x	x	x	x	x			x
26	x	x	x	x	x	x	x	x	x	x	x	x	x			x	x	x	x	x	x			x	x	x	x	x	x			x
28	x	x	x	x	x	x	x	x	x	x	x	x	x			x	x	x	x	x	x	x	x	x	x	x	x	x	x			x
30	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x			x
32	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

Usage of mixed DIMM sizes

All installed memory DIMMs do not have to be the same size, but it is a preferred practice that the following groups of DIMMs be kept the same size:

- Slots 1 - 4
- Slots 5 - 8
- Slots 9 - 12
- Slots 13 - 16
- Slots 17 - 20 (IBM Flex System p460 Compute Node only)
- Slots 21 - 24 (IBM Flex System p460 Compute Node only)
- Slots 25 - 28 (IBM Flex System p460 Compute Node only)
- Slots 29 - 32 (IBM Flex System p460 Compute Node only)

4.7 Active Memory Expansion

The optional Active Memory Expansion feature is a POWER7 technology that allows the effective maximum memory capacity to be much larger than the true physical memory. Applicable to AIX V6.1 or later, this innovative compression and decompression of memory content using processor cycles allows memory expansion of up to 100%.

This situation allows an AIX V6.1 or later partition to do more work with the same physical amount of memory, or a server to run more partitions and do more work with the same physical amount of memory.

Active Memory Expansion uses processor resources to compress and extract memory contents. The trade-off of memory capacity for processor cycles can be an excellent choice, but the degree of expansion varies, based on how compressible the memory content is, and having adequate spare processor capacity available for the compression and decompression. Tests in IBM laboratories using sample workloads showed excellent results for many workloads in terms of memory expansion per additional processor used. Other test workloads had more modest results.

Clients have a great deal of control over Active Memory Expansion usage. Each individual AIX partition can turn on or turn off Active Memory Expansion. Control parameters set the amount of expansion wanted in each partition to help control the amount of processor used by the Active Memory Expansion function. An IBM Public License (IPL) is required for the specific partition that is turning on or off memory expansion. After being turned on, monitoring capabilities in standard AIX performance tools are available, such as **lparstat**, **vmstat**, **topas**, and **svmon**.

Figure 4-18 represents the percentage of processor used to compress memory for two partitions with various profiles. The green curve corresponds to a partition that has spare processing power capacity. The blue curve corresponds to a partition constrained in processing power.

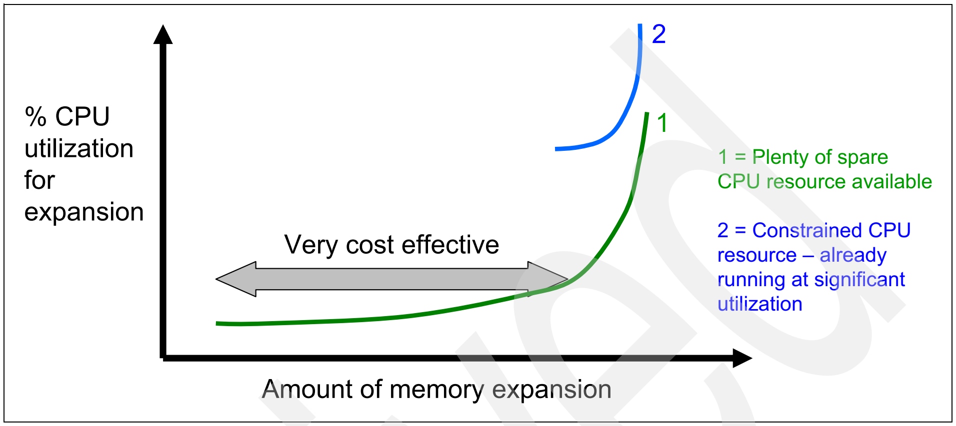

Figure 4-18 Processor usage versus memory expansion effectiveness

Both cases show a *knee of the curve* relationship for processor resources required for memory expansion:

- Busy processor cores do not have resources to spare for expansion.
- The more memory expansion that is done, the more processor resources are required.

The knee varies, depending on how compressible the memory contents are. This situation demonstrates the need for a case by case study to determine whether memory expansion can provide a positive return on investment (ROI). To help you perform this study, a planning tool is included with AIX V6.1 Technology Level 4 or later. You can use this planning tool to sample actual workloads and estimate both how expandable the partition memory is and how much processor resources are needed. Any Power System model runs the planning tool.

Figure 4-19 shows an example of the output returned by this planning tool. The tool outputs various real memory and processor resource combinations to achieve the wanted effective memory and proposes one particular combination. In this example, the tool proposes to allocate 58% of a processor core, to benefit from 45% extra memory capacity.

```
Active Memory Expansion Modeled Statistics:
-----------------------
Modeled Expanded Memory Size :   8.00 GB

Expansion      True Memory       Modeled Memory       CPU Usage
 Factor        Modeled Size      Gain                 Estimate
---------    --------------    -----------------    -----------
  1.21          6.75 GB         1.25 GB [ 19%]          0.00
  1.31          6.25 GB         1.75 GB [ 28%]          0.20
  1.41          5.75 GB         2.25 GB [ 39%]          0.35
  1.51          5.50 GB         2.50 GB [ 45%]          0.58
  1.61          5.00 GB         3.00 GB [ 60%]          1.46

Active Memory Expansion Recommendation:
---------------------
The recommended AME configuration for this workload is to configure the LPAR
with a memory size of 5.50 GB and to configure a memory expansion factor of
1.51.  This will result in a memory expansion of 45% from the LPAR's current
memory size.  With this configuration, the estimated CPU usage due to Active
Memory Expansion is approximately 0.58 physical processors, and the
estimated overall peak CPU resource required for the LPAR is 3.72 physical
processors.
```

Figure 4-19 Output from the AIX Active Memory Expansion planning tool

For more information about this topic, see the white paper *Active Memory Expansion: Overview and Usage Guide*, available at the following website:

http://www.ibm.com/systems/power/hardware/whitepapers/am_exp.html

4.8 Storage

The Power Systems compute nodes have an onboard SAS controller that can manage up to two, non-hot-pluggable internal drives. Both 2.5-inch hard disk drives (HDDs) and 1.8-inch solid-state drives (SSDs) are supported. The drives attach to the cover of the server, as shown in Figure 4-20. Even though the p460 is a full-wide server, it has the same storage options as the p260 and the p24L.

Figure 4-20 The IBM Flex System p260 Compute Node showing the hard disk drive location on the top cover

4.8.1 Storage configuration impact to memory configuration

The type of local drives, HDDs or SSDs, used impacts the form factor of your memory DIMMs:

- If HDDs are chosen, then only Very Low Profile (VLP) DIMMs can be used because of internal spacing. There is not enough room for the 2.5-inch drives to be used with Low Profile (LP) DIMMs (currently the 2 GB and 16 GB sizes). Verify your memory choice to make sure that it is compatible with the local storage configuration.
- The usage of SSDs does not have the same limitation, and LP (or VLP) DIMMs can be used with SSDs.

4.8.2 Local storage and cover options

Local storage options are shown in Table 4-8. None of the available drives are hot-swappable. If you use local drives, you need to order the appropriate cover with connections for your drive type. The maximum number of drives that can be installed in any Power Systems compute node is two. SSD and HDD drives cannot be mixed.

As you see in Figure 4-20 on page 95, the local drives (HDD or SDD) are mounted to the top cover of the system. When ordering your Power Systems compute nodes, choose which cover is appropriate for your system (SSD, HDD, or no drives).

Table 4-8 Local storage options

Feature code	Part number	Description
2.5-inch SAS HDDs		
7069	None	Top cover with HDD connectors for the p260 and the p24L
7066	None	Top cover with HDD connectors for the p460 (full-wide)
8274	42D0627	300 GB 10K RPM non-hot-swap 6 Gbps SAS
8276	49Y2022	600 GB 10K RPM non-hot-swap 6 Gbps SAS
8311	81Y9654	900 GB 10K RPM non-hot-swap 6 Gbps SAS
1.8-inch SSDs		
7068	None	Top cover with SSD connectors for the p260 and the p24L
7065	None	Top Cover with SSD connectors for p460 (full-wide)

Feature code	Part number	Description
8207	74Y9114	177 GB SATA non-hot-swap SSD
No drives		
7067	None	Top cover for no drives on the p260 and the p24L
7005	None	Top cover for no drives on the p460 (full-wide)

4.8.3 Local drive connection

On covers that accommodate drives, the drives attach to an interposer that connects to the system board when the cover is properly installed. This connection is shown in more detail in Figure 4-21.

Figure 4-21 Connector on drive interposer card mounted to server cover

On the system board, the connection for the cover's drive interposer is shown in Figure 4-22.

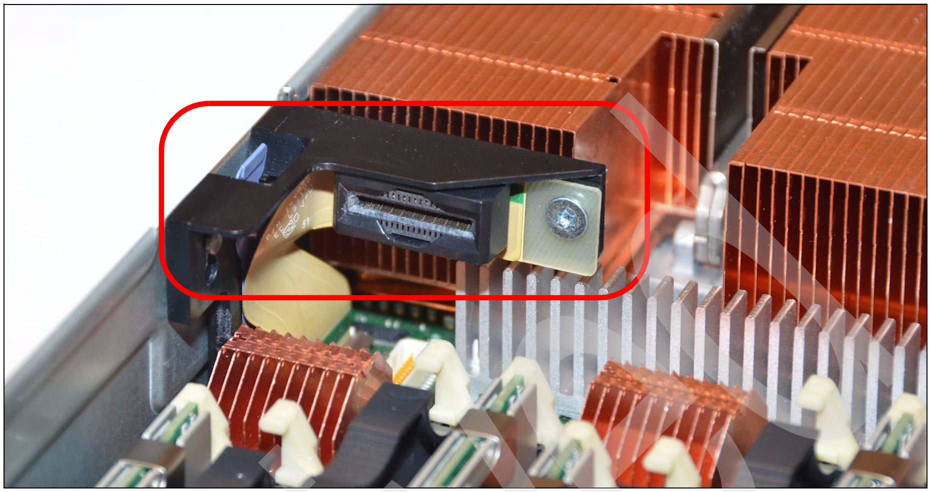

Figure 4-22 Connection for drive interposer card mounted to the system cover (connected to the system board through a flex cable)

4.8.4 RAID capabilities

Disk drives and solid-state drives in the Power Systems compute nodes can be used to implement and manage various types of RAID arrays in operating systems that are on the ServerProven list. For the compute node, you must configure the RAID array by running **`smit sasdam`**, which starts the SAS RAID Disk Array Manager for AIX.

The AIX Disk Array Manager is packaged with the Diagnostics utilities on the Diagnostics CD. Run **`smit sasdam`** to configure the disk drives for use with the SAS controller. The diagnostics CD can be downloaded in ISO file format from the following website:

http://www14.software.ibm.com/webapp/set2/sas/f/diags/download/

For more information, see "Using the Disk Array Manager" in the Systems Hardware Information Center at the following website:

http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/p7ebj/sasusingthesasdiskarraymanager.htm

Tip: Depending on your RAID configuration, you might need to create the array before you install the operating system in the compute node. Before you can create a RAID array, you must reformat the drives so that the sector size of the drives changes from 512 bytes to 528 bytes.

If you later decide to remove the drives, delete the RAID array before you remove the drives. If you decide to delete the RAID array and reuse the drives, you might need to reformat the drives so that the sector size of the drives changes from 528 bytes to 512 bytes.

4.9 I/O adapters

The networking subsystem of the IBM Flex System Enterprise Chassis is designed to provide increased bandwidth and flexibility. The new design also allows for more ports on the available expansion adapters, which allow for greater flexibility and efficiency with your system's design.

4.9.1 I/O adapter slots

There are two I/O adapter slots on the p260 and the p24L. The p460 has four I/O adapter slots.

The key differences from IBM BladeCenter are:

- The I/O adapter slots on IBM Flex System nodes are identical in shape (form factor).
- The I/O adapters for the Power Systems compute nodes have their own connector that plugs into the IBM Flex System Enterprise Chassis midplane.
- There is no onboard network capability in the Power Systems compute nodes other than the Flexible Service Processor (FSP) NIC interface.

We describe the reference codes associated with the physical adapter slots in more detail in "Assigning physical I/O" on page 301.

Slot 1 requirements: You must have an EN4054 4-port 10Gb Ethernet Adapter (Feature Code #1762) or EN2024 4-port 1Gb Ethernet Adapter (Feature Code #1763) card installed in slot 1 of the Power Systems compute nodes.

Similarly, you must have an EN4093 10Gb Scalable Switch (Feature Code #3593), EN2092 1Gb Ethernet Switch (Feature Code #3598) or EN4091 10Gb Ethernet Pass-thru Switch (Feature Code #3700) installed in bay 1 of the chassis.

A typical I/O adapter is shown in Figure 4-23.

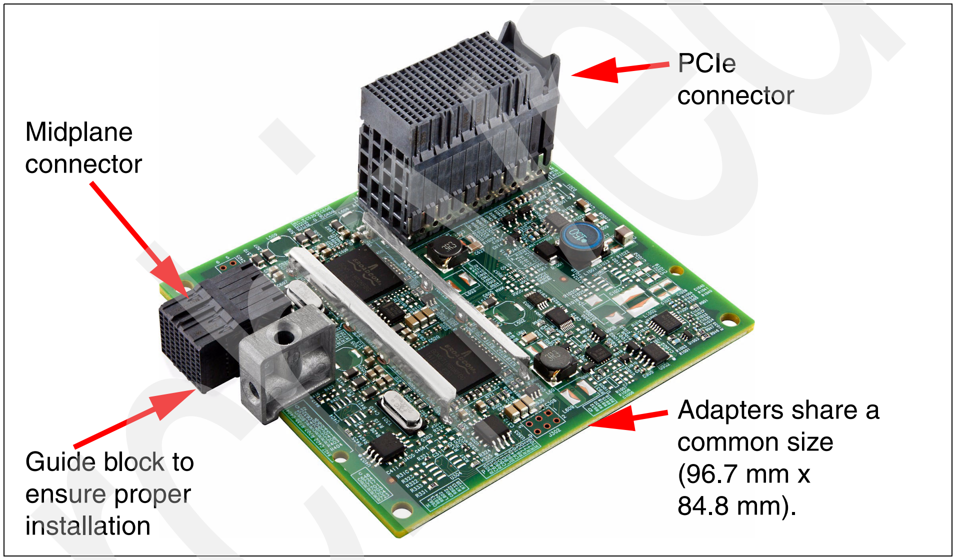

Figure 4-23 The underside of the IBM Flex System EN2024 4-port 1Gb Ethernet Adapter

Note the large connector, which plugs into one of the I/O adapter slots on the system board. Also, notice that it has its own connection to the midplane of the Enterprise Chassis. If you are familiar with IBM BladeCenter systems, several of the expansion cards connect directly to the midplane (such as the CFFh adapters) and others do not (such as the CIOv and CFFv adapters).

4.9.2 PCI hubs

The I/O is controlled by two (IBM Flex System p260 Compute Node) or four (IBM Flex System p460 Compute Node) P7-IOC I/O controller hub chips. This configuration provides additional flexibility when assigning resources within Virtual I/O Server (VIOS) to specific Virtual Machine/LPARs.

4.9.3 Available adapters

Table 4-9 shows the available I/O adapter cards for Power Systems compute nodes.

Table 4-9 Supported I/O adapter for Power Systems compute nodes

Feature Code	Description
1762	IBM Flex System EN4054 4-port 10Gb Ethernet Adapter
1763	IBM Flex System EN2024 4-port 1Gb Ethernet Adapter
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter
1761	IBM Flex System IB6132 2-port QDR InfiniBand Adapter

4.9.4 Adapter naming convention

Figure 4-24 shows the naming structure for the I/O adapters.

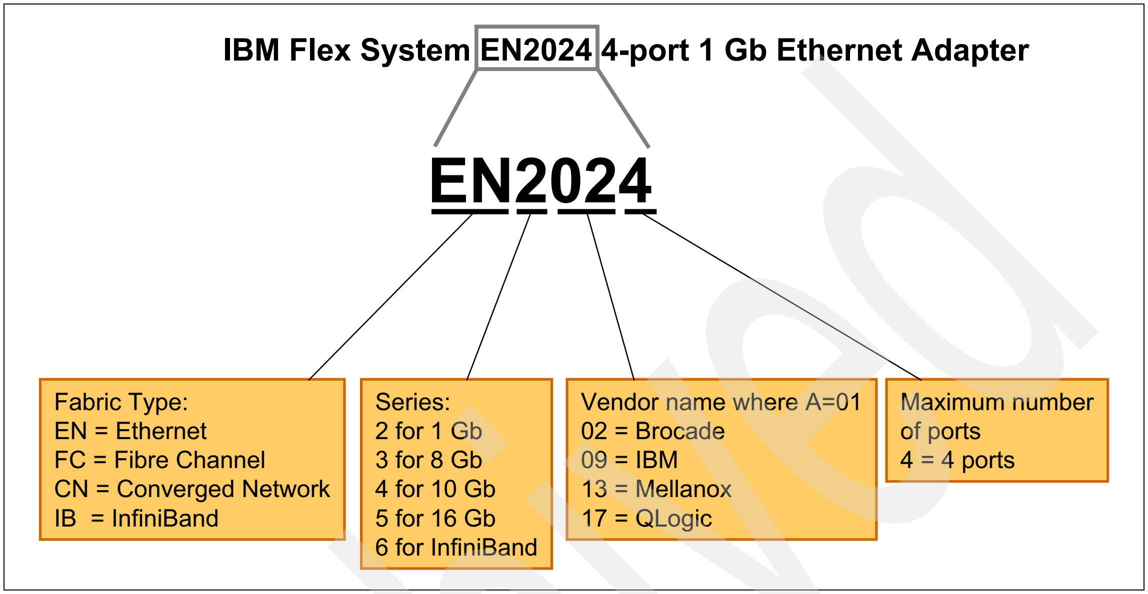

Figure 4-24 Naming structure for the I/O expansion cards

4.9.5 IBM Flex System EN4054 4-port 10Gb Ethernet Adapter

The IBM Flex System EN4054 4-port 10Gb Ethernet Adapter from Emulex enables the installation of four 10 Gb ports of high-speed Ethernet into an IBM Power Systems compute node. These ports interface to chassis switches or pass-through modules, enabling connections within or external to the IBM Flex System Enterprise Chassis.

The firmware for this four port adapter is provided by Emulex, while the AIX driver and AIX tool support are provided by IBM.

Table 4-10 lists the ordering part number and feature code.

Table 4-10 Ordering part number and feature code

Feature Code	Description
1762	EN4054 4-port 10Gb Ethernet Adapter

The IBM Flex System EN4054 4-port 10Gb Ethernet Adapter has the following features and specifications:

- On-board flash memory: 16 MB for FC controller program storage
- Uses standard Emulex SLI drivers
- Interoperates with existing FC SAN infrastructures (switches, arrays, SRM tools (including Emulex utilities), SAN practices, and so on)
- Provides 10 Gb MAC features, such as MSI-X support, jumbo frames (8 K bytes) support, VLAN tagging (802.1Q, PER priority pause / priority flow control), and advanced packet filtering
- No host operating system changes are required. NIC and HBA functionality (including device management, utilities, and so on) are not apparent to the host operating system

Figure 4-26 on page 105 shows the IBM Flex System EN4054 4-port 10Gb Ethernet Adapter.

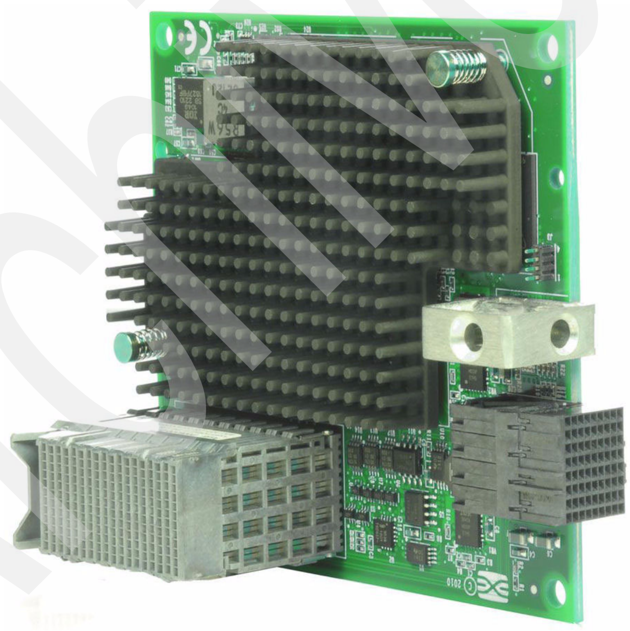

Figure 4-25 The EN2024 4-port 1Gb Ethernet Adapter for IBM Flex System

4.9.6 IBM Flex System EN2024 4-port 1Gb Ethernet Adapter

The IBM Flex System EN2024 4-port 1Gb Ethernet Adapter is a quad-port network adapter from Broadcom that provides 1 Gb per second, full duplex, Ethernet links between a compute node and Ethernet switch modules installed in the chassis. The adapter interfaces to the compute node using the PCIe bus.

Table 4-11 lists the ordering part number and feature code.

Table 4-11 Ordering part number and feature code

Feature Code	Description
1763	EN2024 4-port 1Gb Ethernet Adapter

The IBM Flex System EN2024 4-port 1Gb Ethernet Adapter has the following features:

- Connection to 1000BASE-X environments using Ethernet switches
- Compliance with US and international safety and emissions standards
- Full-duplex (FDX) capability, enabling simultaneous transmission and reception of data on the Ethernet local area network (LAN)
- Preboot Execution Environment (PXE) support
- Wake on LAN support
- MSI and MSI-X capabilities
- Receive Side Scaling (RSS) support
- NVRAM, a programmable, 4 Mb flash module
- Host data transfer: PCIe Gen 2 (one lane)

Figure 4-26 shows the IBM Flex System EN2024 4-port 1Gb Ethernet Adapter.

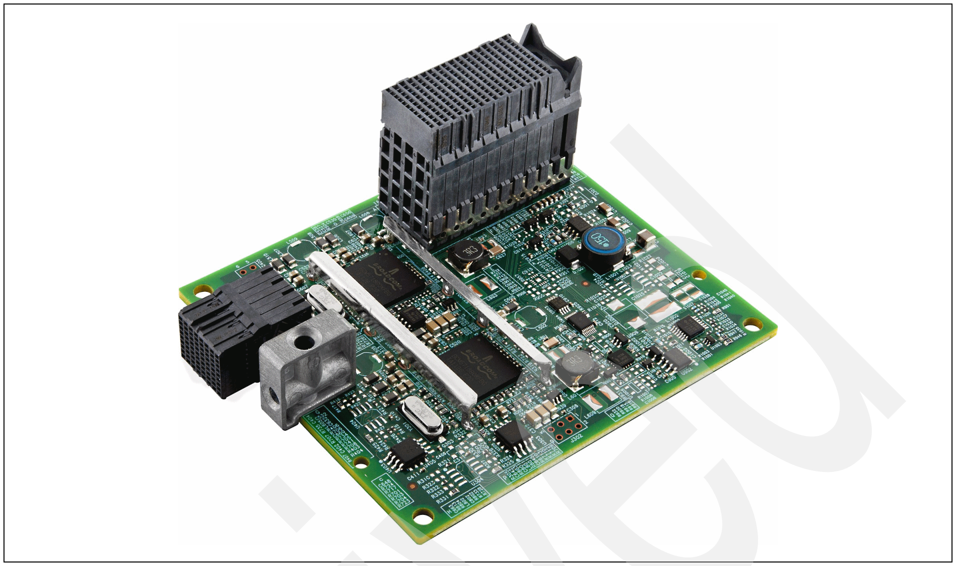

Figure 4-26 The EN2024 4-port 1Gb Ethernet Adapter for IBM Flex System

4.9.7 IBM Flex System FC3172 2-port 8Gb FC Adapter

The IBM Flex System FC3172 2-port 8Gb FC Adapter from QLogic enables high-speed access for IBM Flex System Enterprise Chassis compute nodes to connect to a Fibre Channel storage area network (SAN). This adapter is based on proven QLogic 2532 8Gb ASIC design and works with any of the 8 Gb or 16 Gb IBM Flex System Enterprise Chassis Fibre Channel switch modules.

Table 4-12 lists the ordering part number and feature code.

Table 4-12 Ordering part number and feature code

Feature Code	Description
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter

The IBM Flex System FC3172 2-port 8Gb FC Adapter has the following features:

- Support for Fibre Channel protocol SCSI (FCP-SCSI) and Fibre Channel Internet protocol (FCP-IP)
- Support for point-to-point fabric connection (F-port fabric login)

- Support for Fibre Channel service (classes 2 and 3)
- Configuration and boot support in UEFI

The IBM Flex System FC3172 2-port 8Gb FC Adapter has the following specifications:

- Bandwidth: 8 Gb per second maximum at half-duplex and 16 Gb per second maximum at full-duplex per port
- Throughput: 3200 MBps (full-duplex)
- Support for both FCP-SCSI and IP protocols
- Support for point-to-point fabric connections: F-Port Fabric Login
- Support for Fibre Channel Arbitrated Loop (FCAL) public loop profile: Fibre Loop-(FL-Port)-Port Login
- Support for Fibre Channel services class 2 and 3
- Support for FCP SCSI initiator and target operation
- Support for full-duplex operation
- Copper interface AC coupled

Figure 4-27 shows the IBM Flex System FC3172 2-port 8Gb FC Adapter.

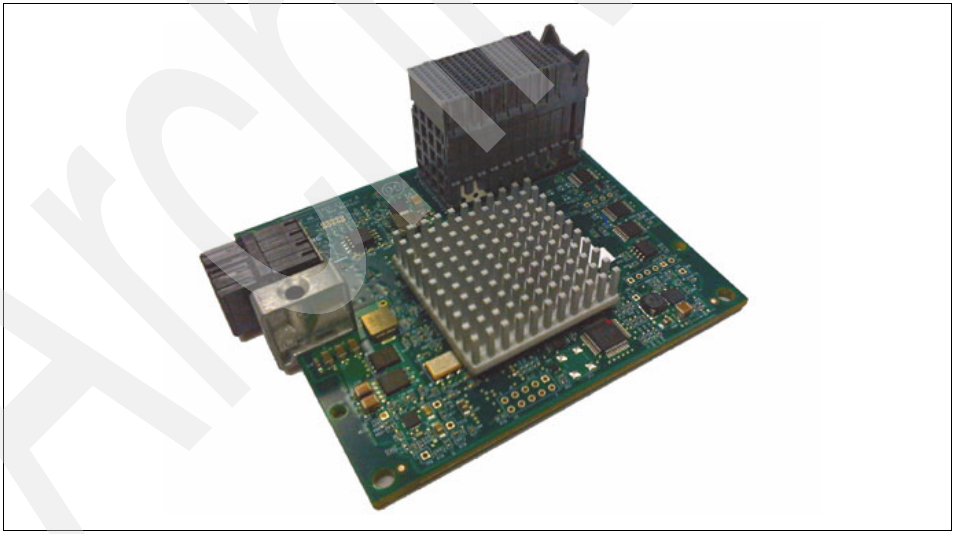

Figure 4-27 The FC3172 2-port 8Gb FC Adapter for IBM Flex System

4.9.8 IBM Flex System IB6132 2-port QDR InfiniBand Adapter

The IBM Flex System IB6132 2-port QDR InfiniBand Adapter from Mellanox provides the highest performing and most flexible interconnect solution for servers used in Enterprise Data Centers, High-Performance Computing, and Embedded environments.

Table 4-13 lists the ordering part number and feature code.

Table 4-13 Ordering part number and feature code

Feature Code	Description
1761	IB6132 2-port QDR InfiniBand Adapter

The IBM Flex System IB6132 2-port QDR InfiniBand Adapter has the following features and specifications:

- ConnectX2 based adapter
- Virtual Protocol Interconnect (VPI)
- InfiniBand Architecture Specification V1.2.1 compliant
- IEEE Std. 802.3 compliant
- PCI Express 2.0 (1.1 compatible) through an x8 edge connector up to 5 GTps
- Processor offload of transport operations
- CORE-Direct application offload
- GPUDirect application offload
- Unified Extensible Firmware Interface (UEFI)
- Wake on LAN (WoL)
- RDMA over Converged Ethernet (RoCE)
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- TCP/UDP/IP stateless offload
- RoHS-6 compliant

Figure 4-26 on page 105 shows the IBM Flex System IB6132 2-port QDR InfiniBand Adapter.

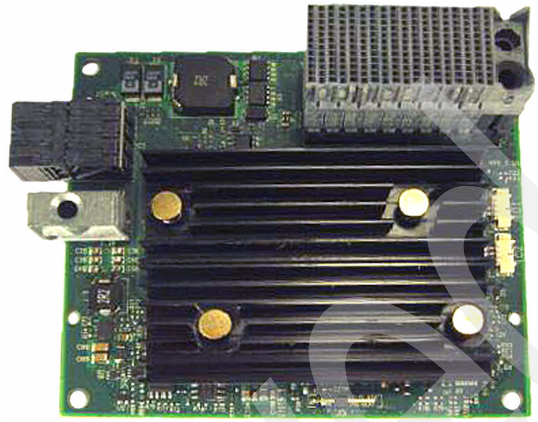

Figure 4-28 The IB6132 2-port QDR InfiniBand Adapter for IBM Flex System

4.10 System management

There are several advanced system management capabilities built into Power Systems compute nodes. A Flexible Support Processor handles most of the server-level system management. It has features, such as system alerts and Serial-over-LAN capability, that we describe in this section.

4.10.1 Flexible Support Processor

A Flexible Support Processor (FSP) provides out-of-band system management capabilities, such as system control, runtime error detection, configuration, and diagnostic tests. Generally, you do not interact with the Flexible Support Processor directly, but by using tools, such as IBM Flex System Manager, Chassis Management Module, and the external IBM Systems Director Management Console.

The Flexible Support Processor provides a Serial-over-LAN interface, which is available using the Chassis Management Module and the **`console`** command.

The IBM Flex System p460 Compute Node, even though it is a full-wide system, has only one Flexible Support Processor.

4.10.2 Serial over LAN (SOL)

The Power Systems compute nodes do not have an on-board video chip and do not support keyboard, video, and mouse (KVM) connections. Server console access is obtained by a SOL connection only. 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 or that are attached to the IBM Flex System Manager. SOL provides console redirection for both System Management Services (SMS) and the server operating system. The SOL feature redirects server serial-connection data over a LAN without requiring special cabling by routing the data using the Chassis Management Module network interface. The SOL connection enables Power Systems compute nodes to be managed from any remote location with network access to the Chassis Management Module.

SOL offers the following advantages:

- Remote administration without KVM (headless servers)
- Reduced cabling and no requirement for a serial concentrator
- Standard Telnet/SSH interface, eliminating the requirement for special client software

The Chassis Management Module CLI provides access to the text-console command prompt on each server through a SOL connection, enabling the Power Systems compute nodes to be managed from a remote location.

4.10.3 Anchor card

The anchor card, shown in Figure 4-29, contains the smart vital product data chip that stores system-specific information. The pluggable anchor card provides a means for this information to be transferable from a faulty system board to the replacement system board. Before the service processor knows what system it is on, it reads the smart vital product data chip to obtain system information.

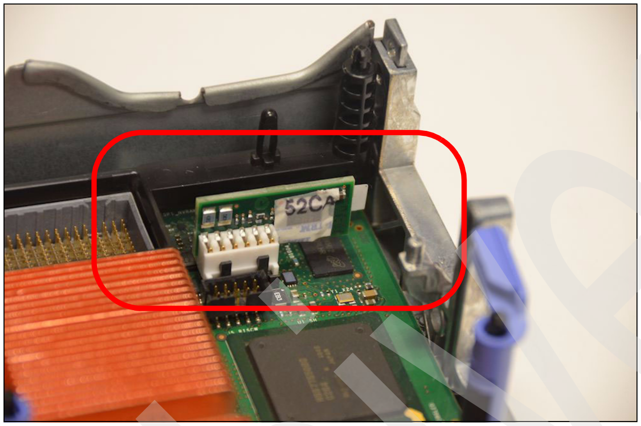

Figure 4-29 Anchor card

The vital product data chip includes information such as machine type, model, and serial number.

4.11 Integrated features

As stated in 4.1.1, “IBM Flex System p260 Compute Node” on page 64, 4.1.2, “IBM Flex System p460 Compute Node” on page 66, and 4.1.3, “IBM Flex System p24L Compute Node” on page 68, the integrated features are as follows:

- Flexible Support Processor
- IBM POWER7 processors
- SAS RAID-capable controller
- USB port

4.12 IBM EnergyScale

IBM EnergyScale technology provides functions that help you to understand and dynamically optimize the processor performance versus processor power and system workload, to control IBM Power Systems power and cooling usage.

The BladeCenter Advanced Management Module and IBM Systems Director Active Energy Manager use EnergyScale technology, enabling advanced energy management features to conserve power and improve energy efficiency. Intelligent energy optimization capabilities enable the POWER7 processor to operate at a higher frequency for increased performance and performance per watt, or reduce frequency to save energy. This feature is called Turbo-Mode.

4.12.1 IBM EnergyScale technology

This section describes the design features and the hardware and software requirements of IBM EnergyScale.

IBM EnergyScale consists of the following elements:

- A built-in EnergyScale device (formerly known as the Thermal Power Management Device (TPMD))
- Power executive software, which includes IBM Systems Director Active Energy Manager, an IBM Systems Directors plug-in, and BladeCenter Advanced Management Module

IBM EnergyScale functions include the following elements:

- Energy trending

  EnergyScale provides the continuous collection of real-time server energy consumption data. This function enables administrators to predict power consumption across their infrastructure and to react to business and processing needs. For example, administrators might use such information to predict data center energy consumption at various times of the day, week, or month.

- Thermal reporting

  IBM Systems Director Active Energy Manager displays measured ambient temperature and calculated exhaust heat index temperature. This information helps identify data center hot spots that require attention.

- Power Saver Mode

  Power Saver Mode reduces the processor frequency and voltage by a fixed amount, reducing the energy consumption of the system, while delivering predictable performance. This percentage is predetermined to be within a safe operating limit and is not user configurable. The server is designed for a fixed frequency drop of 50% from nominal. Power Saver Mode is not supported during boot or reboot operations, although it is a persistent condition that is sustained after the boot, when the system starts executing instructions.

- Dynamic Power Saver Mode

  Dynamic Power Saver Mode varies processor frequency and voltage, based on the use of the POWER7 processors. This setting is configured from BladeCenter Advanced Management Module or IBM Director Active Energy Manager. Processor frequency and usage are inversely proportional for most workloads, implying that, as the frequency of a processor increases, its use decreases, given a constant workload. Dynamic Power Saver Mode takes advantage of this relationship to detect opportunities to save power, based on measured real-time system usage.

  When a system is idle, the system firmware lowers the frequency and voltage to Power Saver Mode values. When fully used, the maximum frequency varies, depending on whether you favor power savings or system performance. If energy savings are preferred and the system is fully used, the system can reduce the maximum frequency to 95% of nominal values. If performance is favored over energy consumption, the maximum frequency is at least 100% of nominal.

  Dynamic Power Saver Mode is mutually exclusive with Power Saver mode. Only one of these modes can be enabled at one time.

- Power capping

  Power capping enforces a limit, specified by you, on power usage. Power capping is not a power-saving mechanism. It enforces power caps by throttling the processors in the system, degrading performance significantly. The idea of a power cap is to set a limit that is not expected to be reached, but that frees up margined power in the data center. The margined power is the amount of extra power allocated to a server during installation in a data center. It is based on those server environmental specifications that usually are never reached because server specifications are always based on maximum configurations and worst case scenarios. The energy cap is set and enabled in BladeCenter Advanced Management Module and in IBM Systems Director Active Energy Manager.

- Soft power capping

  Soft power capping extends the allowed energy capping range further, beyond a region that can be guaranteed in all configurations and conditions. When an energy management goal is to meet a particular consumption limit, soft power capping is the mechanism to use.

- Processor core nap

  The IBM POWER7 processor uses a low-power mode called *nap* that stops processor execution when there is no work to be done by that processor core. The latency of exiting nap falls within a partition dispatch (context switch), such that the IBM POWER Hypervisor™ uses it as a general purpose idle state. When the operating system detects that a processor thread is idle, it yields control of a hardware thread to the POWER Hypervisor. The POWER Hypervisor immediately puts the thread into nap mode. Nap mode allows the hardware to clock-off most of the circuits inside the processor core. Reducing active energy 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 mode until they are licensed, and they return to core nap mode when unlicensed again.

- Processor folding

  Processor folding is a consolidation technique that dynamically adjusts, over the short term, the number of processors available for dispatch to match the number of processors demanded by the workload. As the workload increases, the number of processors made available increases. As the workload decreases, the number of processors made available decreases. Processor folding increases energy savings during periods of low to moderate workload, because unavailable processors remain in low-power idle states longer.

- EnergyScale for I/O

  IBM POWER processor-based systems automatically power off pluggable, Peripheral Component Interconnect (PCI) adapter slots that are empty or not being used. System firmware automatically scans all pluggable PCI slots at regular intervals, looking for the ones that meet the criteria of not being in use and powering them off. This support is available for all POWER processor-based servers and the expansion units they support.

4.12.2 EnergyScale device

The EnergyScale device dynamically optimizes processor performance, depending on the processor power and system workload.

The IBM POWER7 chip provides significant improvement in power and performance over the IBM POWER6 chip. POWER7 has more internal hardware, and power and thermal management functions to interact with:

- More hardware: Eight cores versus two cores, four threads versus two threads per core, and an asynchronous processor core chipset
- Advanced idle power management functions
- Advanced dynamic power management (DPM) functions in all hardware (processor cores, processor core chip, chip-level nest unit level, and chip level)
- Advanced actuators and control
- Advanced accelerators

The new EnergyScale device has a more powerful microcontroller, more analog/digital (A/D) channels, and more buses to handle increased workload, link traffic, and new power and thermal functions.

4.13 Operating system support

The p260 and the p460 are designed to run AIX, VIOS, IBM i, and Linux. The p24L is designed to run Linux only.

A summary of the supported operating systems is in 5.1.2, “Software planning” on page 119.

4.14 Warranty and maintenance agreements

The Power Systems compute nodes have a three year limited on-site warranty. Upgrades to the base warranty are available. An upgraded warranty provides a faster response time for repairs, on-site repairs for most work, and after-hours and weekend repairs.

Details about warranty options and our terms and conditions are at the following website:

http://www.ibm.com/support/warranties/

4.15 Software support and remote technical support

IBM also offers technical assistance to help solve software-related challenges. Our team assists with configuration, how-to questions, and setup of your servers. Information about these options is at the following website:

http://ibm.com/services/us/en/it-services/tech-support-and-maintenance-services.html

5

Planning

In this chapter, we describe the steps you should take before ordering and installing Power Systems compute nodes as part of an IBM Flex System solution.

We cover the following topics in this chapter:

- Planning your system: An overview
- Network connectivity
- SAN connectivity
- Configuring redundancy
- Dual VIOS
- Power planning
- Cooling
- Planning for virtualization

5.1 Planning your system: An overview

One of the initial tasks for your team is to plan for the successful implementation of your Power Systems compute node. This planning includes ensuring that the primary reasons for acquiring the server are effectively planned for. Consider the overall uses for the server, the planned growth of your applications, and the operating systems in your environment. Correct planning of these issues ensures that the server meets the needs of your organization.

We describe the following topics in this section:

- Hardware planning
- Software planning

5.1.1 Hardware planning

Important topics to consider in your planning activities are:

- Network connectivity

  On Power Systems compute nodes, several models of expansion cards are available (as described in 4.9, "I/O adapters" on page 99). Make sure that you choose the best expansion cards for your environment, either chassis, switches, or networking in general, to avoid compatibility issues. Purchasing Ethernet expansion cards requires the purchase of an Ethernet switch in the chassis.

- Fibre Channel and storage area network (SAN) connectivity

  The same considerations described for network connectivity also apply to Fibre Channel and storage area network (SAN) connectivity.

- Hard disk drives (HDDs) and solid-state drives (SSDs)

  If you choose to use your Power Systems compute node with internal disks, your memory choices are affected. SAS and SATA HDD options are available, and solid-state drives (SSDs). Very Low Profile (VLP) memory DIMMs are required if HDDs are chosen (as described in 4.8, "Storage" on page 95). If Low Profile (LP) memory options are chosen, only SSD disks can be used. Choosing the disk type that best suits your needs involves evaluating the size, speed, and price of the options.

- Memory

  Your Power Systems compute node supports a wide spread of memory configurations. The memory configuration depends on whether you have internal disks installed, as described "Hard disk drives (HDDs) and solid-state drives (SSDs)" on page 118). Mixing both types of memory is not recommended. Active memory expansion (AME) is available on POWER7, as is Active Memory Sharing (AMS) when using PowerVM Enterprise Edition. AMS is described in detail in several Redbooks publications, two of which are *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 and *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

- Processor

  Several processor options are available for both the IBM Flex System p260 Compute Node and the IBM Flex System p460 Compute Node (described in 4.5.1, "Processor options for Power Systems compute nodes" on page 77). Evaluate the processor quantity and speed options to determine what processor configuration most closely matches your needs. IBM provides a measurement (called **`rperf`**) that can be used to compare the relative performance of POWER systems in absolute values. The charts can be found at the following website:

  `http://www.ibm.com/systems/power/hardware/reports/system_perf.html`

- Optical media

  The IBM Flex System Enterprise Chassis and the Enterprise Chassis do not provide CD-ROM or DVD-ROM devices as do the previous BladeCenter chassis versions. If you require a local optical drive, use an external USB drive.

5.1.2 Software planning

Determine the primary uses for your Power Systems compute node and how it is set up. Will you be using full system partition, or a virtualized environment which includes virtual servers (formerly named logical partitions, LPARs) and workload partitions (WPARs)?

Operating system support

The IBM POWER7 processor-based systems support three families of operating systems:

- AIX
- IBM i
- Linux

In addition, the Virtual I/O Server can be installed in special virtual servers that provide support to the other operating systems for using features such as virtualized I/O devices, PowerVM Live Partition Mobility, or PowerVM Active Memory Sharing.

For details about the software available on IBM Power Systems servers, see the IBM Power Systems Software™ website at:

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

> **Note:** The p24L supports Virtual I/O Server (VIOS) and Linux only.

The p260 and p460 support the following operating systems and versions.

Virtual I/O Server

The supported versions are Virtual I/O Server 2.2.1.4, or later.

IBM regularly updates the Virtual I/O Server code. For information about the latest update, see the Virtual I/O Server website at:

http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/

IBM AIX V5.3

The supported versions are IBM AIX 5L™ V5.3 with the 5300-12 Technology Level with Service Pack 6, or later (the planned availability is 29 June 2012) AIX 5L V5.3 Service Extension is required.

A partition using AIX 5L V5.3 runs in POWER6 or POWER6+ compatibility mode.

IBM periodically releases maintenance packages (service packs or technology levels) for the AIX 5L operating system. Information about downloading and obtaining the CD-ROM for these packages is on the Fix Central web page at:

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

The Service Update Management Assistant can help you automate the task of checking and downloading operating system files, and is part of the base operating system. For more information about the **suma** command functionality, go to the following web page:

http://www14.software.ibm.com/webapp/set2/sas/f/genunix/suma.html

AIX V6.1

The supported versions are:

- AIX V6.1 with the 6100-07 Technology Level, with Service Pack 3 with APAR IV14283
- AIX V6.1 with the 6100-07 Technology Level, with Service Pack 4, or later (the planned availability 29 June 2012)
- AIX V6.1 with the 6100-06 Technology Level with Service Pack 8, or later (the planned availability is 29 June 2012)

For information about AIX V6.1 maintenance and support, go to the Fix Central website at:

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

AIX V7.1

The supported versions are:

- AIX V7.1 with the 7100-01 Technology Level with Service Pack 3 with APAR IV14284
- AIX V7.1 with the 7100-01 Technology Level with Service Pack 4, or later (the planned availability is 29 June 2012)
- AIX V7.1 with the 7100-00 Technology Level with Service Pack 6, or later (the planned availability is 29 June 2012)

For information about AIX V7.1 maintenance and support, go to the Fix Central website at:

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

IBM i

The supported versions are:

- IBM i 6.1 with i 6.1.1 machine code, or later
- IBM i 7.1, or later

Virtual I/O Server is required to install IBM i in a Virtual Server on IBM Flex System p260 Compute Node and IBM Flex System p460 Compute Node. All I/O must be virtualized.

For a detailed guide about installing and operating IBM i with Power Based compute nodes, go to the following website:

http://ibm.com/systems/resources/systems_power_hardware_blades_i_on_blade_readme.pdf

Linux

Linux is an open source operating system that runs on numerous platforms from embedded systems to mainframe computers. It provides a UNIX like implementation in many computer architectures.

At the time of this writing, the supported versions of Linux on POWER7 processor technology-based servers are as follows:

- SUSE Linux Enterprise Server 11 Service Pack 2 for POWER, with current maintenance updates available from Novell to enable all planned functionality
- Red Hat Enterprise Linux 5.7 for POWER, or later
- Red Hat Enterprise Linux 6.2 for POWER, or later

Linux operating system licenses are ordered separately from the hardware. You can obtain Linux operating system licenses from IBM to be included with your POWER7 processor technology-based servers, or from other Linux distributors.

> **Important:** For systems ordered with the Linux operating system, IBM ships the most current version available from the distributor. If you require another version than the one shipped by IBM, you must obtain it by downloading it from the Linux distributor's website. Information concerning access to a distributor's website is on the product registration card delivered to you as part of your Linux operating system order.

For information about the features and external devices supported by Linux, go to the following website:

`http://www.ibm.com/systems/p/os/linux/`

For information about SUSE Linux Enterprise Server, go to the following website:

`http://www.novell.com/products/server`

For information about Red Hat Enterprise Linux Advanced Servers, go to the following website:

`http://www.redhat.com/rhel/features`

> **Important:** Be sure to update your system with the latest Linux on Power service and productivity tools from the IBM website at:
>
> `http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html`

Full system partition planning

In the full system partition installation, you have several AIX version options, as explained in "Operating system support" on page 119.

When you install AIX V6.1 TL7 and AIX V7.1 TL1, you can virtualize through WPARs, as described in 8.3.1, “Installing AIX” on page 364. (Older versions of AIX 5L V5.2 and V5.3 on lower TL levels can run WPARS within a Virtual Server running AIX V7.)

Also, Linux installations are supported on the Power Systems compute node. Supported versions are listed in “Operating system support” on page 119.

> **Important:** Methods for installing these operating systems are described in Chapter 8, “Operating system installation” on page 317.

Virtualized environment planning

If you decide to implement a virtualized environment, you can create AIX and Linux virtual servers on the Power Systems compute node with or without a VIOS. If you choose not to use VIOS, the number of virtual servers is limited by the amount of expansion cards in the Power Systems compute node. If you choose to use VIOS, you are able to virtualize the limited number of expansion cards to create client virtual servers. Remember that you must use VIOS 2.2.1.

Integrated Virtualization Manager (IVM) and Systems Director Management Console (SDMC) are not supported in the Power Systems compute node environment, so you need an IBM Flex System Manager attached to your Power Systems compute node Flexible Service Processor (FSP) to create virtual servers and perform virtualization. The process for attaching the Power Systems compute node to the IBM Flex System Manager is described in 8.1.1, “Firmware update methods” on page 318.

> **Important:** PowerVM provides several types of licensing, called editions. Only Standard and Enterprise Editions are supported for Power Systems compute nodes. Be sure to evaluate the options available in each of those editions and purchase the correct license for what you are implementing.

If you plan to use advanced features such as Live Partition Mobility or Active Memory Sharing, the Enterprise Edition is required. Information about these features is at the following website:

http://ibm.com/systems/power/software/virtualization/editions/

As mentioned in 5.1.1, “Hardware planning” on page 118, **rperf** reports can be used to check processor values and equivalences.

Implementing a dual VIOS solution is the best way to achieve a high availability (HA) environment. This environment allows for maintenance on one VIOS without disrupting the clients, and avoids depending on just one VIOS to do all of the work functions.

Dual VIOS: If you want a dual VIOS environment, external disk access is required, as the two internal disks are connected to the same SAS/SATA controller. The two internal disks are used for the rootvg volume group on one VIOS only.

5.2 Network connectivity

Network connectivity in Power Systems compute nodes is provided by the I/O adapters installed in the nodes. The adapters are functionally similar to the CFFh cards used in BladeCenter servers.

The Ethernet adapters currently supported by compute nodes are listed in Table 5-1. For more details about the supported expansion cards, see 4.9, “I/O adapters” on page 99.

Table 5-1 Supported Ethernet adapters

Feature Code	Supported Ethernet adapters
1762	IBM Flex System EN4054 4-port 10Gb Ethernet Adapter
1763	IBM Flex System EN2024 4-port 1Gb Ethernet Adapter

5.2.1 Ethernet switch module connectivity

There are a number of I/O modules that can be used to provide network connectivity. These modules include Ethernet switch modules that provide integrated switching capabilities for the chassis, and pass-through modules that make internal compute node ports available external to the chassis. Using the Ethernet switch modules might provide required or enhanced functions, and simplified cabling. However, in some circumstances (for example, specific security policies or certain network requirements), it is not possible to use integrated switching capabilities, so pass-through modules are required.

Make sure that the external interface ports of the switches selected are compatible with the physical cabling used or planned to be used in your data center. Also, make sure that the features and functions required in the network are supported by the proposed switch modules.

Detailed information about I/O module configuration can be found in *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984.

The available Ethernet switches and pass-through modules are listed in Table 5-2.

Table 5-2 Available switch options for the chassis

Feature Code	Description
3593	IBM Flex System Fabric EN4093 10Gb Scalable Switch
3598	IBM Flex System EN2092 1Gb Ethernet Scalable Switch
3700	IBM Flex System EN4091 10Gb Ethernet Pass-thru

Table 5-3 lists the common selection considerations that might be useful when selecting a switch module.

Table 5-3 Switch module selection criteria

Suitable switch module →	**EN2092 1Gb Ethernet Switch**	**EN4093 10Gb Scalable Switch**
Requirement		
Gigabit Ethernet to nodes/10 Gb Ethernet uplinks	Yes	Yes
10 Gb Ethernet to nodes/10 Gb Ethernet uplinks	Yes	Yes
Basic Layer 2 switching (VLAN and port aggregation)	Yes	Yes
Advanced Layer 2 switching: IEEE-based features (Spanning Tree Protocol and QoS)	Yes	Yes
Layer 3 IPv4 switching (forwarding, routing, and ACL filtering)	Yes	Yes
Layer 3 IPv6 switching (forwarding, routing, and ACL filtering)	Yes	Yes
10 Gb Ethernet CEE/FCoE	No	Yes
Switch stacking	No	Yes
vNIC support	No	Yes
IBM VMready®	Yes	Yes

5.2.2 VLANs

Virtual LANs (VLANs) are commonly used in the Layer 2 network to split up groups of network users into manageable broadcast domains, to create a logical segmentation of workgroups, and to enforce security policies among logical segments. VLAN considerations include the number and types of VLANs supported, VLAN tagging protocols supported, and specific VLAN configuration protocols implemented.

All IBM Flex System switch modules support the 802.1Q protocol for VLAN tagging.

Another usage of 802.1Q VLAN tagging is to divide one physical Ethernet interface into several logical interfaces that belong to more than one VLAN. A compute node can send and receive tagged traffic from several VLANs on the same physical interface. This task can be done with network adapter management software (the same used for NIC teaming). Each logical interface appears as a separate network adapter in the operating system, with its own set of characteristics, such as IP addresses, protocols, and services.

Having several logical interfaces can be useful in cases when an application requires more than two separate interfaces and you do not want to dedicate a whole physical interface to it (for example, not enough interfaces or low traffic). It might also help to implement strict security policies for separating network traffic using VLANs, while having access to server resources from other VLANs, without needing to implement Layer 3 routing in the network.

To be sure that the application deployed supports logical interfaces, check the application documentation for possible restrictions applied to the NIC teaming configurations, especially in the case of a clustering solutions implementation.

For more information about Ethernet switch modules, see *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984.

5.3 SAN connectivity

SAN connectivity in the Power Systems compute nodes is provided by the expansion cards. The list of SAN Fibre Channel (FC) adapters currently supported by the Power Systems compute nodes is listed in Table 5-4. For more details about the supported expansion cards, see 4.9, “I/O adapters” on page 99.

Table 5-4 Supported FC adapters

Feature Code	Description
1764	IBM Flex System FC3172 2-port 8Gb FC Adapter

Important: At the time of writing, the FC3052 2-port 8Gb FC Adapter and FC5022 2-port 16Gb FC Adapter were not supported by the Power Systems compute nodes.

The SAN and Fibre Channel I/O modules are installed in the IBM Flex System chassis. This installation includes SAN switch modules that provide integrated switching capabilities and pass-through modules that make internal compute node ports available to the outside.

Use SAN switches whenever possible, because you can use this configuration to mix complex configuration and zoning settings inside the chassis or to integrate with your existing SAN configuration.

Ensure that the external interface ports of the switches or pass-through modules selected are compatible with the physical cabling used or planned to be used in your data center. Also, ensure that the features and functions required in the SAN are supported by the proposed switch modules or pass-through modules.

Detailed information about these modules is in Chapter 3, “Systems management”, in *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984. The available switch and pass-through options are listed in Table 5-5.

Table 5-5 SAN switch options for the chassis

Feature Code	Description
3770	IBM Flex System FC5022 16Gb SAN Scalable Switch
3595	IBM Flex System FC3171 8Gb SAN Switch
3591	IBM Flex System FC3171 8Gb SAN Pass-thru

5.4 Configuring redundancy

Your environment might require continuous access to your network services and applications. Providing highly available (HA) network resources is a complex task that involves the integration of multiple hardware and software components. One HA component is to provide network infrastructure availability.

This availability is required for both network and SAN connectivity.

5.4.1 Network redundancy

Network infrastructure availability can be achieved by implementing certain techniques and technologies. Most of these items are widely used standards, but several are specific to the IBM Flex System Enterprise Chassis. This section describes the most common technologies that can be implemented in an IBM Flex System environment to provide a highly available network infrastructure.

In general, a typical LAN infrastructure consists of server NICs, client NICs, and network devices, such as Ethernet switches and that cables that connect them. The potential failures in a network include port failures (both on switches and servers), cable failures, and network device failures.

To provide high availability and redundancy:

- Avoid or minimize single points of failure, that is, provide redundancy for network equipment and communication links. The IBM Flex System Enterprise Chassis has built-in redundancy:
  - Two or four ports on I/O expansion cards on each compute node
  - Two separate communication paths to I/O modules through dual midplane connections
  - Two I/O module bays per dual port for device redundancy

  For a sample connection topology between I/O adapters and I/O modules, see Chapter 3, “Systems management”, of *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984.

  Implement technologies that provide automatic failover in the case of any failure. This implementation can be done by using certain feature protocols that are supported by network devices, together with server-side software.

  Consider implementing the following technologies, which can help you to achieve a higher level of availability in an IBM Flex System network solution (depending on your network architecture):

  - Spanning Tree Protocol
  - Layer 2 failover (also known as Trunk Failover)

– Virtual Link Aggregation Groups (VLAG)
– Virtual Router Redundancy Protocol (VRRP)
– Routing protocol (such as RIP or OSPF)

Redundant network topologies

The IBM Flex System Enterprise Chassis can be connected to the enterprise network in several ways, as shown in Figure 5-1.

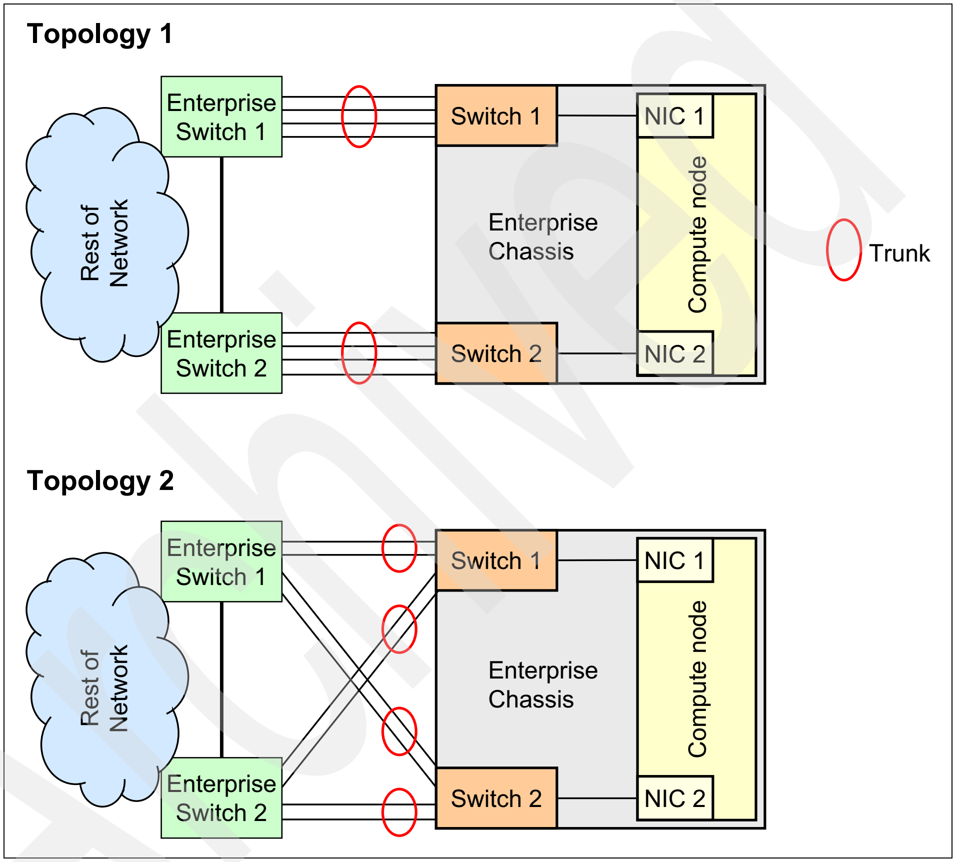

Figure 5-1 IBM Flex System redundant LAN integration topologies

Topology 1 in Figure 5-1 has each switch module in the chassis directly connected to one of the enterprise switches through aggregation links, using external ports on the switch. The specific number of external ports used for link aggregation depends on your redundancy requirements, performance considerations, and real network environments. This topology is the simplest way to integrate IBM Flex System into an existing network, or to build a new one.

Topology 2 in Figure 5-1 on page 129 has each switch module in the chassis with two direct connections to two enterprise switches. This topology is more advanced, and it has a higher level of redundancy, but certain specific protocols such as Spanning Tree or Virtual Link Aggregation Groups must be implemented. Otherwise, network loops and broadcast storms can cause the meltdown of the network.

Spanning Tree Protocol

Spanning Tree Protocol is a 802.1D standard protocol used in Layer 2 redundant network topologies. When multiple paths exist between two points on a network, Spanning Tree Protocol or one of its enhanced variants can prevent broadcast loops and ensure that the switch uses only the most efficient network path. Spanning Tree Protocol is also used to enable automatic network reconfiguration in case of failure. For example, enterprise switches 1 and 2, together with switch 1 in chassis, create a loop in a Layer 2 network (see Topology 2 in Figure 5-1 on page 129). We must use Spanning Tree Protocol in that case as a loop prevention mechanism (because a Layer 2 network cannot operate in a loop).

Assume that the link between enterprise switch 2 and chassis switch 1 is disabled by Spanning Tree Protocol to break a loop, so traffic is going through the link between enterprise switch 1 and chassis switch 1. If there is a link failure, Spanning Tree Protocol reconfigures the network and activates the previously disabled link. The process of reconfiguration can take tenths of a second, and the service is unavailable during this time.

Whenever possible, plan to use trunking with VLAN tagging for interswitch connections, which can help you achieve higher performance by increasing interswitch bandwidth, and higher availability by providing redundancy for links in the aggregation bundle.

STP modifications, such as Port Fast Forwarding or Uplink Fast, might help improve STP convergence time and the performance of the network infrastructure. Additionally, several instances of STP might run on the same switch simultaneously, on a per-VLAN basis (that is, each VLAN has its own copy of STP to load-balance traffic across uplinks more efficiently).

For example, assume that a switch has two uplinks in a redundant loop topology, and several VLANs are implemented. If single STP is used, then one of these uplinks is disabled and the other carries traffic from all VLANs. However, if two STP instances are running, then one link is disabled for one set of VLANs while carrying traffic from another set of VLANs, and vice versa. Both links are active, thus enabling more efficient use of available bandwidth.

Layer 2 failover

Depending on the configuration, each compute node can have one IP address per each Ethernet port, or it can have one virtual NIC consisting of two or more physical interfaces with one IP address. This configuration is known as NIC teaming technology. From an IBM Flex System perspective, NIC Teaming is useful when you plan to implement high availability configurations with automatic failover if there are internal or external uplink failures.

We can use only two ports on a compute node per virtual NIC for high availability configurations. One port is active, and the other is standby. One port (for example, the active port) is connected to the switch in I/O bay 1, and the other port (for example, the standby port) is to be connected to the switch in I/O bay 2. If you plan to use an Ethernet expansion card for high availability configurations, then the same rules apply. Active and standby ports need to be connected to a switch on separate bays.

If there is an internal port or link failure of the active NIC, the teaming driver switches the port roles. The standby port becomes active and the active port becomes standby. This action is done quickly, within a few seconds. After restoring the failed link, the teaming driver can perform a failback or can do nothing, depending on the configuration.

Look at topology 1 in Figure 5-1 on page 129. Assume that NIC Teaming is on, the compute node NIC port connected to switch 1 is active, and the other node is on standby. If something goes wrong with the internal link to switch 1, then the teaming driver detects the status of NIC port failure and performs a failover. But what happens if external connections are lost (the connection from chassis switch 1 to Enterprise Switch 1 is lost)? The answer is that nothing happens because the internal link is still on and the teaming driver does not detect any failure. So the network service becomes unavailable.

To address this issue, the Layer 2 Failover technique is used. Layer 2 Failover can disable all internal ports on the switch module if there is an upstream links failure. A disabled port means no link, so the NIC Teaming driver performs a failover. This special feature is supported on the IBM Flex System and BladeCenter switch modules. Thus, if Layer 2 Failover is enabled and you lose connectivity with Enterprise Switch 1, then the NIC Teaming driver performs a failover and the service is available through Enterprise Switch 2 and chassis switch 2.

Layer 2 Failover is used with NIC active or standby teaming. Before using NIC Teaming, verify whether it is supported by the operating system and applications.

> **Important:** To avoid possible issues when you replace a failed switch module, do *not* use automatic failback for NIC teaming. A newly installed switch module has no configuration data, and it can cause service disruption.

Virtual Link Aggregation Groups (VLAGs)

In many data center environments, downstream switches connect to upstream devices which consolidate traffic, as shown in Figure 5-2.

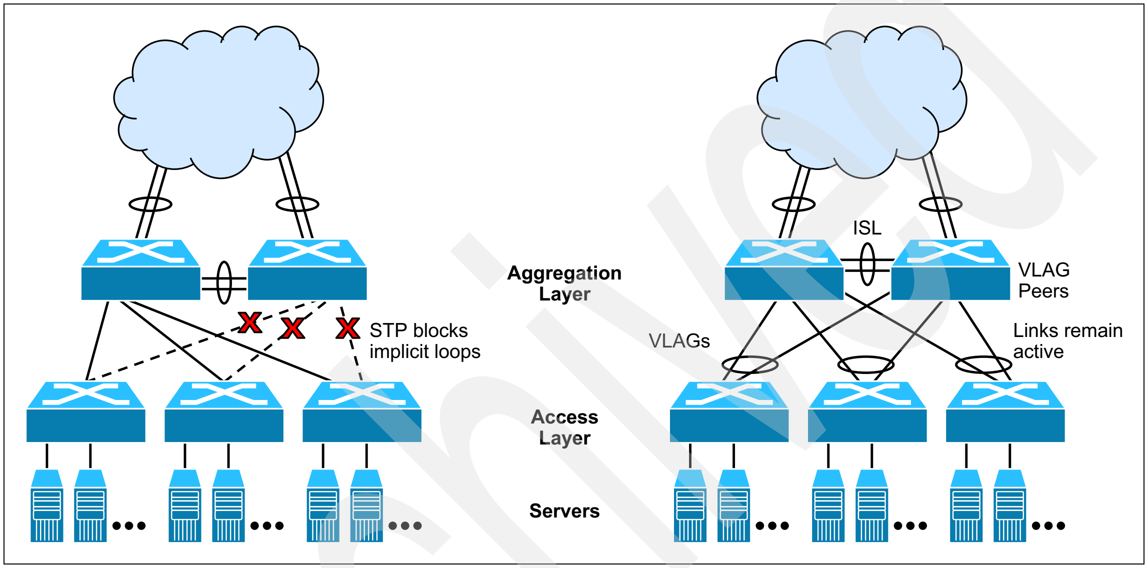

Figure 5-2 Typical switching layers with STP versus VLAG

A switch in the access layer might be connected to more than one switch in the aggregation layer to provide network redundancy. Typically, the Spanning Tree Protocol is used to prevent broadcast loops, blocking redundant uplink paths. This setup has the unwanted consequence of reducing the available bandwidth between the layers by as much as 50%. In addition, STP might be slow to resolve topology changes that occur during a link failure, which can result in considerable MAC address flooding.

Using Virtual Link Aggregation Groups (VLAGs), the redundant uplinks remain active, using all the available bandwidth. Using the VLAG feature, the paired VLAG peers appear to the downstream device as a single virtual entity for establishing a multiport trunk. The VLAG-capable switches synchronize their logical view of the access layer port structure and internally prevent implicit loops. The VLAG topology also responds more quickly to link failure and does not result in unnecessary MAC address flooding.

VLAGs are also useful in multi-layer environments for both uplink and downlink redundancy to any regular LAG-capable device, as shown in Figure 5-3.

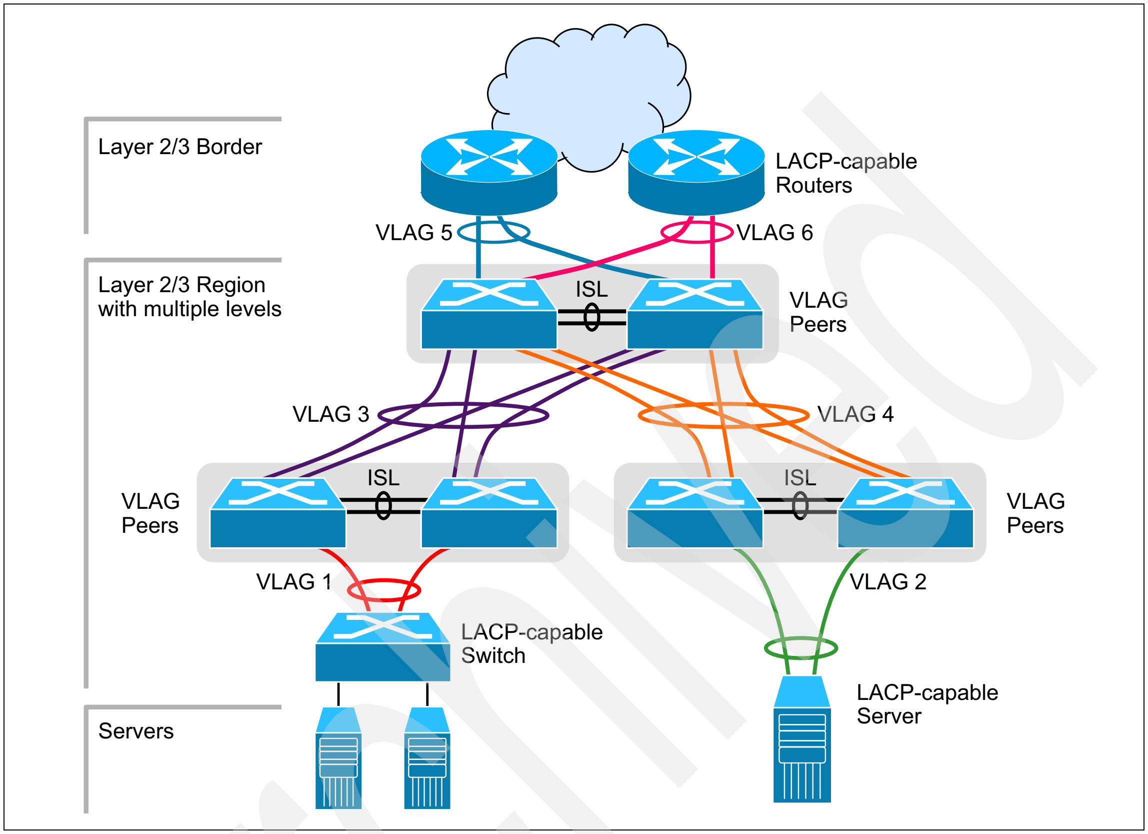

Figure 5-3 VLAG with multiple layers

5.4.2 SAN and Fibre Channel redundancy

SAN infrastructure availability can be achieved by implementing certain techniques and technologies. Most of them are widely used standards. This section describes the most common technologies that can be implemented in an IBM Flex System environment to provide high availability for a SAN infrastructure.

In general, a typical SAN infrastructure consists of storage FCs, client FCs, and SAN devices, such as SAN switches and the cables that connect them. The potential failures in a SAN include port failures (both on the switches and in storage), cable failures, and device failures.

Consider the scenario of dual-FC, dual-SAN switch redundancy, connected with storage attached through a SAN for an p460. In this scenario, the OS has four paths to each storage, and the behavior of the multipathing driver might vary, depending on the storage and switch type. This scenario is one of the best scenarios for high availability. The two adapters prevent an adapter fault, the two switches prevent the case of a switch fault or firmware upgrade, and, as the SAN has two paths to each storage device, the worst scenario is the failure of the complete storage. Figure 5-4 shows this scenario.

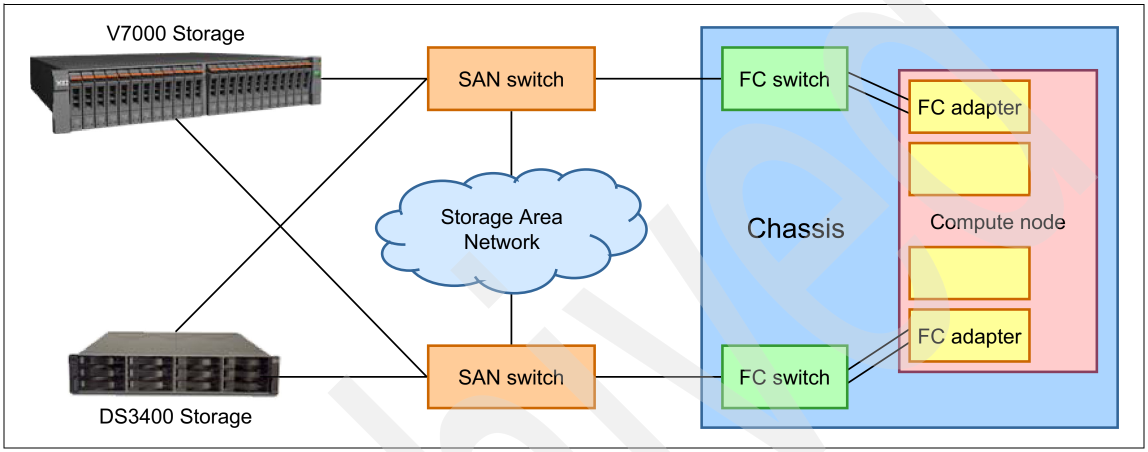

Figure 5-4 Dual-FC and dual-SAN switch redundancy connection, as applied to the IBM Flex System p460 Compute Node

This configuration might be improved by adding multiple paths from each Fibre Channel switch in the chassis to the SAN, which protects against a cable failure.

Another scenario for the p260 is one in which the redundancy in the configuration is provided by the Fibre Channel switches in the chassis. There is no hardware redundancy on the compute node, as it has only two expansion cards, with one used for Ethernet access and the other for Fibre Channel access. For this reason, in the case of Fibre Channel or Ethernet adapter failure on the compute node, redundancy is maintained. Figure 5-5 shows this scenario.

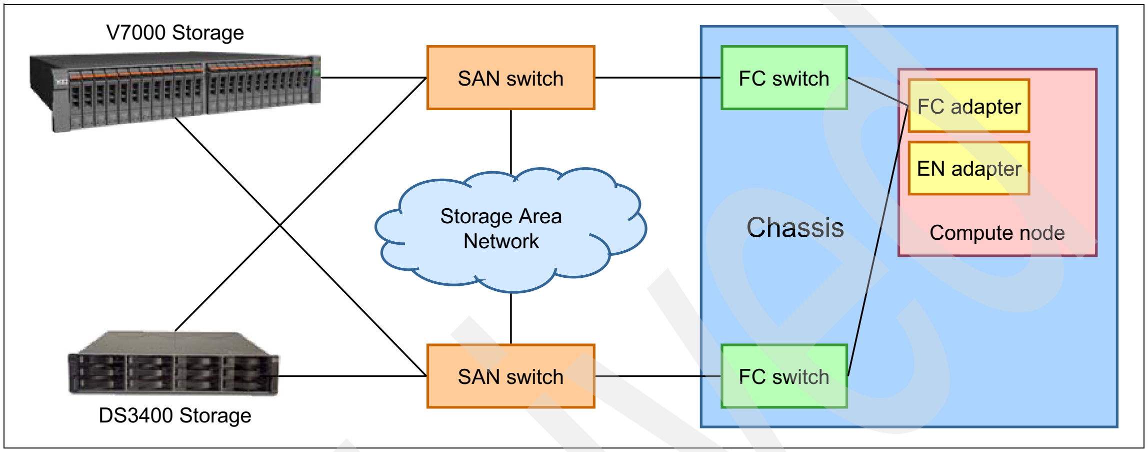

Figure 5-5 Dual-SAN switch connection with the IBM Flex System p260 Compute Node

5.5 Dual VIOS

Dual VIOS is supported in the Power Systems compute node. This configuration has certain limitations based on the computer node and the hardware installed on the computer node.

To have dual VIOS support, you need:

- A system managed by IBM Flex System Manager.
- An IBM Flex System p460 Compute Node. The p260 does not have enough expansion cards for dual VIOS.
- Three expansion cards (one Ethernet and two FC for each VIOS).

The p260 supports only two expansion cards. Internal disks are attached to only one PCI bus, so only one VIOS can manage internal disks. Therefore, to have dual VIOS, you need two Fibre Channel adapters and one Ethernet adapter, which is one more than the two supported on the p260.

Dual VIOS: The dual VIOS environment is not currently supported in the p260 or in the p24L. This feature might be added in future versions based on certain adapter configurations.

5.5.1 Dual VIOS on Power Systems compute nodes

One of the capabilities that is available with Power Systems compute nodes managed by an IBM Flex System Manager is the ability to implement dual Virtual I/O Servers in the same way as SDMC- or HMC-managed systems are able to.

Previously, with IVM-managed systems, the VIOS/IVM installed itself on virtual server number one, then additional virtual servers can be created by using IVM. No additional VIOS virtual servers are allowed to be created under this management method.

With IBM Flex System Manager, the creation of virtual servers and the type of operating system environment that they support can occur before any operating system installation. The only limitation from a dual VIOS perspective is the availability of disk and network physical resources. Physical resource assignment to a virtual server is made at the level of the expansion card slot or controller slot (physical location code). Individual ports and internal disks cannot be individually assigned. This type of assignment is not unique to Power Systems compute nodes and is a common practice for all Power platforms.

A dual VIOS environment setup requires the creation of the two virtual servers, both of which are set for a VIOS environment. After the virtual servers are created with the appropriate environment setting and physical resources assigned to support independent disk and network I/O, then the VIOS operating systems can be installed.

When planning a dual VIOS environment on a computer node, your hardware configuration requires two virtual servers, both of which must have a physical Ethernet connection and disk resources available. The following examples describe several of the possible hardware configurations to support a dual VIOS environment. These examples are not intended to be all-inclusive.

IBM Flex System p460 Compute Node

A typical basic configuration is 16 GB of memory, and a single internal disk. To support a dual VIOS environment, the following additional options are required:

- One Ethernet adapter (an EN4054 4-port 10Gb Ethernet Adapter or an EN2024 4-port 1Gb Ethernet Adapter), with two ports assigned to each VIOS. Two Ethernet adapters can be used to reduce single points of failure.

- Two Fibre Channel adapters (using FC3172 2-port 8Gb FC Adapters)
- One IBM Flex System Enterprise Chassis, with at least one Ethernet switch or pass-through node and one Fibre Channel switch or pass-through node.

As mentioned earlier, the four ports are assigned in pairs to each of the two VIOS virtual servers if only one Ethernet adapter is used, or each Ethernet adapter on the p460 is assigned to each VIOS if two Ethernet adapters are used. Each FC Card on the p460 is assigned to each VIOS.

Both VIOS servers in this example boot from SAN. The SAS controller and internal drive can be owned only by one VIOS and, in this example, could not be used.

This example for the p460, while not all-inclusive, provides the basics for a dual VIOS environment. Memory requirements for additional virtual servers beyond the base order amounts are not considered and need to be evaluated before ordering either model.

The actual steps of creating a dual VIOS are not covered here; however, the result of this type of configuration performed on a p460 is shown in Figure 5-6.

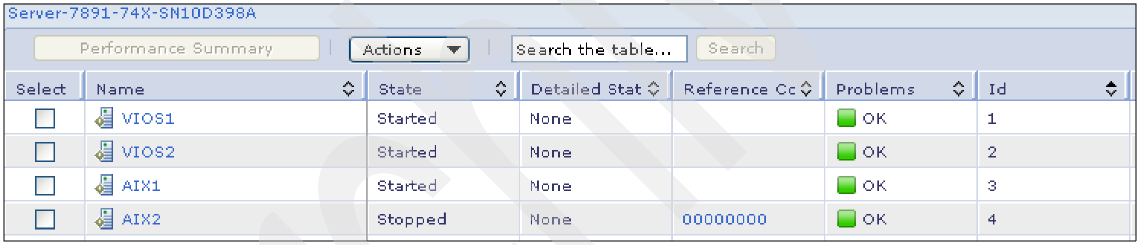

Figure 5-6 Dual VIOS configuration on an IBM Flex System p460 Compute Node

With the two virtual I/O servers installed, the normal methods of creating a Share Ethernet Adapter (SEA) failover for virtual networking and redundant paths for the client virtual server disks (N_Port ID Virtualization (NPIV) and virtual SCSI (vSCSI)) can be configured.

5.6 Power planning

When planning the power consumption for your Power Systems compute node, you must consider the server estimated power consumptions highs and lows based on the power supply features installed in the chassis and tools, such as Active Energy Manager. You can use these features to manage, measure, and monitor your energy consumption.

5.6.1 Power Systems compute node power supply features

The peak power consumption is 632 W for the IBM Flex System p260 Compute Node and 1264 W for the IBM Flex System p460 Compute Node. Power is provided by the chassis power supply modules. The maximum measured value is the worst-case power consumption expected from a fully populated server under an intensive workload. It also takes into account component tolerance and non-ideal operating conditions. Power consumption and heat load vary greatly by server configuration and use.

Use the IBM Systems Energy Estimator to obtain a heat output estimate based on a specific configuration. The Estimator is available at the following website:

http://www-912.ibm.com/see/EnergyEstimator

5.6.2 Power Systems compute node PDU and UPS planning

Planning considerations for your IBM Flex System configuration depend on your geographical location. Your need for power distribution units (PDUs) and uninterruptible power supply (UPS) units varies based on the electrical power that feeds your data center (AC or DC, 220 V or 110 V, and so on). These specifications define the PDUs, UPS units, cables, and support you need.

For information about planning your PDU and UPS configurations, see the *IBM Flex System Power Guide*, available at the following address:

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4401

The chassis power system is designed for efficiency using data center power, and consists of 3-phase, 60A Delta 200 VAC (North America) or 3-phase 32A wye 380-415 VAC (international). The Chassis may also be fed from single phase 200-240VAC supplies, if required.

Power cabling: 32A at 380-415V 3-phase (international)

As shown in Figure 5-7, one 3-phase 32A wye PDU (WW) can provide power feeds for two chassis. In this case, an appropriate 3-phase power cable is selected for the Ultra-Dense Enterprise PDU+, which then splits the phases, supplying one phase to each of the three PSUs within each chassis. One 3-phase 32A wye PDU can power two fully populated chassis within a rack. A second PDU may be added for power redundancy from an alternative power source, if the chassis is configured N+N.

Figure 5-7 shows a typical configuration with a 32A 3-phase wye supply at 380-415 VAC (often termed “WW” or “International”) N+N.

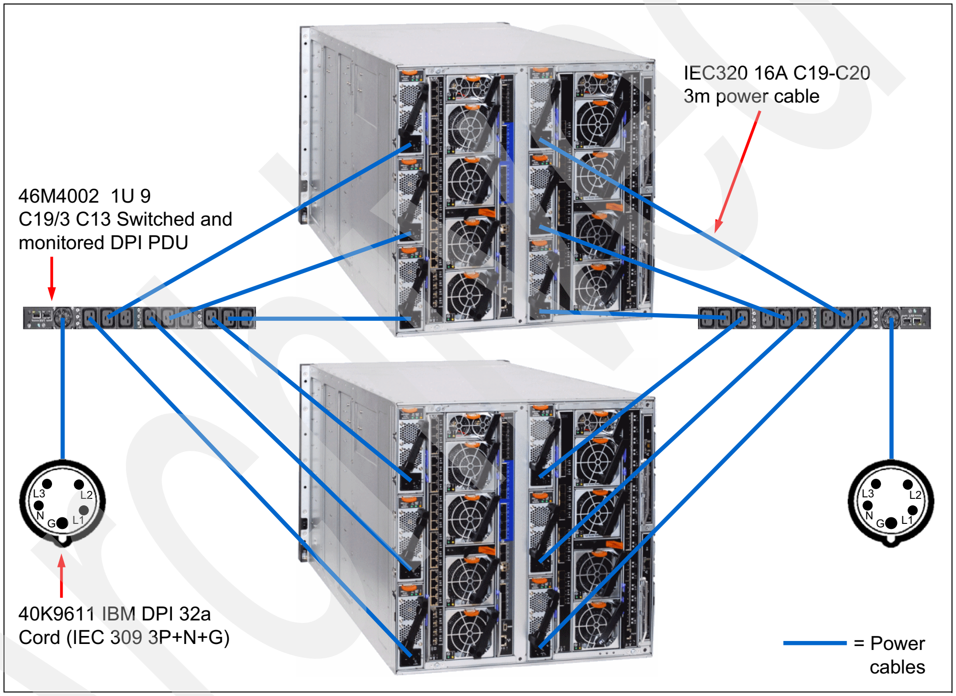

Figure 5-7 Example power cabling 32A at 380-415V 3-phase (international)

The maximum number of Enterprise Chassis that may be installed with a 42 U rack is four, so this configuration requires a total of four 32A 3-phase wye feeds into the rack, to provide for a fully redundant N+N configuration.

Power cabling: 60A at 208V 3-phase (North America)

In North America, this configuration requires four 60A 3-phase delta supplies at 200 - 208 VAC, so an optimized 3-phase configuration is shown in Figure 5-8.

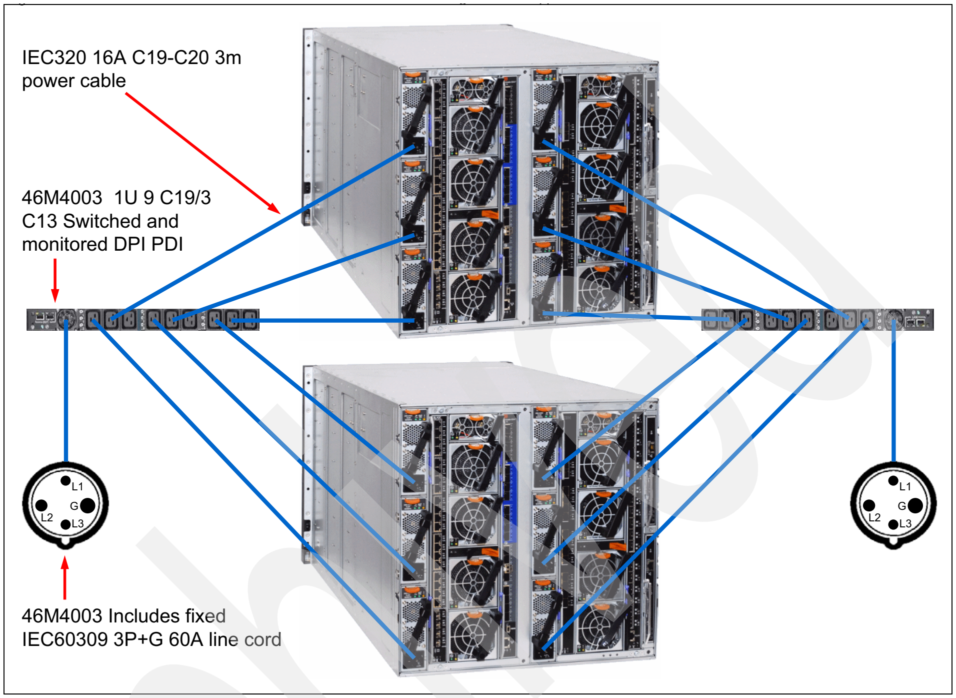

Figure 5-8 Example power cabling 60 A at 208 V 3-phase configuration

5.6.3 Chassis power supplies

The number of power supplies required depend on the number of nodes installed within a chassis and the level of redundancy required. When installing additional nodes, the power supplies are installed from the bottom upwards.

A maximum of six power supplies can be installed in the IBM Flex System Enterprise Chassis. The power supplies are 80 PLUS Platinum certified and are 2500 W output, rated at 200 VAC, with oversubscription to 3538 W output at 200 VAC. The power supplies also contain two independently powered 40 mm cooling fans.

80 PLUS is a performance specification for power supplies used within servers and computers. To meet the 80 PLUS standard, the power supply must have an efficiency of 80% or greater, at 20, 50, and 100 percent of rated load with a Power Factor (PF) of 0.09 or greater. The standard has several grades, such as Bronze, Silver, Gold, and Platinum. Further information about 80 PLUS is at the following website:

http://www.80PLUS.org

Figure 5-9 shows the location of the power supply bays at the rear of the Enterprise Chassis.

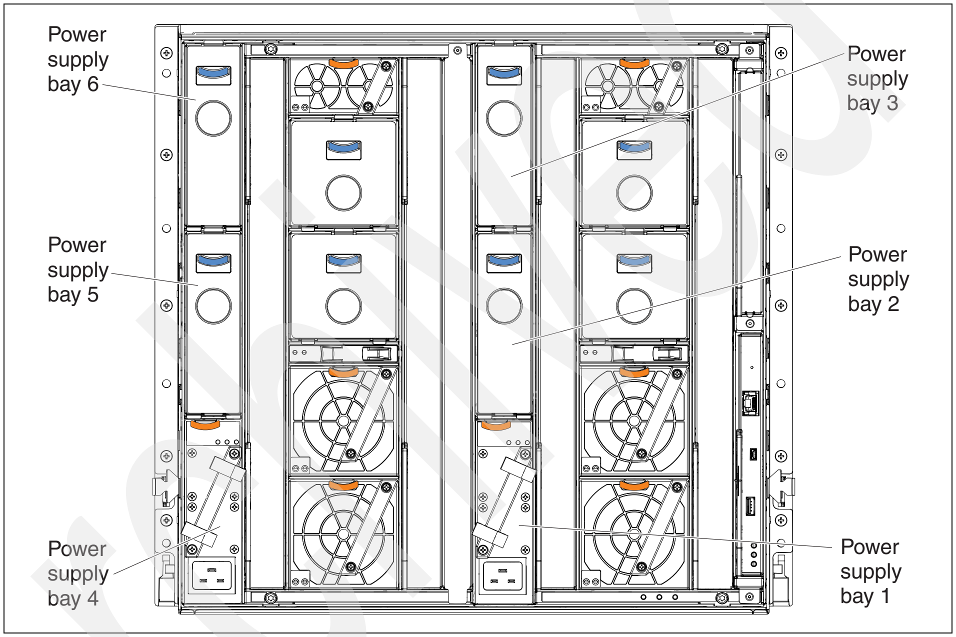

Figure 5-9 Power supply location

The Enterprise Chassis allows configurations of power supplies to provide N+N or N+1 redundancy. Three (or six, for redundancy) power supplies allow for a balanced, three-phase configuration.

All power supply modules are combined into a single power domain in the chassis, which distributes power to each compute node and I/O module through the Enterprise Chassis midplane. The midplane is a highly reliable design with no active components. Each power supply is designed to provide fault isolation.

Power monitoring of the AC and DC signals from the power supplies allows the Chassis Management Module to accurately monitor these signals.

The integral power supply fans are not dependent upon the power supply being functional. Rather, they operate and are powered independently from the midplane.

For detailed information about the power supply features of the chassis, see the *IBM Flex System Power Guide*, available at the following website:

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4401

5.6.4 Power supply management policies

There are two areas of power supply management that can be set using the Chassis Management Module: Power Redundancy and Power Capping/Limiting.

These settings can be accessed as shown in Figure 5-10. The Power Modules and Management window is shown.

AEM: More granular power controls can be set using Active Energy Manager, as described in 5.7.2, “Active Energy Manager” on page 151.

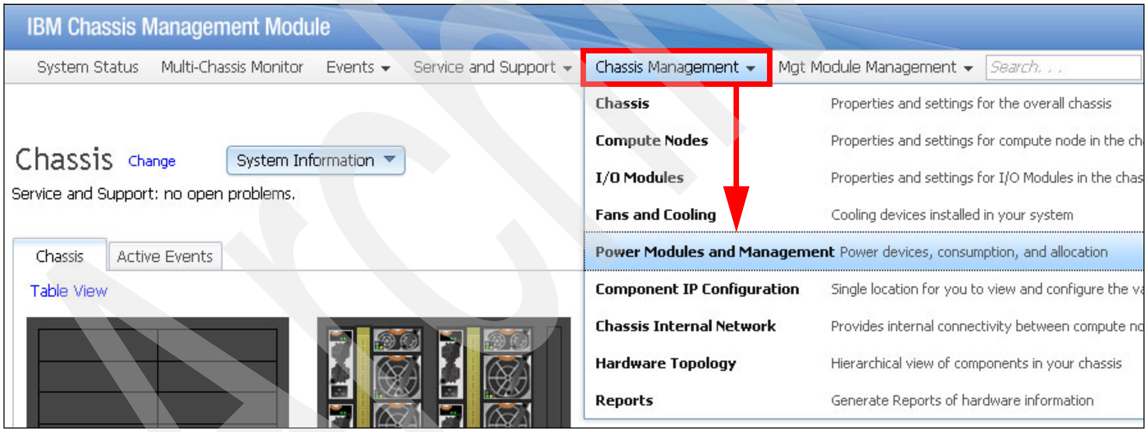

Figure 5-10 Accessing the Power Management options in the Chassis Management Module

Power redundancy settings

There are five power management redundancy settings available for selection:

- AC Power Source Redundancy

  Intended for dual AC power sources into the chassis. Maximum input power is limited to the capacity of two power modules. This approach is the most conservative one, and is best used when all four power modules are installed. When the chassis is correctly wired with dual AC power sources, one AC power source can fail without affecting compute node server operation. Some compute nodes may not be allowed to power on if doing so would exceed the policy power limit.

- AC Power Source Redundancy with Compute Node Throttling Allowed

  Similar to the AC Power Source Redundancy. This policy allows higher input power, and capable compute nodes may be allowed to throttle down if one AC power source fails.

- Power Module Redundancy

  Intended for a single AC power source in the chassis where each power module is on its own dedicated circuit. Maximum input power is limited to one less than the number of power modules when more than one power module is present. One power module can fail without affecting compute node operation. Multiple power module failures can cause the chassis to power off. Some compute nodes may not be allowed to power on if doing so would exceed the policy power limit.

- Power Module Redundancy with Compute Nodes Throttling Allowed

  Similar to Power Module Redundancy. This policy allows higher input power; however, capable compute nodes may be allowed to throttle down if one power module fails.

- Basic Power Management

  Maximum input power is higher than other policies and is limited only by the nameplate power of all the power modules combined. This approach is the least conservative one, because it does not provide any protection for an AC power source or power module failure. If any single power supply fails, the compute node or chassis operation might be affected.

The power redundancy options are shown in Figure 5-11. We clicked **Change** (next to Power Redundancy with Compute Node Throttling Policy) to show the power redundancy options.

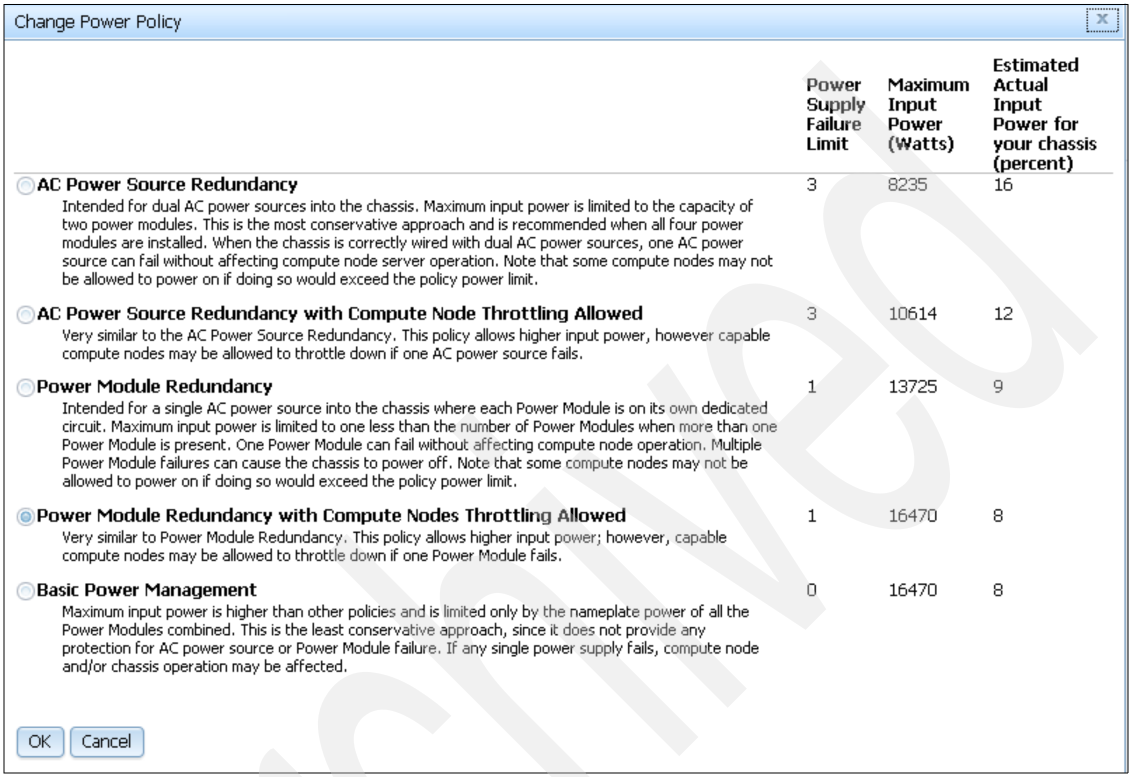

Figure 5-11 Changing the redundancy

5.6.5 Power limiting and capping policies

Simple power capping policies can be set to limit the amount of power consumed by the chassis. Two policy options are available, which you can configure with the Chassis Management Module:

- No Power Capping: The maximum input power is determined by the active Power Redundancy policy.
- Static Capping: Sets an overall chassis limit on the maximum input power. In a situation where powering on a component could cause the limit to be exceeded, the component cannot power on. No capping is the default. Static capping can be set as a percentage, starting with a minimum value. If there is insufficient power available to power on a compute node, the compute node does not come online.

The power capping options can be set as shown in Figure 5-12.

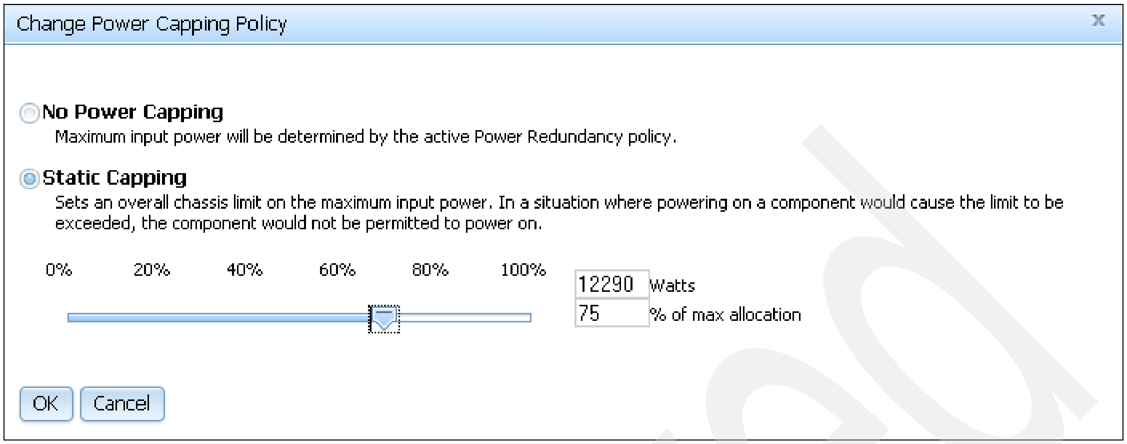

Figure 5-12 Setting power capping in the Chassis Management Module

5.6.6 Chassis power requirements

It is expected that the initial configuration (based on the IBM PureFlex System configuration that is ordered), plus any additional nodes, contains the necessary number of power supplies.

You need to know the number of power supplies needed to support the number of Power Systems compute nodes in the IBM Flex System Enterprise Chassis when adding a Power Systems compute node to an existing chassis. In addition, you must know the relationship between the number of Power Systems compute nodes and the number of power supplies in the chassis.

N+N redundancy

Table 5-6 lists the minimum number of power supplies required in a chassis to support the designated number of compute nodes, assuming average half-wide powers of 500 W, 600 W, and 700 W per compute node with a throttling policy enabled. If there is a a power fault, all compute nodes in the chassis must throttle the power to the average power indicated in the Fault column.

Table 5-6 Minimum number of N+N power supplies required to support half-wide compute nodes

Number of half-wide compute nodes	500 W ITE		600 W ITE		700 W ITE	
	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]
14	6	412 W	6	412 W	Not supported	

Number of half-wide compute nodes	500 W ITE		600 W ITE		700 W ITE	
	Number of power supplies[a]	**Fault[b]**	**Number of power supplies[a]**	**Fault[b]**	**Number of power supplies[a]**	**Fault[b]**
13	6	444 W	6	444 W	Not supported	
12	6	481 W	6	481 W	6	481 W
11	6	500 W	6	525 W	6	525 W
10	4	335 W	6	577 W	6	577 W
9	4	380 W	4	380 W	6	650 W
8	4	437 W	4	437 W	6	700 W
7	4	499 W	4	499 W	4	499 W
6	4	500 W	4	583 W	4	583 W
5	4	500 W	4	600 W	4	700 W
4	2	305 W	4	600 W	4	700 W
3	2	407 W	2	407 W	2	407 W
2	2	500 W	2	600 W	2	611 W
1	2	500 W	2	600 W	2	700 W

a. Theoretical number. Might require unrealistic throttle levels.
b. If there is a power fault, all compute nodes in the chassis must throttle the power to the average power indicated in the Fault column.

Table 5-7 shows similar information for full-wide compute nodes of 1000 W, 1200 W, and 1400 W with a throttling policy enabled. The number of supplies indicated in the table is theoretical, and might not represent a practical configuration. For example, although a specific number of supplies may be adequate to power the indicated configuration under normal operation, it may require the compute nodes to throttle the power to unrealistic or impossible levels during a fault to keep the system running.

Table 5-7 Minimum number of N+N power supplies required to support full-wide compute nodes

Number of full-wide compute nodes	1000 W ITE		1200 W ITE		1400 W ITE	
	Number of power supplies[a]	**Fault[b]**	**Number of power supplies[a]**	**Fault[b]**	**Number of power supplies[a]**	**Fault[b]**
7	6	824 W	6	824 W	Not supported	

Number of full-wide compute nodes	1000 W ITE		1200 W ITE		1400 W ITE	
	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]
6	6	962 W	6	962 W	6	962 W
5	4	670 W	6	1154 W	6	1154 W
4	4	874 W	4	874 W	6	1400 W
3	4	1000 W	4	1166 W	4	1166 W
2	2	610 W	4	1200 W	4	1400 W
1	2	1000 W	2	1200 W	2	1222 W

a. Theoretical number. Might require unrealistic throttle levels.
b. If there is a power fault, all compute nodes in the chassis must throttle the power to the average power indicated in the Fault column.

N+1 redundancy

Table 5-8 lists the minimum number of power supplies required in a chassis to support the designated number of compute nodes, assuming average half-wide powers of 400 W, 500 W, and 600 W per compute node. If there is a loss of one power supply, all compute nodes in the chassis must throttle the power to the average power indicated in the Fault column.

Table 5-8 Minimum number of N+1 power supplies required to support half-wide compute nodes

Number of half-wide compute nodes	500 W ITE		600 W ITE		700 W ITE	
	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]
14	4	418 W	4	418 W	5	594 W
13	4	450 W	4	450 W	5	640 W
12	4	488 W	4	488 W	4	488 W
11	4	500 W	4	532 W	4	532 W
10	3	339 W	4	585 W	4	585 W
9	3	385 W	3	385 W	4	658 W
8	3	442 W	3	442 W	3	442 W
7	3	500 W	3	505 W	3	505 W
6	2	500 W	3	589 W	3	589 W

Number of half-wide compute nodes	500 W ITE		600 W ITE		700 W ITE	
	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]
5	3	500 W	3	600 W	3	700 W
4	2	305 W	3	600 W	3	700 W
3	2	407 W	2	407 W	2	407 W
2	2	500 W	2	600 W	2	611 W
1	2	500 W	2	600 W	2	700 W

a. Theoretical number. Might require unrealistic throttle levels.
b. If there is a power fault, all compute nodes in the chassis must throttle the power to the average power indicated in the Fault column.

Table 5-9 shows similar information for full-wide compute nodes with average powers of 1000 W, 1200 W, and 1400 W. The number of supplies indicated in the table are theoretical, and might not represent a practical configuration. For example, although a specific number of supplies might be adequate to power the indicated configuration under normal operation, it may require the blades to throttle the power to unrealistic or impossible levels during a fault to keep the system running.

Table 5-9 Minimum number of N+1 power supplies required to support full-wide compute nodes

Number of full-wide compute nodes	1000 W ITE		1200 W ITE		1400 W ITE	
	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]	Number of power supplies[a]	Fault[b]
7	4	836 W	4	836 W	5	1188 W
6	4	976 W	4	976 W	4	976 W
5	3	678 W	4	1170 W	4	1170 W
4	3	884 W	3	884 W	3	884 W
3	3	1000 W	3	1178 W	3	1178 W
2	2	610 W	3	1200 W	3	1400 W
1	2	1000 W	2	1200 W	2	1222 W

a. Theoretical number. Might require unrealistic throttle levels.
b. If there is a power fault, all compute nodes in the chassis must throttle the power to the average power indicated in the Fault column.

5.7 Cooling

The flow of air within the Enterprise Chassis follows a front to back cooling path; cool air is drawn in at the front of the chassis and warm air is exhausted to the rear.

There are two cooling zones for the nodes: a left zone and a right zone.

The cooling is scaled up as required, based upon which node bays are populated. The number of cooling fans required for a given number of nodes is described further in this section.

Air is drawn in both through the front node bays and the front airflow inlet apertures, at the top and bottom of the chassis.

When a node is not inserted in a bay, an airflow damper closes in the midplane, meaning that absolutely no air is drawn in through an unpopulated bay. When a node is inserted into a bay, the damper is opened mechanically by insertion of the node, allowing for cooling of the node in that bay.

5.7.1 IBM Flex System Enterprise Chassis fan population

The fans are populated dependent on nodes installed. To support the base configuration and up to four half-wide nodes (or two full-wide nodes), a chassis ships with four 80 mm fans and two 40 mm fans preinstalled.

The minimum configuration of 80 mm fans is four, which provide cooling for a maximum of four half-wide nodes, as shown in Figure 5-13. This configuration is the base configuration.

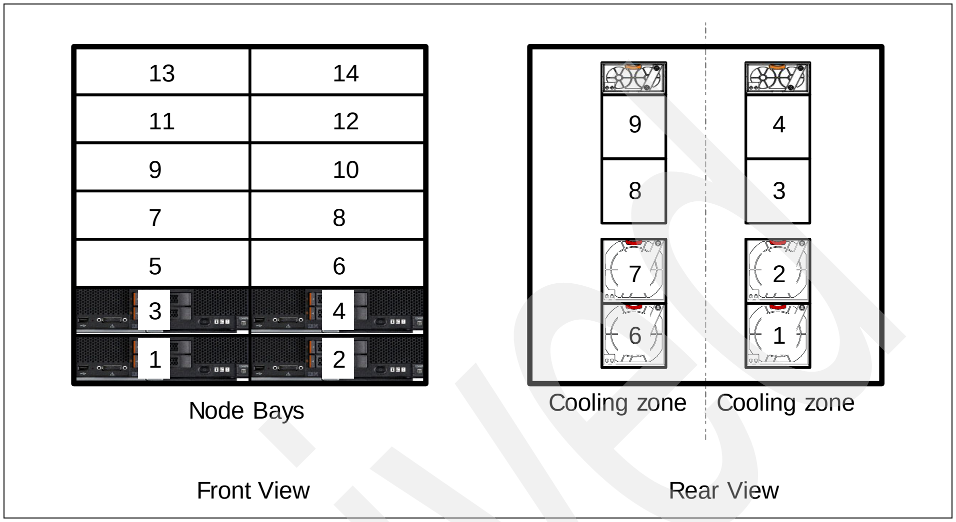

Figure 5-13 Four 80 mm fan modules support a maximum of four half-wide nodes installed

Six installed 80 mm fans support four more half-wide nodes within the chassis, to a maximum of eight, as shown in Figure 5-14.

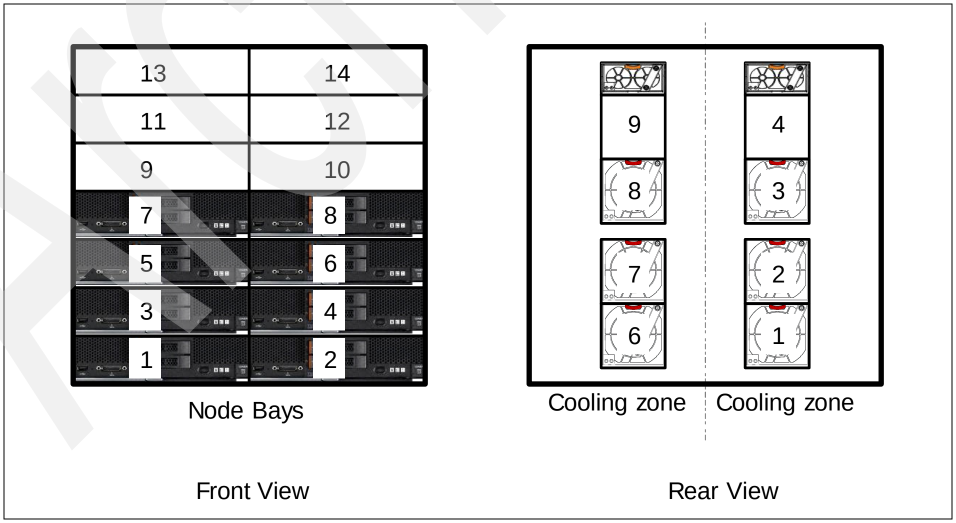

Figure 5-14 Six 80 mm fan modules support a maximum of eight half-wide nodes

To cool more than eight half-wide (or four full-wide) nodes, all the fans must be installed, as shown in Figure 5-15.

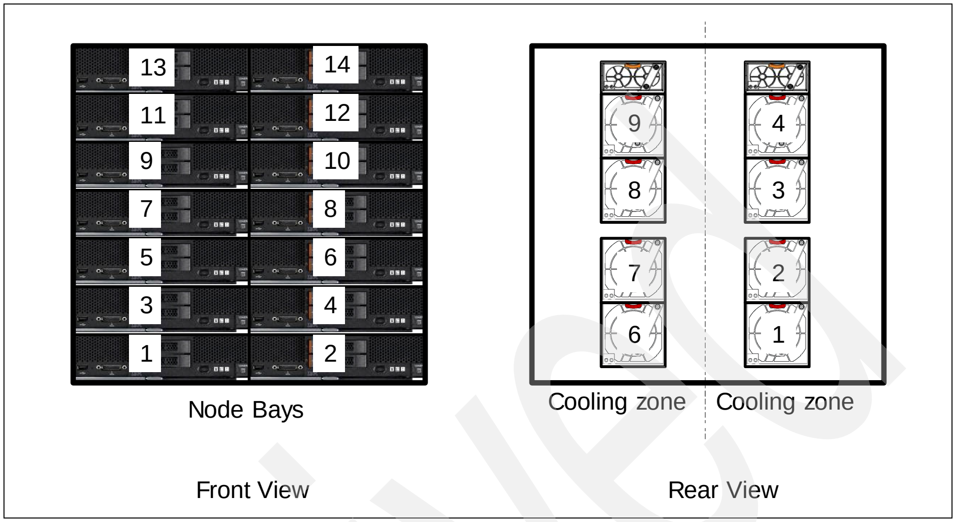

Figure 5-15 Eight 80 mm fan modules support from 7 to 14 nodes

5.7.2 Active Energy Manager

Active Energy Manager measures, monitors, and manages the energy components that are built into IBM Systems, enabling a cross-platform management solution. Active Energy Manager extends the scope of energy management so that facility providers have a complete view of energy consumption in the data center.

Active Energy Manager is an IBM Flex System Manager or SDMC extension that supports the following endpoints:

- IBM Flex System
- IBM BladeCenter
- Power Systems
- System x
- IBM System z®
- IBM System Storage® and non-IBM platforms[1]

In addition, Active Energy Manager can collect information from select facility providers, including Liebert SiteScan from Emerson Network Power and SynapSense.

[1] IBM storage systems and non-IBM platforms can be monitored through IBM or non-IBM (Raritan and Eaton) Power Distribution Unit (PDU) support.

Active Energy Manager uses agent-less technology, so no agents must be installed on the endpoints.

Monitoring and management functions apply to all IBM systems that are enabled for IBM Systems Director Active Energy Manager.

- Monitoring functions include power trending, thermal trending, IBM and non-IBM PDU support, support for facility providers, energy thresholds, and altitude input.
- Management functions include power capping and power savings mode.

Active Energy Manager also provides a source of energy management data that can be used by Tivoli enterprise solutions, such as IBM Tivoli Monitoring. For more information about IBM Tivoli Monitoring, go to the following website:

http://www.ibm.com/software/tivoli/products/monitor/

5.7.3 Supported environment

The p260 and p460 and the Enterprise Chassis comply with ASHRAE Class A3 specifications.

Here are the specifications for the supported operating environment:

- 5 - 40 °C (41 - 104 °F) at 0 - 914 m (0 - 3,000 ft)
- 5 - 28 °C (41 - 82 °F) at 914 - 3,050 m (3,000 - 10,000 ft)
- Relative humidity: 8 - 85%
- Maximum altitude: 3,050 m (10,000 ft)

5.8 Planning for virtualization

The new Power Systems compute nodes provide features available in high-end POWER servers, such as virtualization, when connected to the IBM Flex System Manager. You can use virtualization to create and manage virtual servers and take full advantage of the PowerVM virtualization features, such as IBM Micro-Partitioning®, AMS, NPIV, and Live Partition Mobility.

To partition your Power Systems compute node, it must be attached to the IBM Flex System Manager. The process for connecting your Power Systems compute node to both nodes is described in 7.1.3, “Preparing to use the IBM Flex System Manager for partitioning” on page 284.

The key element for planning your partitioning is knowing the hardware you have in your Power Systems compute node, as that hardware is the only limit that you have for your virtual servers. Adding VIOS to the equation solves much those limitations.

Support for IVM: IVM is not supported on the Power Systems compute nodes in IBM Flex System.

5.8.1 Virtual servers without VIOS

Preconfigured VIOS: The VIOS that is preconfigured must be removed when working with a Power Systems compute node that is part of an initial order.

Partitions on a Power Systems compute node without VIOS might be available on certain configurations, as described in the configuration examples that follow:

- Sample Configuration 1:

  One IBM Flex System p260 Compute Node, with one IBM Flex System EN2024 4-port 1Gb Ethernet Adapter, 32 GB of memory, internal disks, and an IBM Flex System FC3172 2-port 8Gb FC Adapter.

  In this sample, you can create two partitions:

  Virtual Server 1 consists of:

  – One processor
  – 16 GB
  – Internal disks
  – One port on the IBM Flex System EN2024 4-port 1Gb Ethernet Adapter
  – AIX operating system

  Virtual Server 2 consists of:

  – One processor
  – 16 GB
  – SAN-attached disks through the IBM Flex System FC3172 2-port 8Gb FC Adapter
  – One port on the IBM Flex System EN2024 4-port 1Gb Ethernet Adapter
  – Linux operating system

- Sample Configuration 2:

  One IBM Flex System p460 Compute Node, with two IBM Flex System EN2024 4-port 1Gb Ethernet Adapters, 32 GB of memory, and two IBM Flex System FC3172 2-port 8Gb FC Adapters.

  In this sample, you can create two partitions:

  Virtual Server 1 consists of:

  – One processor
  – 10 GB
  – SAN-attached disks through the IBM Flex System FC3172 2-port 8Gb FC Adapter
  – One IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking
  – AIX operating system

- Virtual Server 2 consists of:

  – One processor
  – 22 GB
  – SAN-attached disks through the IBM Flex System FC3172 2-port 8Gb FC Adapter
  – One IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking
  – AIX operating system

- Sample Configuration 3:

  One IBM Flex System p460 Compute Node, with 1 IBM Flex System EN2024 4-port 1Gb Ethernet Adapter, 32 GB of memory, internal disks, and two IBM Flex System FC3172 2-port 8Gb FC Adapter.

  In this sample, you can create three partitions:

  Virtual Server 1 consists of:

  – 0.5 processor
  – 8 GB
  – Internal disks
  – One port from an IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking
  – AIX operating system

Virtual Server 2 consists of:

- 0.5 processor
- 10 GB
- SAN-attached disks through an IBM Flex System FC3172 2-port 8Gb FC Adapter
- One port of the IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking
- Linux operating system

Virtual Server 3 consists of:

- One processor
- 14 GB
- SAN-attached disks through an IBM Flex System FC3172 2-port 8Gb FC Adapter
- One port of the IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking
- AIX operating system

> **Important:** Configurations shown in the following samples are not the only configurations supported. You can use several combinations of expansion cards and memory; the limitations are disk and network access.

5.8.2 Virtual server with VIOS

You can use IBM Flex System Manager management to install a dual VIOS, as described in 5.5, “Dual VIOS” on page 135. Setting up a VIOS environment is the key to overcoming the hardware limitations you might have on your Power Systems compute node. This environment supports:

- Up to 160 Virtual Servers on p260
- Up to 320 Virtual Servers on p460
- Up to 120 Virtual Servers on p24L

As explained in “Virtualized environment planning” on page 123, you can install only one VIOS server on the IBM Flex System p260 Compute Node. You can install two VIOS servers on the IBM Flex System p460 Compute Nodes, depending on your hardware configuration.

VIOS can solve many of the hardware limitations (buses, cards, disk, and memory) you find when creating virtual servers on your Power Systems compute node. For more information, see Chapter 7, “Virtualization” on page 275.

Sample configurations for VIOS installations are:

- Sample Configuration 1:

  One IBM Flex System p460 Compute Node, with two IBM Flex System EN2024 4-port 1Gb Ethernet Adapters, 32 GB of memory, and two IBM Flex System FC3172 2-port 8Gb FC Adapters.

  For this sample, you can create two VIOS servers:

  VIOS Server 1 consists of:

  – One processor
  – 4 GB
  – SAN-attached disks through the IBM Flex System FC3172 2-port 8Gb FC Adapter
  – One IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking

  Virtual Server 2 consists of:

  – One processor
  – 4 GB
  – SAN-attached disks through the IBM Flex System FC3172 2-port 8Gb FC Adapter
  – One IBM Flex System EN2024 4-port 1Gb Ethernet Adapter for networking

The VIOS virtual servers should be configured for redundant access to storage and the network.

Additional AIX, Linux, or IBM i client virtual servers can now be configured by using resources from the VIO virtual servers, with the assurance that the loss of a VIOS does not result in a client losing access to storage or the network.

6

Management setup

The IBM Flex System Enterprise Chassis brings a whole new approach to management. This approach is based on a global management appliance, the IBM Flex System Manager (FSM), which you can use to view and manage functions for all of your Enterprise Chassis components. These components include the Chassis Management Module, I/O modules, computer nodes, and storage.

Methods of system management, such as direct access to the Chassis Management Module and the I/O modules, are still supported. Setup and access to each system management method are described in 6.2, "Chassis Management Module" on page 161.

You can manage your Enterprise Chassis more proficiently with the Chassis Management Module and IBM Flex System Manager.

We cover the following topics in this chapter:

- IBM Flex System Enterprise Chassis security
- Chassis Management Module
- Management network
- IBM Flex System Manager
- FSM initial setup
- Basic management of Power Systems compute nodes
- IBM Flex System Manager options and tasks

6.1 IBM Flex System Enterprise Chassis security

The focus of IBM on smarter computing is evident in the improved security measures implemented in the IBM Flex System Enterprise Chassis. Today's world of computing demands tighter security standards and native integration with computing platforms. For example, the virtualization movement increased the need for a high degree of security, as more mission-critical workloads are consolidated to fewer and more powerful servers. The IBM Flex System Enterprise Chassis takes a new approach to security with a ground-up chassis management design to meet new Trusted Computing Group (TCG) security standards.

Here are additional security enhancements and features in the Enterprise Chassis:

- Single sign-on (central user management).
- End-to-end audit logs.
- Secure boot with Trusted Platform Module (TPM) and Core Root of Trust Measurement (CRTM), also known as BIOS Bootblock.
- Intel processor-based compute nodes using Intel Trusted Execution Technology (TXT) technology.
- Signed firmware updates to ensure authenticity.
- Secure communications.
- Certificate authority and management.
- Chassis and compute node detection and provisioning.
- Role-based access control.
- Security policy management.
- Management protocols that are the same as the protocols supported in the BladeCenter AMM for compatibility with earlier versions.
- Non-secure protocols are disabled by default in the Chassis Management Module, with lock settings to prevent inadvertent or malicious enabling.
- Supports up to 84 local Chassis Management Module user accounts.
- Supports up to 32 simultaneous sessions.
- Planned support for Dynamic Root of Trust Measurement (DRTM).

The Enterprise Chassis ships with secure settings by default, with two security policy settings supported:

- Secure: This setting is the default one. It ensures a secure chassis infrastructure that supports:
  – Strong password policies with automatic validation and verification checks
  – Updated passwords that replace the default passwords after initial setup
  – Secure communication protocols, such as SSH, SSL, and HTTPS
  – Certificates to establish secure and trusted connections for applications that run on the management processors
- Legacy: This setting provides flexibility in chassis security. It provides:
  – Weak password policies with minimal controls
  – Manufacturing default passwords that do not have to be changed
  – Unencrypted communication protocols, such as Telnet, SNMP v1, TCP command mode, CIM-XML, FTP server, and TFTP server

The centralized security policy makes the Enterprise Chassis easy to configure. In essence, all components run the same security policy that is provided by the Chassis Management Module. This configuration ensures that all I/O modules run with a hardened attack surface, as shown in Figure 6-1.

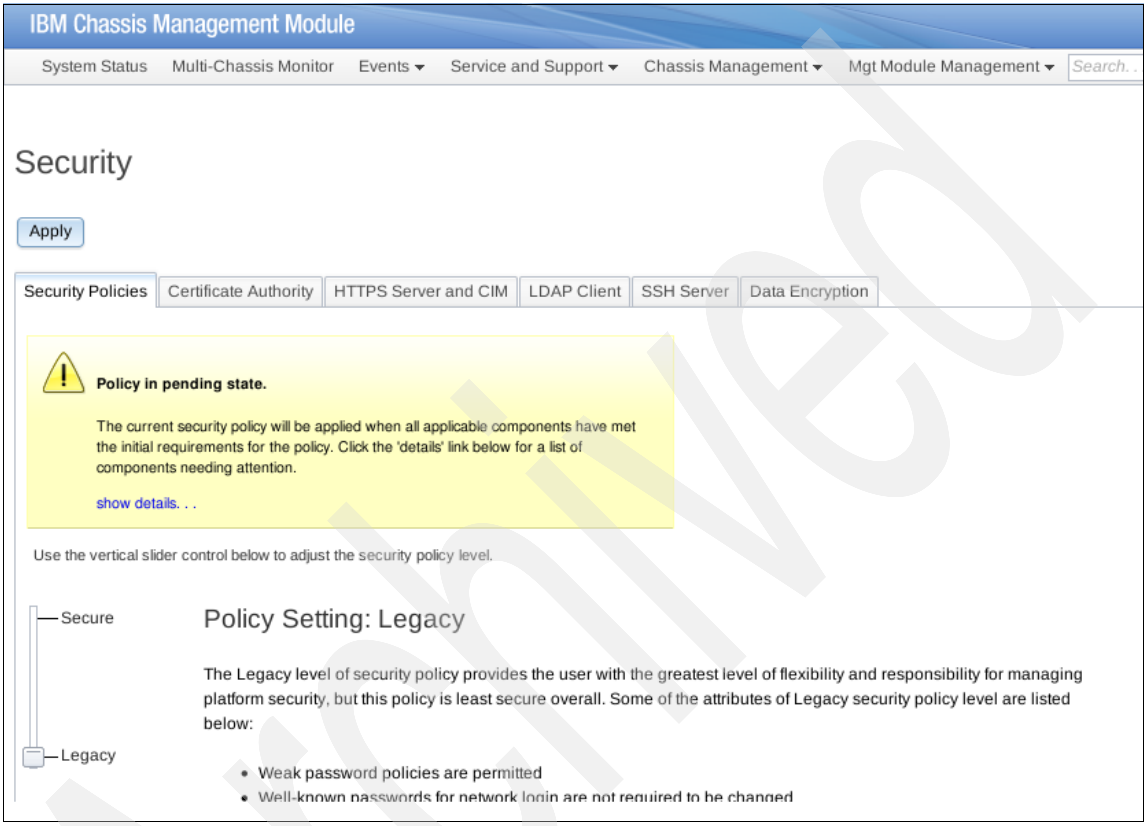

Figure 6-1 The basic security configuration of Enterprise Chassis, as shipped

Figure 6-2 shows a sample configuration of HTTPS access to the Chassis Management Module.

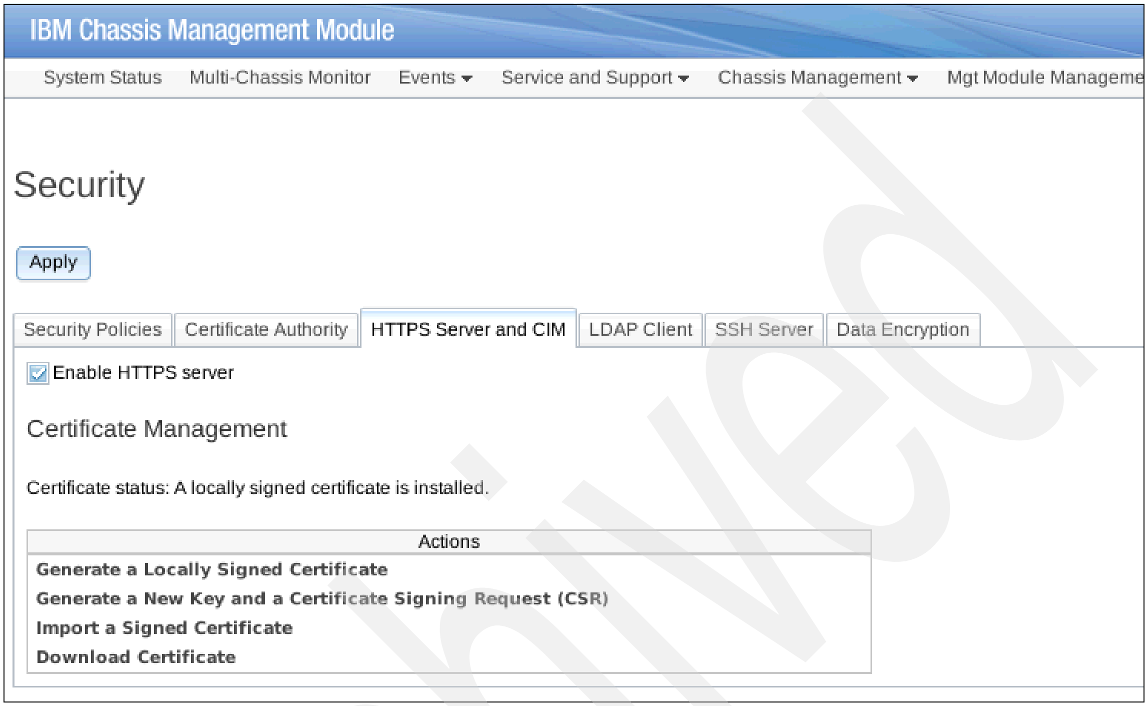

Figure 6-2 HTTPS setup

6.2 Chassis Management Module

The Chassis Management Module manages hardware elements within a single chassis. As such, the Chassis Management Module is central to chassis management and is required in the Enterprise Chassis.

The Chassis Management Modules are inserted in the back of the chassis, and are vertically oriented. When looking at the back of the chassis, the Chassis Management Module bays are on the far right. The Chassis Manager tab in FSM shows this configuration clearly, as shown in Figure 6-3.

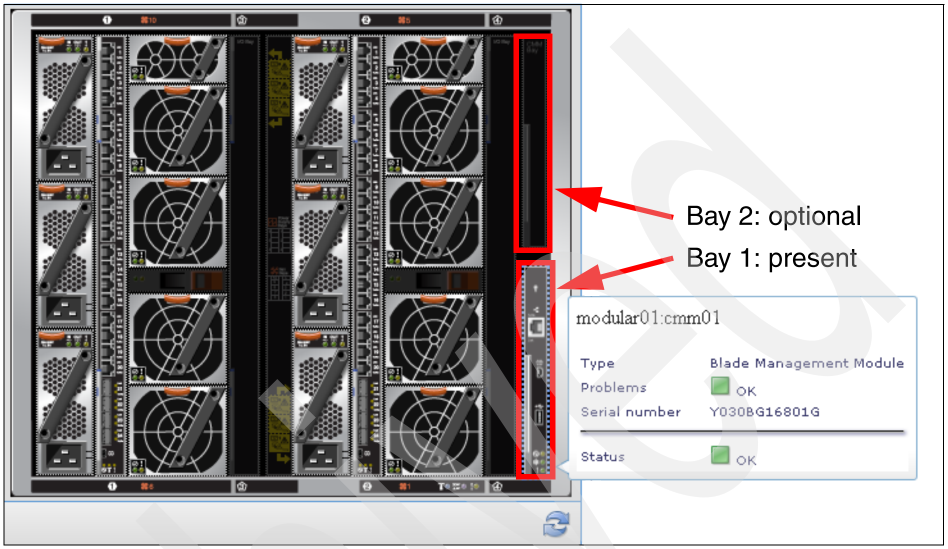

Figure 6-3 FSM chassis manager table view that shows the Chassis Management Module bays

The Chassis Management Module can also be seen in the Chassis Management Module GUI, as shown in Figure 6-4.

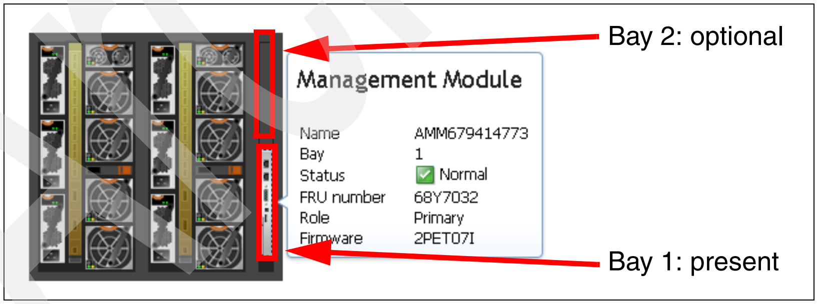

Figure 6-4 View of the Chassis Management Module in the GUI

The following section describes the usage models and features of the Chassis Management Module.

For a hardware overview of the CMM, see *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984.

6.2.1 Overview of the Chassis Management Module

The Chassis Management Module (CMM) is a hot-swap module that provides system management functions for all devices installed in the Enterprise Chassis. An Enterprise Chassis comes with at least one CMM and supports module redundancy. Only one module is active at a time.

Systems management of either a single chassis or multiple chassis can be performed directly through the CMM or through the IBM Flex System Manager. Using the CMM is only useful for base management tasks. We describe IBM Flex System Manager in 6.4, "IBM Flex System Manager" on page 192.

Through an embedded firmware stack, the CMM implements functions to monitor, control, and provide external user interfaces to manage all chassis resources. These functions include:

- Functionality similar to Advanced Management Module (AMM) in BladeCenter
- Single chassis management
- The Chassis Management Module initial setup wizard
- An IMMv2-based hardware engine
- Optional redundancy
- A Chassis Management Module service advisor
- Compatibility with earlier versions of BladeCenter CLI and SNMP
- Support for IBM Feature on Demand
- Support for IBM Fabric Manager

The CMM automatically detects installed modules in the Enterprise Chassis and stores vital product data (VPD) on them.

Chassis Management Module default network information

The CMM does not have a fixed static IPv6 address by default. Initial access to the CMM in an IPv6 environment can be done by either using the IPv4 IP address or the IPv6 link-local address. The IPv6 link-local address is automatically generated, based on the MAC address of the Chassis Management Module. By default, the Chassis Management Module is configured to respond to Dynamic Host Configuration Protocol (DHCP) first, before using a static IPv4 address.

If a DHCP response is not received within 3 minutes of the CMM Ethernet port being connected to the network, the CMM uses the factory default IP address and subnet mask. During this 3-minute interval, the CMM is not accessible.

The CMM has the following default settings:

- IP address: 192.168.70.100
- Subnet: 255.255.255.0
- User ID: USERID (all capital letters)
- Password: PASSW0RD (all capital letters, with a zero instead of the letter O)

The default information, MAC address, and IPv6 link-local address is available in the network access card attached to all new CMMs, as shown in Figure 6-5.

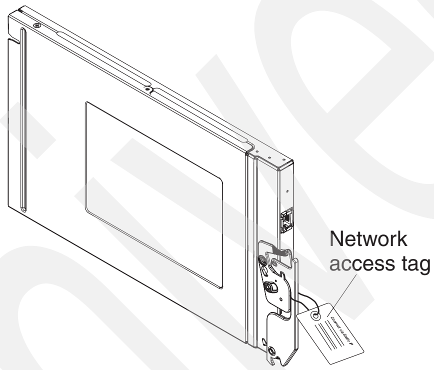

Figure 6-5 Chassis Management Module network access tag

Initial configuration of the Chassis Management Module

The CMM is accessible by using a browser and either the DHCP-obtained address or the default IP settings. The default security setting is *secure*, so HTTPS is required to connect to the CMM.

To perform the initial configuration, complete the following steps:

1. Open a browser and go to the IP address of the CMM, either the DHCP-obtained address or the default IP settings. The Login window opens, as shown in Figure 6-6.

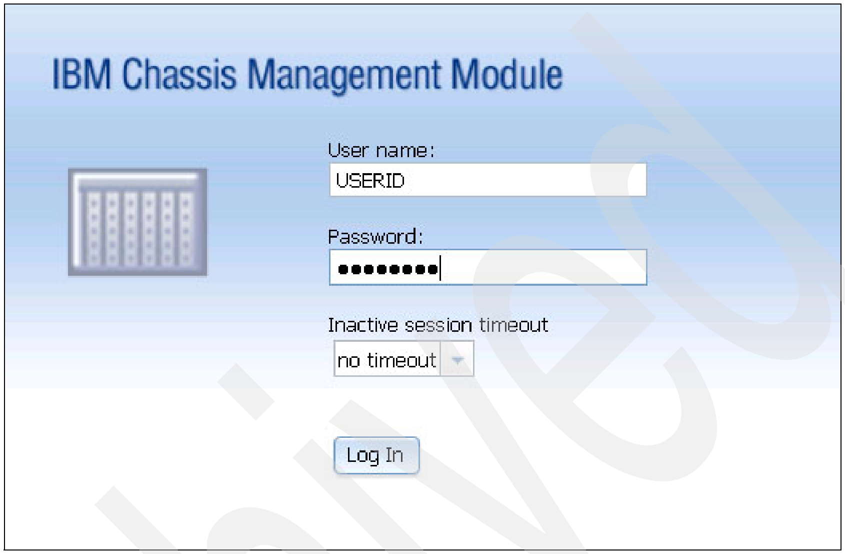

Figure 6-6 CMM login

2. Log in with the default user ID and password, and a window opens that shows the system status information and a graphic view of the chassis (see Figure 6-11 on page 169). Across the top of this window are groups that can be expanded for specific functions. The initial setup wizard is contained in the Mgt. Module Group in the Configuration function, as shown in Figure 6-7.

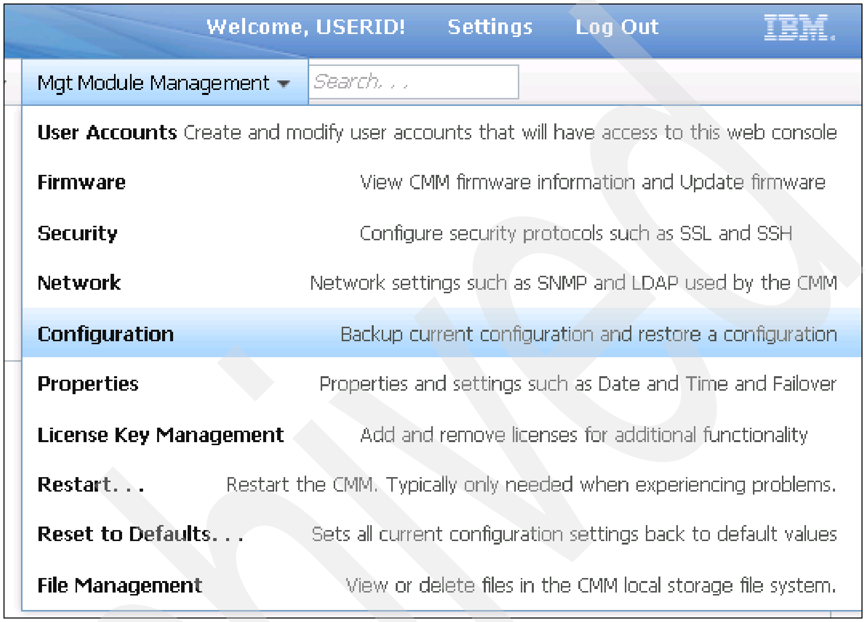

Figure 6-7 CMM management group

3. Several options are shown that can be use to manage the Chassis Management Module configuration. For this first time connection, click **Initial Setup Wizard**, as shown in Figure 6-8.

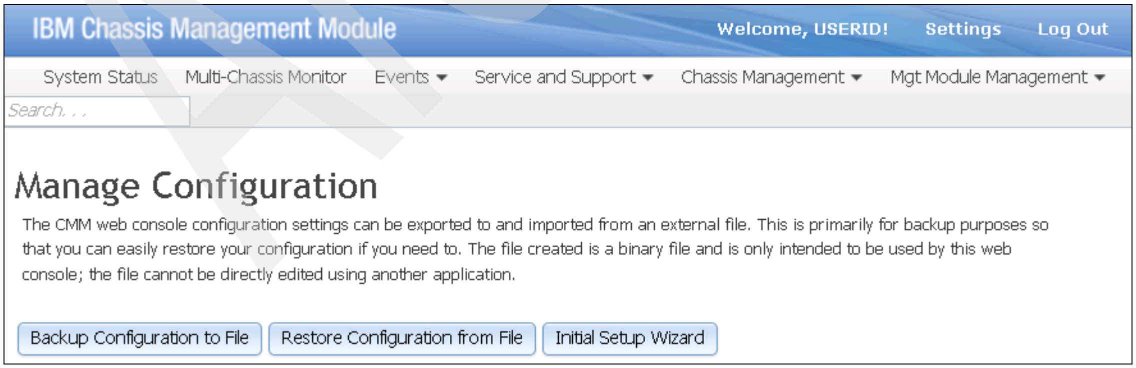

Figure 6-8 CMM configuration management

When the wizard starts, the first window shows the steps that are performed on the left side of the window, and a basic description of the steps in the main field.

Figure 6-9 shows the Welcome window of the setup wizard. This wizard is similar to other IBM wizards. Navigation buttons for the wizard are in the lower left of each window.

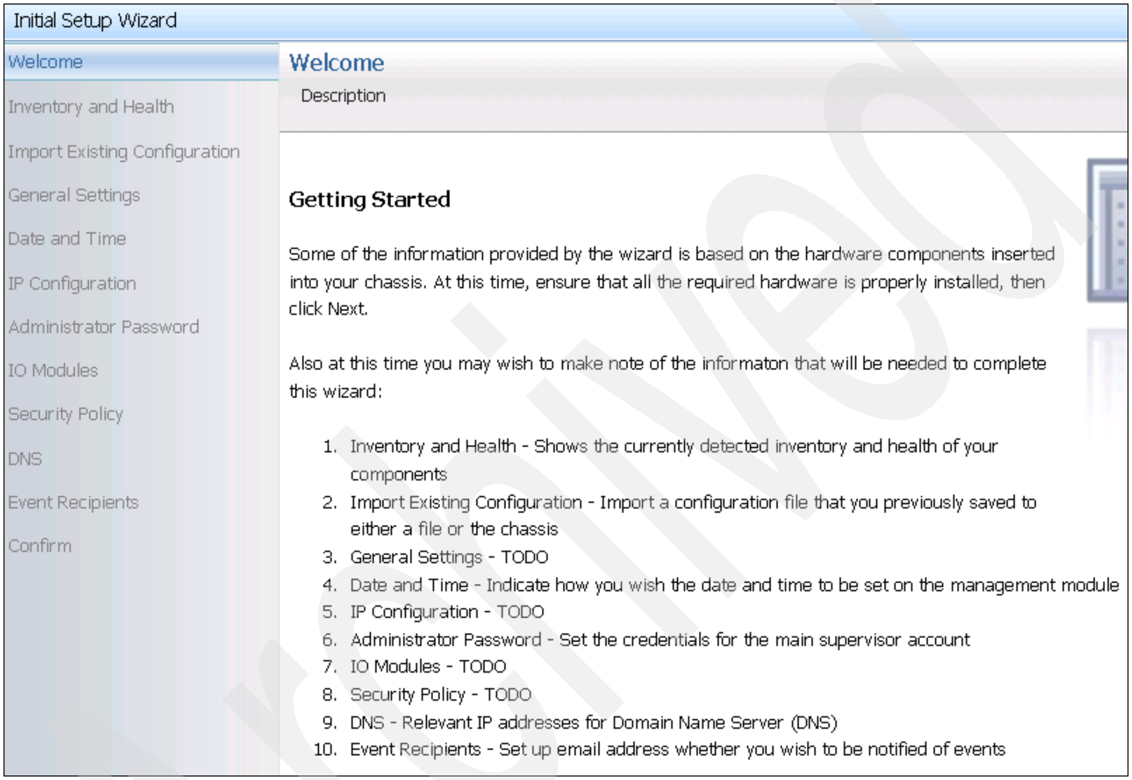

Figure 6-9 Chassis Management Module initial setup wizard Welcome window

4. Proceed through each step of the wizard by clicking **Next**, entering the information as required.
5. Click **Finish** to complete the last step of the wizard.

For further details about using the Chassis Management Module, see 6.2.2, “CMM functions” on page 168.

6.2.2 CMM functions

The Chassis Management Module web interface has a menu structure at the top of each page that gives you easy access to most functions, as shown in Figure 6-10.

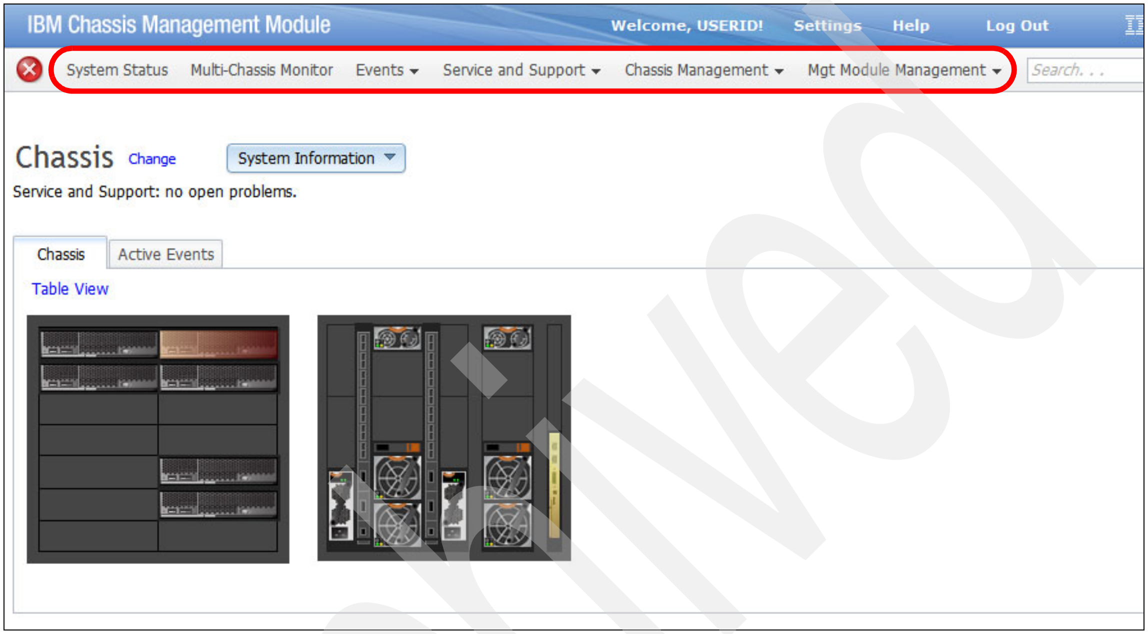

Figure 6-10 CMM home page

The menus are as follows:

- System Status tab
- Multi-Chassis Monitor tab
- Events tab
- Service and Support tab
- Chassis Management tab
- Mgmt Module Management tab

System Status tab

The System Status tab shows the System Status window, which is the default window when you enter the CMM web interface (Figure 6-11). You can also access this window by clicking **System Status**. This window shows a graphical systems view of a selected chassis, active events, and general systems information.

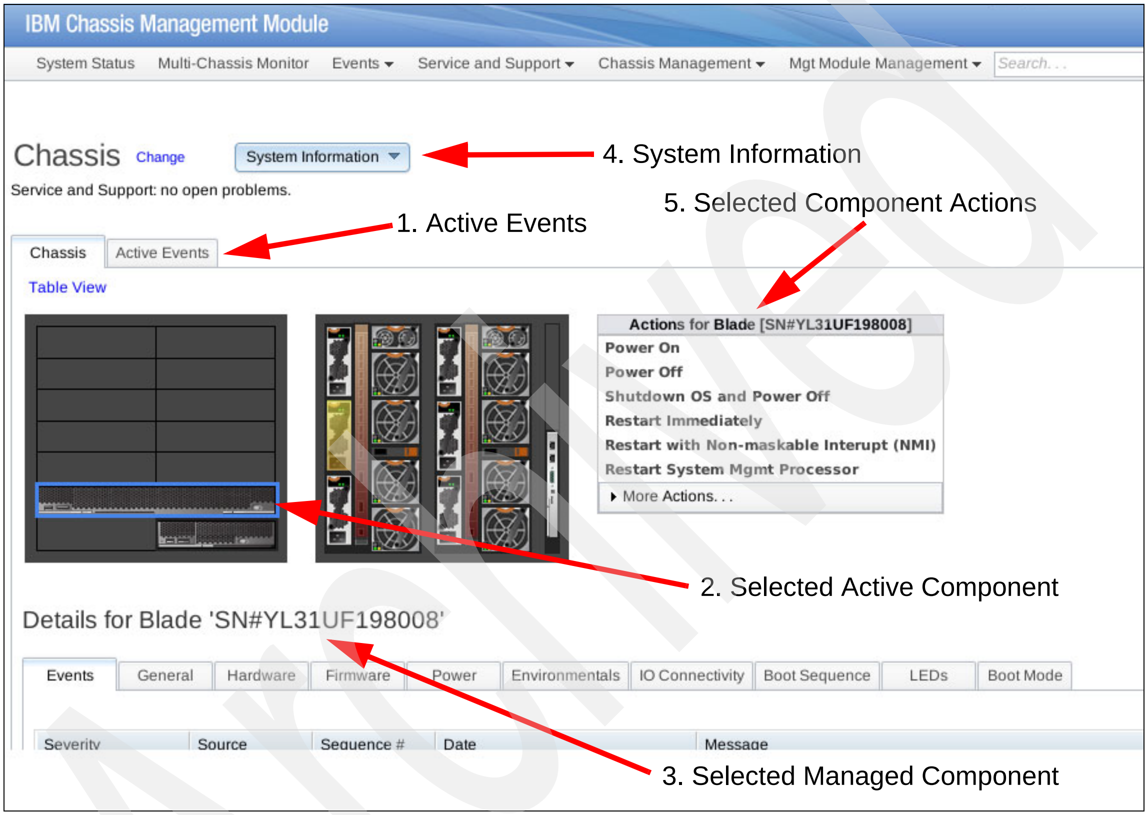

Figure 6-11 System Status tab

The graphical view of the chassis is active, so changes are reflected immediately. The available selections are (with the numbers that match callouts in Figure 6-11):

1. Active Events: Shows the active events from the chassis, including errors, informational events, and so on. Click an event to display detailed information.

2. Selected active components: All major components of the chassis can be clicked for more information. Select a component of interest (in Figure 6-11 on page 169, we select the IBM Flex System p460 Compute Node), and a dialog box opens with information about that component, for example, serial number, name, bay, and so on. You can power a component on or off in this dialog box, or view other details about the component.
3. Selected managed component: With a component selected, this dialog box shows several tabs with additional information, for example, events, hardware, firmware, and LEDs. Other components, such as I/O modules and computer nodes, and power options, are available.
4. System Information. This selection shows a quick view of chassis information in a drop-down menu.
5. Selected Component Actions: I/O modules and computer nodes activate the actions menu, from which you can power on and off, restart, and perform several other options. The More actions link provides additional component-specific actions.

Multi-Chassis Monitor tab

In the Multi-Chassis Monitor tab, you can manage and monitor other IBM Flex System chassis, as shown in Figure 6-12. Click the chassis name to show details about that chassis. Click the link to start the CMM interface for that chassis.

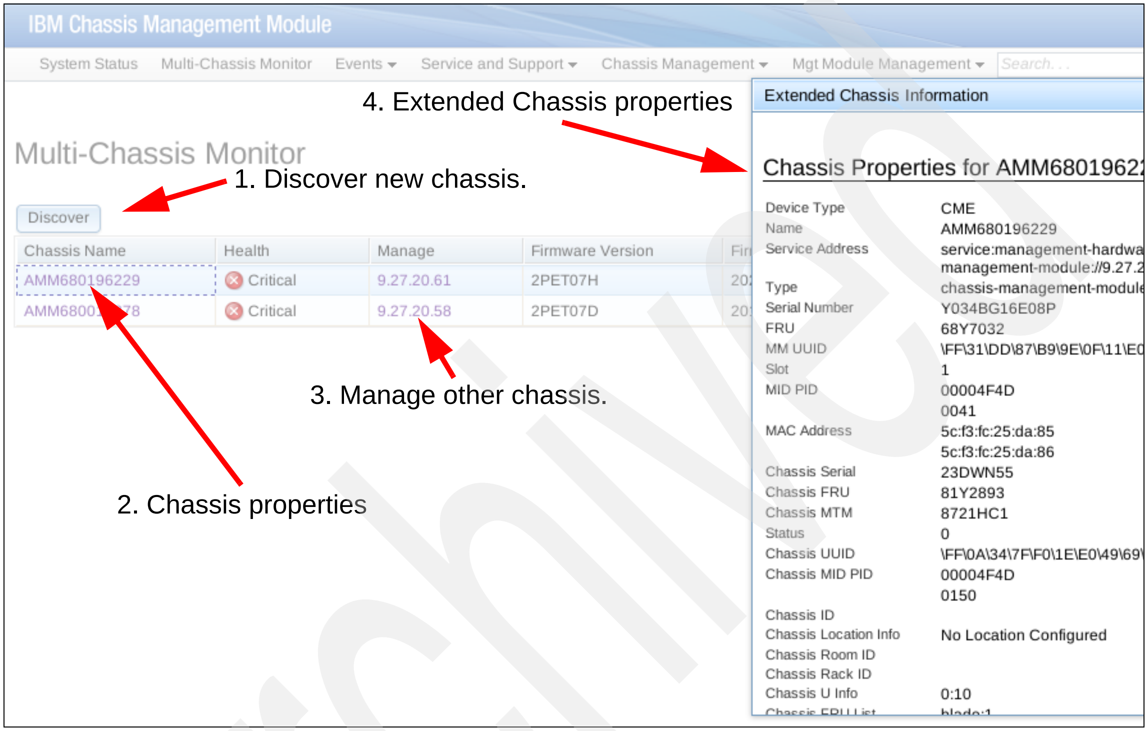

Figure 6-12 Multi-Chassis Monitor tab

The following selections are available (with the numbers that match callouts in Figure 6-12):

1. Discover new chassis: Discover other chassis in the network.
2. Chassis properties: The grid marked in point 2 of Figure 6-12 shows a quick view of the other chassis discovered by the Chassis Management Module, listing the name, health, management IP address, firmware version, and firmware release day. Click the chassis name, and an Events Log dialog box opens.
3. Manage other chassis: With this option, you can manage other chassis directly from this grid. Click a chassis IP address, and another tab opens, where you can manage the chassis.

4. Chassis properties: After you select a chassis from the Chassis Information tab, click the Events Log dialog box and the Chassis Properties tab opens, showing your IP address, location, computer nodes, and so on.

Events tab

This tab (Figure 6-13) has two options, Event Log (shown in Figure 6-14), and Event Recipients, which provide options to send an SNMP alert or send an email using Simple Mail Transfer Protocol (SMTP).

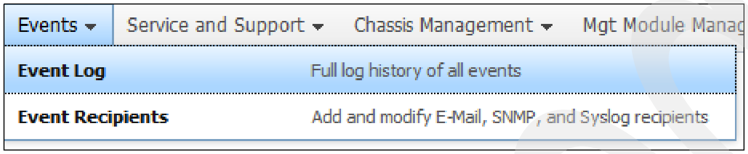

Figure 6-13 Events tab

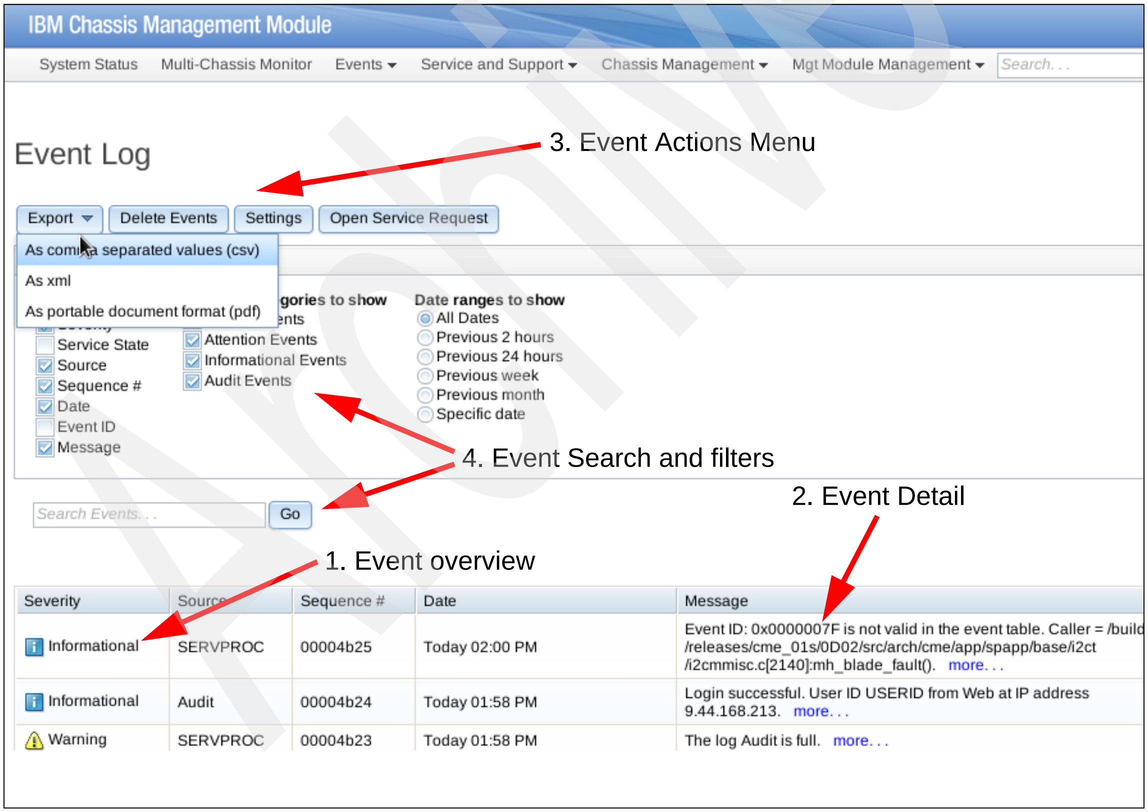

Figure 6-14 Event Logs window

The callouts shown in Figure 6-14 on page 172 are described as follows:

1. Event overview: This grid shows general information about the event listing, including severity, source, sequence, date, and message.
2. Event detail: Click the **More** link to show detailed information about the selected event.
3. Event actions menu: Several options are available to manage logs:
   a. You can use the Export option to export your event log in various formats (.csv, XML, or PDF).
   b. Use the Delete Events option to delete all selected items, with the additional option of selecting audit, systems, or both.
   c. With the Settings option, you can add a log event when a log is 75% full.
   d. The Open Service Request option is enabled when you select one of the events from the grid, followed by a prompt for a small description.
4. Event search and filters: The event grid can become large over time. With Event Search, you can search by keyword and use several filters, as shown in Figure 6-14 on page 172.

Service and Support tab

The Service and Support tab is used for reviewing detected problems, troubleshooting, opening a service request, and for updating chassis settings.

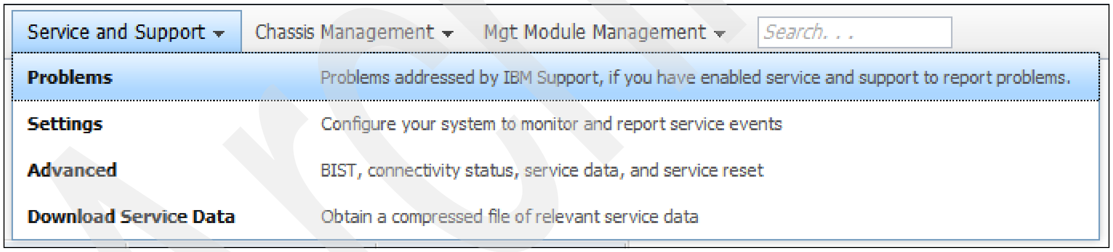

Figure 6-15 Service and Support tab

The Service and Support menu has four menu items:

1. Problems: Shows a grid of detected problems. You can open a service request directly to IBM.
2. Settings: Use this menu item to configure the chassis, enter contact information, country, proxy access, and so on.

3. Advanced Status: This menu item provides advanced service information and additional service tasks. You might be directed by IBM Support staff to review or perform tasks in this section.
4. Download Service Data: Using this menu item, you can download Chassis Management Module data, send management module data to an email recipient (SMTP must be set up first), and download blade data.

Chassis Management tab

This menu is used for reviewing or changing the properties of the components of the chassis. The menu is shown in Figure 6-16.

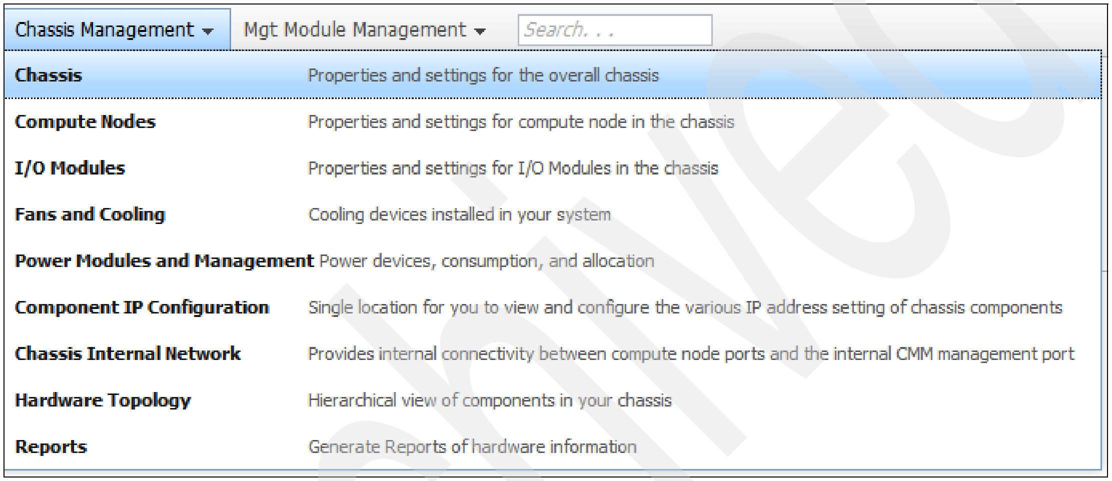

Figure 6-16 Chassis Management tab

Chassis tab

Clicking **Chassis** from the menu shows a window where you can view or change chassis-level data (Figure 6-17).

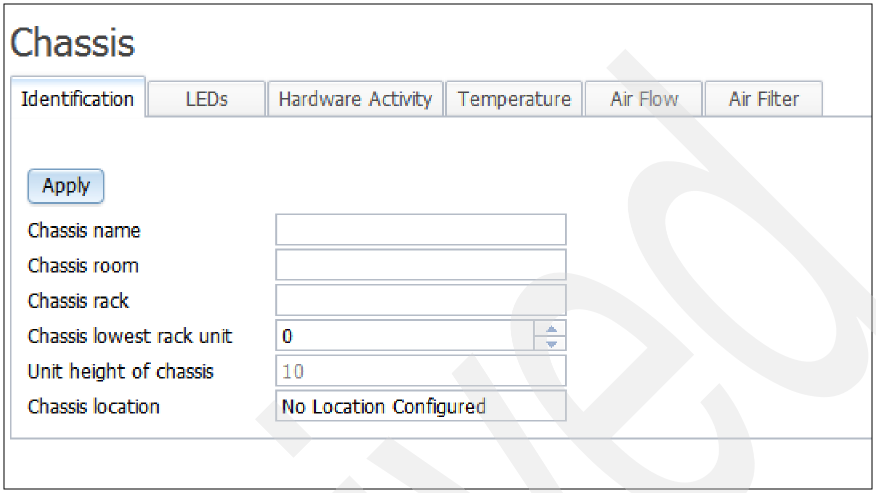

Figure 6-17 Chassis tab

Compute Nodes tab

Clicking **Compute Nodes** from Figure 6-16 on page 174 shows a window that lists the servers installed in the chassis. Clicking one of the names in the Device Name column opens a window with details about that server, as shown in Figure 6-18.

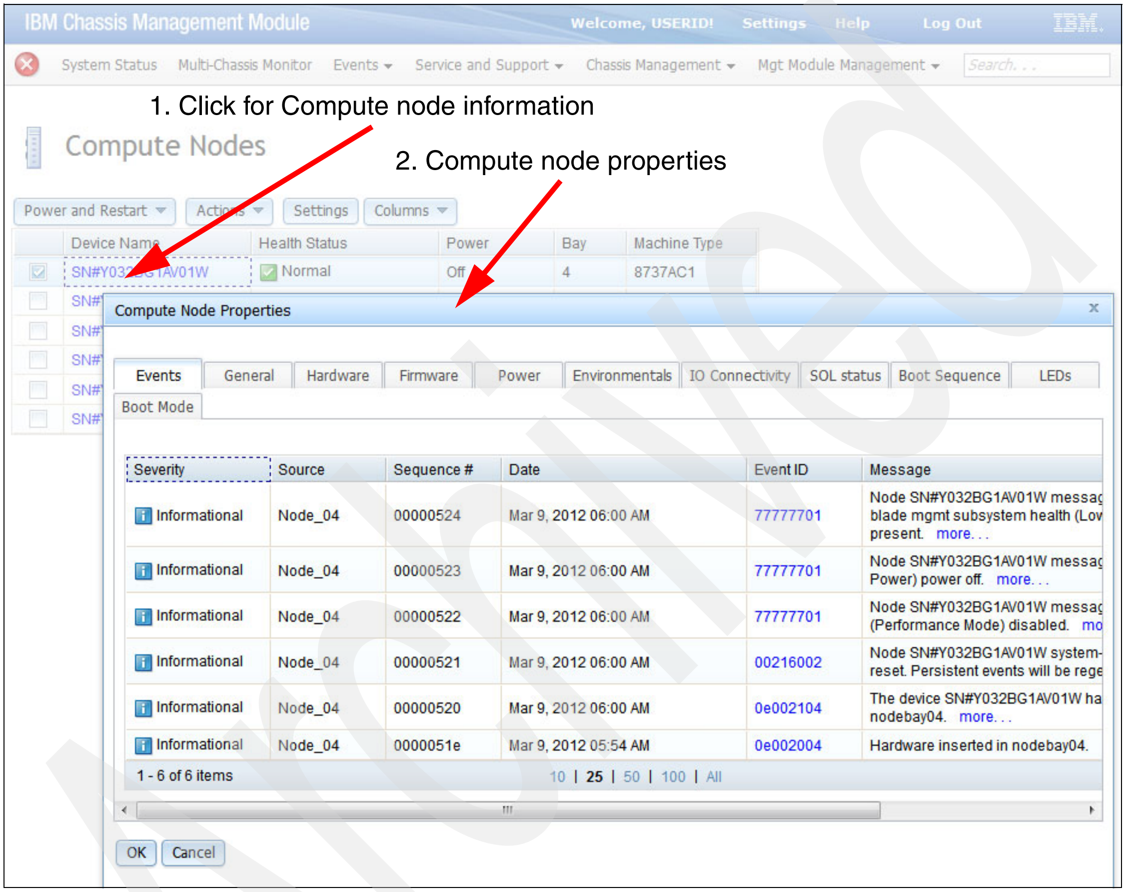

Figure 6-18 Computer Nodes tab

I/O Modules tab

The I/O Modules window is similar to the Compute Nodes window. A grid opens and shows the I/O modules. Clicking a module name opens other panes with the properties of that module (Figure 6-19).

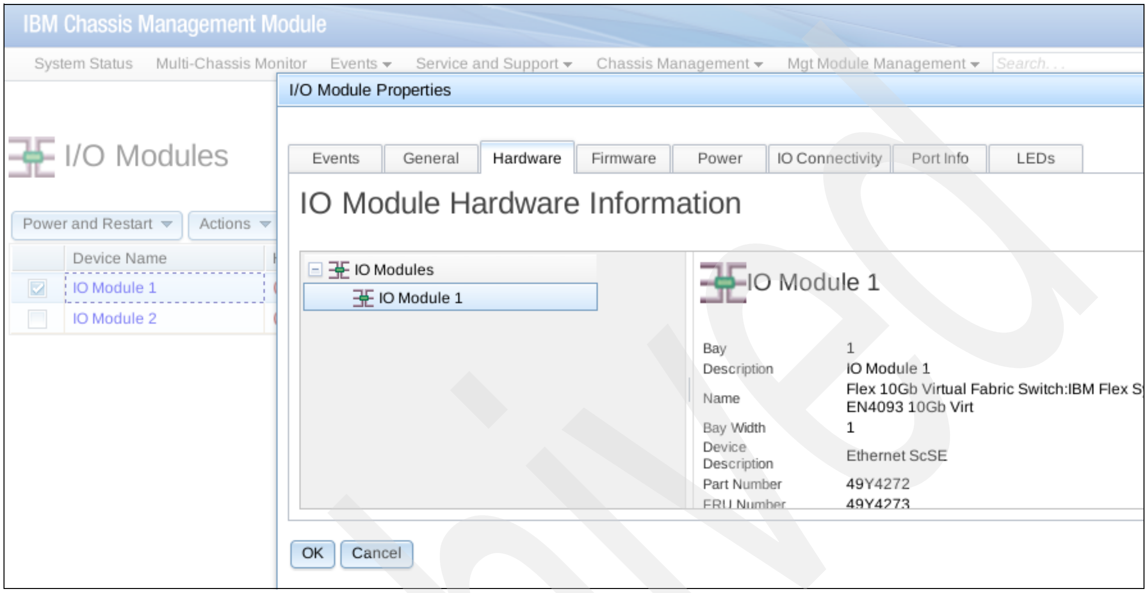

Figure 6-19 I/O Modules tab

Fans and Cooling tab

The Fans and Cooling window shows the fans and their operational status. Select an item in the list to show information about it (Figure 6-20).

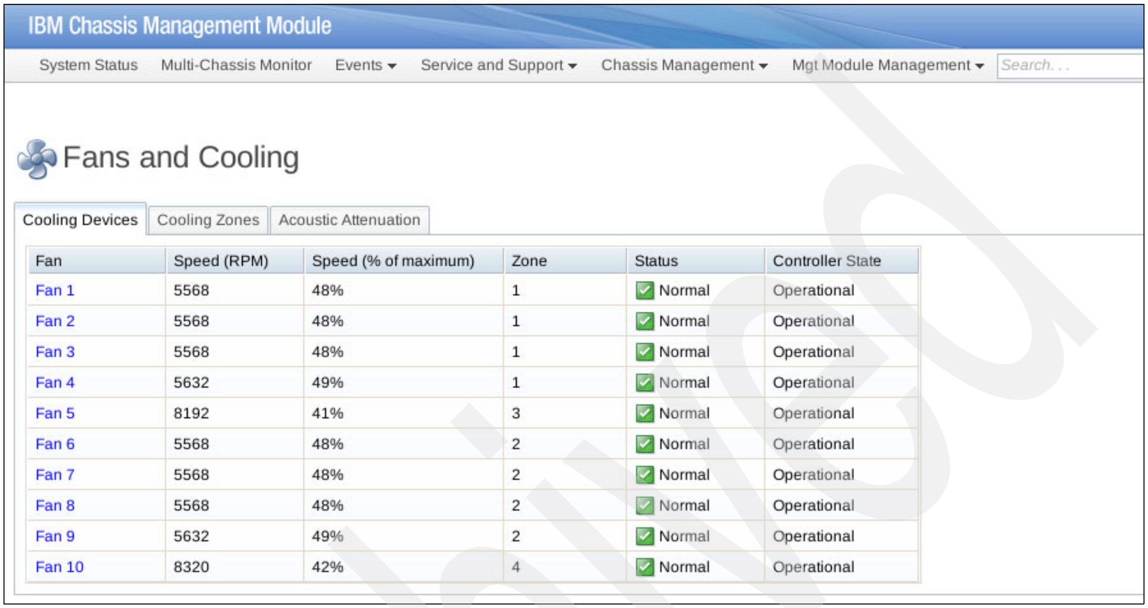

Figure 6-20 Fans and Cooling tab

Power Modules and Management tab

In the Power Modules and Management window (Figure 6-21), you can manage the power subsystem.

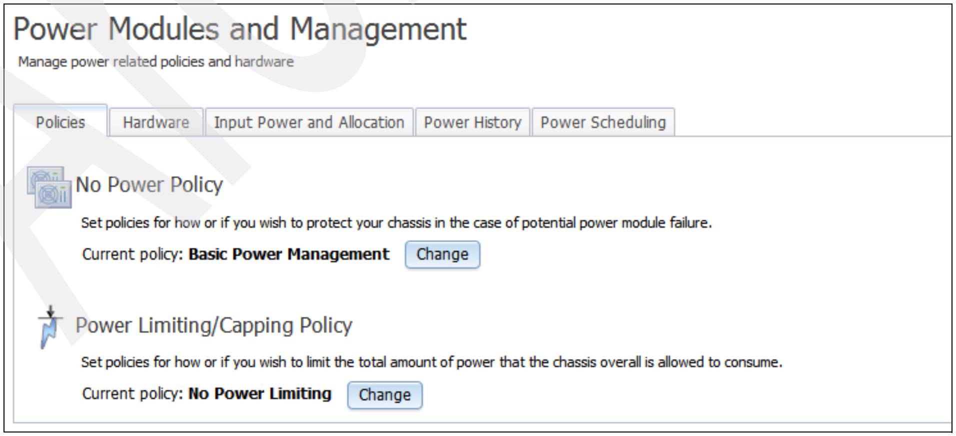

Figure 6-21 Power Modules and Management window

The Power Modules and Management window has the following features:

- The Policies tab shows the power polices that are currently enabled. If you click **Change** in Figure 6-21 on page 178, you can modify the current policy in the window that opens (Figure 6-22).

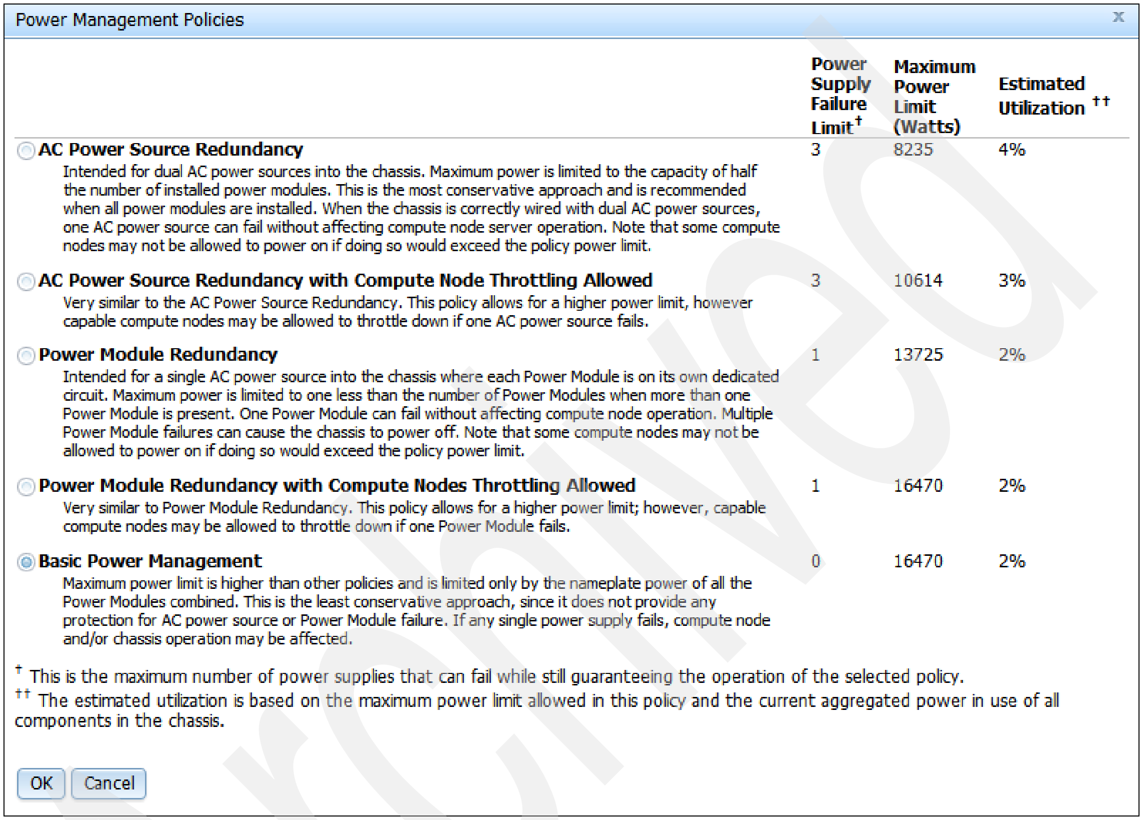

Figure 6-22 Power Management Policies window

- The Hardware tab lists the power supply details (Figure 6-23).

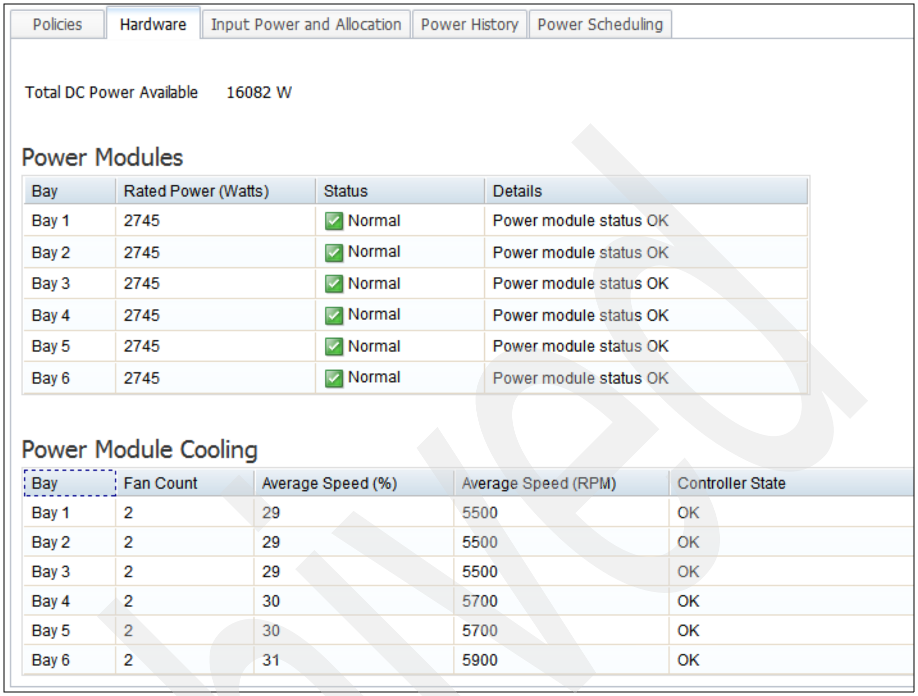
Policies | Hardware | Input Power and Allocation | Power History | Power Scheduling

Total DC Power Available 16082 W

Power Modules

Bay	Rated Power (Watts)	Status	Details
Bay 1	2745	Normal	Power module status OK
Bay 2	2745	Normal	Power module status OK
Bay 3	2745	Normal	Power module status OK
Bay 4	2745	Normal	Power module status OK
Bay 5	2745	Normal	Power module status OK
Bay 6	2745	Normal	Power module status OK

Power Module Cooling

Bay	Fan Count	Average Speed (%)	Average Speed (RPM)	Controller State
Bay 1	2	29	5500	OK
Bay 2	2	29	5500	OK
Bay 3	2	29	5500	OK
Bay 4	2	30	5700	OK
Bay 5	2	30	5700	OK
Bay 6	2	31	5900	OK

Figure 6-23 Hardware tab

- The Input Power and Allocation tab shows charts and details of energy use on the chassis. Figure 6-24 shows an example of one of these charts.

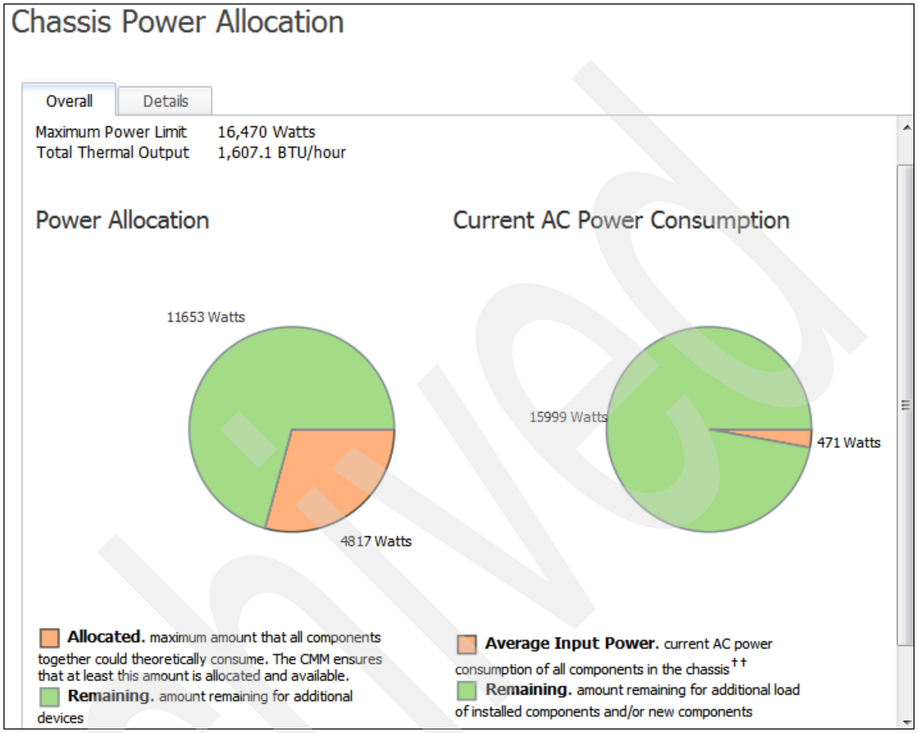

Figure 6-24 Chassis power allocation window

- The Power History tab graphically shows historical power consumption at selected intervals.
- You can use the Power Scheduling tab to configure schedules to power off or on or power cycle one or more compute nodes based on location in the chassis, serial number, or machine-type-model number.

Component IP configuration

This menu item lists all the components and IP configuration (if available) of the chassis.

Chassis internal networking

This menu item shows the status of the internal network and permits modifications. You can enable or disable the internal network here.

Hardware topology

You can use this menu item to view all the hardware installed in the chassis, right down to the individual component level, such as a DIMM. Figure 6-25 shows an example.

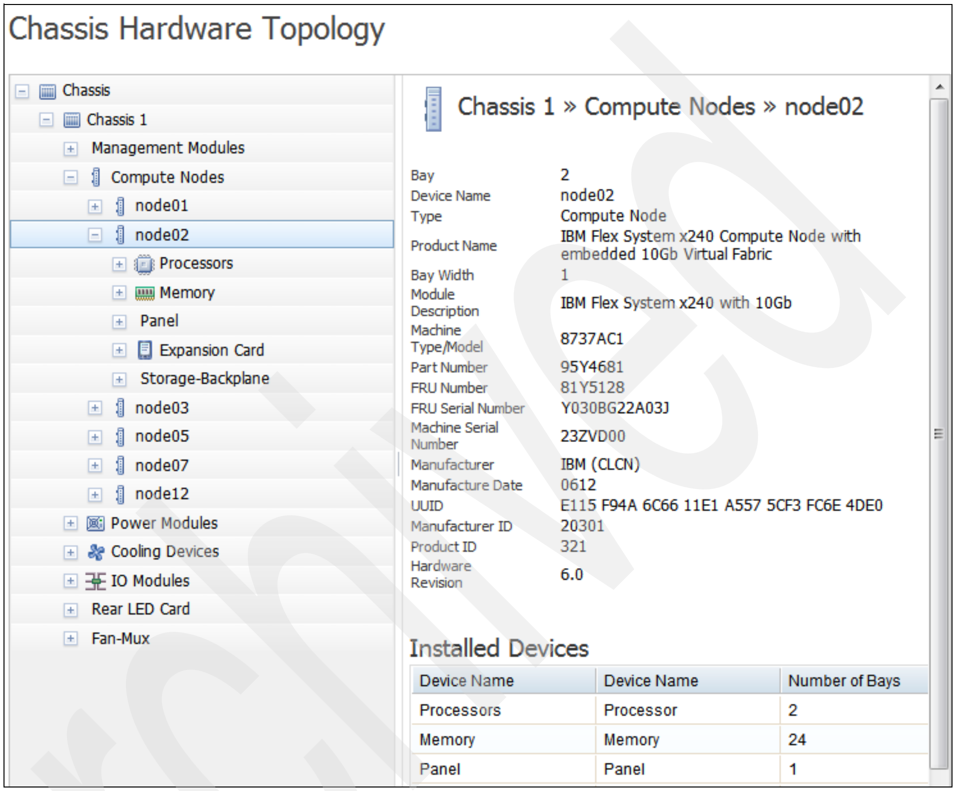

Figure 6-25 Hardware topology

Reports

This menu item shows reports that list all MAC addresses or unique IDs used by components in the chassis.

Mgmt Module Management tab

This tab, shown in Figure 6-26, has options for performing user management tasks, firmware upgrades, security management, network management, and so on. The tab is shown in Figure 6-26.

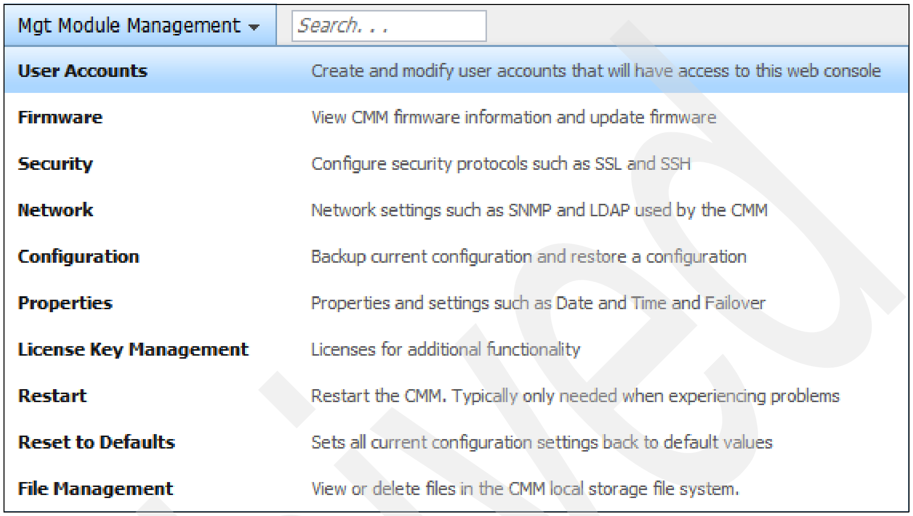

Figure 6-26 Mgmt Module Management tab

User Accounts

This option provides access to user accounts and permission groups, for which you can add users, change passwords, and create groups for access to specific functions. Click a user name to view additional information, as shown in Figure 6-27.

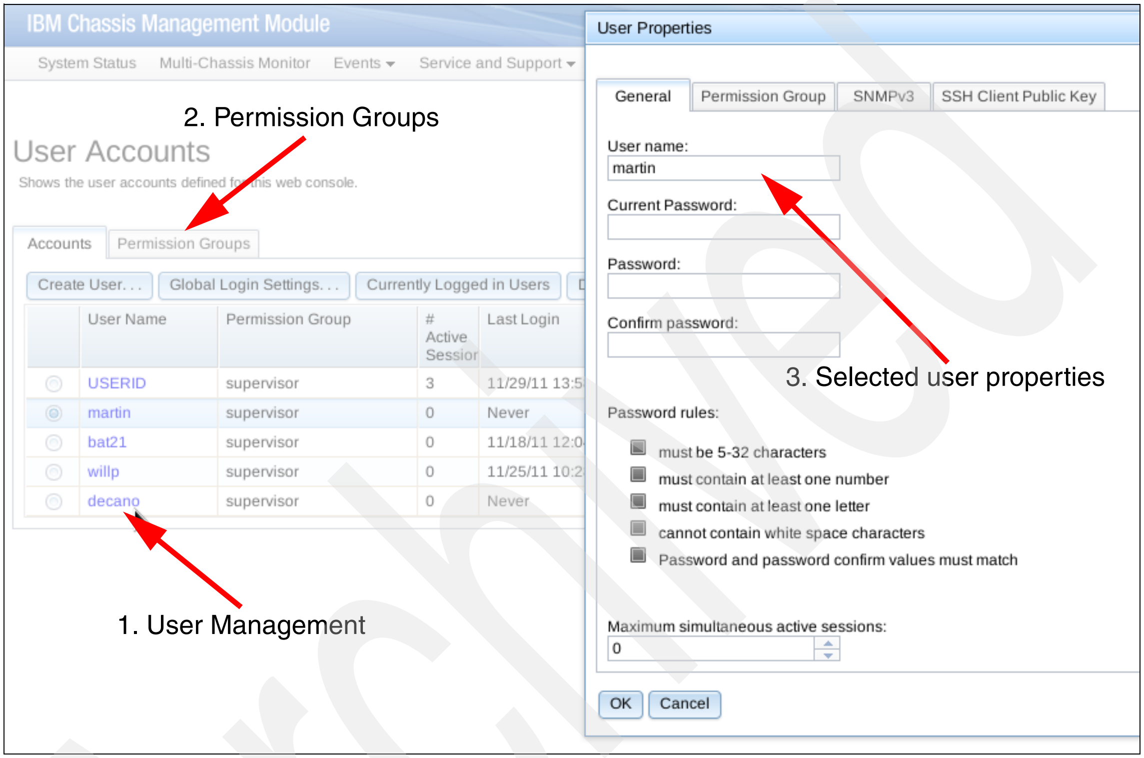

Figure 6-27 User and group management

Firmware menu

Enables firmware upgrades and views of current firmware state.

Security

You can use this menu to configure security policies and set up a certificate authority (CA), enable HTTPS or SSH access, and configure an LDAP for logins. Figure 6-28 shows the Security Policies tab of this window.

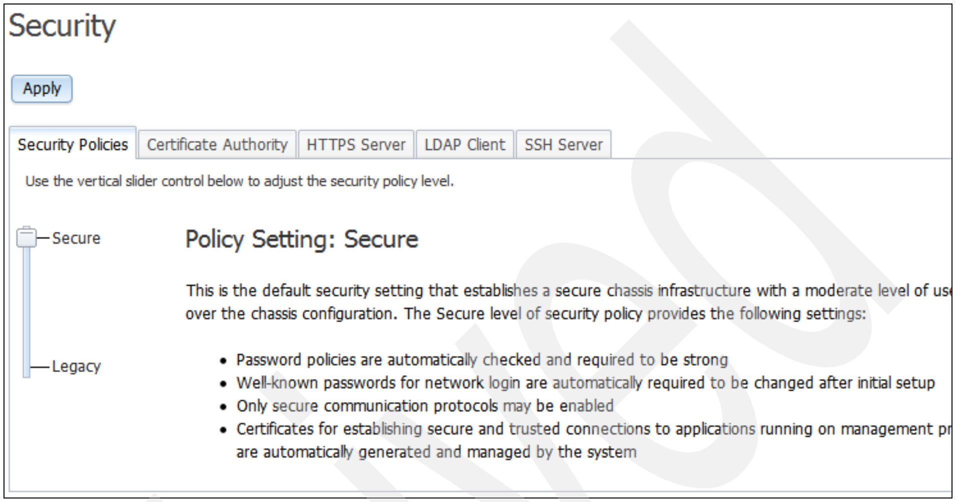

Figure 6-28 Security Policies tab

Network

All network setup options for your chassis are available in this menu. The Chassis Management Module supports IPv4 and IPv6. You can set up SMTP, Domain Name System (DNS), Lightweight Directory Access Protocol (LDAP), Hypertext Transfer Protocol (HTTP), and so on, as shown on Figure 6-29.

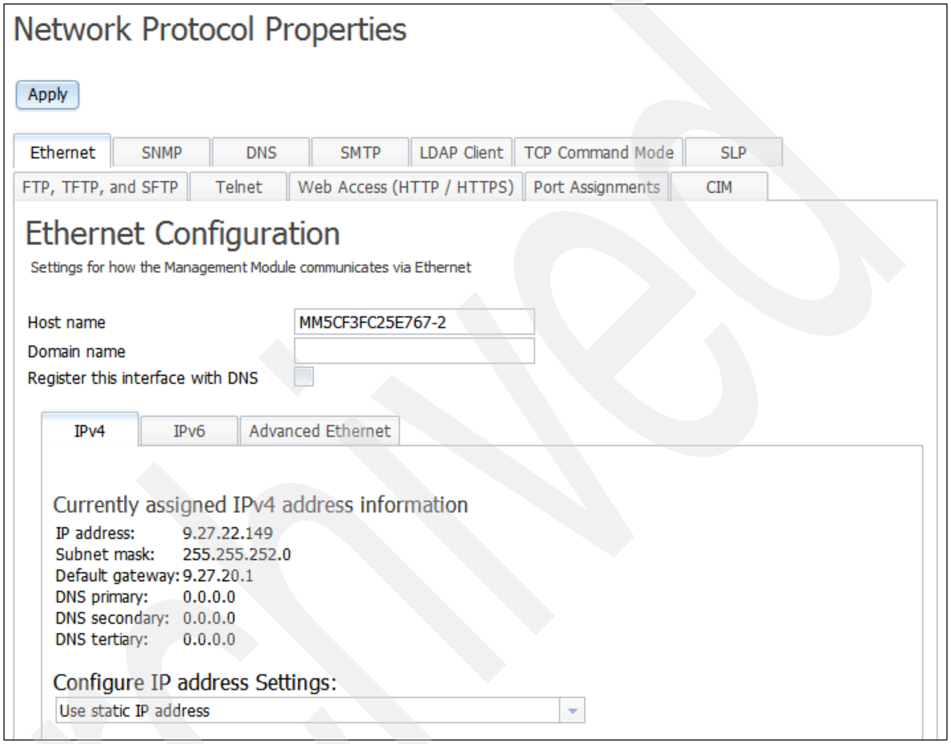

Figure 6-29 Network Protocol Properties window

Configuration

You can use this menu to back up and restore your Chassis Management Module configuration. Use the Initial Setup Wizard to walk you through these setup steps.

Properties

You can use this window to set up your Chassis Management Module name, time and date, and standby Chassis Management Module management details. Figure 6-30 shows an example.

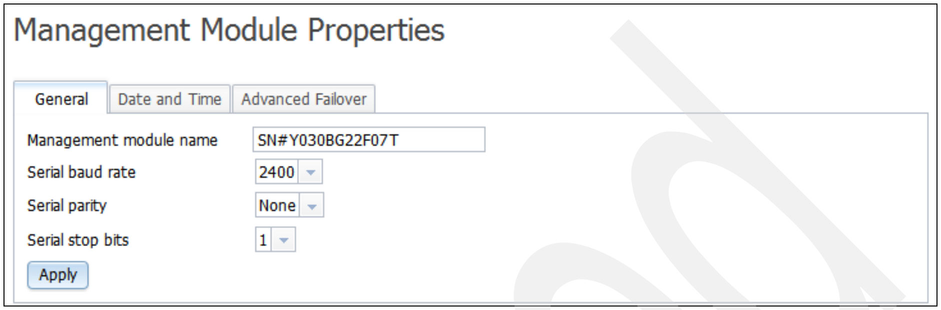

Figure 6-30 Management Module Properties

License Key Management

You can use this window to manage all of your chassis licensed features. Figure 6-31 shows an example.

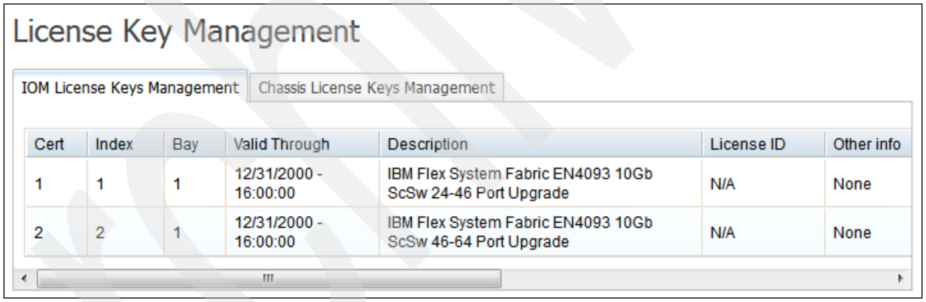

Figure 6-31 License Key Management

6.2.3 Accessing the node through the CMM

This section describes how to access the Power Systems compute node through the CMM. This option is useful when the FSM is not available.

Before you begin, you need the IP address of the Chassis Management Module. You can access the CMM using SSH or a browser. The browser method is described here.

Important: When a Power Systems compute node is discovered and managed by a Flex System Manager, Serial Over LAN (SOL) must be disabled. Given that the Power Systems compute node is ordered as part of one of the IBM PureFlex System configurations, it is discovered and managed by a Flex System Manager. Therefore, the Chassis Management Module access is disabled in most cases. This section is intended to address situations where an additional Power Systems compute node is ordered for an existing Enterprise Chassis. In that case, the Power Systems compute node is not discovered by an Flex System Manager and can be managed and installed using the Chassis Management Module.

To access the node through the CMM, complete the following steps:

1. Open a browser and point it to the following URL (where *system_name* is the host name or IP address of the Chassis Management Module):

   `https://system_name`

   The window in Figure 6-32 opens.

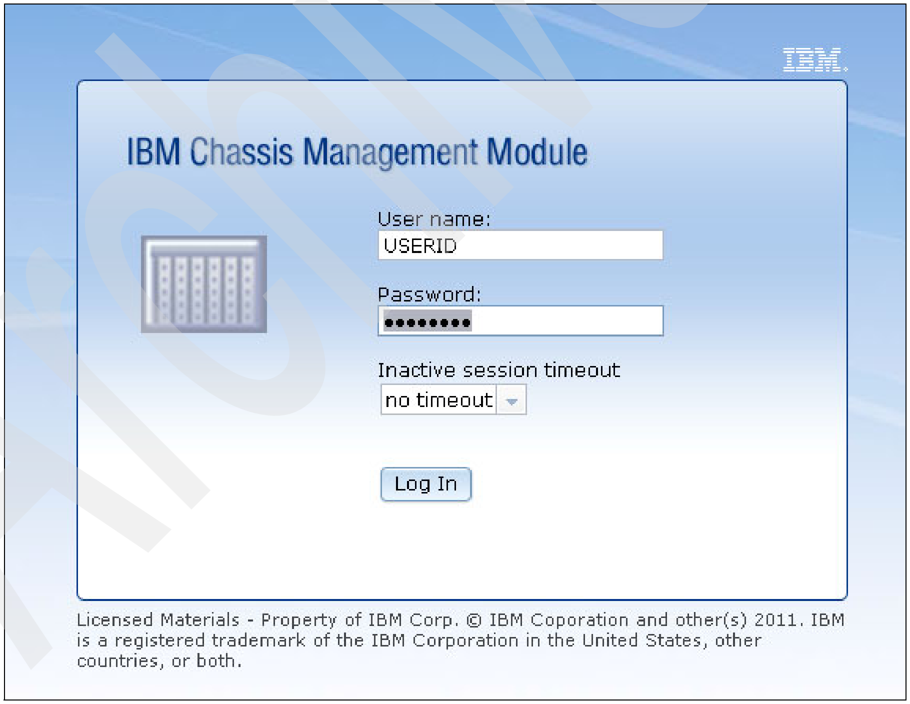

Figure 6-32 Chassis Management Module login window

2. Log in with your user ID and password. The System Status window of the Chassis Management Module opens, as shown in Figure 6-33, with the Chassis tab active. If not, click **System Status** from the menu bar at the top of the window.

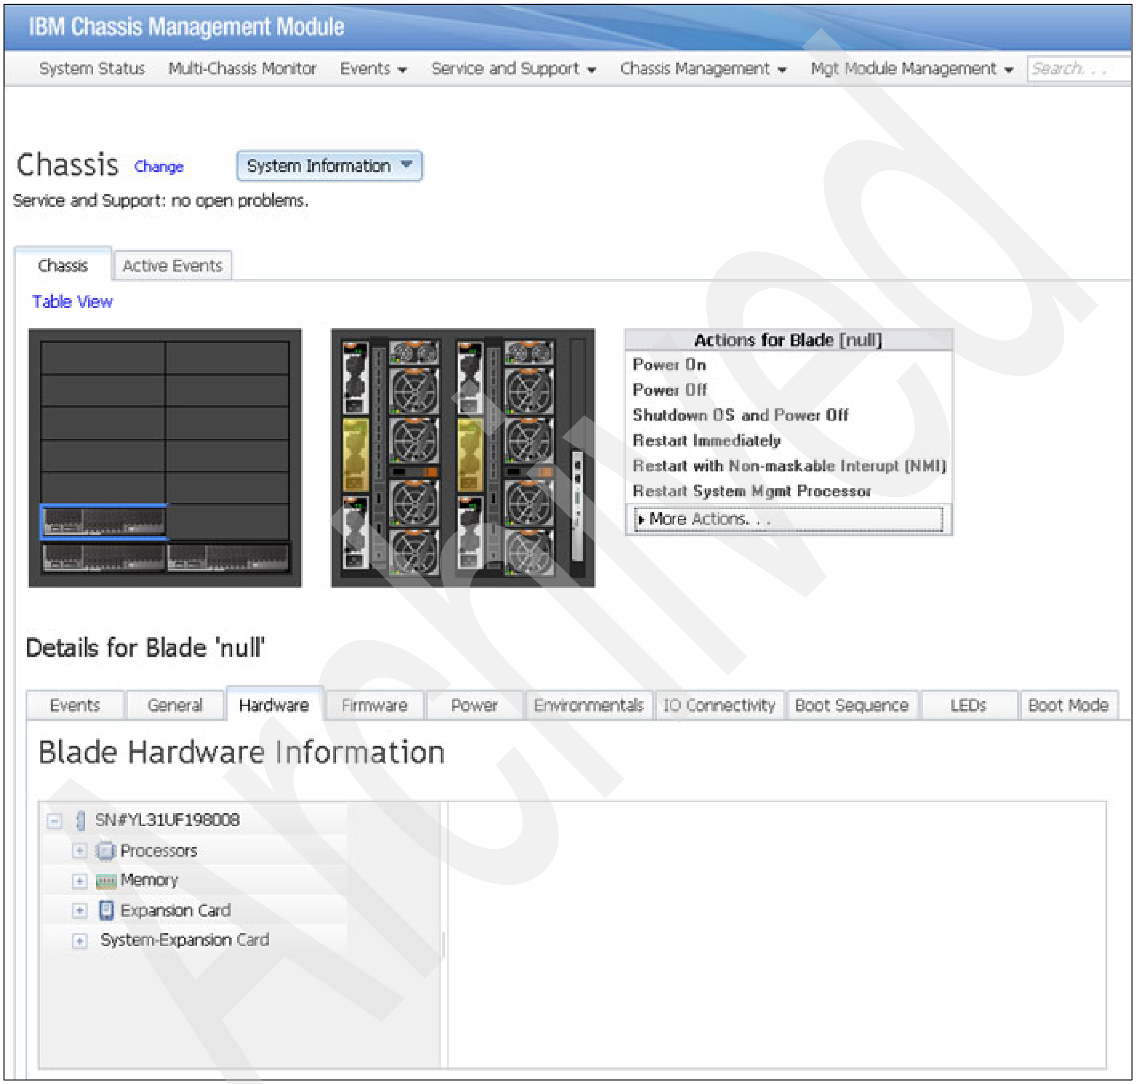

Figure 6-33 Chassis Management Module with node in bay 3 selected

3. Select the Power Systems compute node image of the chassis. Figure 6-33 shows the node in bay 3 selected. The Actions menu to the right of the graphics is useful when working with the node.

4. Click **More Actions** → **Launch Blade Console** to access the option to launch a console (Figure 6-34).

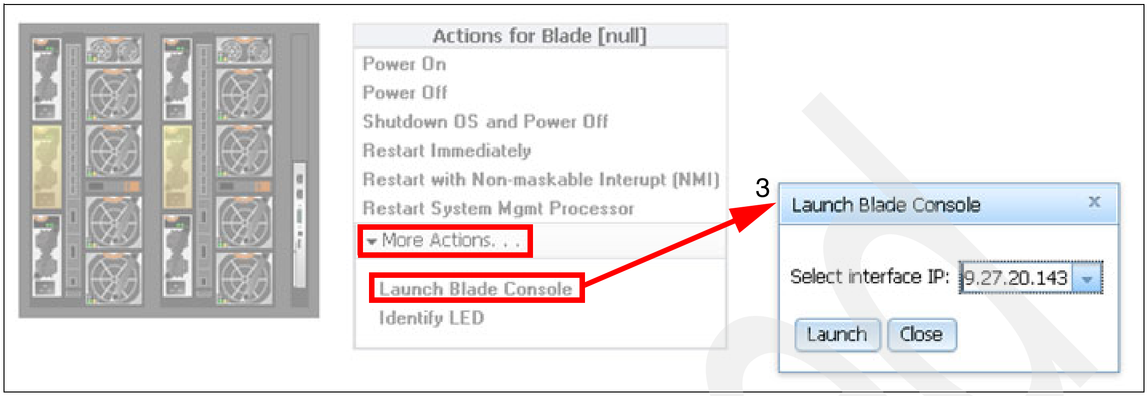

Figure 6-34 Launch console on Power Systems compute node from Chassis Management Module

5. Enter the IP address of the node or select it from the menu.
6. Power on the node using the **Power On** option in the Actions menu. The resulting progress indicator is shown in Figure 6-35.

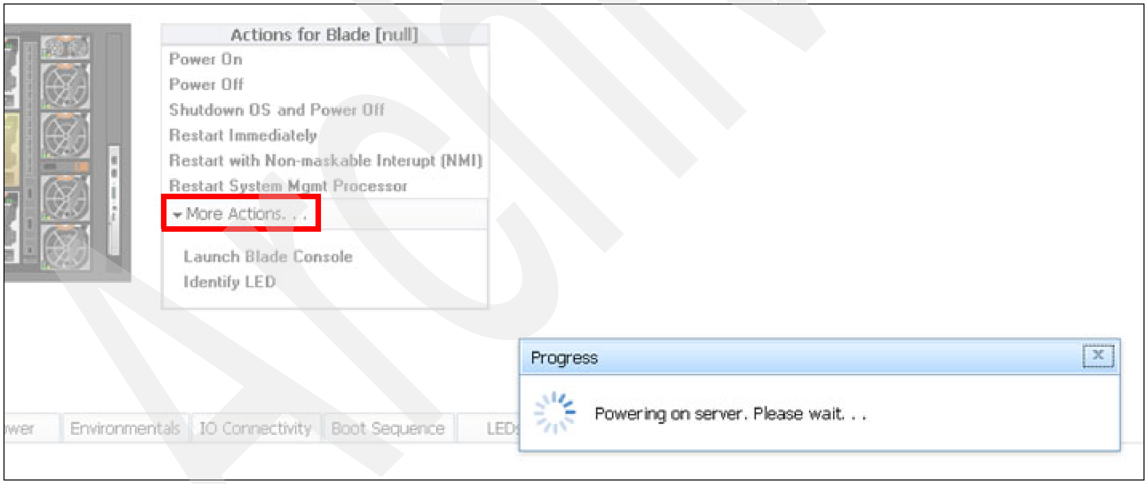

Figure 6-35 Power on the Power Systems compute node

You interact with the node as it boots. You can enter SMS to install the node, or allow it to boot to an already installed operating system, from which you can log in to the console.

6.3 Management network

In an IBM Flex System Enterprise Chassis, you can configure separate management and data networks.

The management network is a private and secure Gigabit Ethernet network used to complete management-related functions throughout the chassis, including management tasks related to the compute nodes, switches, and the chassis itself.

The management network is shown in Figure 6-36 (it is the blue line). It connects the CMM to the compute nodes, the switches in the I/O bays, and the FSM. The FSM connection to the management network is through a special Broadcom 5718-based management network adapter (Eth0). The management networks in multiple chassis can be connected together through the external ports of the CMMs in each chassis via a GbE Top-of-Rack switch.

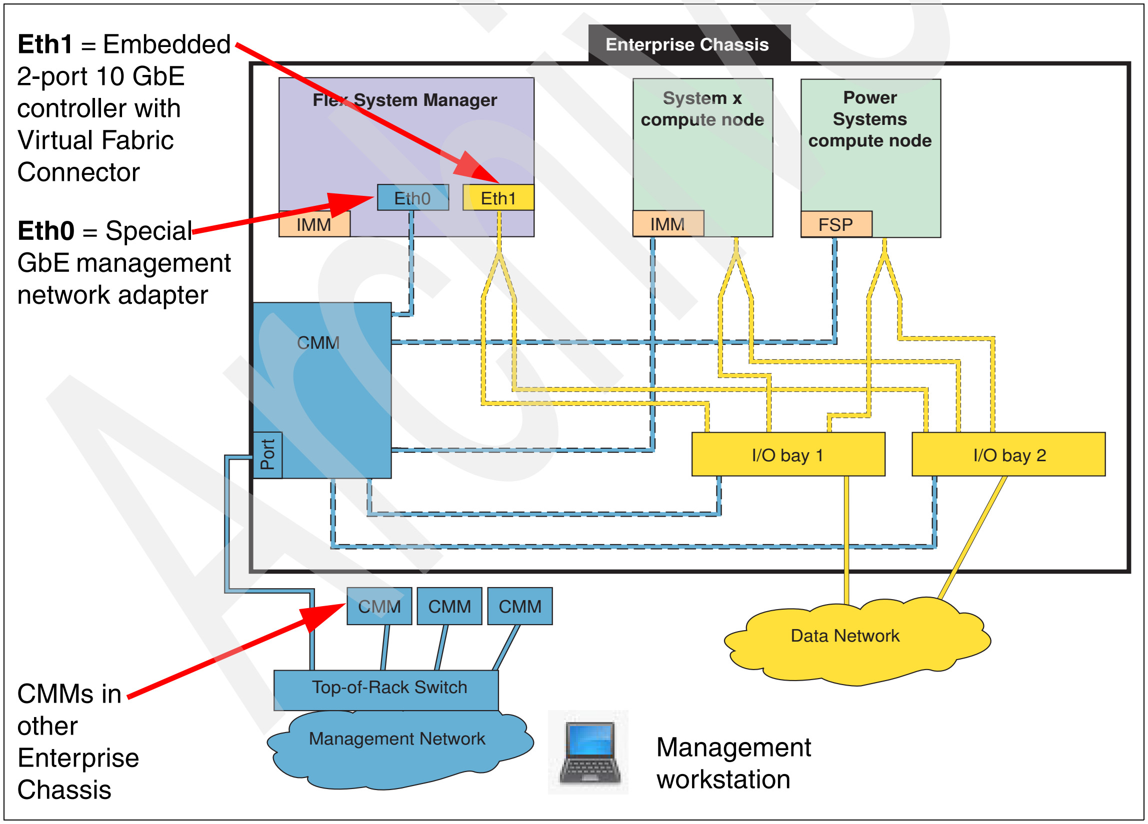

Figure 6-36 Separate management and production data networks

The yellow line in the Figure 6-36 on page 191 shows the production data network. The FSM also connects to the production network (Eth1) so that it can access the Internet for product updates and other related information.

Important: The management node console can be connected to the data network for convenience of access.

One of the key functions that the data network supports is discovery of operating systems on the various network endpoints. Discovery of operating systems by the FSM is required to support software updates on an endpoint, such as a compute node. You can use the FSM Checking and Updating Compute Nodes wizard to discover operating systems as part of the initial setup.

6.4 IBM Flex System Manager

The IBM Flex System Manager (FSM) is a high performance, scalable system management appliance, based on the IBM Flex System x240 Compute Node. FSM hardware has systems management software preinstalled, and you can configure, monitor, and manage FSM resources in up to four chassis.

The following list describes the high-level features and functions of the IBM Flex System Manager:

- Supports a comprehensive, pre-integrated system that is configured to optimize performance and efficiency.
- Automated processes triggered by events simplify management and reduce manual administrative tasks.
- Centralized management reduces the skills and the number of steps it takes to manage and deploy a system.
- Enables comprehensive management and control of energy utilization and costs.
- Automates responses for a reduced need for manual tasks (custom actions / filters, configure, edit, relocate, and automation plans).
- Full integration with server views, including virtual server views enables efficient management of resources.

The preinstallation contains a set of software components that are responsible for performing certain management functions. These components must be activated using the available IBM Feature on Demand (FoD) software entitlement licenses, and they are licensed on a per-chassis basis. You need one license for each chassis you plan to manage.

The management node comes standard without any entitlement licenses, so you must purchase a license to enable the required FSM functionality.

As described in Chapter 2, “IBM PureFlex System” on page 15, there are two versions of IBM Flex System Manager: base and advanced.

The IBM Flex System Manager base feature set offers the following functionality:

- Supports up to four managed chassis
- Supports up to 5,000 managed elements
- Auto-discovery of managed elements
- Overall health status
- Monitoring and availability
- Hardware management
- Security management
- Administration
- Network management (Network Control)
- Storage management (Storage Control)
- Virtual machine lifecycle management (VMControl Express)

The IBM Flex System Manager advanced feature set offers all capabilities of the base feature set plus:

- Image management (VMControl Standard)
- Pool management (VMControl Enterprise)

6.4.1 Hardware overview

The FSM, from a hardware point of view, is a locked-down compute node with a specific hardware configuration designed for optimal performance of the preinstalled software stack. The FSM looks similar to the Intel based x240. However, there are slight differences between the system board designs, which make these two hardware nodes not interchangeable.

Figure 6-37 shows a front view of the FSM.

Figure 6-37 IBM Flex System Manager

Figure 6-38 shows the internal layout and major components of the FSM.

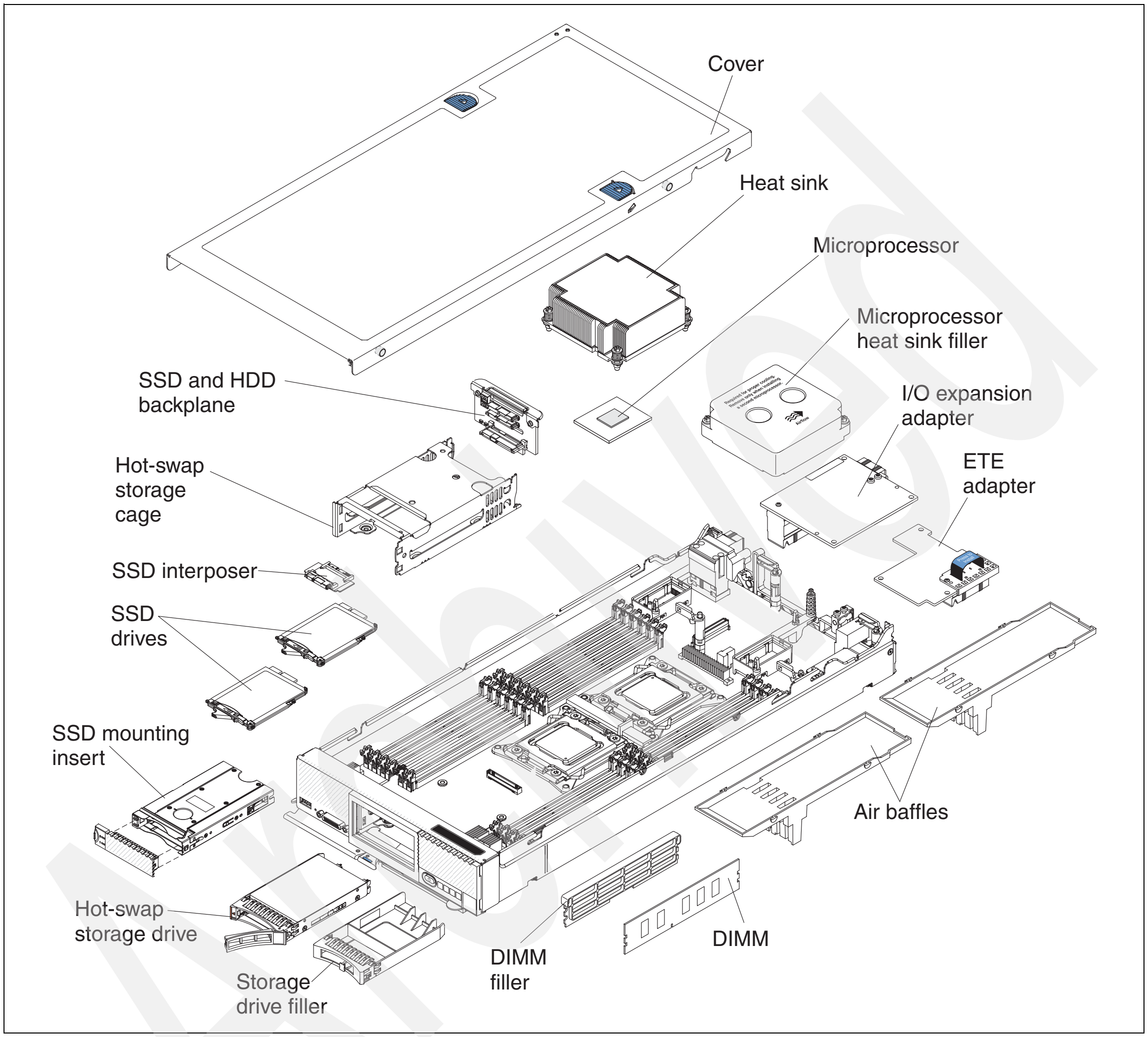

Figure 6-38 Exploded view of the IBM Flex System Manager node showing major components

The FSM comes preconfigured with the components described in Table 6-1.

Table 6-1 Features of the IBM Flex System Manager node (8731)

Feature	Description
Processor	1x Intel Xeon Processor E5-2650 8C 2.0GHz 20MB Cache 1600MHz 95W

Feature	Description
Memory	8x 4GB (1x 4GB, 1Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333MHz LP RDIMM
SAS Controller	One LSI 2004 SAS Controller
Disk	1x IBM 1TB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD 2x IBM 200GB SATA 1.8" MLC SSD (configured in a RAID 1 configuration)
Integrated NIC	Embedded dual-port 10 Gb Virtual Fabric Ethernet controller (Emulex BE3) Dual-port 1 GbE Ethernet controller on a management adapter (Broadcom 5718)
Systems Management	Integrated Management Module v2 (IMMv2) Management network adapter

Figure 6-39 shows the internal layout of the FSM.

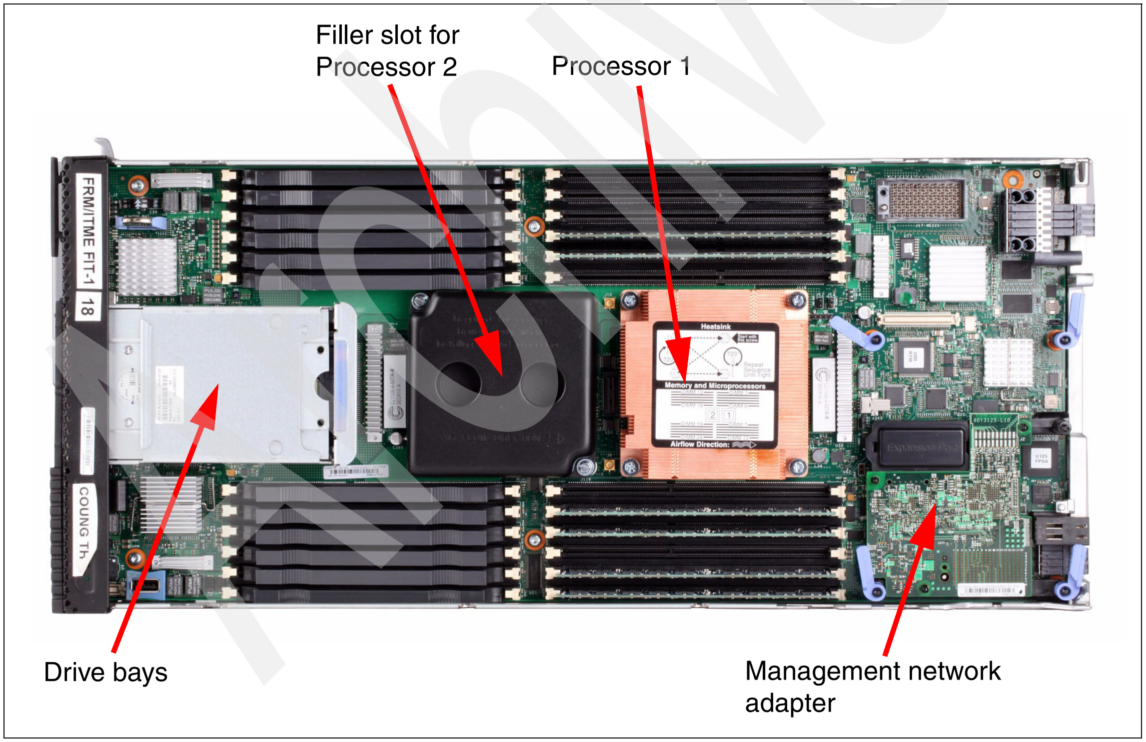

Figure 6-39 Internal view showing the major components of IBM Flex System Manager

Front controls

The FSM has similar controls and LEDs as the IBM Flex System x240 Compute Node. Figure 6-40 shows the front of an FSM with the location of the control and LEDs.

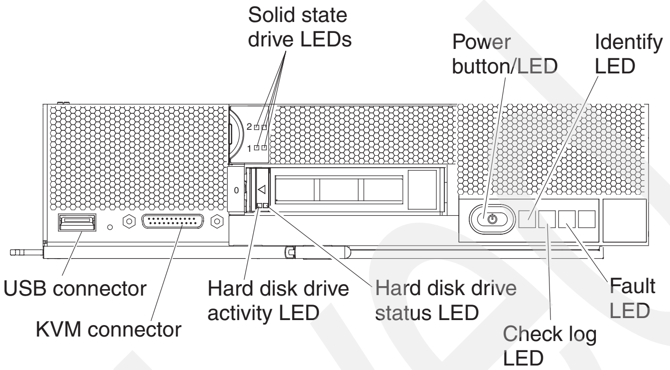

Figure 6-40 FSM front panel - controls and LEDs

Storage

The FSM ships with 2x IBM 200GB SATA 1.8" MLC SSD and 1x IBM 1TB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD drives. The 200 GB SSD drives are configured as a RAID 1 pair, providing roughly 200 GB of usable space. The 1 TB SATA drive is not part of a RAID group.

The partitioning of the disks is listed in Table 6-2.

Table 6-2 Detailed SSD and HDD disk partitioning

Physical disk	Virtual disk size	Description
SSD	50 MB	Boot disk
SSD	60 GB	OS/Application disk
SSD	80 GB	Database disk
HDD	40 GB	Update repository
HDD	40 GB	Dump space
HDD	60 GB	Spare disk for OS/Application
HDD	80 GB	Spare disk for database
HDD	30 GB	Service Partition

Management network adapter

The management network adapter is a standard feature of the FSM and provides a physical connection into the private management network of the chassis. The adapter is shown in Figure 6-38 on page 195 as the everything-to-everything (ETE) adapter.

The management network adapter contains a Broadcom 5718 Dual 1GbE adapter and a Broadcom 5389 8-port L2 switch. This card is one of the features that makes the FSM unique compared to all other nodes supported by the Enterprise Chassis. The management network adapter provides a physical connection into the private management network of the chassis so that the software stack has visibility into both the data and management networks. The L2 switch on this card is automatically set up by the IMMv2 and connects the FSM and the onboard IMMv2 into the same internal private network.

All other nodes supported by the Enterprise Chassis have a connection only into the management network through the management controller (IMMv2 for System x nodes; FSP for POWER nodes), which is not accessible through the operating system.

Console breakout cable

The FSM connects to local video, USB keyboard, and USB mouse devices by connecting the Console Breakout Cable. The Console Breakout Cable connects to a connector on the front bezel of the FSM. The Console Breakout Cable also provides a serial connector. Figure 6-41 shows the Console Breakout Cable.

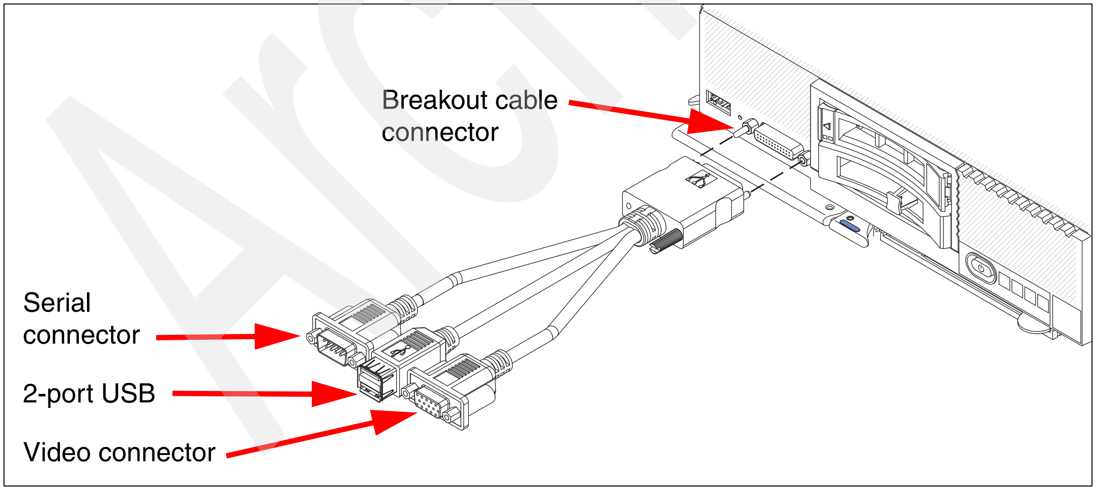

Figure 6-41 The Console Breakout Cable connecting to the IBM Flex System Manager (and x240)

The Console Breakout Cable is included with the chassis.

6.4.2 Software features

The main features of IBM Flex System Manager management software are:

- Monitoring and problem determination
  - A real-time multichassis view of hardware components with overlays for additional information.
  - Automatic detection of issues in your environment through event setup that triggers alerts and actions.
  - Identification of changes that might impact availability.
  - Server resource utilization by a virtual machine or across a rack of systems.
- Hardware management
  - Automated discovery of physical and virtual servers and interconnections, applications, and supported third-party networking.
  - Inventory of hardware components.
  - Chassis and hardware component views.
  - Hardware properties.
  - Component names/hardware identification numbers.
  - Firmware levels.
  - Utilization rates.
- Network management
  - Management of network switches from various vendors.
  - Discovery, inventory, and status monitoring of switches.
  - Graphical network topology views.
  - Support for KVM, pHyp, VMware virtual switches, and physical switches.
  - VLAN configuration of switches.
  - Integration with server management.
  - Per-virtual machine network usage and performance statistics provided to VMControl.
  - Logical views of servers and network devices grouped by subnet and VLAN.
- Storage management
  - Discovery of physical and virtual storage devices.
  - Support for virtual images on local storage across multiple chassis.

- Inventory of physical storage configuration.
- Health status and alerts.
- Storage pool configuration.
- Disk sparing and redundancy management.
- Virtual volume management.
- Support for virtual volume discovery, inventory, creation, modification, and deletion.

- Virtualization management (base feature set)
  - Support for VMware, Hyper-V, KVM, and IBM PowerVM.
  - Create virtual servers.
  - Edit virtual servers.
  - Manage virtual servers.
  - Relocate virtual servers.
  - Discover virtual server, storage, and network resources, and visualize the physical-to-virtual relationships.
- Virtualization management (advanced feature set)
  - Create new image repositories for storing virtual appliances and discover existing image repositories in your environment.
  - Import external, standards-based virtual appliance packages into your image repositories as virtual appliances.
  - Capture a running virtual server that is configured just the way you want, complete with guest operating system, running applications, and virtual server definition.
  - Import virtual appliance packages that exist in the Open Virtualization Format (OVF) from the Internet or other external sources.
  - Deploy virtual appliances quickly to create new virtual servers that meet the demands of your ever-changing business needs.
  - Create, capture, and manage workloads.
  - Create server system pools, where you can consolidate your resources and workloads into distinct and manageable groups.
  - Deploy virtual appliances into server system pools.
  - Manage server system pools, including adding hosts or additional storage space and monitoring the health of the resources and the status of the workloads in them.

- Group storage systems together using storage system pools to increase resource utilization and automation.
- Manage storage system pools by adding storage, editing the storage system pool policy, and monitoring the health of the storage resources.

- Additional features
  - A resource-oriented chassis map provides an instant graphical view of chassis resources, including nodes and I/O modules.
    - A fly-over provides an instant view of individual server's (node) status and inventory
    - A chassis map provides an inventory view of chassis components, a view of active statuses that require administrative attention, and a compliance view of server (node) firmware.
    - Actions can be taken on nodes, such as working with server-related resources, showing and installing updates, submitting service requests, and launching the remote access tools.
  - Remote console.
    - Open video sessions and mount media, such as DVDs with software updates, to their servers from their local workstation.
    - Remote Keyboard, Video, and Mouse (KVM) connections.
    - Remote Virtual Media connections (mount CD, DVD, ISO, and USB media).
    - Power operations against servers (Power On/Off/Restart).
  - Hardware detection and inventory creation.
  - Firmware compliance and updates.
  - Automatic detection of hardware failures.
    - Provides alerts.
    - Takes corrective action.
    - Notifies IBM of problems to escalate problem determination.
  - Health status (such as processor utilization) on all hardware devices from a single chassis view.
  - Administrative capabilities, such as setting up users within profile groups, assigning security levels, and security governance.

6.5 FSM initial setup

FSM is an appliance that is delivered with all the required software preinstalled. When this software stack is started for the first time, a startup wizard starts that steps through the required configuration process, such as licensing agreements and Transmission Control Protocol/Internet Protocol (TCP/IP) configuration for the appliance.

When configuration is complete, the FSM is ready to manage the chassis it is installed in and other chassis (up to four). Now that the chassis is managed, individual components, such as compute nodes and switches, can also be managed.

This section describes how to use the startup wizards and use the chassis management selection and basic POWER based compute node management functions.

> **Important:** At the time of this writing, IBM Flex System Manager is required for any configuration that contains a Power Systems compute node. It is also assumed that IBM Flex System Manager is preconfigured to manage the initial chassis. In that event, the steps in this section are not required unless IBM Flex System Manager is being reinstalled.

6.5.1 Flex System Manager startup wizard

FSM is based on an Intel compute node and has the same options for obtaining an initial console. These options are the IMMv2 remote console or using the supplied dongle and front port on the FSM node to connect directly to the keyboard, display, and mouse or a KVM unit.

To monitor the FSM startup process, connect a console using one of these methods before powering up the FSM node. The steps that follow use the IMMv2 remote console method.

To initiate an IMMv2 remote console session, complete the following steps:

1. Start a browser session, as shown in Figure 6-42, to the IP address of the FSM IMMv2.

Important: The IP address of the IMMv2 of Intel compute nodes can be determined by using the Chassis Management Module or CLI. By default, the interface is set to use DHCP. You can change this setting to a static address using the CMM, a CLI, or a console connected directly to the VGA port on the front of the FSM (accessible with the use of the console breakout cable).

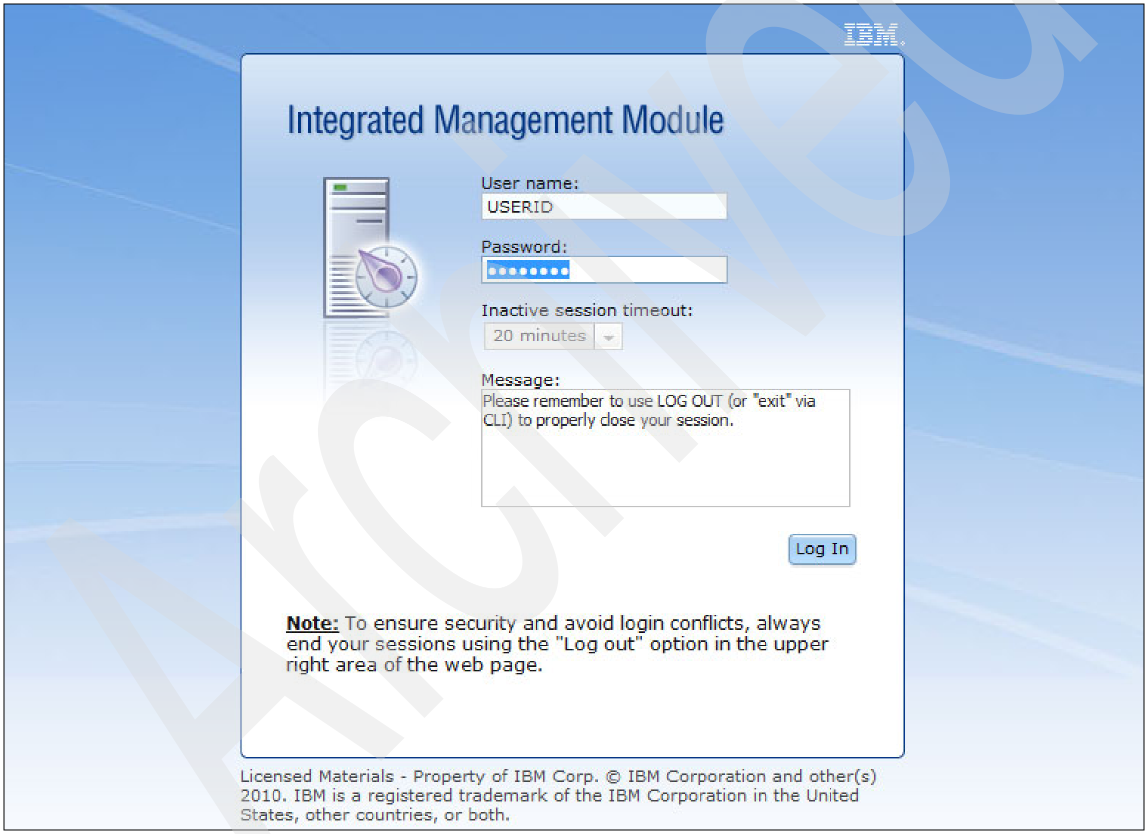

Figure 6-42 IMMv2 login

2. After logging in to the IMMv2, click **Server Management** from the navigation options, as shown in Figure 6-43.

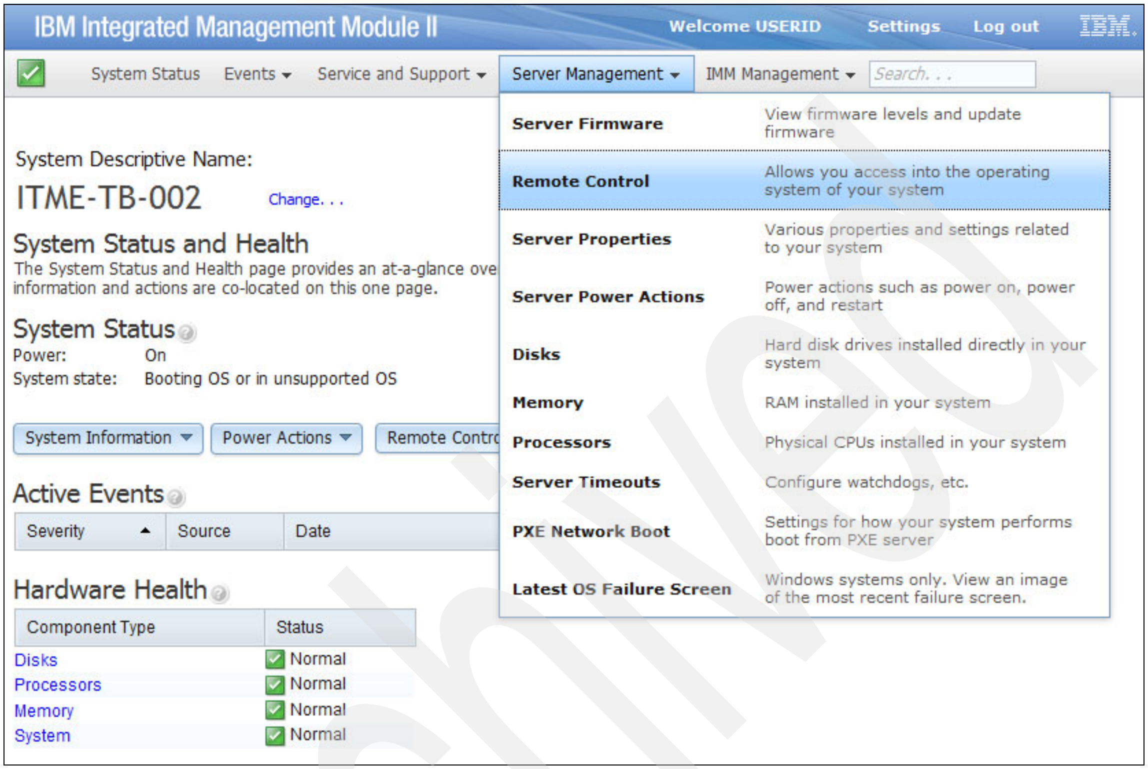

Figure 6-43 Remote control option in IMMv2

3. In the Remote Control window, click **Start remote control in single-user mode**, as shown in Figure 6-44. This action starts a Java applet on the local desktop that is used as a console session to the FSM.

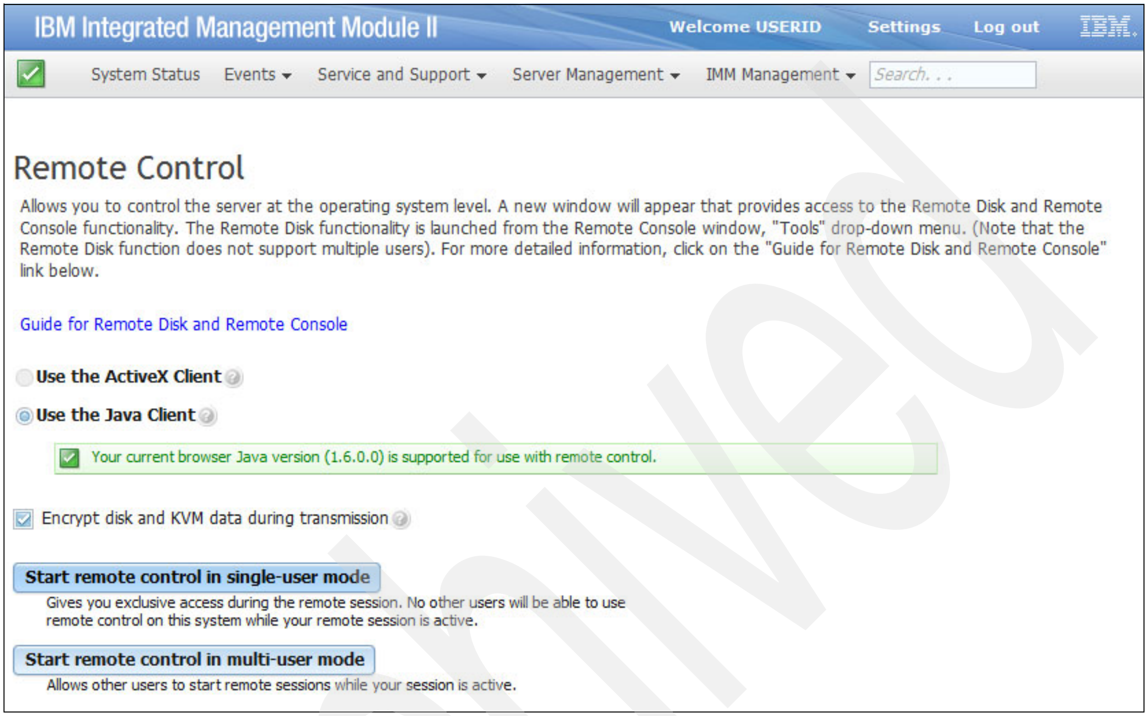

Figure 6-44 Starting remote console from IMMv2

Figure 6-45 shows the Java console window opened to the FSM appliance before power is applied.

Figure 6-45 FSM console in power off state

4. The FSM can be powered in several ways, including the physical power button on the FSM, or from the Chassis Management Module. For this example, using the **Tools/Power/On** option from the remote console menu, as shown in Figure 6-46, is the most convenient.

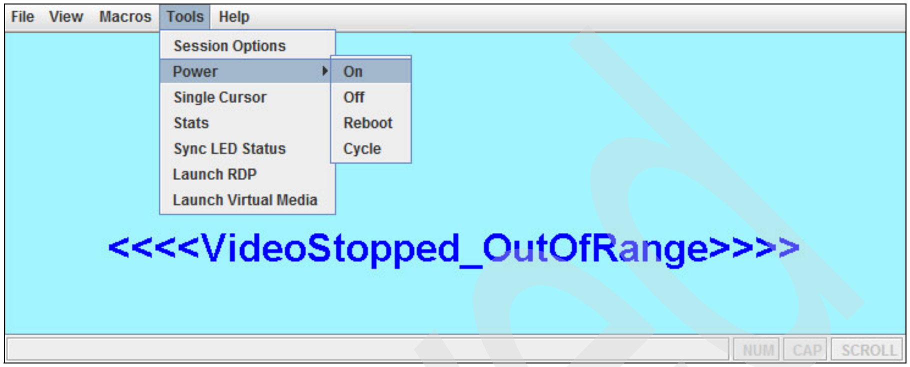

Figure 6-46 Powering on the FSM from the remote console session

As the FSM powers up and boots, the process can be monitored, but no input is accepted until the License Agreement window, shown in Figure 6-47, opens.

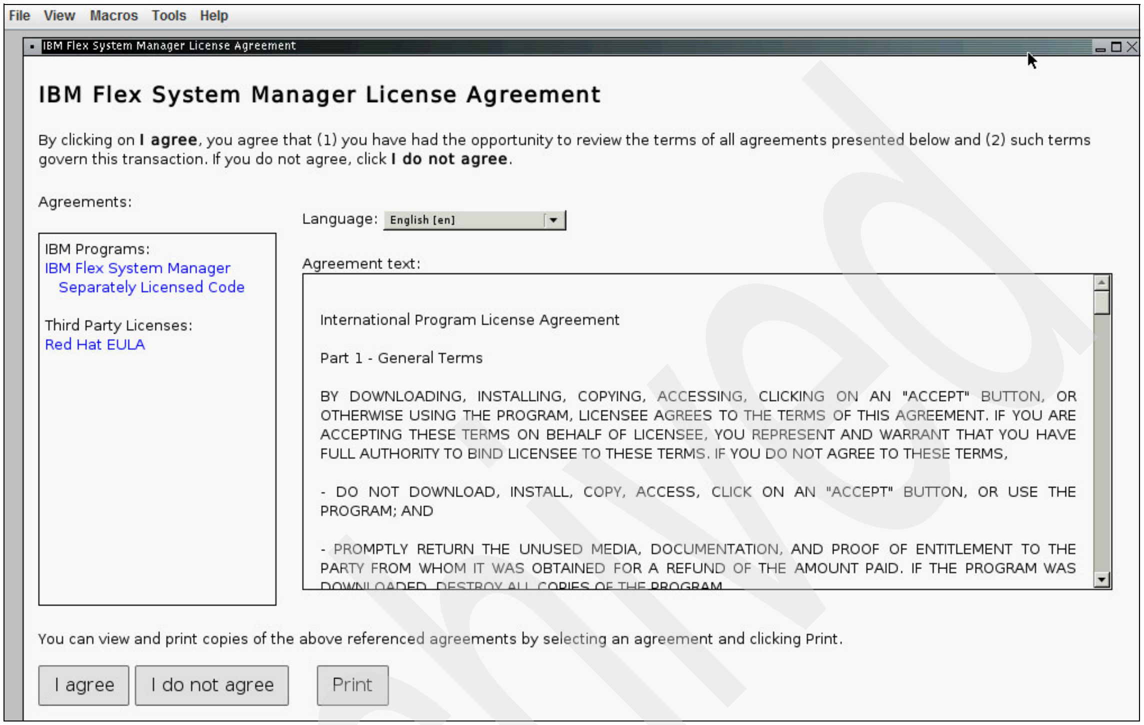

Figure 6-47 FSM license agreement

5. Click **I agree** to continue, and the startup wizard Welcome window opens, as shown in Figure 6-48.

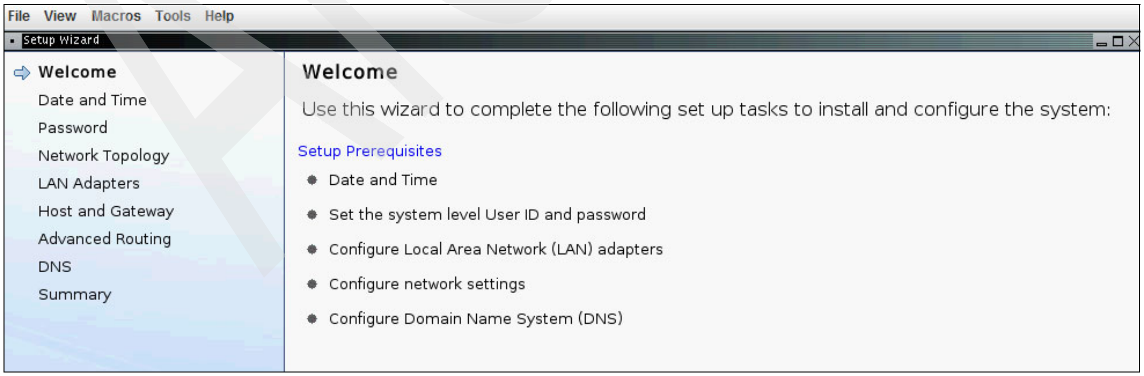

Figure 6-48 FSM Welcome window

6. Click **Data and Time** from the wizard menu to open the window shown in Figure 6-49. Set the time, date, time zone, and Network Time Protocol server, as needed.

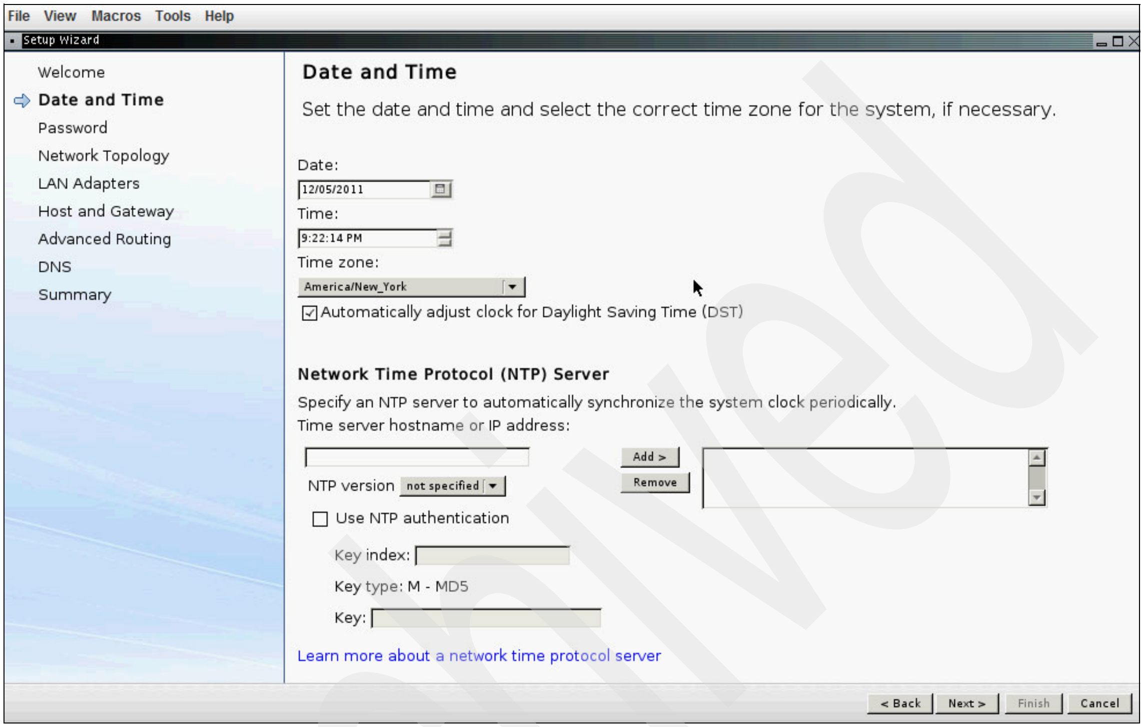

Figure 6-49 Setting the FSM date and time

Click **Next**.

7. Create a user ID and password for accessing the GUI and CLI. User ID and password maintenance, including creating additional user IDs, is available in IBM Flex System Manager after the startup wizard completes. Figure 6-50 shows the creation of user ID *USERID* and entering a password.

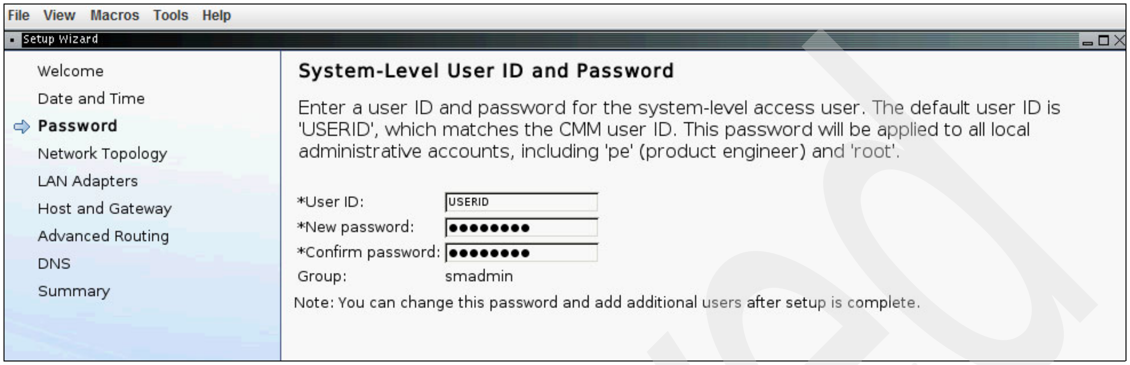

Figure 6-50 FSM system level user ID and password step

Click **Next** to continue.

8. Network topology options include separate networks for management and data, or a single network for both data and management traffic from the chassis. The preferred practice is to have separate management and data networks. To simplify this example, a combined network is configured, using the topology on the right side of Figure 6-51.

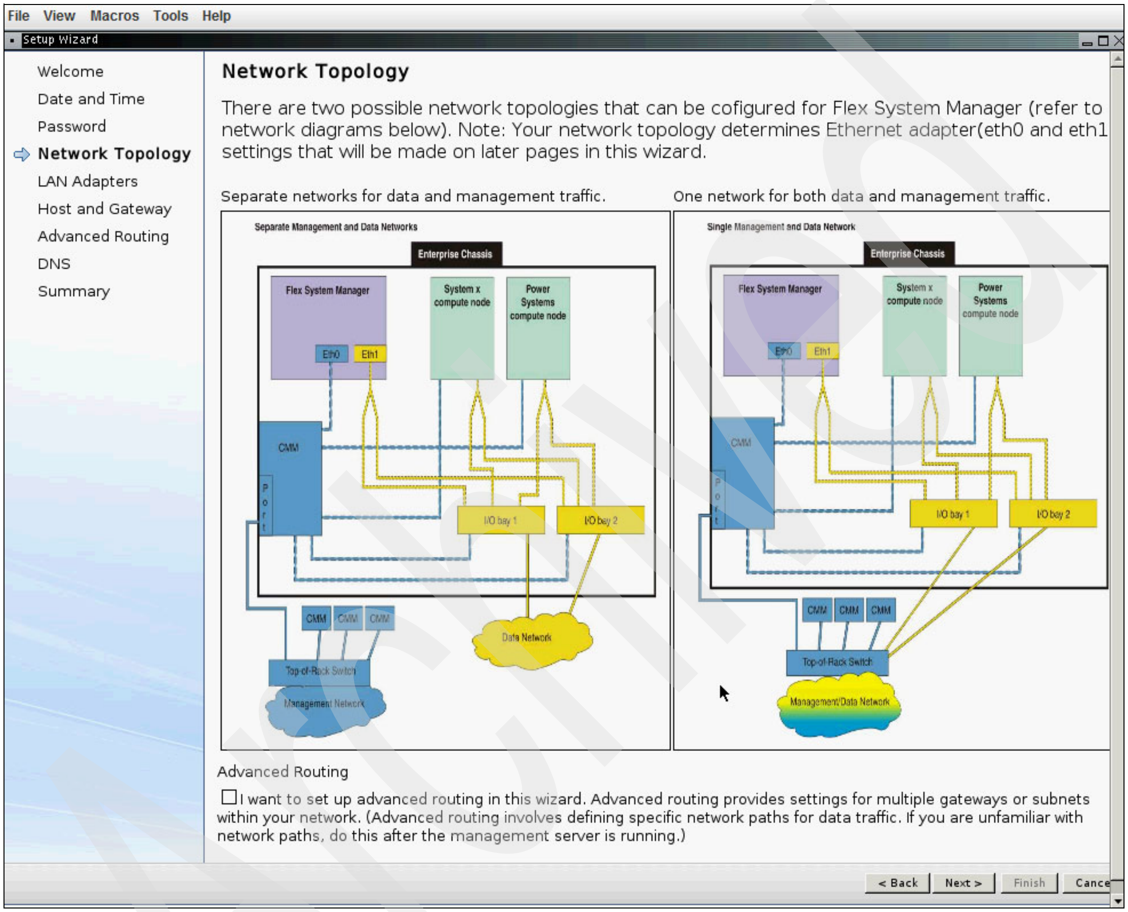

Figure 6-51 FSM network topology options

Click **Next** to continue to the actual network configuration.

9. The LAN adapter configuration is shown in Figure 6-52 on page 211. There are two adapters listed. The first is from the FSM management network adapter, which allows FSM to communicate on the chassis management network. Traffic from this adapter flows through the Chassis Management Module and uses the CMM physical connection to the network.

The second LAN adapter represents one of the integrated Ethernet ports or LAN on motherboard (LOM). Traffic from this adapter flows through the Ethernet switch in the first I/O switch bay of the chassis, and is used as a separate data connection to the FSM. The radio button for the first adapter is preselected (Figure 6-52).

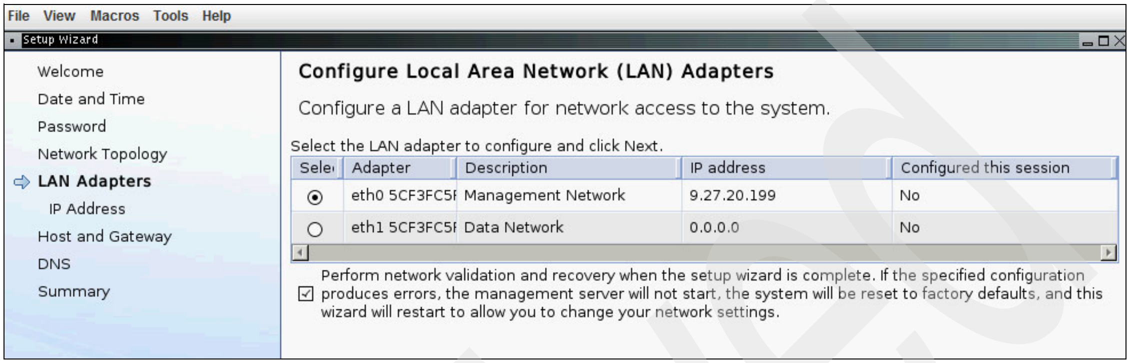

Figure 6-52 FSM LAN adapter configuration

Click **Next** to continue.

10. The Configure IP Address window, shown in Figure 6-53, allows the selection of DHCP or static IP options for IPv4 and IPv6 addressing. Select the wanted options, enter the information as required, and then click **Next**.

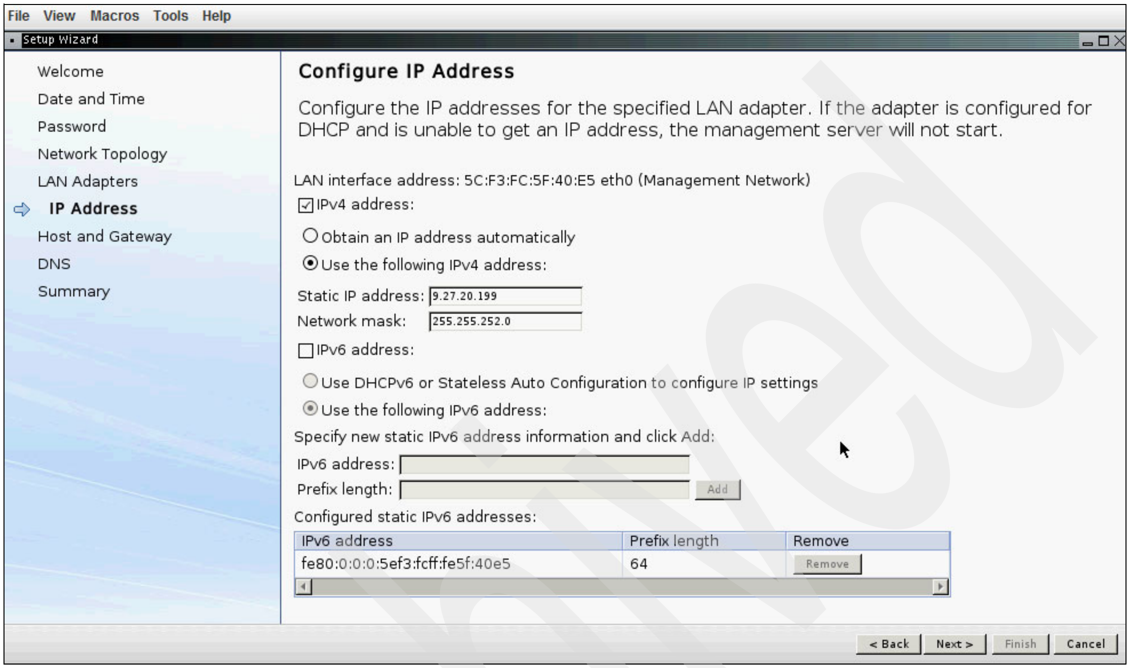

Figure 6-53 FSM IP address assignment

After completing the previous step, the wizard cycles back to the Initial LAN Adapter window and preselects the next adapter in the list, as shown in Figure 6-54.

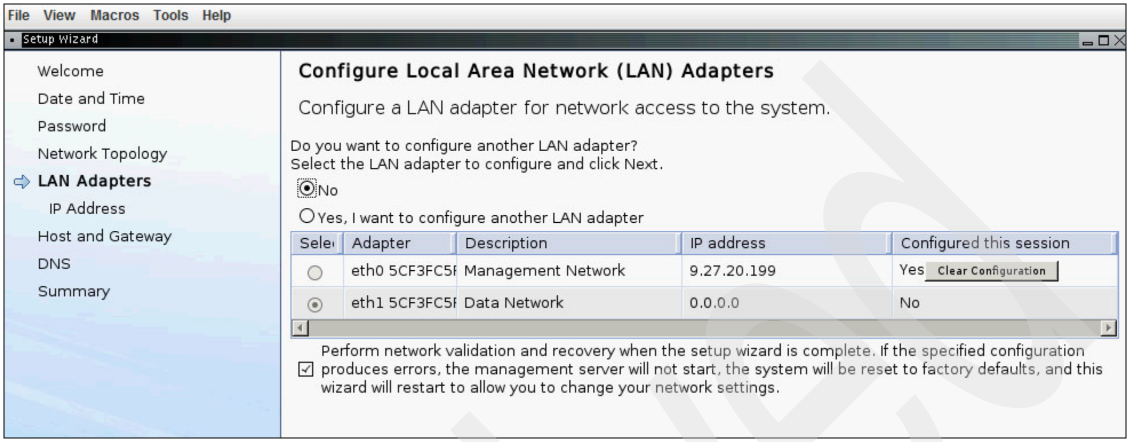

Figure 6-54 FSM LAN adapter configuration continue option

In our example, we are using a combined network topology and a single adapter, so additional IP addresses are not needed.

11. Select **No** by the question, “Do you want to configure another LAN adapter?”, as shown in figure Figure 6-54.

    Click **Next** to continue.

12. After IP address assignment, the host name and gateway are configured as shown in Figure 6-55. Enter the wanted host name, domain name, and default gateway address. Note the IP address and the default gateway adapter in the drop-down menu and ensure that the values are correct.

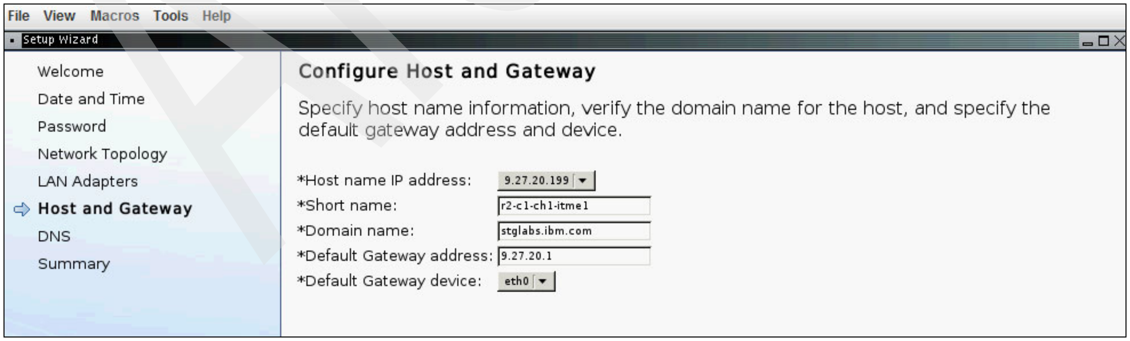

Figure 6-55 FSM host name and gateway configuration

Click **Next** to continue.

> **Important:** It is expected that the host name of the FSM is available on the domain name server.

13. You can enable the use of DNS services and add the address of one or severs and a domain suffix search order.

    Enter the information, as shown in Figure 6-56, and click Next.

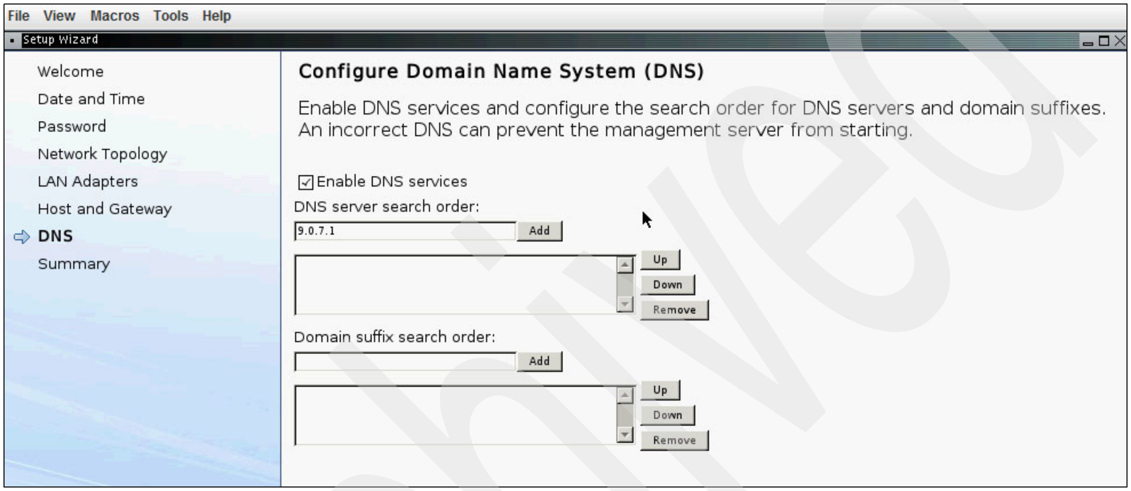

Figure 6-56 FSM DNS services configuration

14. The final step of the setup wizard is shown in Figure 6-57. This windows shows a summary of all configured options.

    To change a selection, click **Back**. If no changes are needed, click **Finish**.

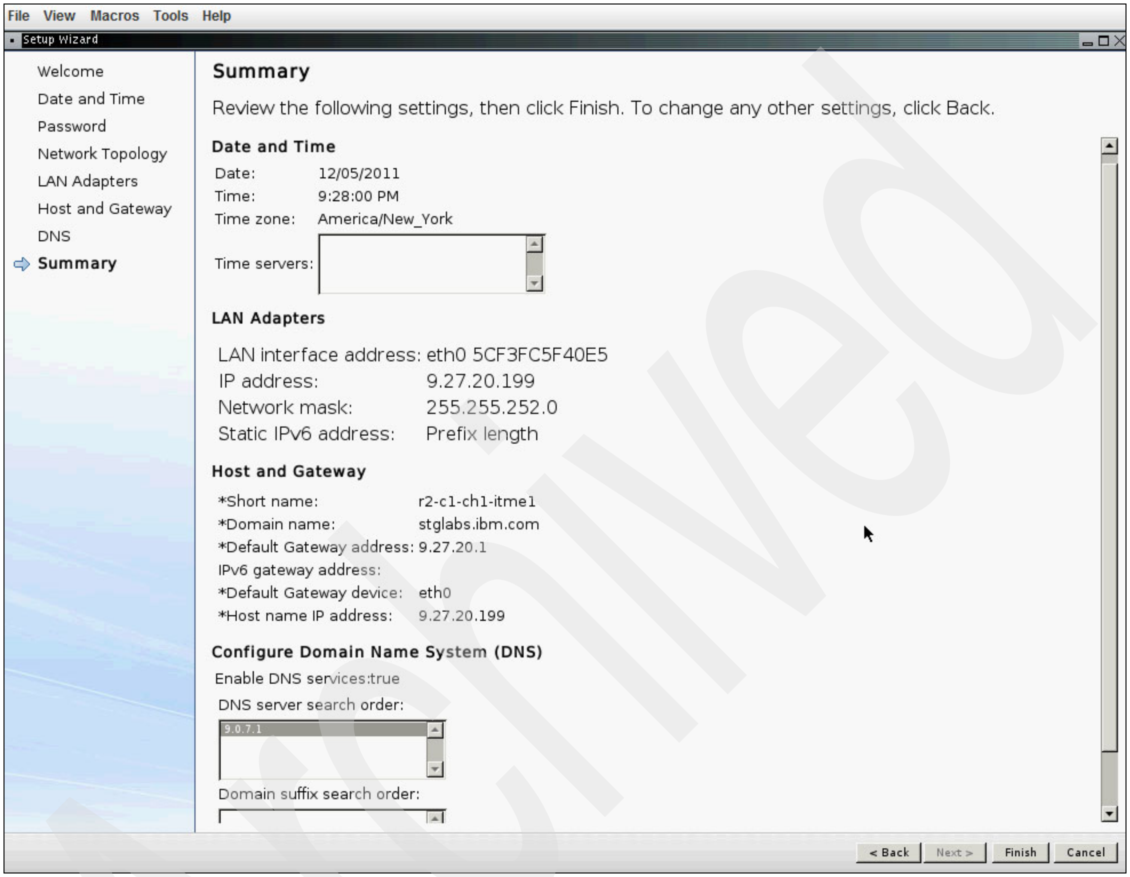

Figure 6-57 FSM startup wizard summary window

After **Finish** is clicked, the final configuration and setup proceeds automatically without any further input, as shown in Figure 6-58 through Figure 6-61 on page 217.

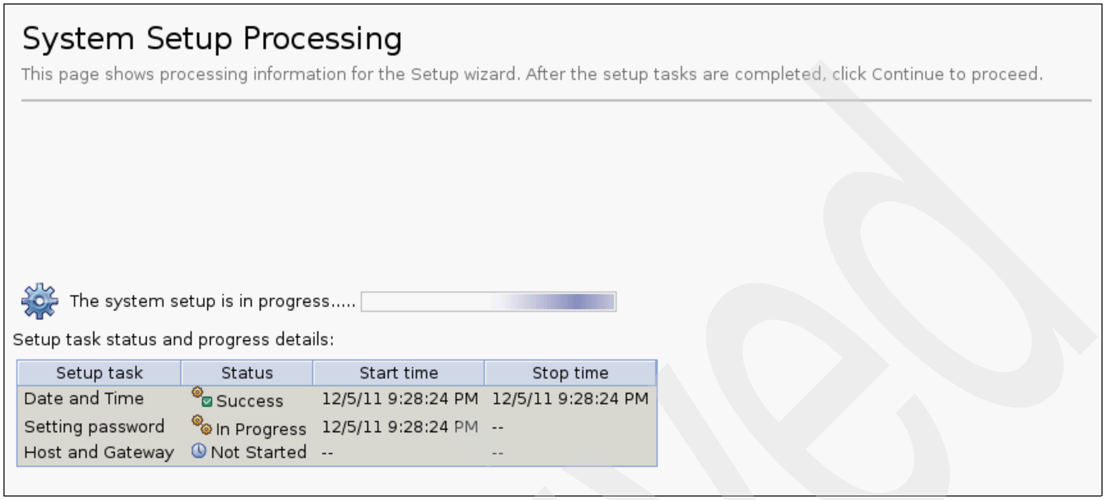

Figure 6-58 FSM system setup processing status

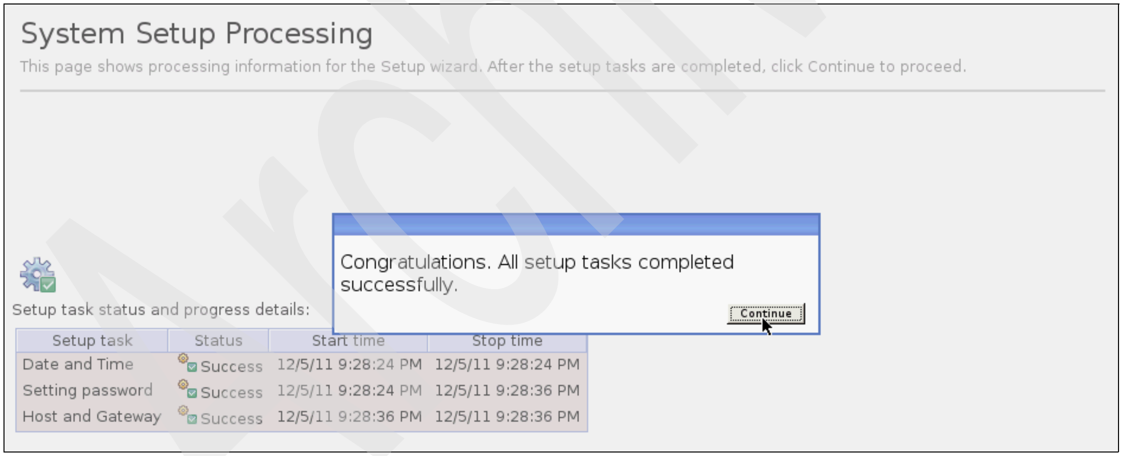

Figure 6-59 FSM system setup processing completion

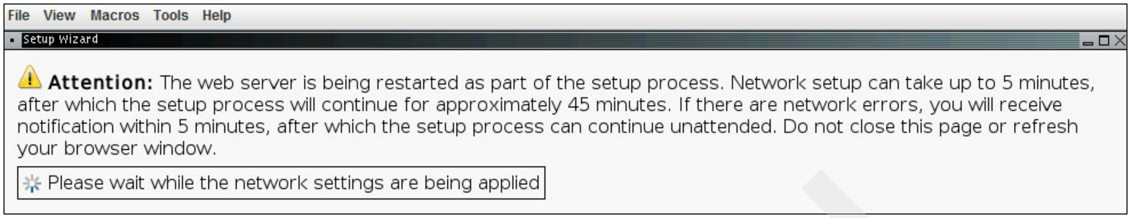

Figure 6-60 FSM startup

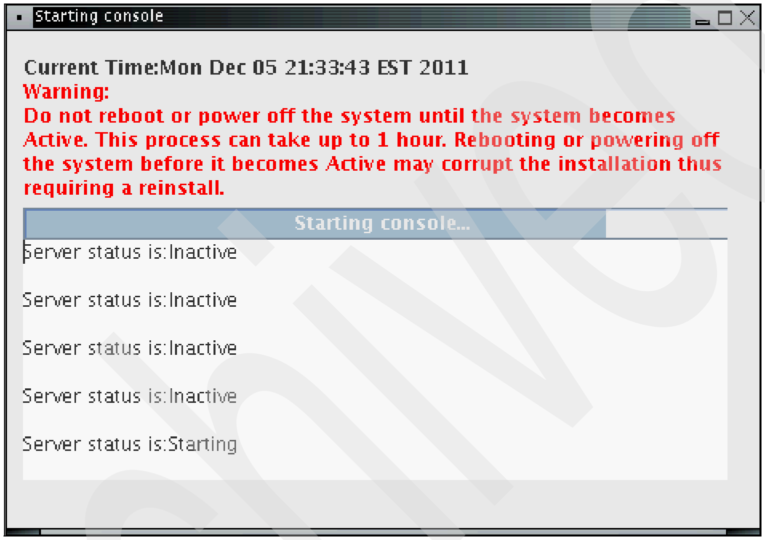

Figure 6-61 FSM startup status

15. With startup completed, the local browser on the FSM also starts. A list of untrusted connection challenges opens.

Accept these challenges by clicking **I Understand the Risks** and **Add Exception**, as shown in Figure 6-62 and Figure 6-63.

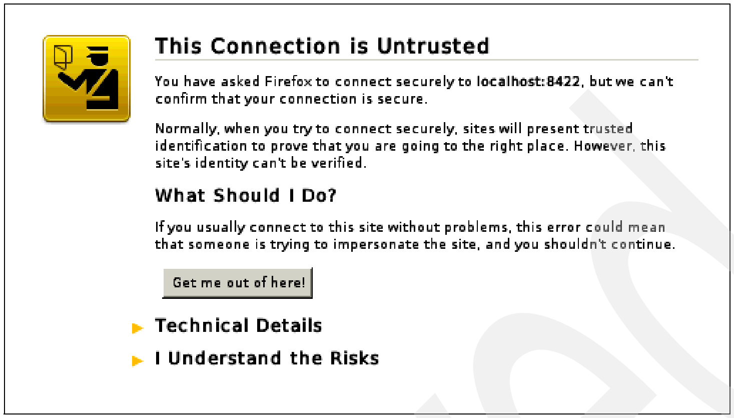

Figure 6-62 FSM browser setup

Figure 6-63 FSM browser exception add

16. Click **Confirm Security Exception**, as shown in Figure 6-64.

Figure 6-64 FSM security exception confirmation

17. With the security exceptions cleared, the Login window of the IBM Flex System Manager GUI opens.

Enter the user ID and credentials that were entered in the startup wizard, and click **Log in**, as shown in Figure 6-65.

Figure 6-65 FSM Login window

A Getting Started window opens and reminds you that initial setup tasks must be completed (Figure 6-66).

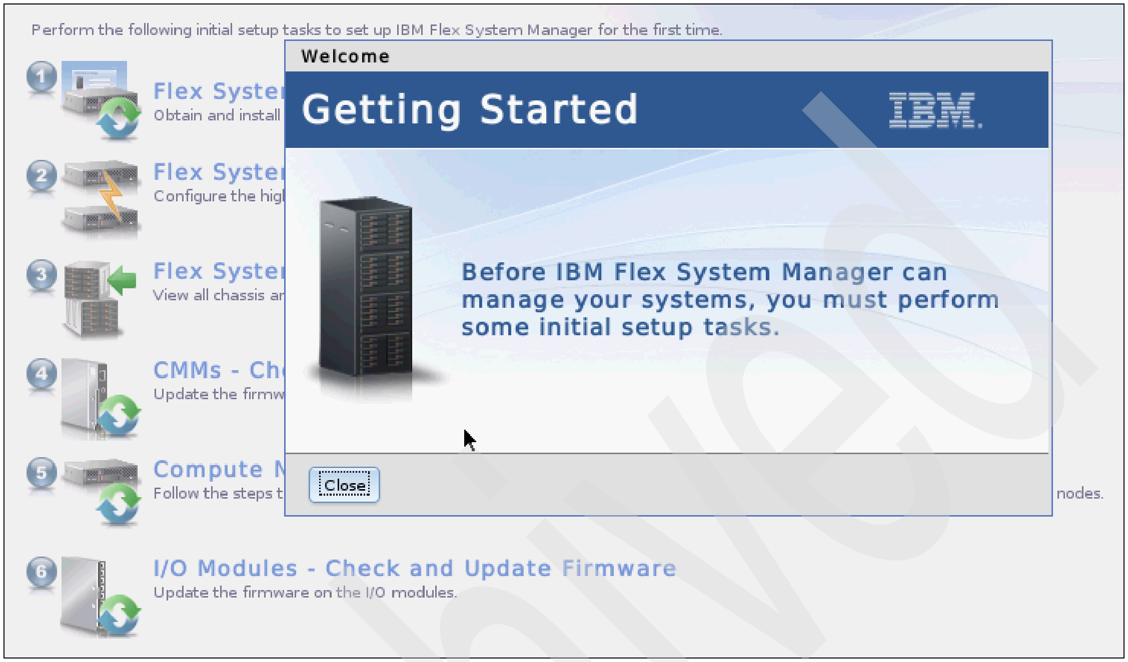

Figure 6-66 FSM Getting Started reminder

The startup wizard and initial login are complete. The FSM is ready for further configuration and use. Our example uses a console from the remote console function of the IMMv2. At this time, a secure browser session can be started to the FSM.

6.5.2 Flex System Manager Update Manager configuration

The IBM Flex System Manager Update Manager is responsible for obtaining and applying chassis, switch, system firmware, and operating system updates from IBM. In addition, the Update Manager is used to update the FSM itself. The updates can be obtained through an Internet connection from the FSM, or they can be downloaded manually from IBM to another workstation, then copied to the FSM by an FTP or SCP connection. After the files are copied to the FSM, they must be imported into the Update Manager.

To download these updates, we must set up the Internet connection, as described in "Direct Internet connection" on page 222.

Direct Internet connection

To set up and test the Internet connection, complete the following steps:

1. Starting from the Home page, click the **Plug-ins** tab. The Plug-ins window lists all of the managers that are available on the FSM, as shown in Figure 6-67.

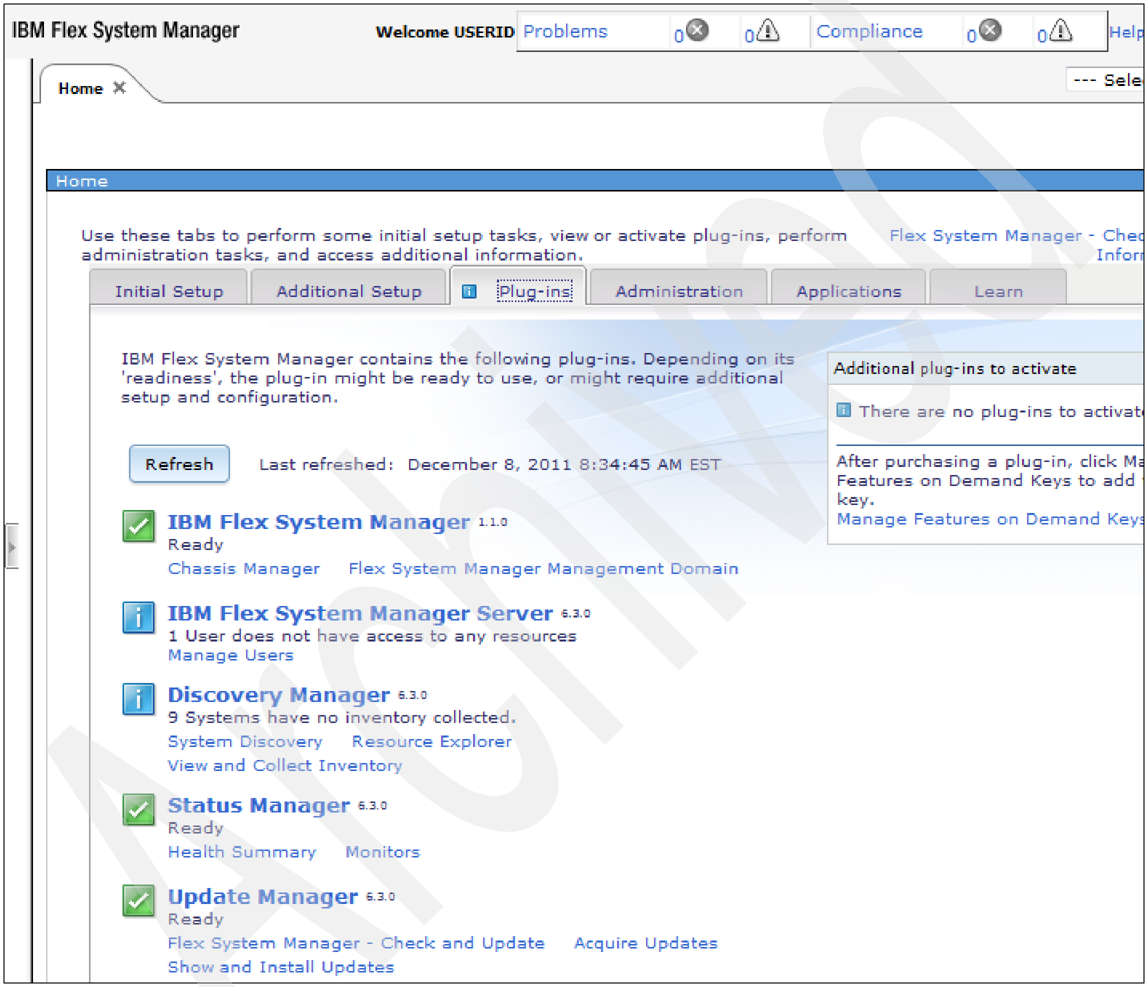

Figure 6-67 FSM list of managers

2. From the list of managers, click **Update Manager** to open the window shown in Figure 6-68.

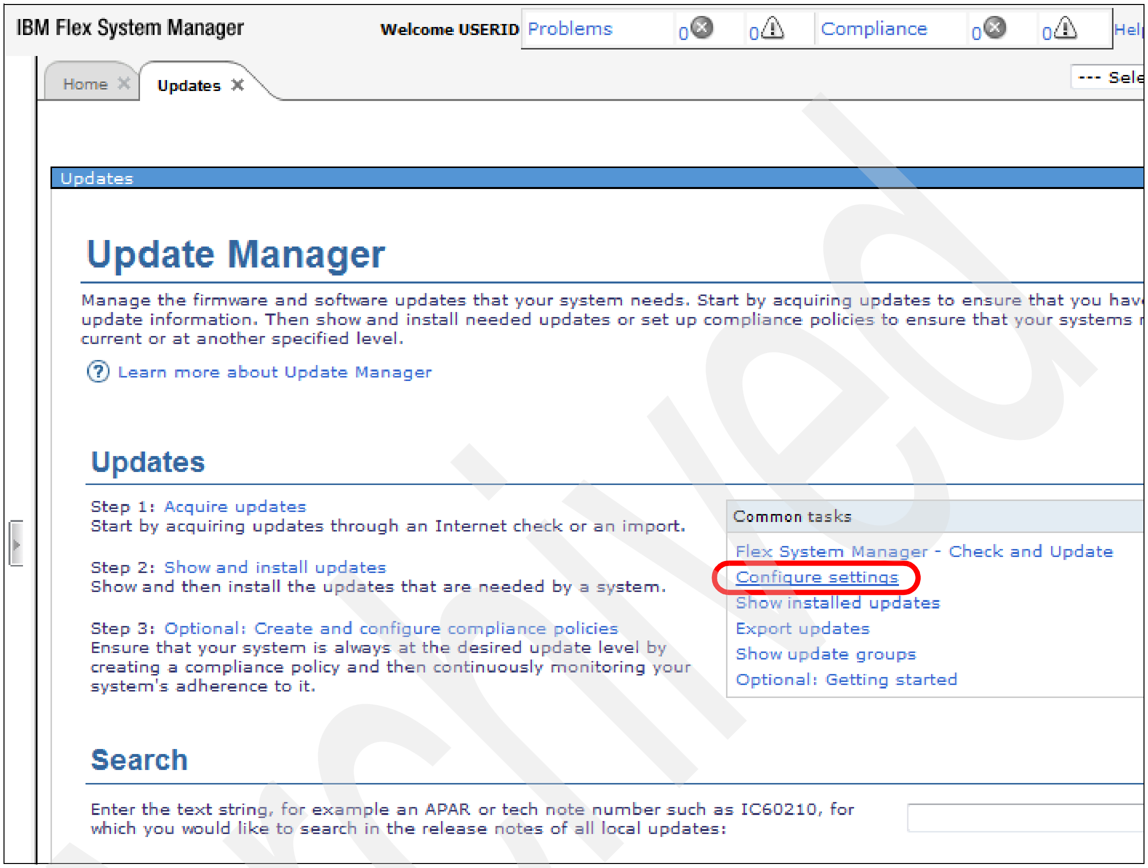

Figure 6-68 FSM Update Manager

3. In the Common task box, click **Configure settings** to open the window shown in Figure 6-69. You can use this window to configure a direct Internet connection, or use the configuration settings to use an existing proxy server.

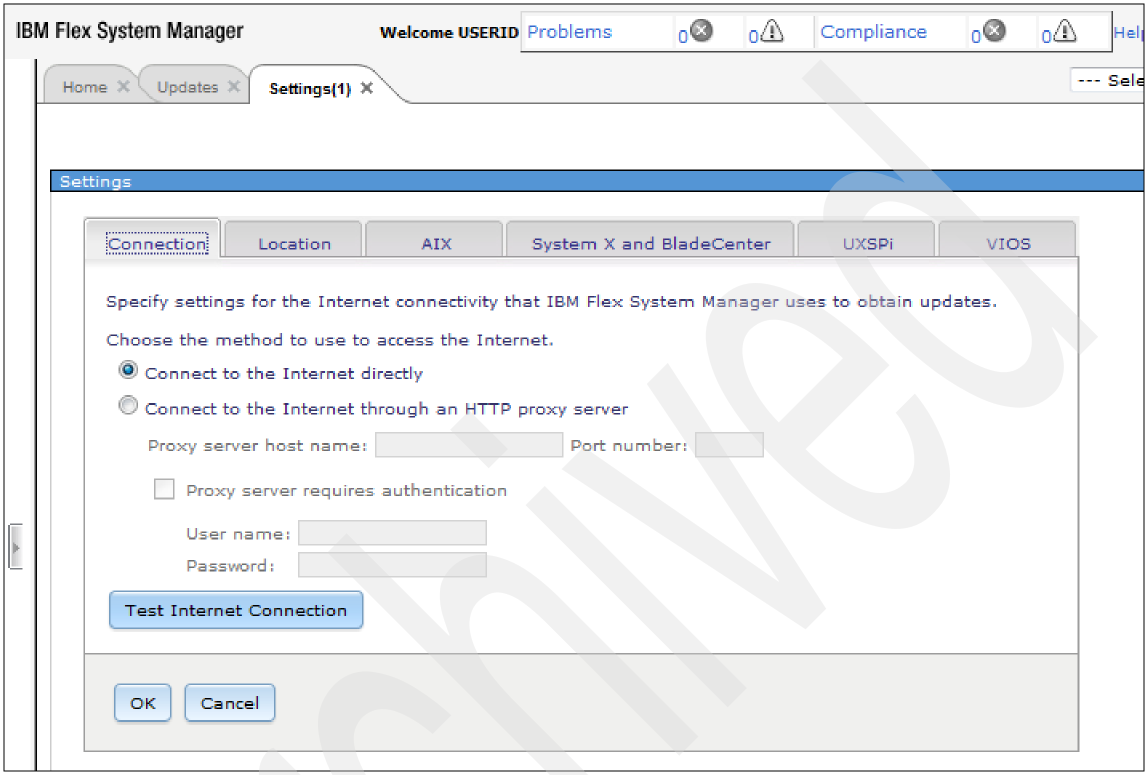

Figure 6-69 FSM Update Manager Internet connection settings

4. With the settings complete, click **Test Internet Connection** to verify the connection.

The test attempts to make a connection to a target IBM server. During the test, a progress indicator opens, as shown in Figure 6-70.

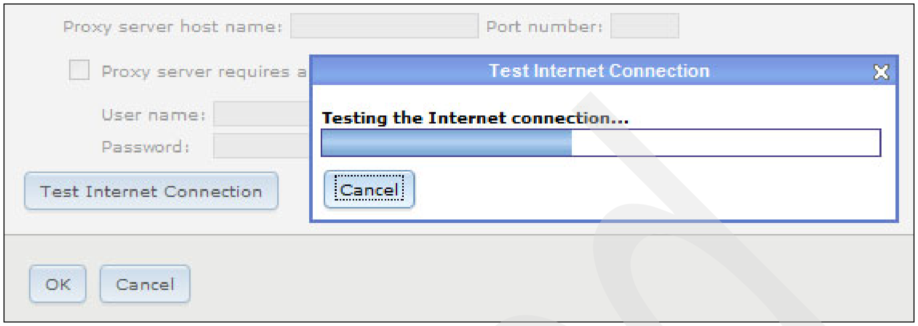

Figure 6-70 FSM testing Internet connection for Update Manager

A successful completion message opens (Figure 6-71).

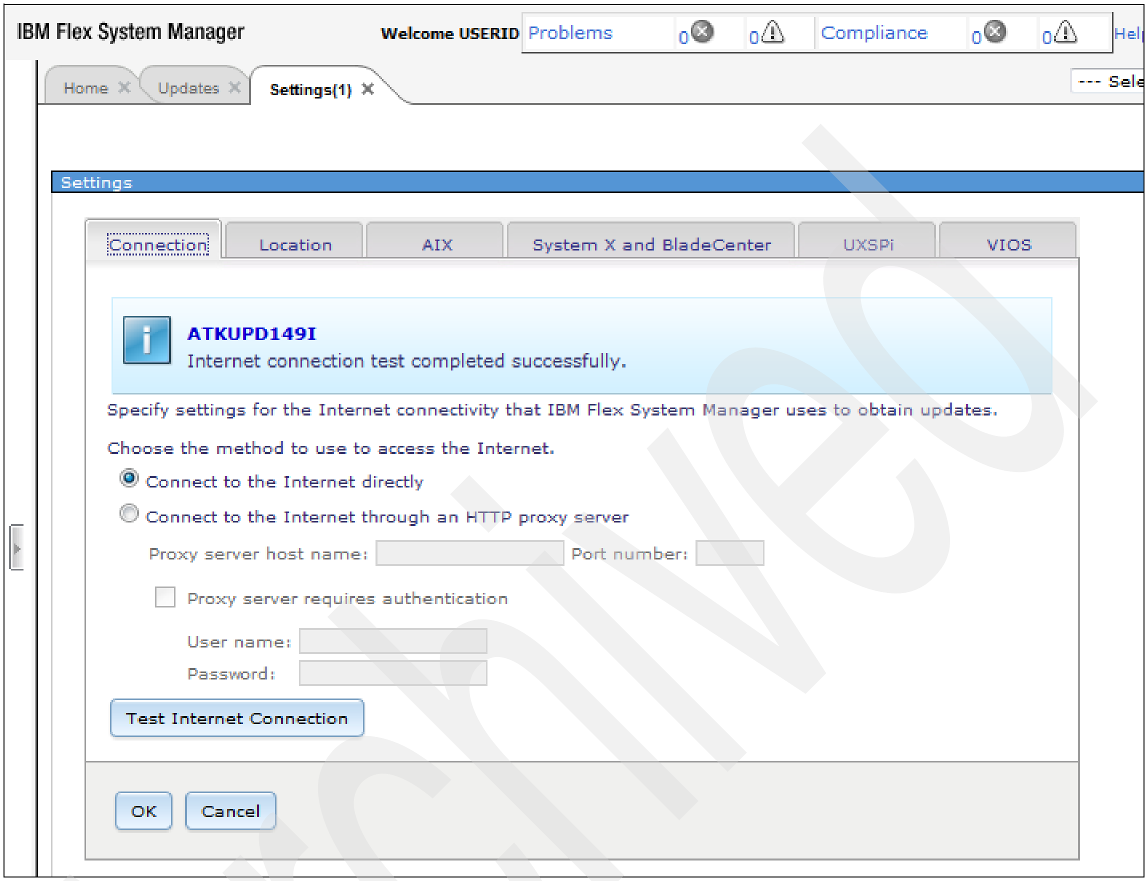

Figure 6-71   Successful Internet connect test for Update Manager

After the test succeeds, the Update Manager can obtain update packages directly from IBM.

If a direct Internet connection is not allowed for the FSM, complete the steps described in “Importing update files” to import the update files into Update Manager.

Importing update files

This section describes how to import files into Update Manager. Because the FSM is an appliance, methods that moves files directly to the FSM are not allowed. The file movement process must be initiated by the FSM from a user login, either from the GUI or from a CLI.

The **scp** command is used to copy the update files from a local workstation to the FSM. The update files on the local workstation are obtained from IBM Fix Central. From an **ssh** login, you have access only to the `/home/`*`userid`* directory. Additional subdirectories can be created and files copied and removed from these subdirectories, but running **cd** to the subdirectory is a restricted operation.

To import the update files using the GUI, complete the following steps:

1. Beginning at the Update Manager window, click **Acquire updates**, as shown in Figure 6-72.

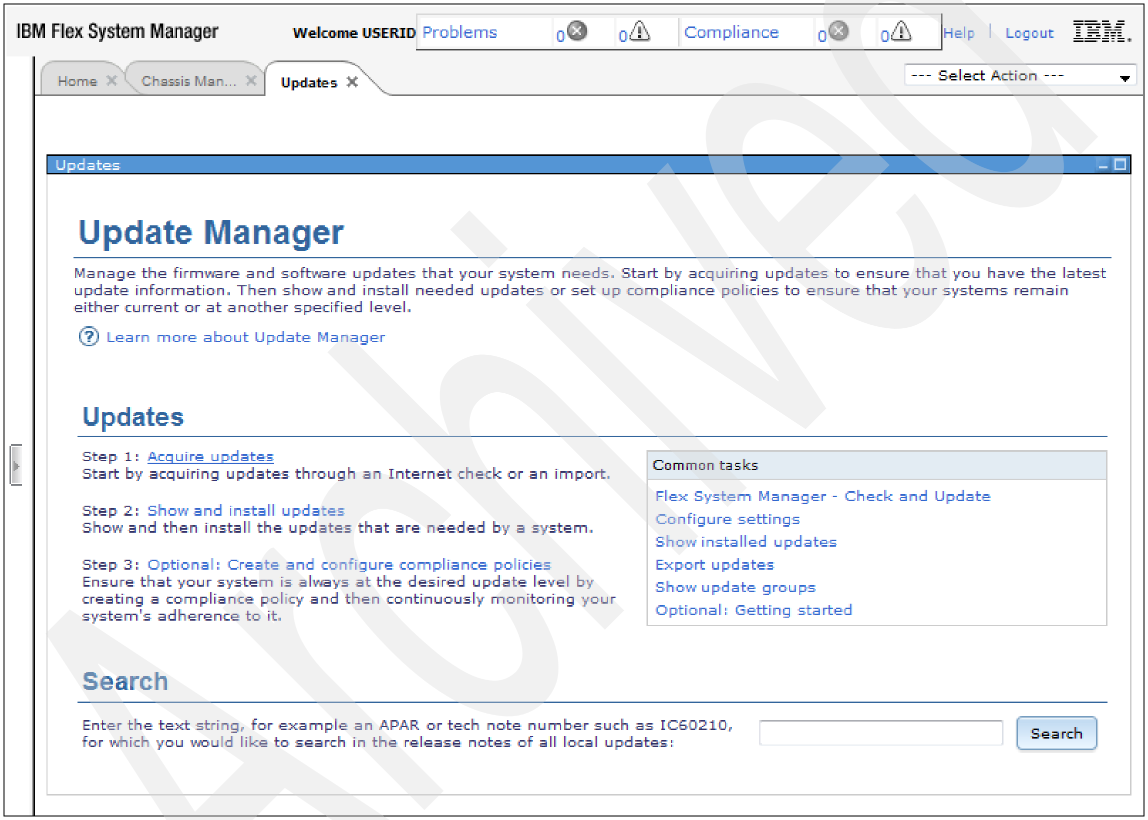

Figure 6-72 FSM Update Manager

Figure 6-73 shows the window that opens. Two options are available: Check for updates using an Internet connection, or Import updates from the file system. For this example, we import the updates.

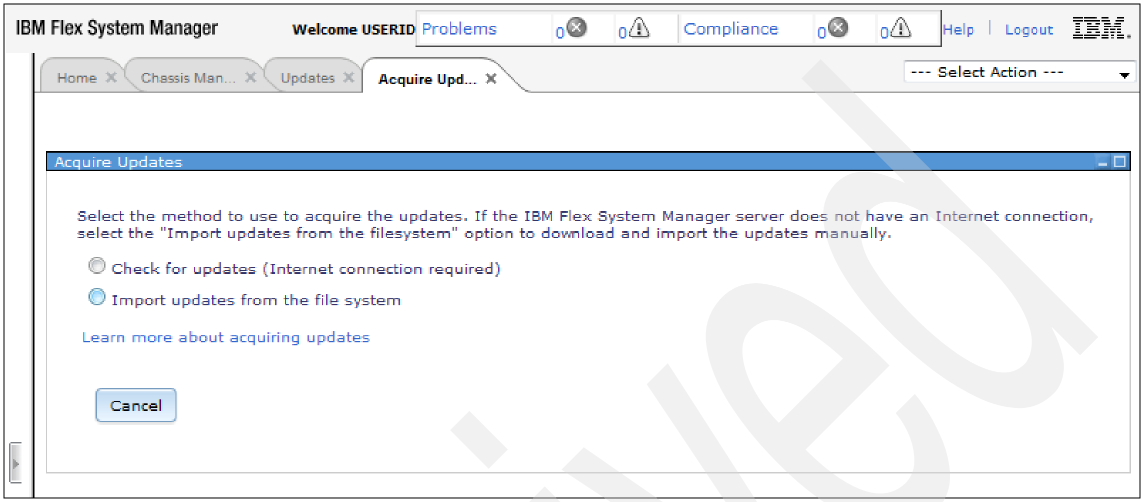

Figure 6-73 Update Manager Acquire Updates selection

2. Select **Import updates from the file system**.

3. Enter the path for the updates that were manually copied to the IBM Flex System Manager, as shown in Figure 6-74.

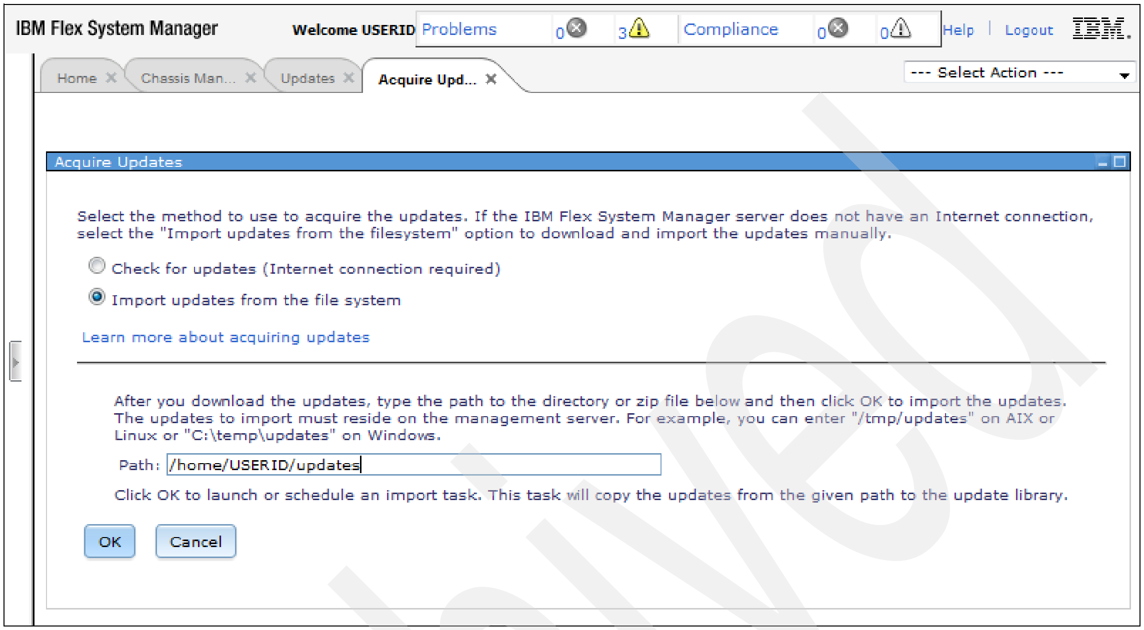

Figure 6-74 Import path to update files

4. Click **OK**, and the IBM Flex System Manager job scheduler opens, as shown in Figure 6-75.

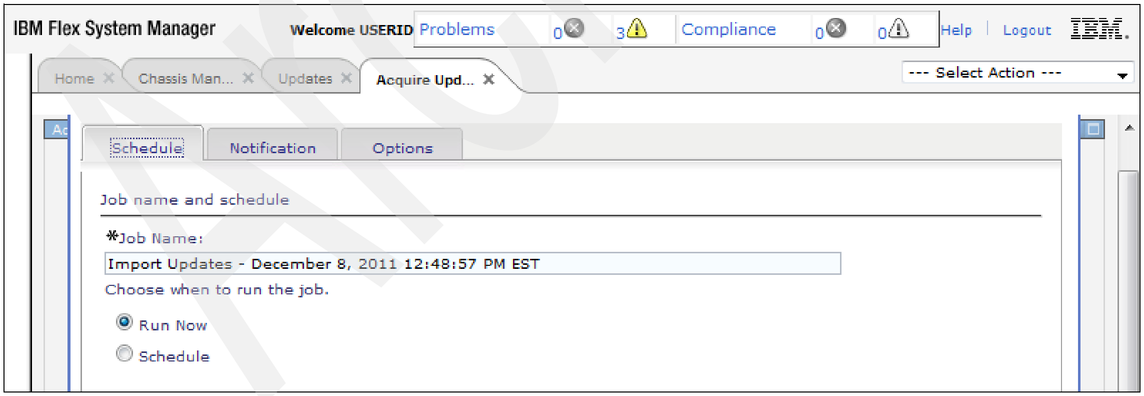

Figure 6-75 Starting the import updates job on FSM

For this example, we left the default **Run Now** selected.

5. Click **OK** at the bottom of the window to start the job.

   When the import updates job starts, the Acquire Updates window refreshes with a message that indicates the new job. The status of the running job can be monitored by clicking **Display Properties**, as shown in Figure 6-76.

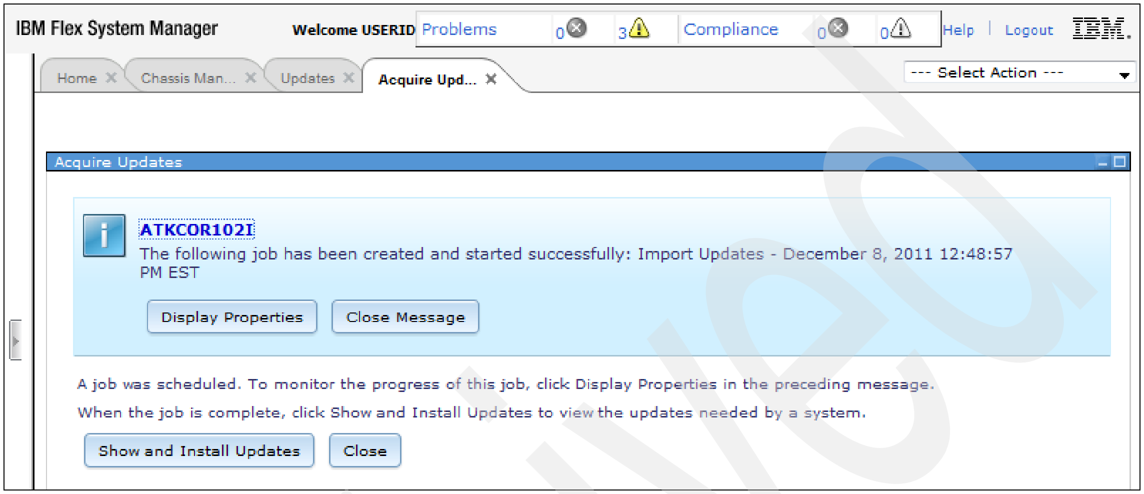

Figure 6-76 Update Manager

6.5.3 Initial chassis management with IBM Flex System Manager

Most tasks in the IBM Flex System Manager can be accomplished by more than one method using the GUI. In this section, we describe one common method.

After the initial setup of the FSM finishes, FSM discovers any available chassis. You can then decide which chassis is managed by the current FSM. To accomplish this task, complete the following steps:

1. Click the **Home** tab.
2. Click the **Initial Setup** tab to open the Initial Setup window.

3. Click **IBM Flex System Manager Domain - Select Chassis to be Managed** (Figure 6-77).

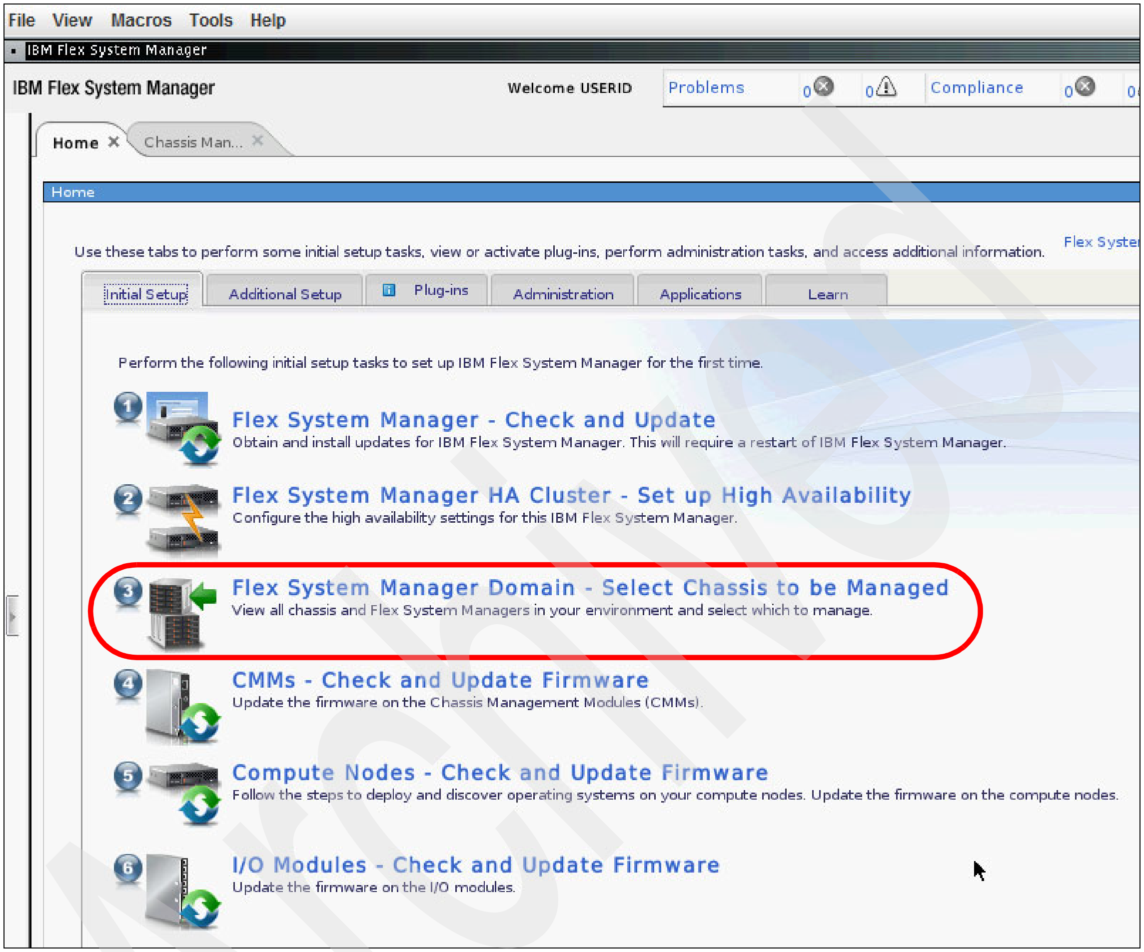

Figure 6-77 FSM initial setup window

A window with a list of available chassis opens, as shown in Figure 6-78.

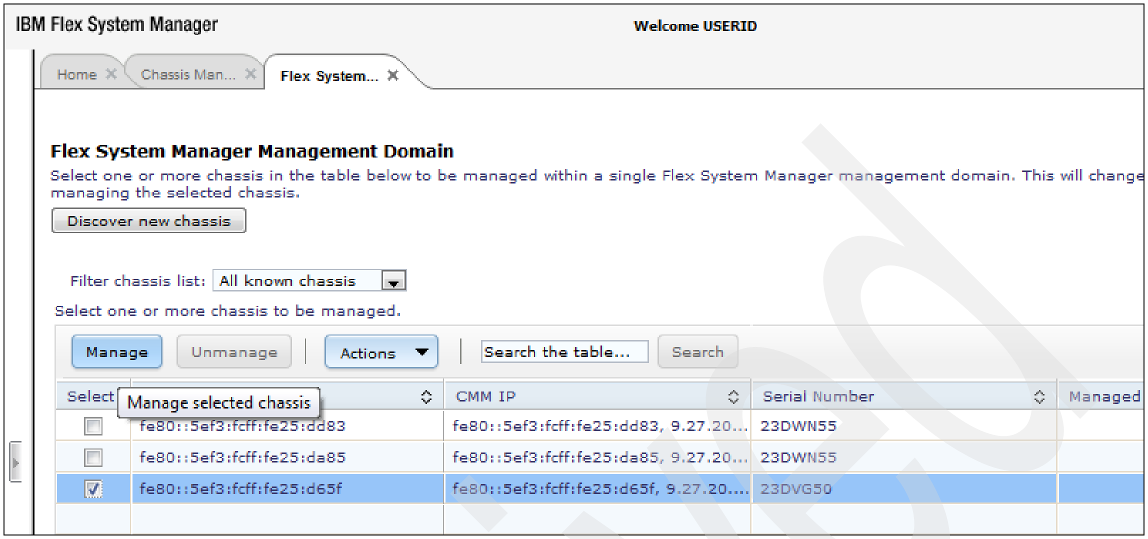

Figure 6-78 FSM chassis selection for management

4. Select the box in front of the wanted chassis.
5. Click **Manage**. The Manage Chassis window opens.

The Manage Chassis window, shown in Figure 6-79, lists the selected chassis. A drop-down menu lists the available IBM Flex System Manager systems.

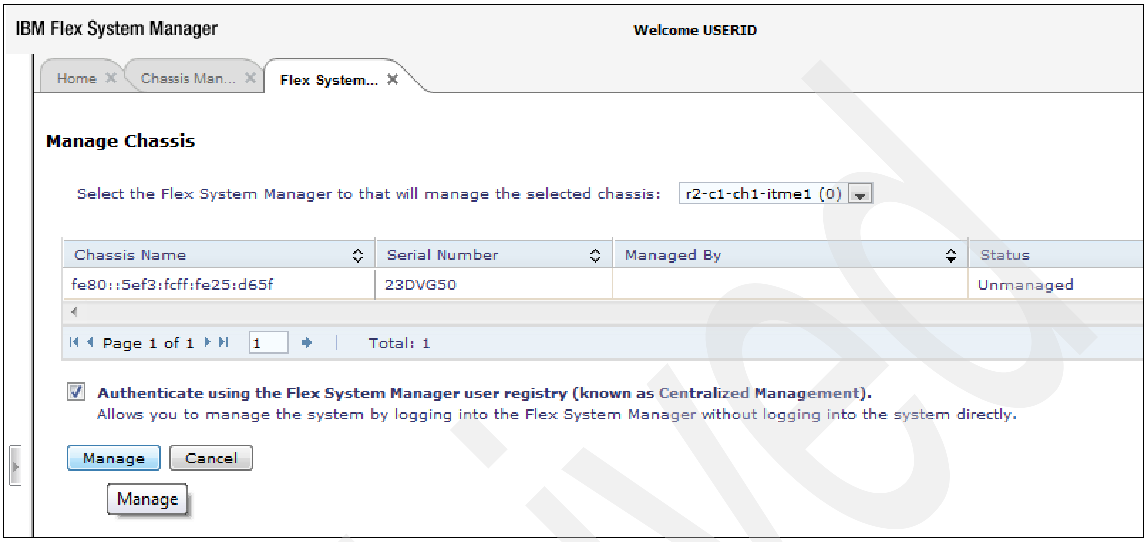

Figure 6-79 FSM - manage chassis options

6. Ensure that the chassis and IBM Flex System Manager selections are correct.
7. Click **Manage**. This action updates the Message column from Waiting to Finalizing, then Managed, as shown in Figure 6-80 and Figure 6-81 on page 234.

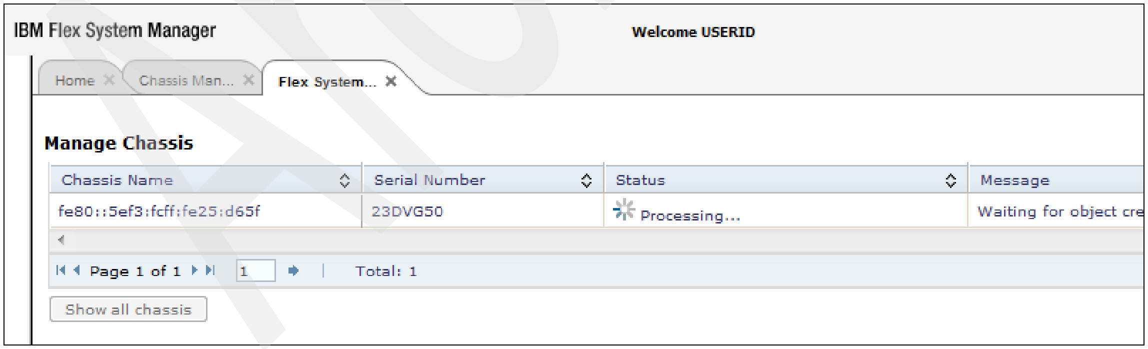

Figure 6-80 FSM manage chassis Step 1

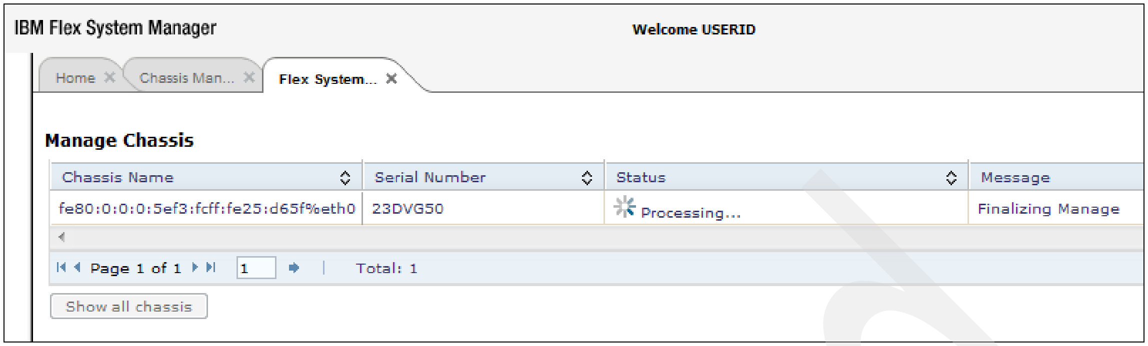

Figure 6-81 FSM manage chassis Step 2

8. After the successful completion of the manage chassis process, click **Show all chassis**, as shown in Figure 6-82.

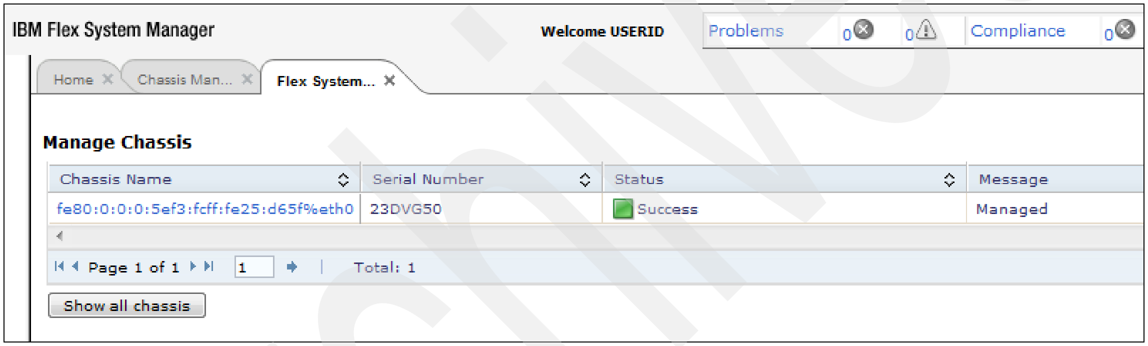

Figure 6-82 FSM manage chassis steps completed

The resulting window is the original IBM Flex System Manager Management Domain window, with the target chassis as the managing IBM Flex System Manager (Figure 6-83).

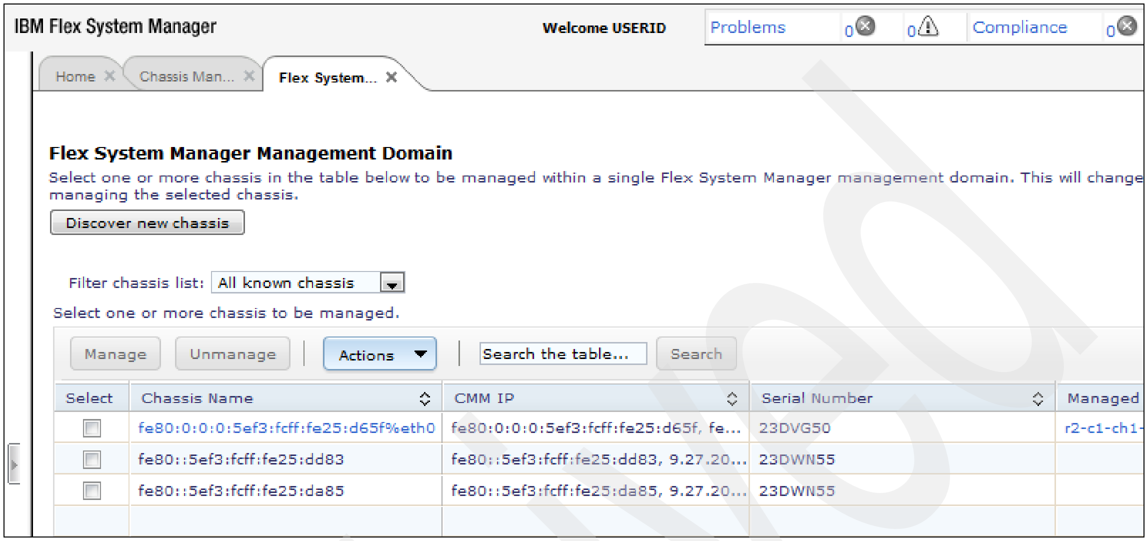

Figure 6-83 FSM with management domain updated

With the Enterprise Chassis now managed by the IBM Flex System Manager, the typical management functions on a Power Systems compute node can be performed.

6.6 Basic management of Power Systems compute nodes

Basic compute node management consists primarily of the following items:

- Discovery/access
- Inventory
- Hardware power on or off
- Virtual server creation
- Creating virtual consoles to virtual servers
- Firmware updates
- Error collection and reporting
- EnergyScale feature enablement
- Mobility

More advanced functions, such as VMControl functionality, are also available in the IBM Flex System Manager, but are not described in this book.

6.6.1 Managing Power Systems resources

The starting point for all of the functions is the Manage Power Systems Resources window. This part of the IBM Flex System Manager GUI can be started by the following steps.

Most operations in the IBM Flex System Manager use the Home page as the starting point. To access Manage Power Systems Resources, complete the following steps:

1. Click **Chassis Manager** (Figure 6-84).

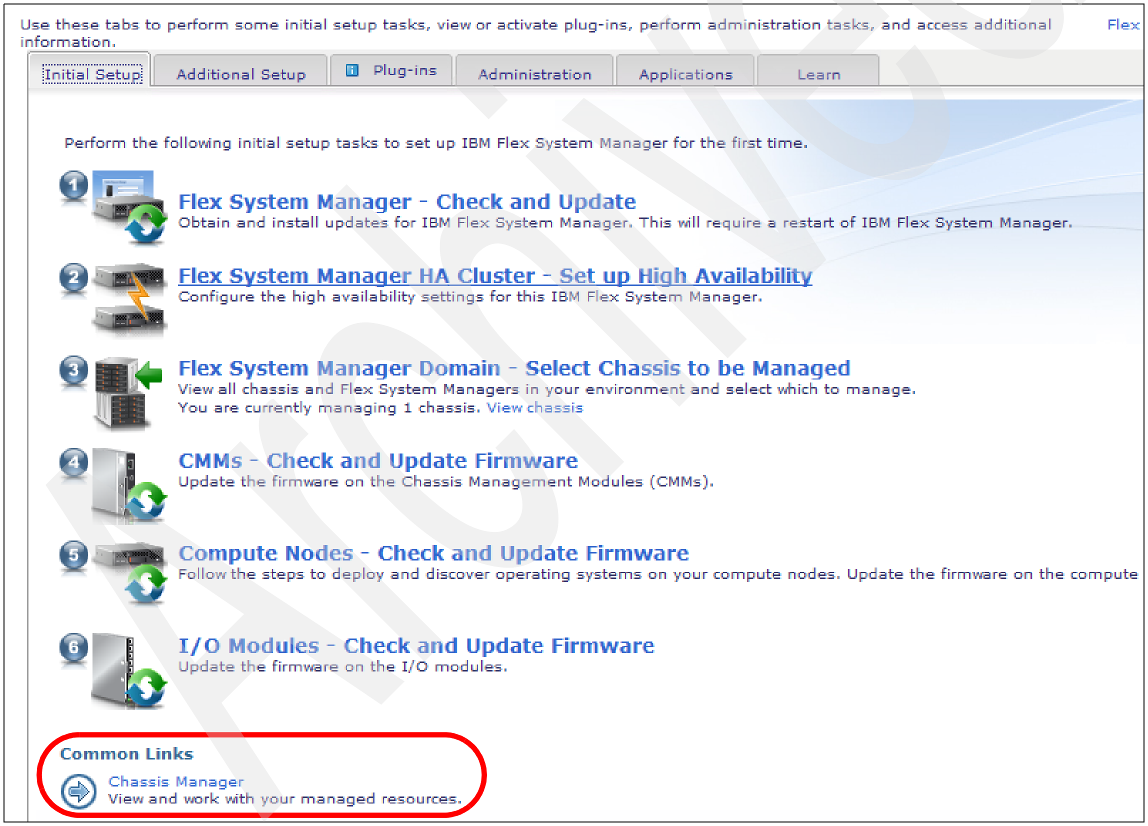

Figure 6-84 FSM Home page

A new tab opens that shows a list of managed chassis (Figure 6-85).

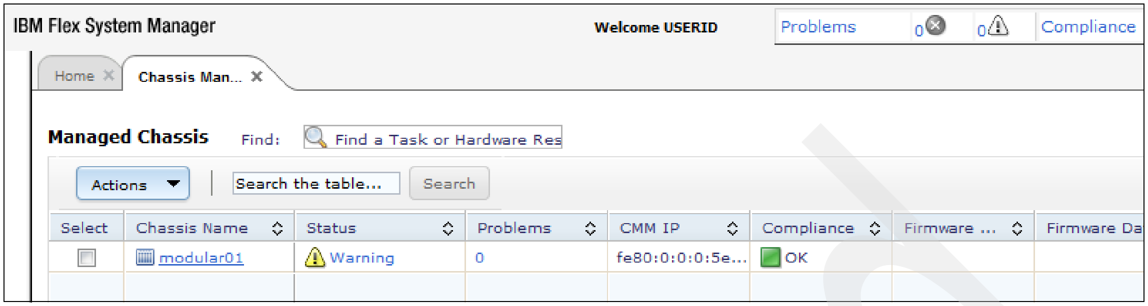

Figure 6-85 FSM Chassis Manager view

2. Click the name of the wanted chassis in the chassis name column (in this case, **modular01**).

   A window with a graphical view of the chassis opens (Figure 6-86).

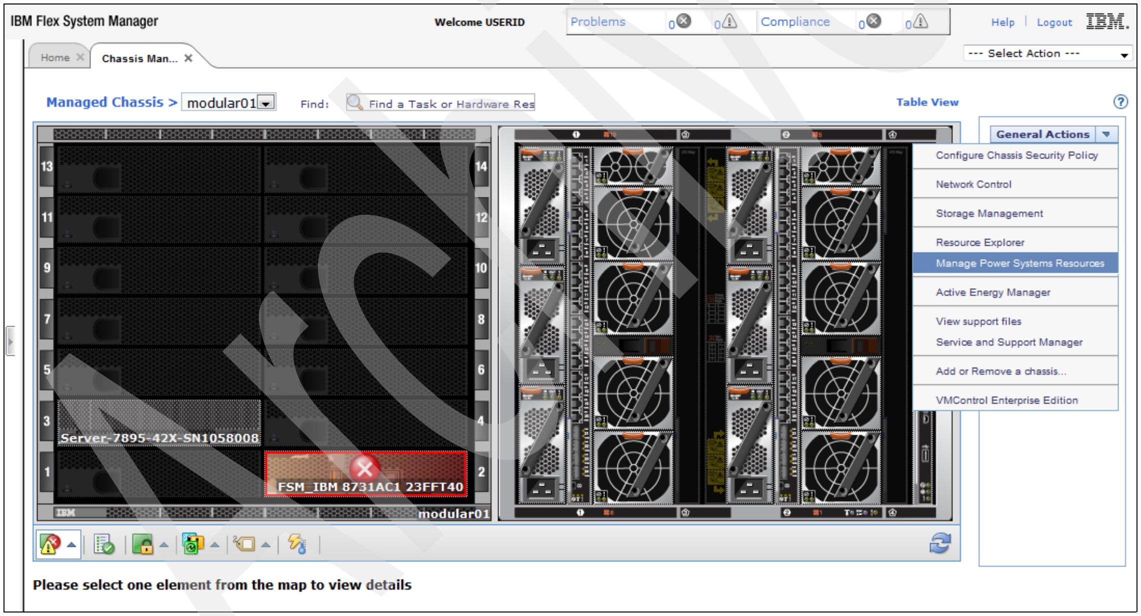

Figure 6-86 FSM Chassis Manager graphical view

3. Click the **General Actions** drop-down menu and click **Manage Power Systems Resources**.

A new tab is created along the top edge of the GUI, and the Manage Power Systems Resources window opens.

Important: Readers who are familiar with the Systems Director Management Console will recognize this part of the IBM Flex System Manager GUI, as it is nearly identical in form and function for both applications.

Requesting access to the Flexible Service Processor

Typically, a Power Systems compute node is automatically discovered, but access must be requested to the Flexible Service Processor on these nodes. The following example shows a discovered node in a No Access state (Figure 6-87).

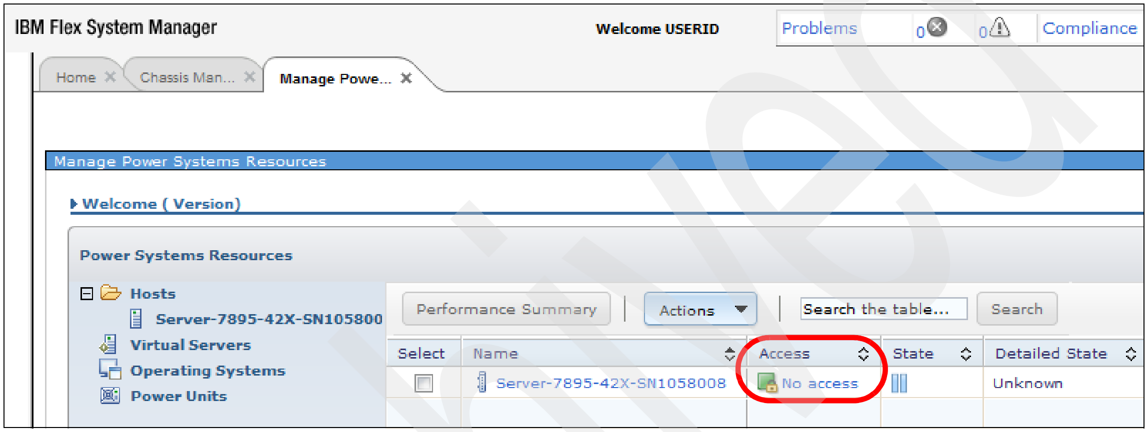

Figure 6-87 FSM Power Systems resources

To request access, complete the following steps:

1. Right-click the wanted server object, as shown in Figure 6-88, and click **Request Access**.

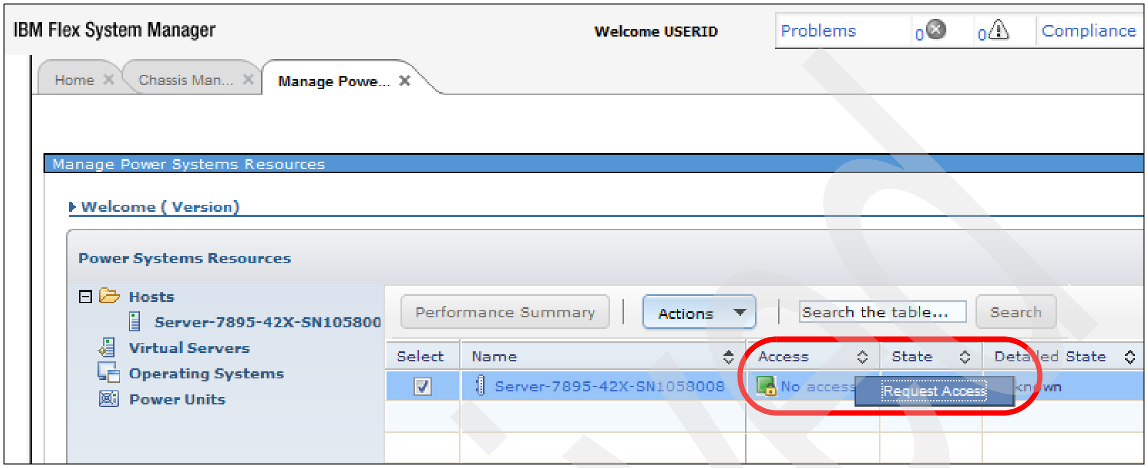

Figure 6-88 Requesting access to a Power Systems compute node

Figure 6-89 shows the next window, which steps you through the process. Notice that the User ID box is prepopulated with the Hardware Management Console (HMC) ID and is disabled. The Password box accepts any password for the HMC user ID and essentially sets the password with this first use.

> **Important:** Remember this password set for initial access, as it is needed if access to the node is requested again.

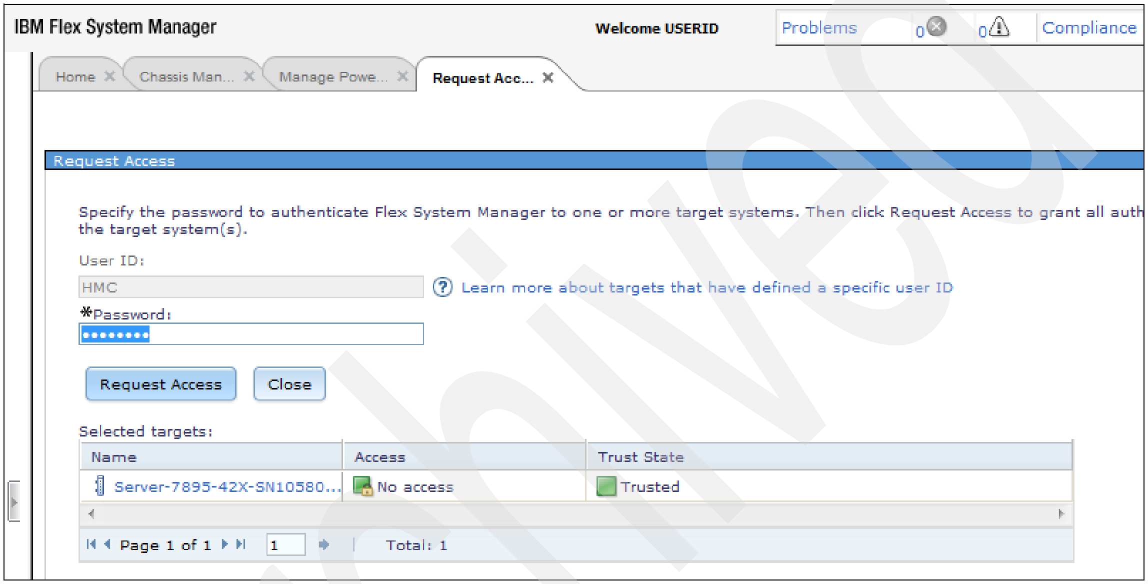

Figure 6-89 Initial password set to Flexible Service Processor

2. After the password is entered, click **Request Access**.

   Figure 6-90 and Figure 6-91 on page 241 show the process steps and the successful completion of the access requested, as indicated by the **OK** in the Access column.

Figure 6-90 Access request in process

Figure 6-91 Completed access request

3. With the access request complete, click **Close** to exit the window and return to the Manage Power Systems Resources window, as shown in Figure 6-92.

   Many of the columns now contain information obtained from this limited communication with the Flexible Service Processor.

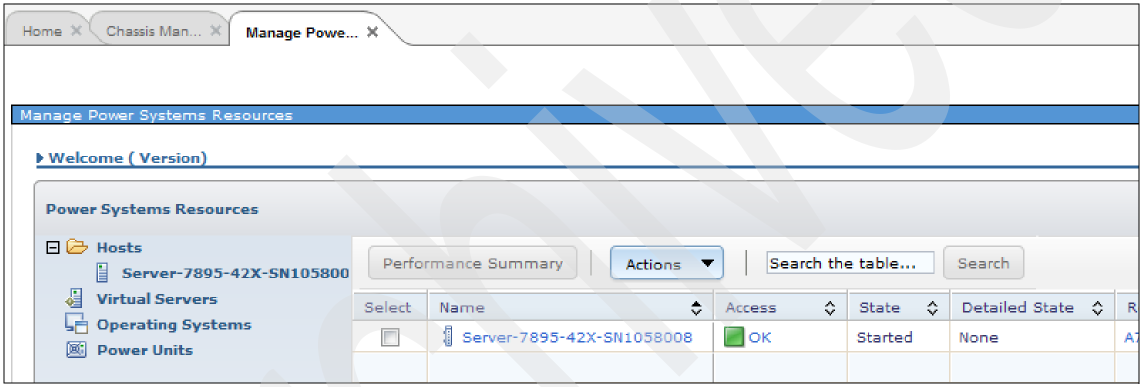

Figure 6-92 Updated Power Systems resources - now with access

Inventory collection

In order for the FSM to accurately manage a Power Systems compute node, inventory information must be collected. To accomplish this task, perform the following steps:

1. Right-click the server object in the list, as shown in Figure 6-93.

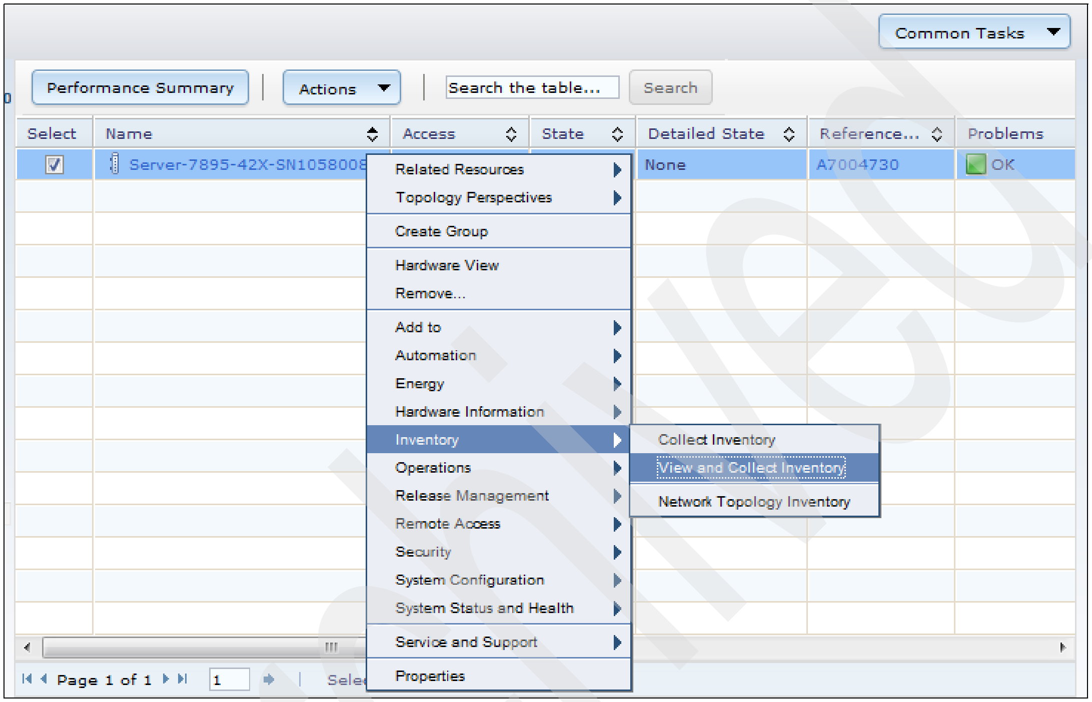

Figure 6-93 Starting inventory request of Power Systems compute node

2. Click **Inventory/View and Collect Inventory** to start the collection.

   In Figure 6-94, notice that, to the right of the Collect Inventory button, a time stamp of the last collection is displayed. In this case, inventory has never been collected for this node.

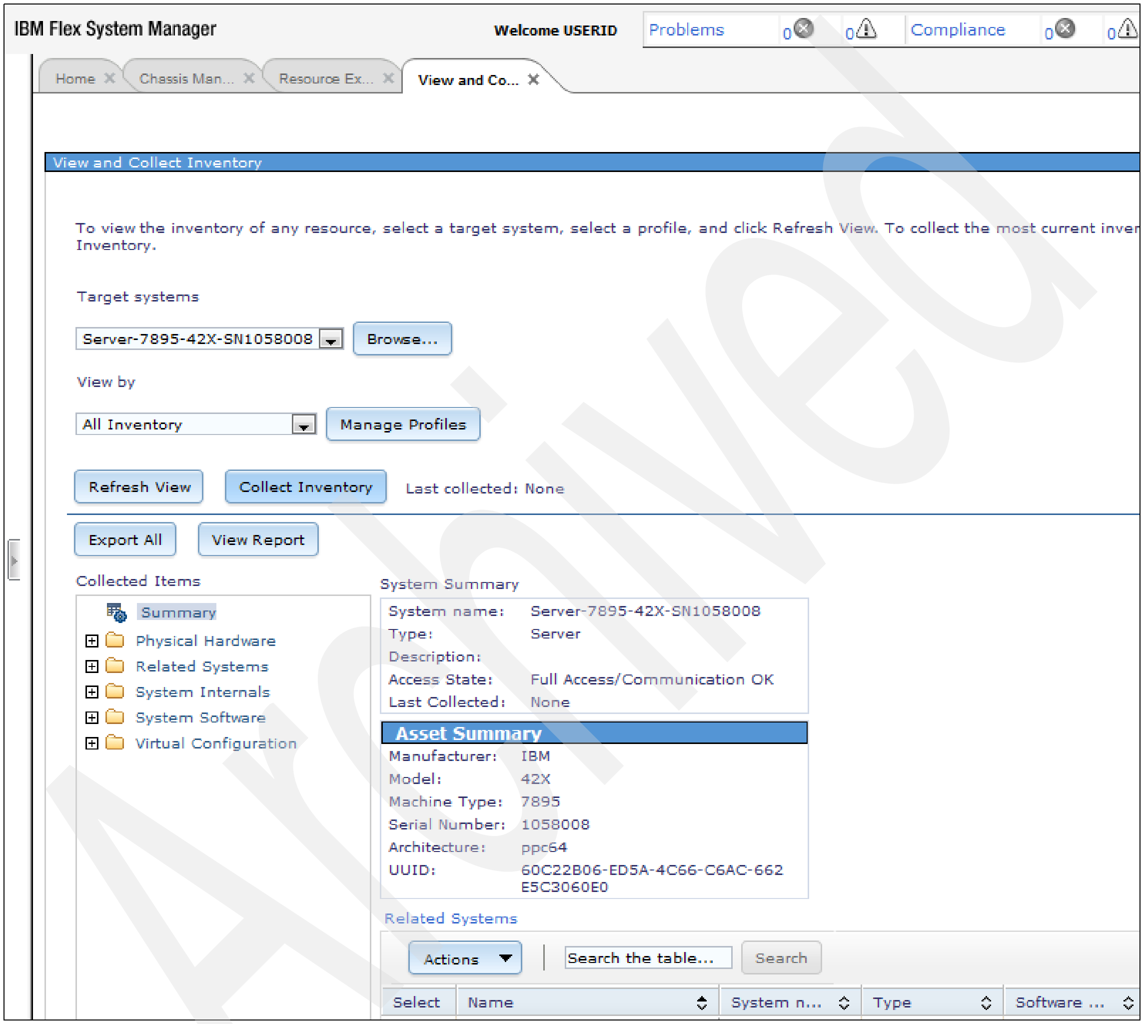

Figure 6-94   Starting inventory collection

3. Click **Collect Inventory** to start the process.

   Nearly all processes in the IBM Flex System Manager application are run as jobs and can be scheduled. The scheduling can be immediate or in the future.

Figure 6-95 shows the job scheduler window that opens when the inventory collection process is started.

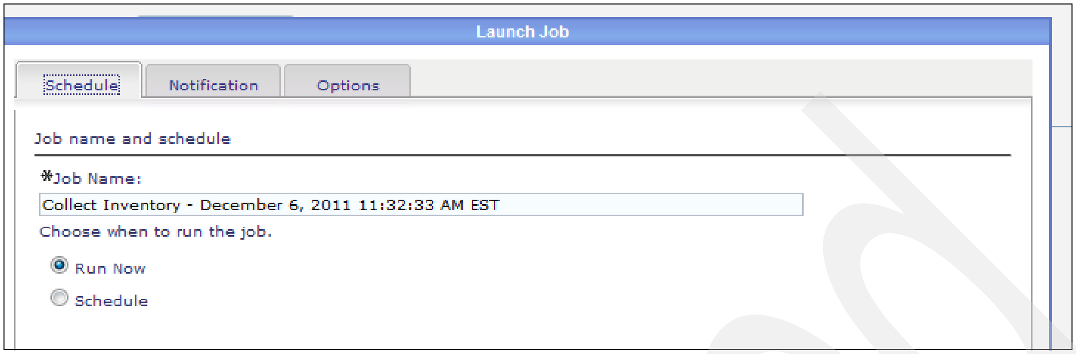

Figure 6-95 Scheduling inventory collection job

4. Select **Run Now** and click **OK** at the bottom of the window.

   When the job starts, a notification is sent to the originating window with options to **Display Properties** or **Close Message** (Figure 6-96).

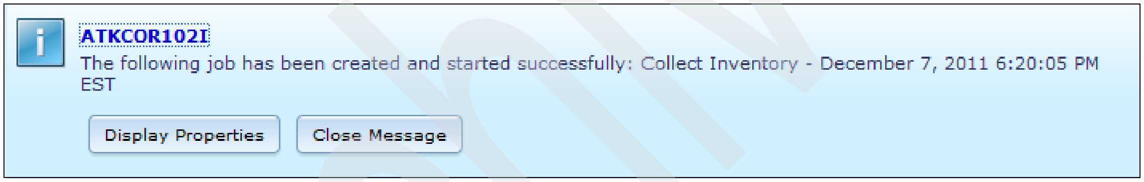

Figure 6-96 Inventory job start notification

Clicking **Display Properties** opens the window shown in Figure 6-97. The job properties window has several tabs that can be used to review additional job details. The General tab shown indicates that the inventory collection job completed without errors.

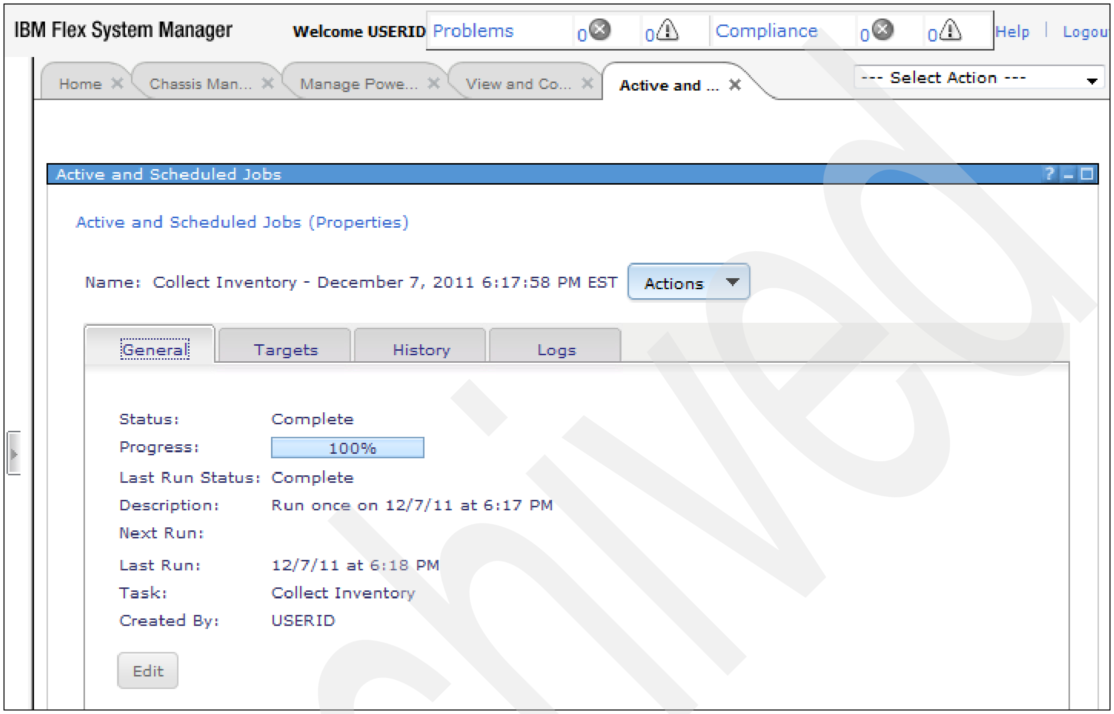

Figure 6-97 Inventory job status

The Active and Scheduled Jobs tab and the View and Collect Inventory tabs near the top of the window can be closed.

Now that you have completed access and inventory collection, you can now use IBM Flex System Manager to manage the node.

6.6.2 Opening a console

Occasionally, it is necessary to open a console session on a virtual server, for example, if there are networking problems that prevent access through the Internet. Each virtual server has only one console that can be open at a time. You can access each console through the IBM Flex System Manager.

To open a virtual console, complete the following steps:

1. Open one of the windows that lists the virtual servers.

   You can accomplish this task in many ways. In the example that follows, we use Resource Explorer. Figure 6-98 shows how to open the console.

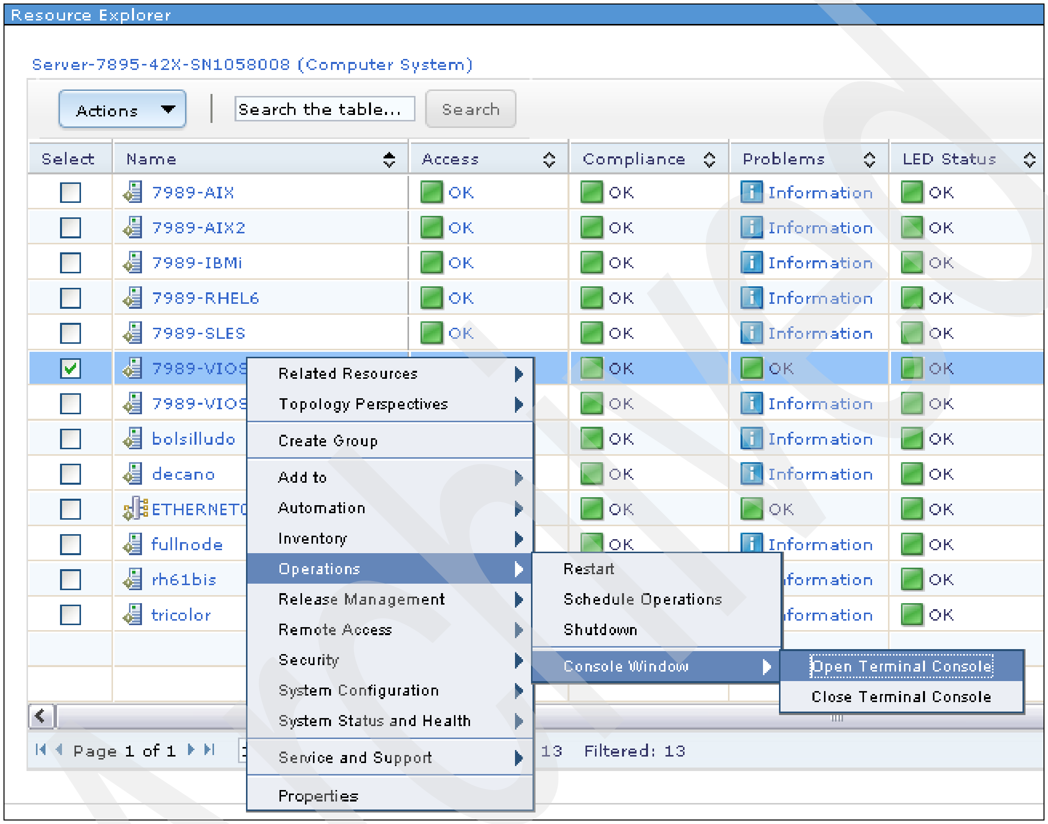

Figure 6-98 Open a console on a virtual server from the FSM

2. Enter the password of the login ID used to access the FSM.
3. Enter the password to open the console.

4. The Terminal Console tab opens and shows a message and an **OK** button. Click **OK** to return to the Resource Explorer tab (or the tab you started the console from) (Figure 6-99).

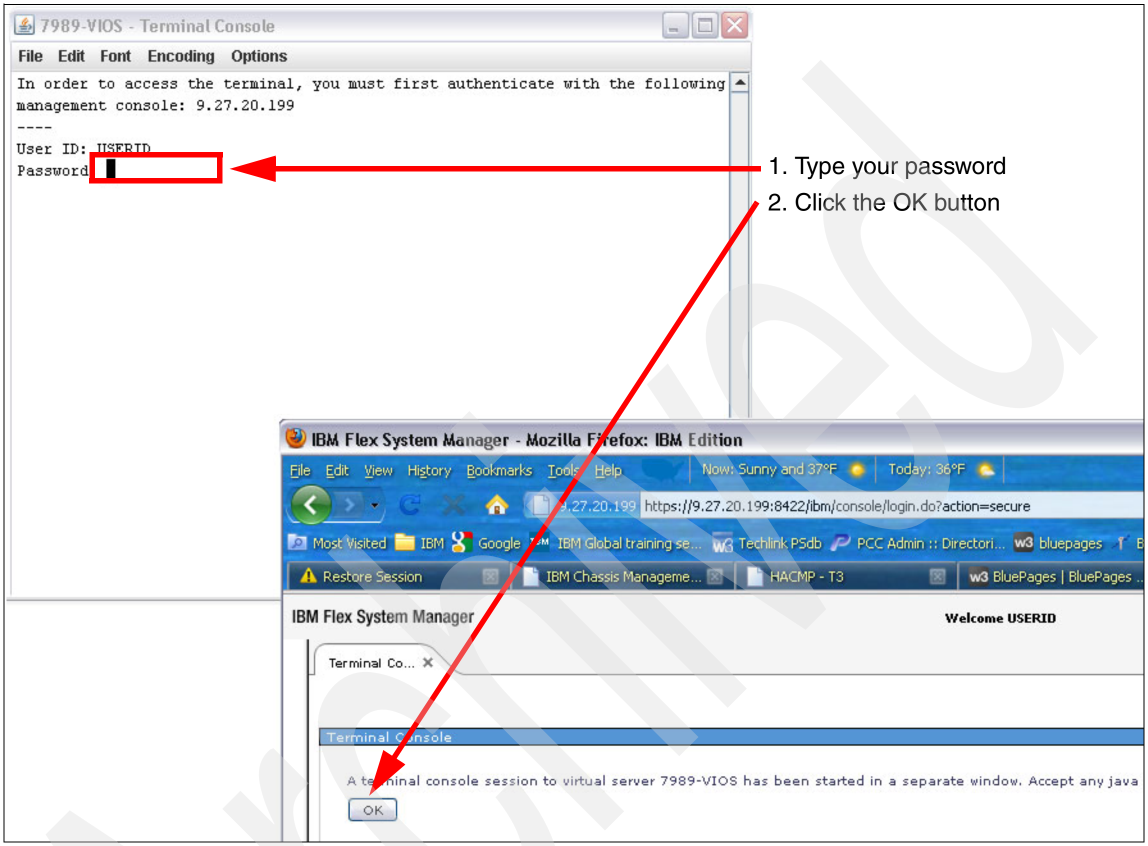

Figure 6-99 Validating with the FSM

If the virtual server ID for the console you are launching is the number 1, the console opens as shown in Figure 6-100.

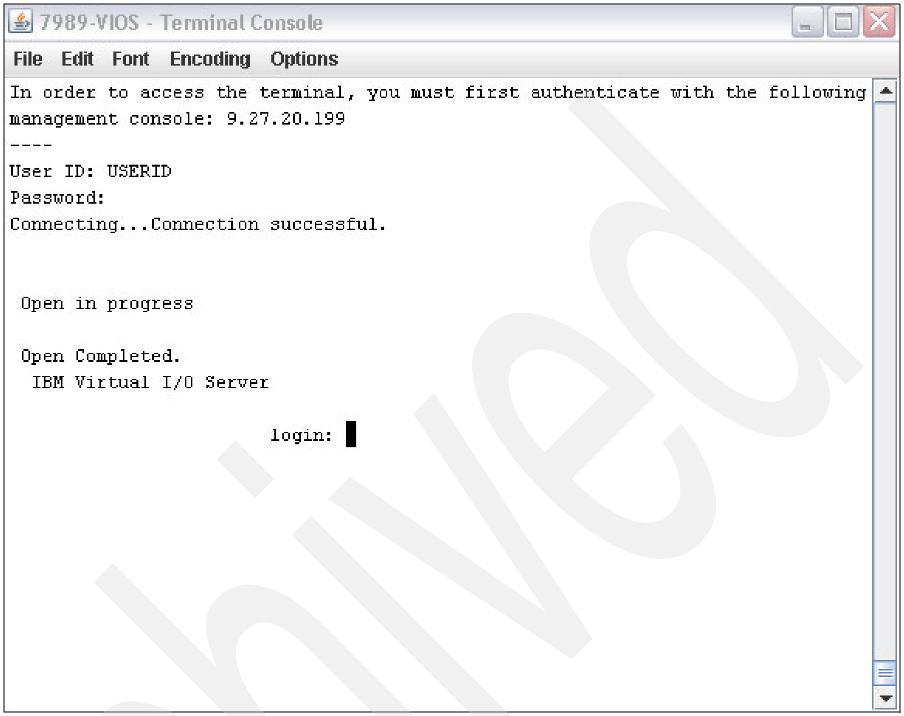

Figure 6-100 Successful launch of virtual server ID 1

If Serial Over LAN (SOL) is not disabled, you receive the error shown in Figure 6-101. To learn the process to disable SOL, see 6.6.3, “Disabling Serial Over LAN (SOL)” on page 249.

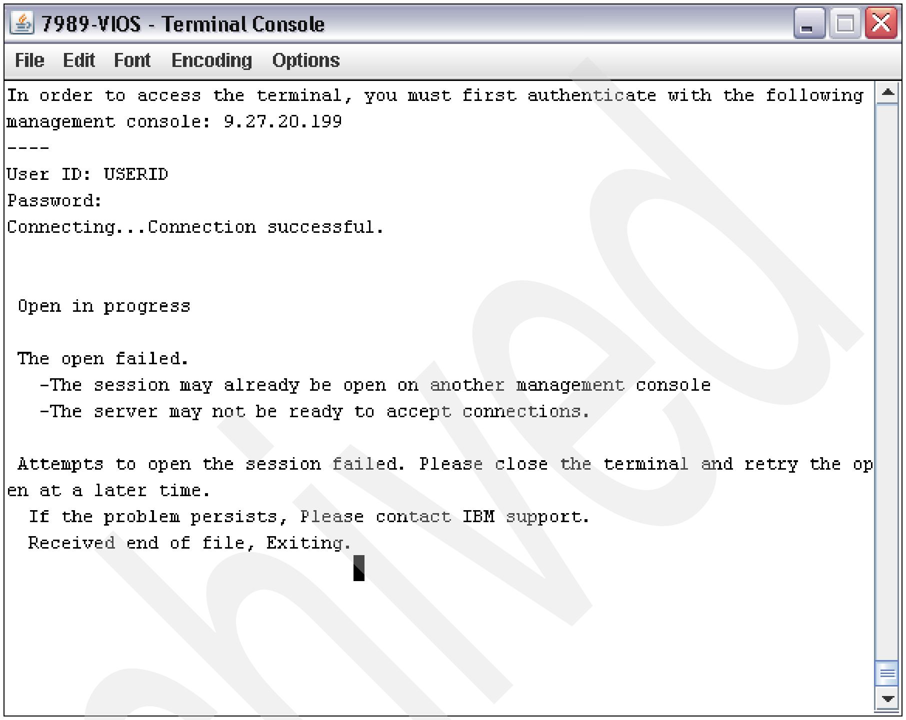

Figure 6-101 Console open failure on virtual server ID 1 with SOL enabled

6.6.3 Disabling Serial Over LAN (SOL)

When a Power Systems compute node is managed by an IBM Flex System Manager, you must disable SOL on the chassis.

> **Important:** There is an option to disable SOL at the individual compute node level.

Accessing the Chassis Management Module using a browser

To access the CMM through a web browser, complete the following steps:

1. Open a browser and point it to the following URL, where *system_name* is the host name or IP address of the Chassis Management Module:

   `https://`*`system_name`*

The Login window opens (Figure 6-102).

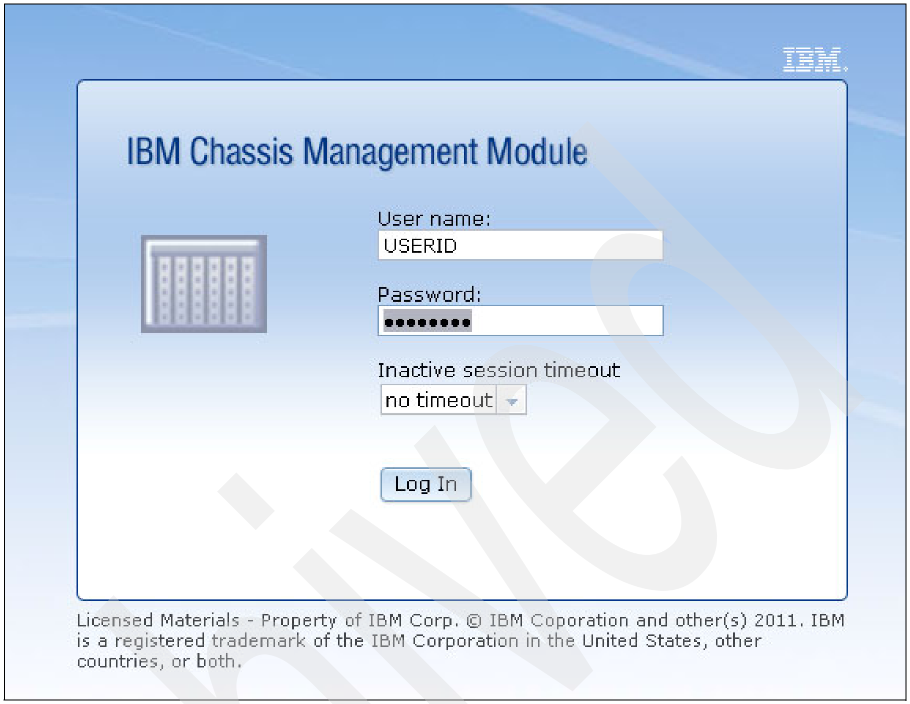

Figure 6-102 Chassis Management Module Login window

2. Log in using a valid user and password. If this is the first time you are logging in to the Chassis Management Module, the System Status window of the Chassis Management Module opens the Chassis tab. If this is not the first time you are logging in, you are returned to the place where you were when you logged off.
3. Click **Chassis Management** → **Compute Nodes** from the menu bar in the CMM interface.

   The Compute Nodes window opens and shows all the compute nodes in the chassis. In our chassis, we had two compute nodes: a Power Systems compute node and the IBM Flex System Manager.

Disabling SOL

To disable SOL on the chassis, complete the following steps, which are also shown in Figure 6-103:

1. Click the **Settings** tab.
2. Click the **Serial Over LAN** tab.
3. Clear the **Serial Over LAN** check box.
4. Click **OK**.

The change takes effect as soon as the window closes.

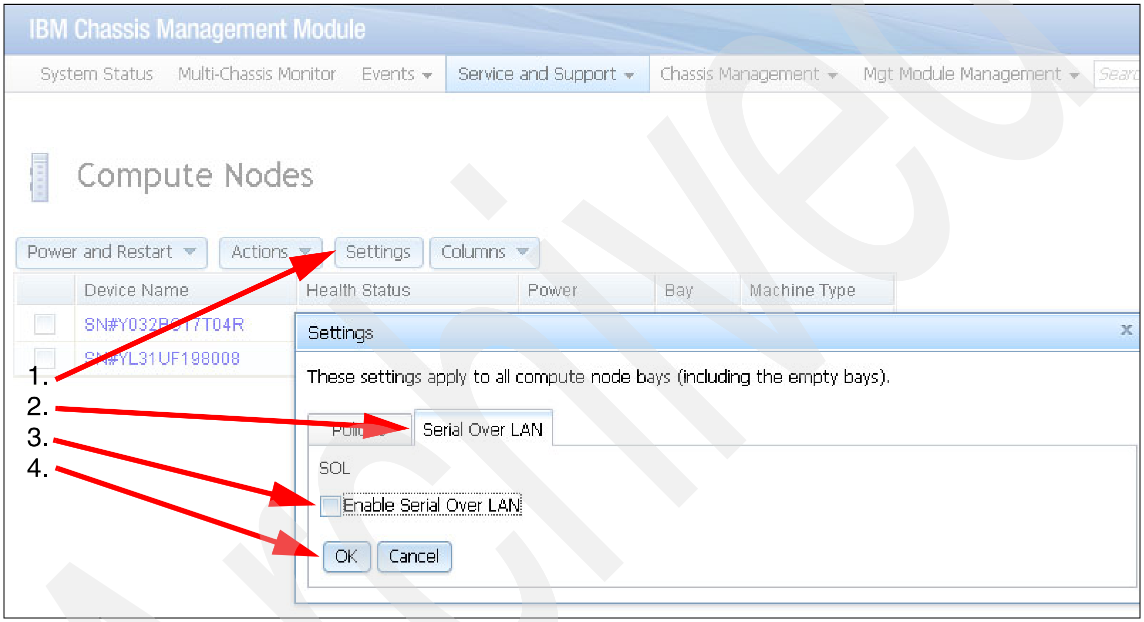

Figure 6-103 Disable SOL for all compute nodes from the Chassis Management Module

6.7 IBM Flex System Manager options and tasks

In this section, we describe a subset of FSM options that you can use to manage your chassis, and the options available in FSM.

6.7.1 Initial setup tab

After logging in, the Home page opens and shows the Initial setup tab (which is selected). This window has options for managing your environment, as shown in Figure 6-104.

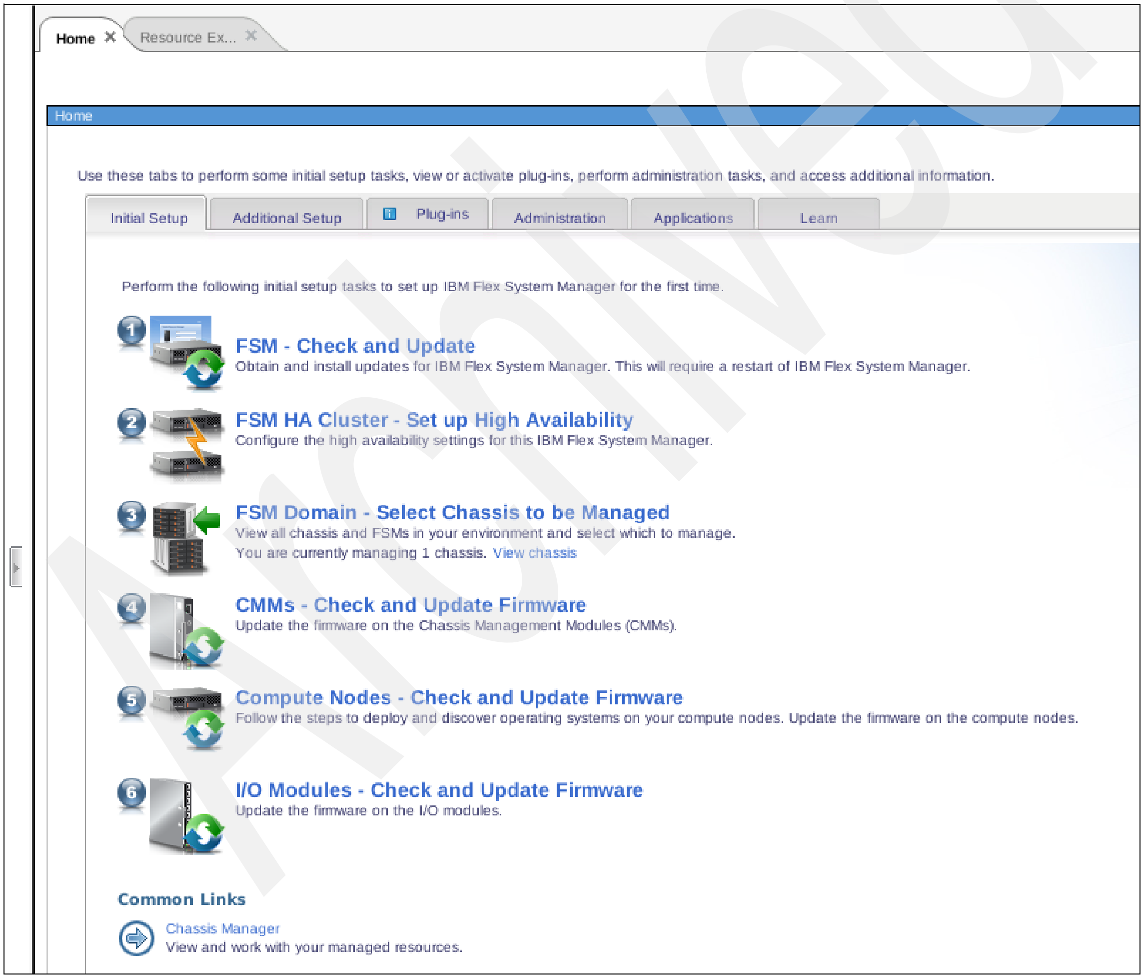

Figure 6-104 Flex System Manager management window

This window provides access to the functions listed in the following sections.

FSM - Check and update

With this option, you can upgrade your IBM Flex System Manager code. If the FSM has access to the Internet, the upgrade codes can be downloaded directly and installed. If you do not have access to the Internet in the FSM, you can download the code package using another system, then upload it manually to the FSM.

FSM HA Cluster - Set up High Availability

Set up high availability between two FSM management nodes.

FSM Domain - Select Chassis to be Managed

You can use this option to manage a chassis. The FSM can manage more chassis than the one where your IBM Flex System Manager is installed. Figure 6-105 shows the chassis management window. From this window, you can discover new chassis and manage discovered chassis.

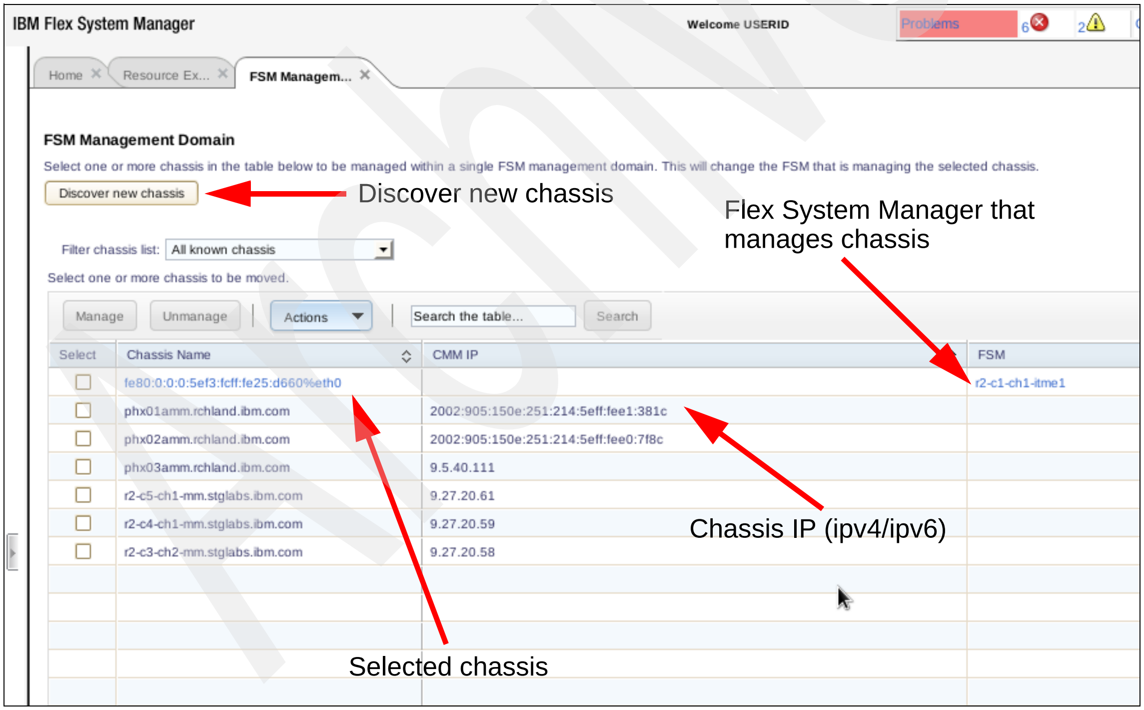

Figure 6-105 Chassis Management

After you select the chassis (as shown in Figure 6-105 on page 253), you can start managing it, and a window similar to Figure 6-106 opens. You see a front and back graphical view of the chassis. Click a component, and click **Action** to select actions applicable to that component (see the area marked “Selected component properties” in Figure 6-106). Actions include restart, power off, access to a command line or console, and so on.

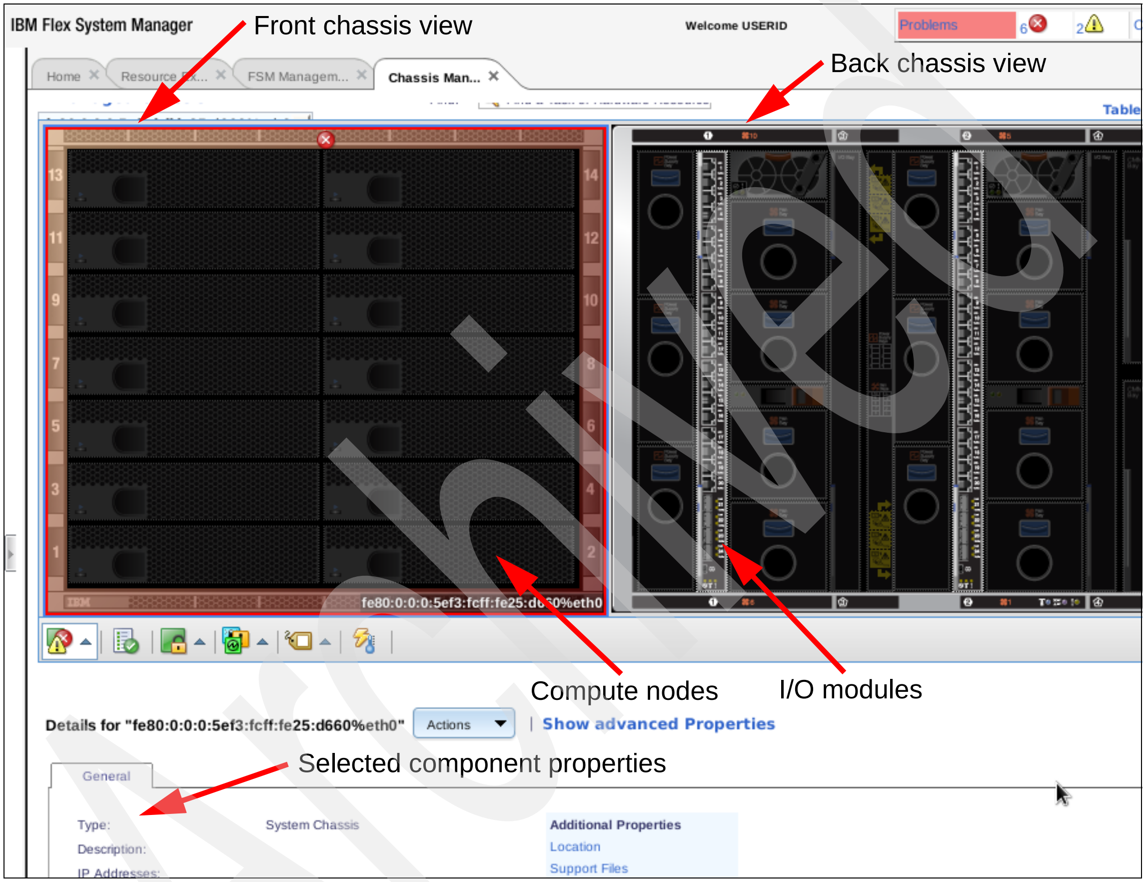

Figure 6-106 Selected Chassis Management

Chassis Management Modules - Check and Update Firmware

If your IBM Flex System Manager is connected to the Internet, you can update your firmware directly from the Internet. If your IBM Flex System Manager is not directly connected to the Internet, you can download the firmware manually from another system, then upload it manually to your IBM Flex System Manager.

Compute Nodes - Check and Upgrade Firmware

After your compute node is discovered, you have several actions that you can take, as shown on Figure 6-107.

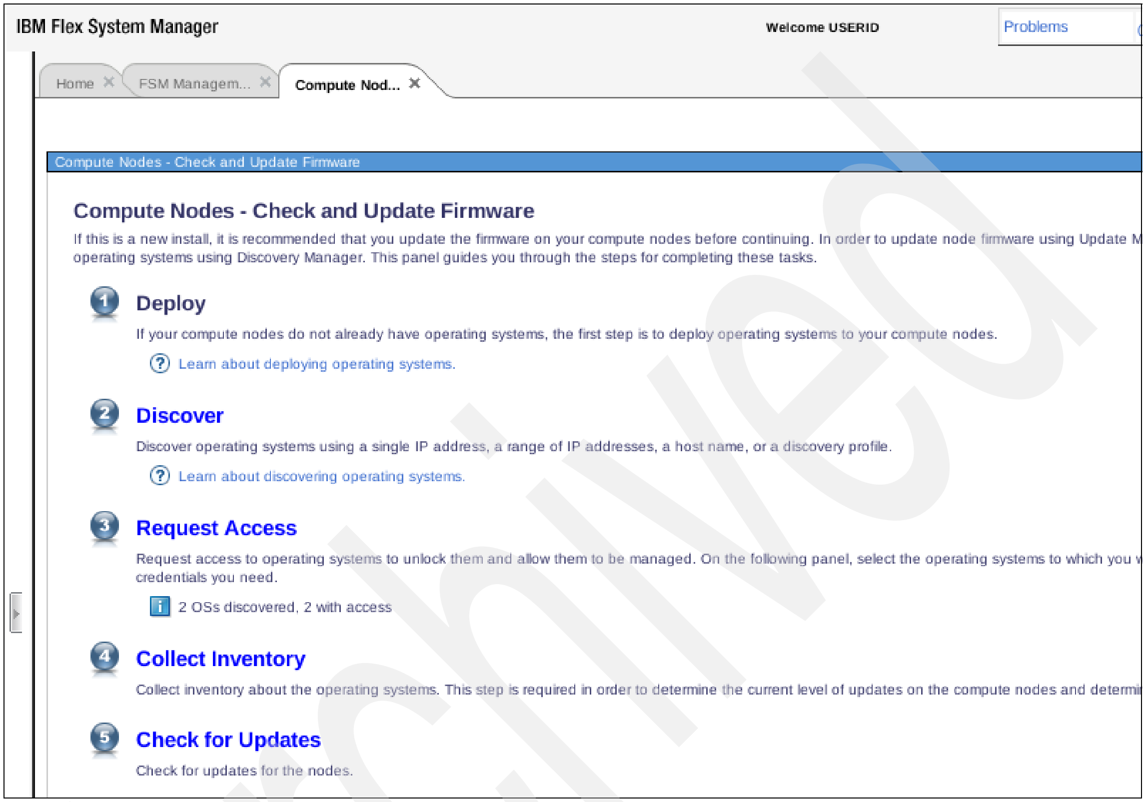

Figure 6-107 Compute Nodes - Management

You can choose the following actions:

- Deploy: Deploy an operating system.
- Discover: Discover operating systems, components, and I/O modules. After systems are discovered, you can request access and start managing them.

- Request Access: After your systems and components are discovered, you can request access to them with this option, as shown in Figure 6-108.

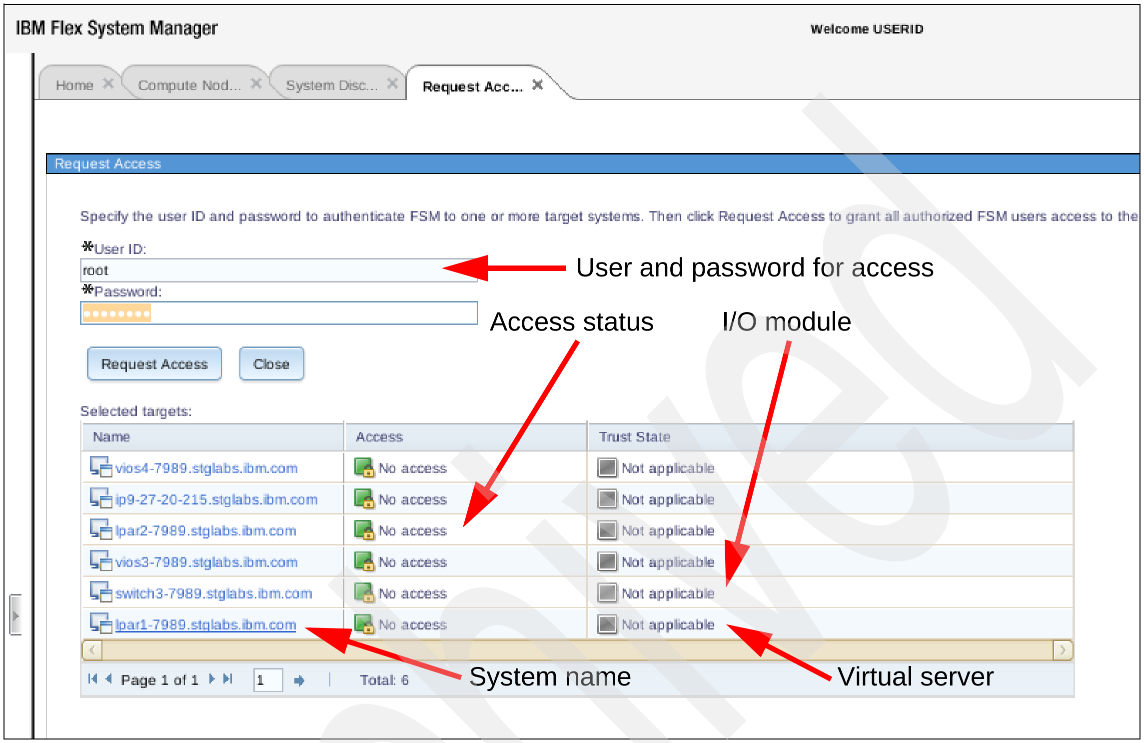

Figure 6-108 Request access to components

- Collect inventory: After you discover your system and request access, you can collect and review the systems inventory. The systems inventory shows you information about hardware and operating systems for the systems you select. There are several filter and export options available, as shown in Figure 6-109.

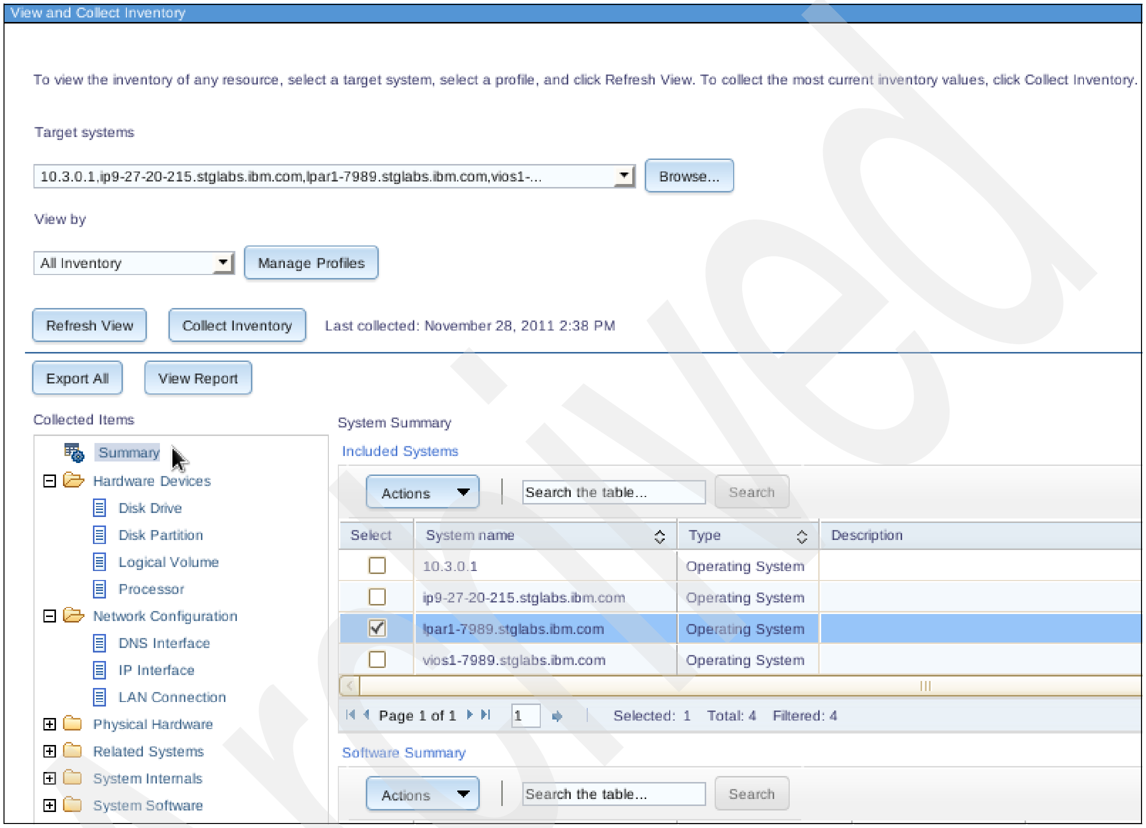

Figure 6-109 Collect inventory management

- Check for Updates: If your FSM is connected to the Internet, you can update your firmware and operating system directly from the Internet. If the FSM is not connected to the Internet, you can download the firmware and operating system manually to another system, and then use that system to upgrade your system firmware and operating system.

I/O modules - Check and Upgrade Firmware

After your I/O modules are discovered, you can perform several operations on them:

- Request Access: In Figure 6-108 on page 256, you can request access to the discovered I/O modules.
- Collect Inventory: After you have access to your I/O modules, you can start collecting inventory on them, as shown in Figure 6-109 on page 257.
- Check for updates: If the FSM is connected to the Internet, you can update your I/O module firmware directly from the Internet. If the FSM is not directly connected to the Internet, you can download it manually to another system and upgrade your I/O module firmware from that system.

6.7.2 Additional setup tab

In this window, you have access to settings such as the IBM Electronic Service Agent™ (ESA) setup, LDAP setup, user setup and more, as shown in Figure 6-110.

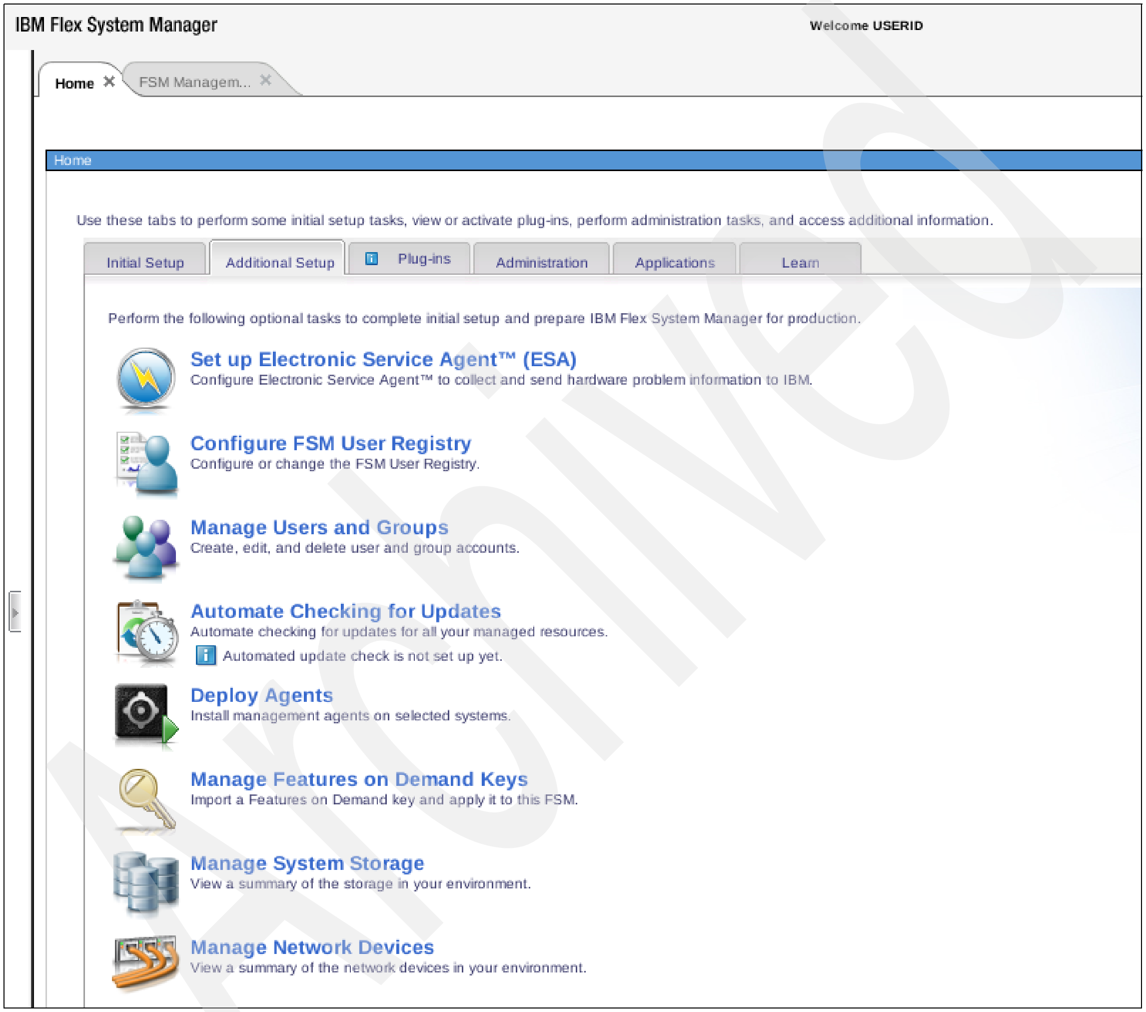

Figure 6-110 Additional Setup window

Set up Electronic Service Agent (ESA)

ESA is an IBM monitoring tool that reports hardware events to a support team automatically. You can use this setting to set up ESA on your IBM Flex System Manager.

Configure FSM User Registry

This setting connects the IBM Flex System Manager to an external LDAP server.

Manage Users and Groups

This setting opens the FSM access management area. From here, you can create, modify, and delete FSM accounts and groups.

Automatic Checking for Updates

Using this setting, your IBM Flex System Manager checks periodically for upgrades through the Internet and informs you when new upgrades are available.

Deploy Agents

When you use this setting, IBM Flex System Manager deploys monitor agents to monitor several items of your compute nodes. You can deploy agents to discovered systems with full access (discovery and collection can be started from this point, before the agent installation).

Manage Features on Demand Keys

IBM Flex System Manager supports additional features and add-ons. To enable them, you must add features keys through this setting.

Manage System Storage

As part of the new total management approach, storage management is integrated into the FSM, as shown in Figure 6-111. After you discover your storage appliance and request access to it through the FSM, you can start managing it.

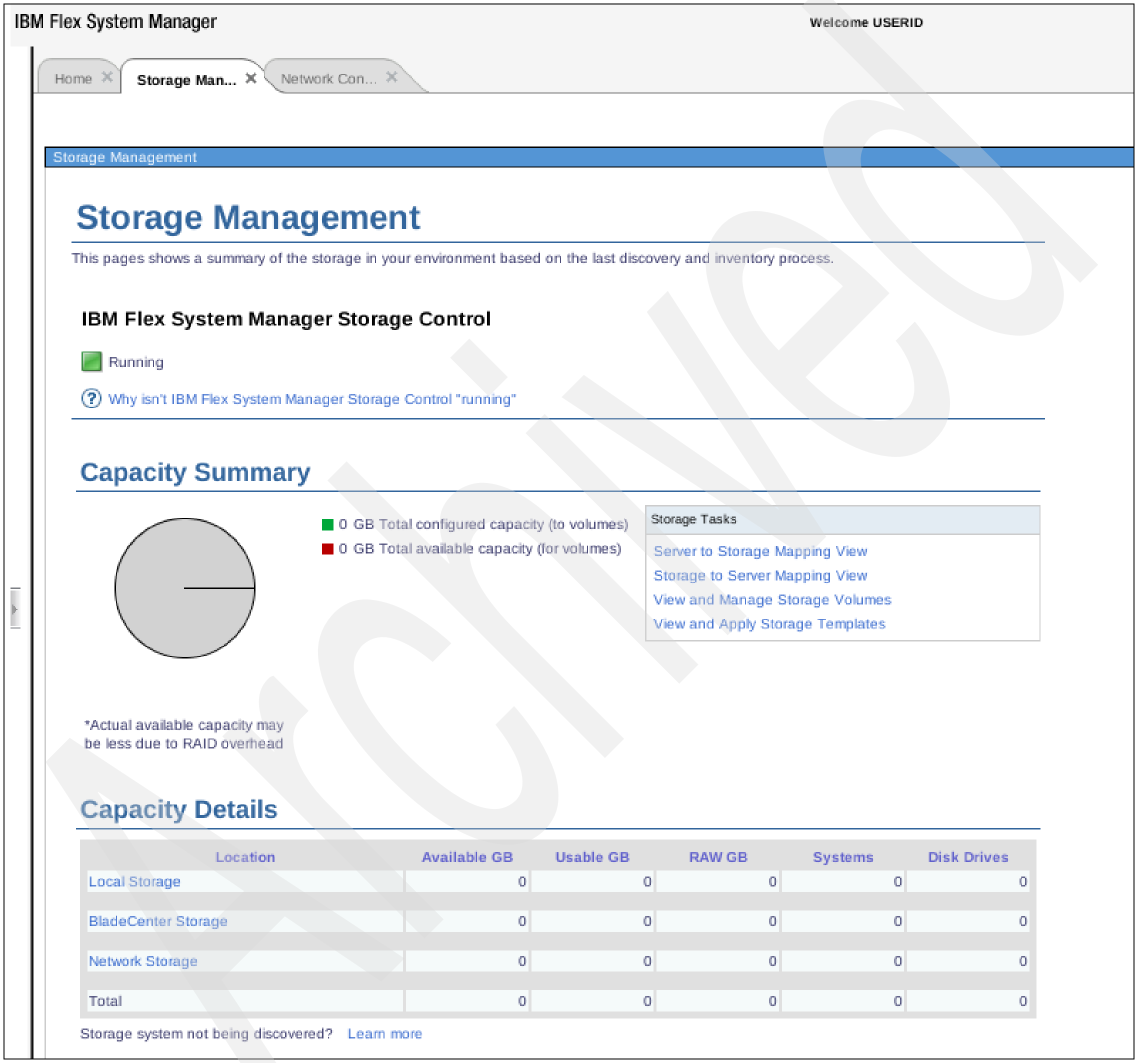

Figure 6-111 Flex System Manager Storage Management

Manage Network Devices

You can use IBM Flex System Manager to manage your network and network devices while the network devices are discovered and have full access. The Network Control window is shown in Figure 6-112.

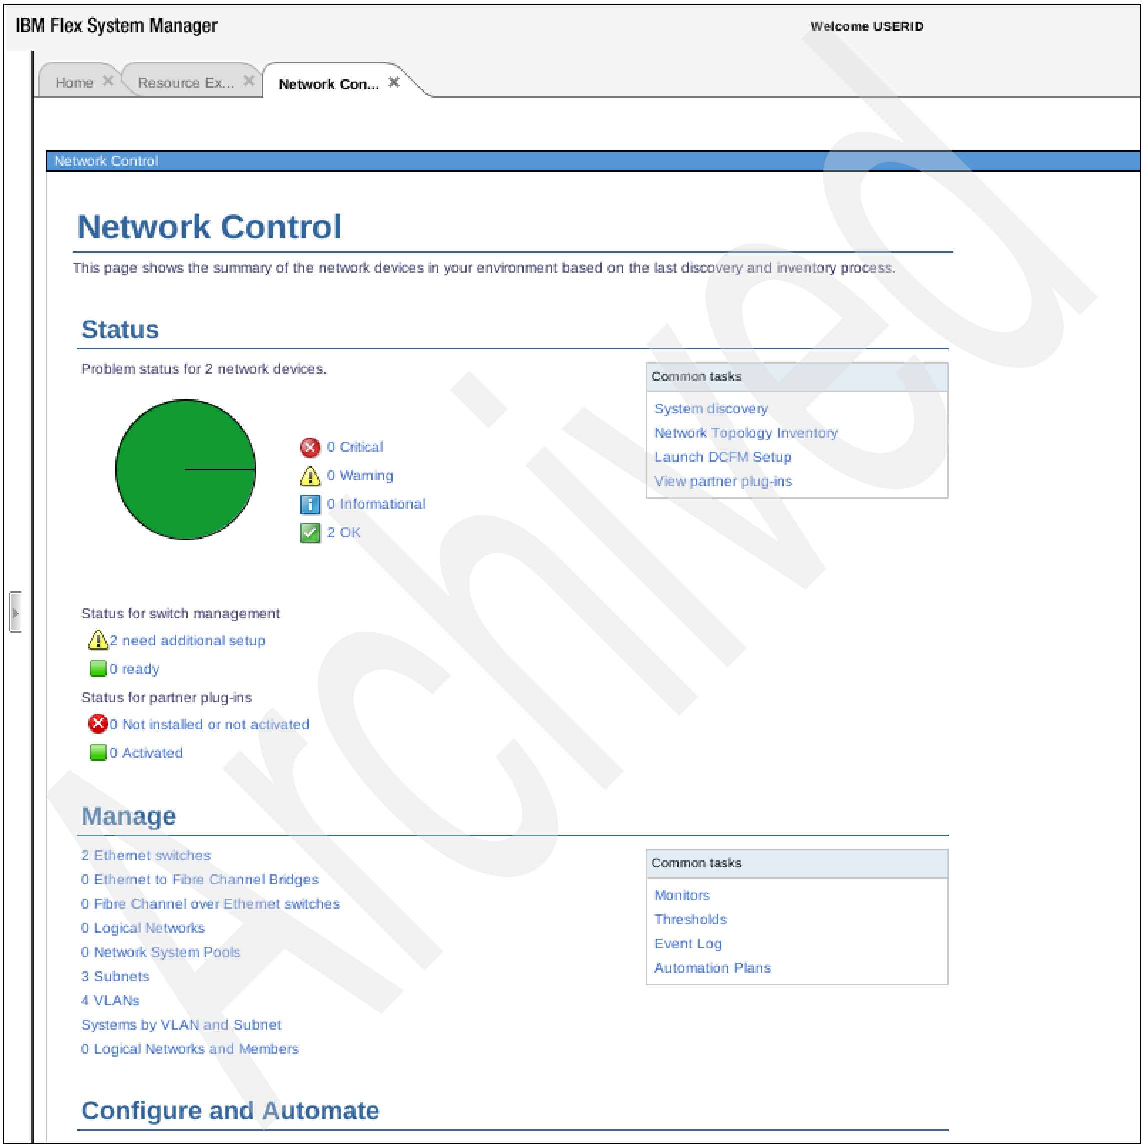

Figure 6-112 Manage Network Devices

6.7.3 Plug-ins tab

The plug-ins tab has options for managing the FSM, managing virtual servers, checking status, managing discovery, and more, as shown in Figure 6-113. Figure 6-113 shows only a portion of the available entries.

Several of the plug-ins require licensing and are included on a trial basis.

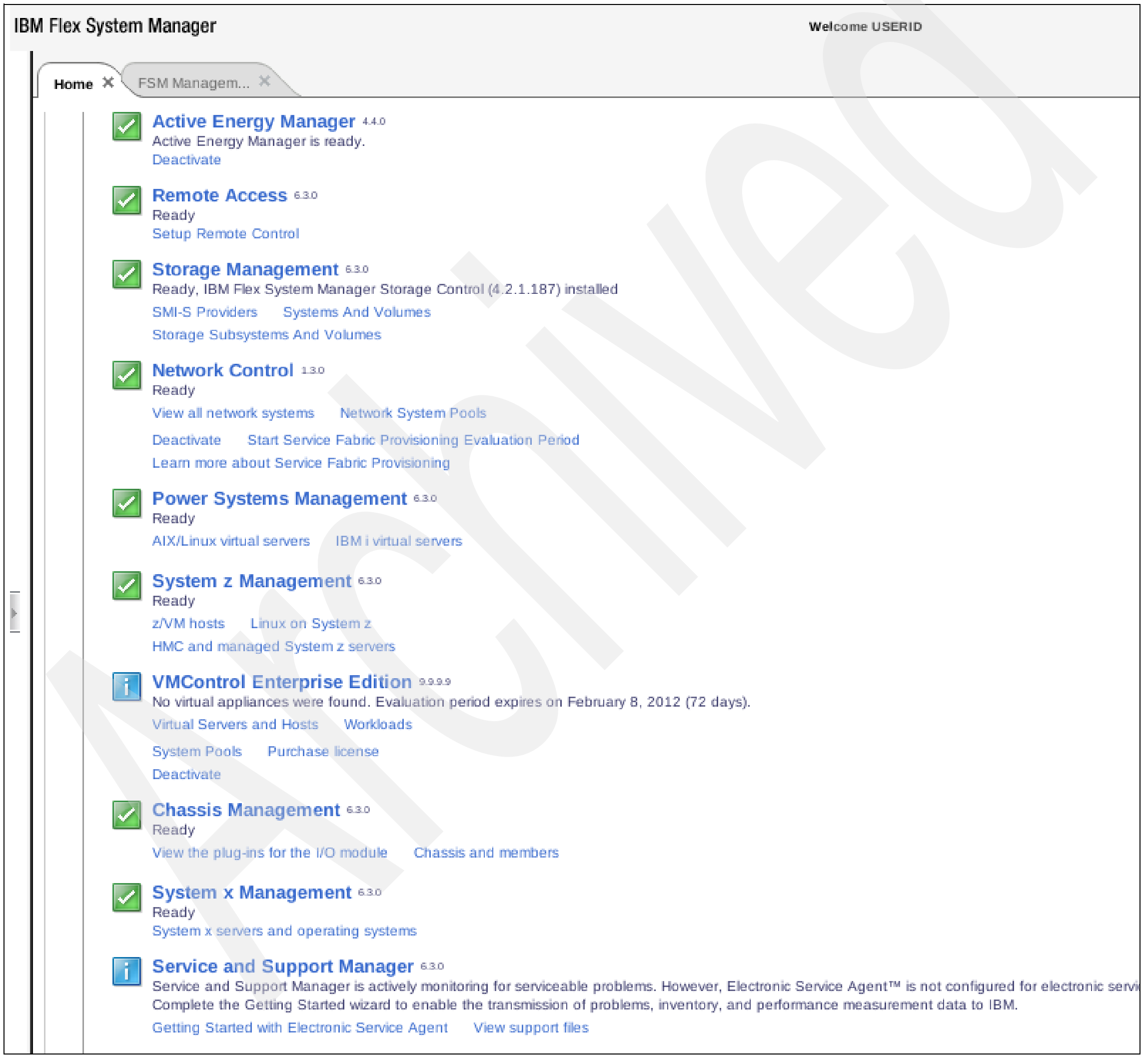

Figure 6-113 Plug-ins tab options

Flex Systems Manager

You can use this tab to monitor and manage the IBM Flex System Manager itself. This tab shows a graphic overview of all resources, indicating the state of selected resources (critical, warning, informational, or OK messages). Below the graphic is general information about the IBM Flex System Manager regarding uptime, processor use, last backup, active events, and so on.

You can also create shortcuts for functions you frequently use for IBM Flex System Manager, chassis management, managing Power System resources, the IBM Flex System Manager management domain, event log, backup and restore, and high availability settings (Figure 6-114).

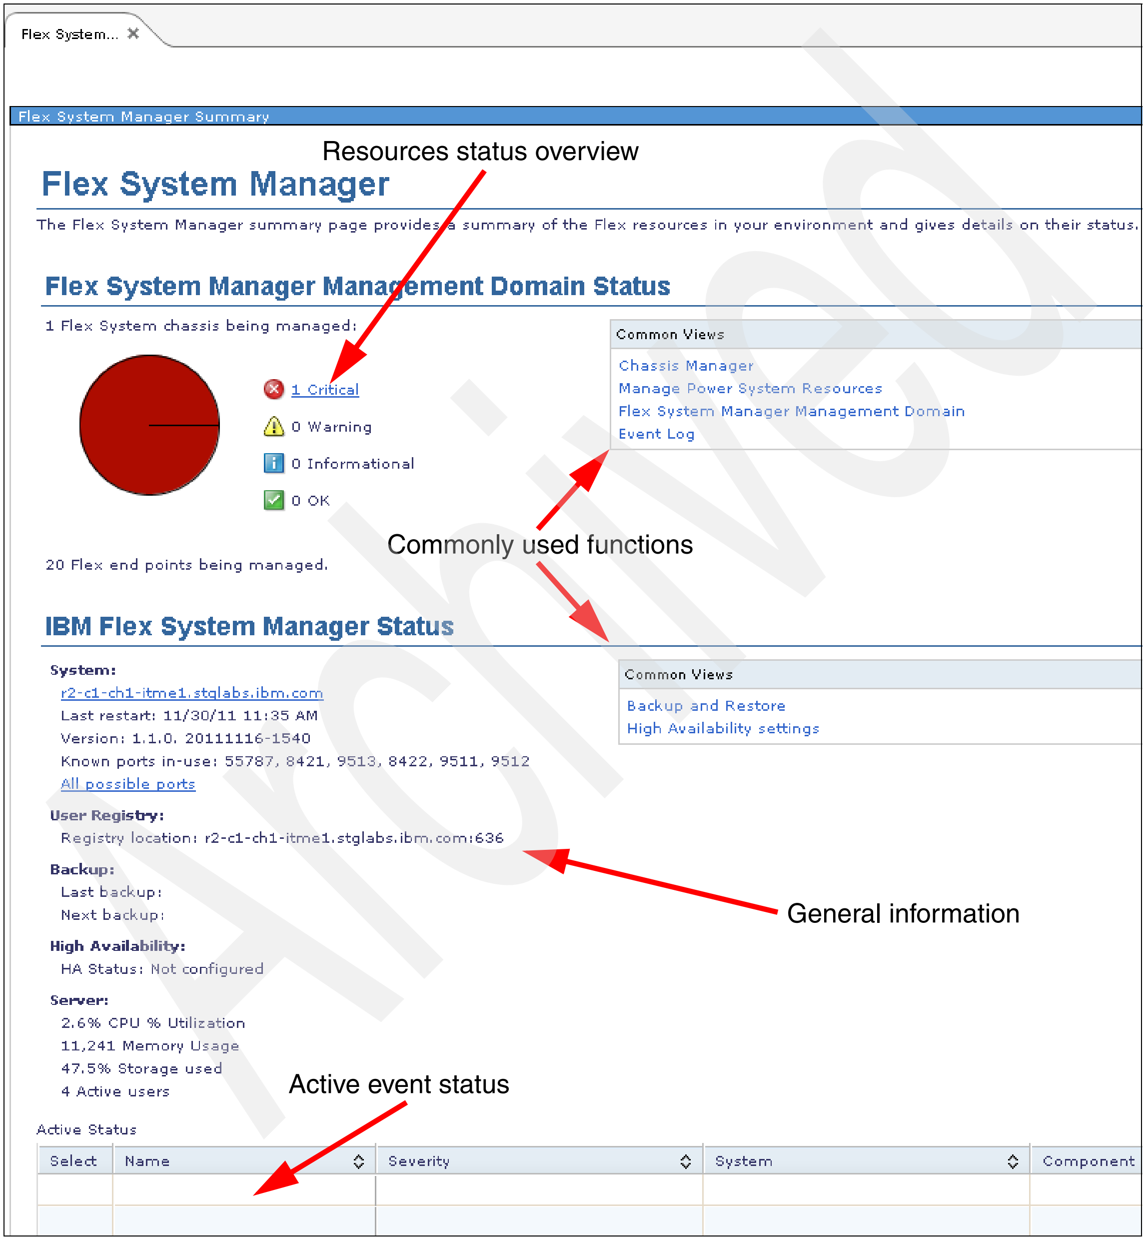

Figure 6-114 Flex System Manager - Management Overview

IBM Flex System Manager Server

This plug-in manages the server side of the IBM Flex System Manager. It shows information about systems discovered, processor use, ports used, and general user information. Shortcuts are available for common management tasks, such as System Discovery, Collect Inventory, Find a task, Find a resource, Resource Explorer, and User/Group management.

Discovery Manager

You can use Discovery Manager to discover and connect to the systems at your site. The Discovery Manager window shows an overview of all discovered systems, which ones you have access to, and which ones you have collected inventory from. It also has options to explore all discovered resources by category (Figure 6-115).

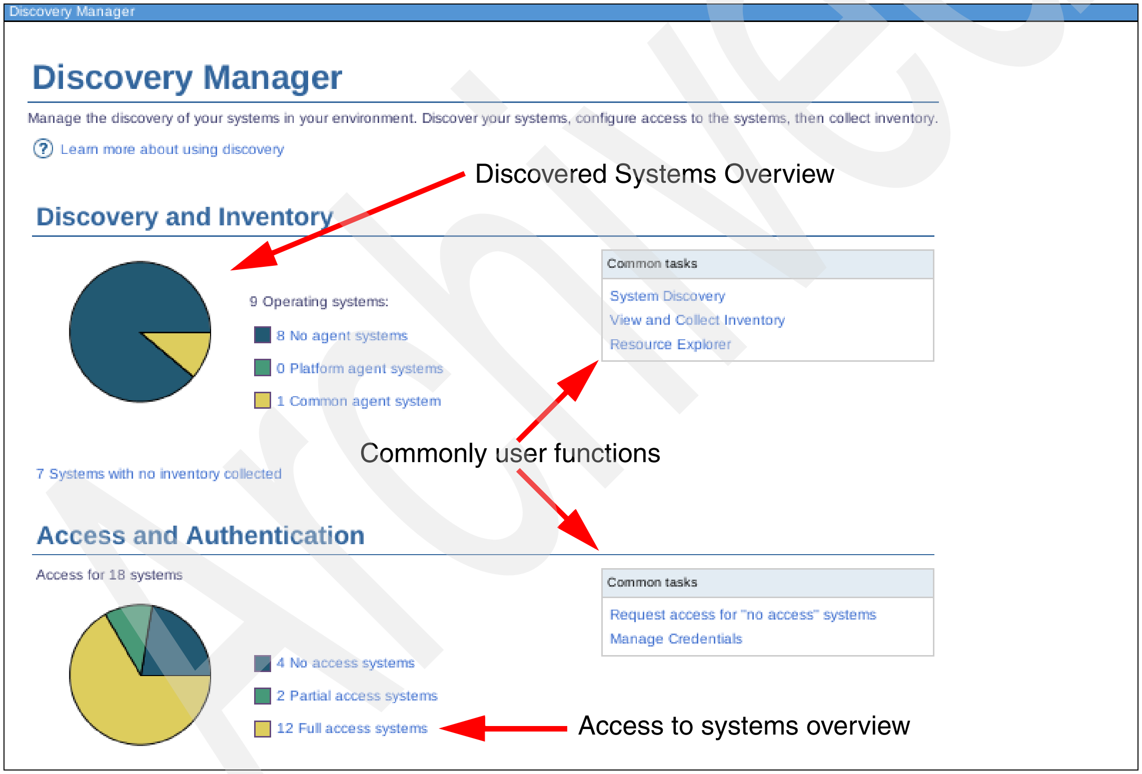

Figure 6-115 Discovery Manager plug-in

Status Manager

This window shows a tactical overview with a pie chart of all resources and systems managed by IBM Flex System Manager, dividing the chart into critical, warning, informational, and OK statuses. As with the other plug-ins, it has quick access menus for frequently used functions, for example, health summary, view problems, monitors, and event logs.

Update Manager

One of the main features of the IBM Flex System Manager is the ability to perform system upgrades and software upgrades on all systems and components that are managed by the IBM Flex System Manager. The Update Manager window is shown in Figure 6-116.

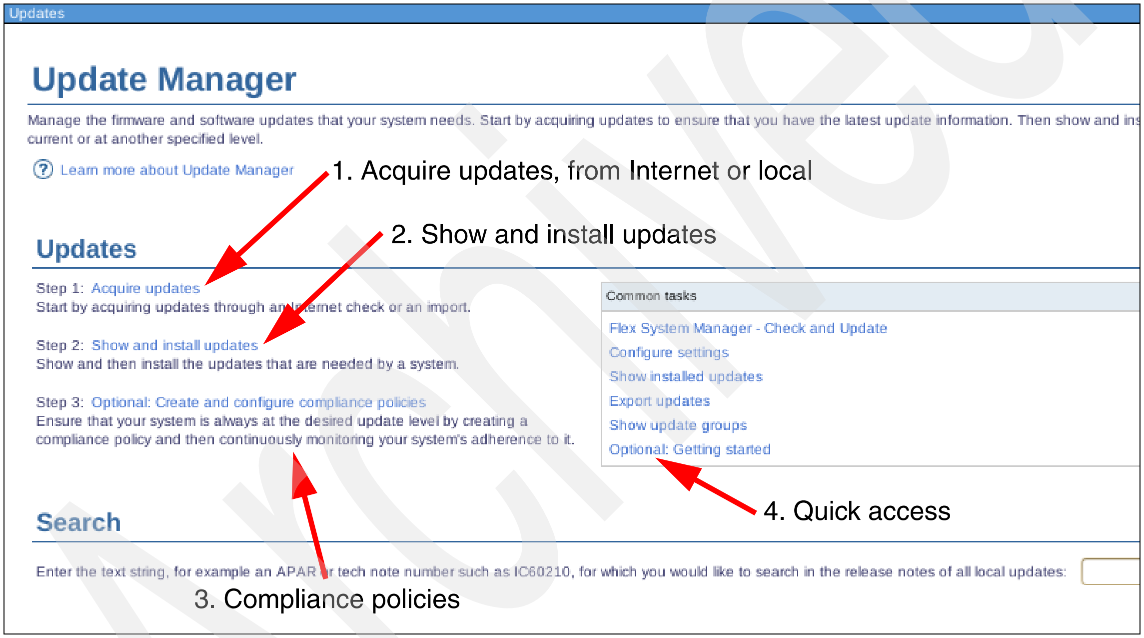

Figure 6-116 Update Manager window

The available features are:

- Acquire updates: If your machine is connected to the Internet, your IBM Flex System Manager can automatically download firmware and software upgrades for its managed systems. If you do not have Internet access, you can download them manually to another system, and upload them to the IBM Flex System Manager using a USB stick, a DVD, TFTP, and so on.

- Show and install updates: After obtaining the updates (using the Internet or manual download), you can view and install them using this setting. A firmware upgrade example for a Power Systems compute node, including captures, is shown in 8.1.5, “Firmware update using IBM Flex System Manager” on page 328.
- Compliance Policies: Using this setting, you can define compliance policies for specific systems to ensure that they are consistently at the wanted update level, and then monitor the adherence of each system to the policy.
- Quick Access: You can use this setting to access shortcuts to the most frequently used firmware actions, such as configure settings, Flex Systems Manager Check and Update, and Show Installed Updates.

Automation Manager

This feature shows an overview of active and scheduled jobs for your IBM Flex System Manager, including information about completed, failed, and upcoming jobs.

Active Energy Manager

This feature manages the energy settings of your IBM Flex System Manager managed systems. Detailed information about this feature is in *Implementing IBM Systems Director Active Energy Manager 4.1.1*, SG24-7780.

Remote Access

You can use this feature to access managed systems using CLI, web access, remote access, or remote control. Access depends on the type of system, for example, I/O modules might have web and CLI access, but not remote. Computer nodes might have CLI, remote control, and web access.

Storage management

For more information about this feature, see “Manage System Storage” on page 261.

Network Control with Service Fabric Provisioning

For more information about this feature, see “Manage Network Devices” on page 262.

Power Systems Management

You can use this feature to assume the role of the Hardware Management Consoles and Systems Director Management Consoles to manage the Power Systems servers in your data center. From here you can create partitions, manage virtual servers, set up dual VIOS in an IBM Flex System environment, and access features, such as live partition mobility. The Power System management overview is shown in Figure 6-117.

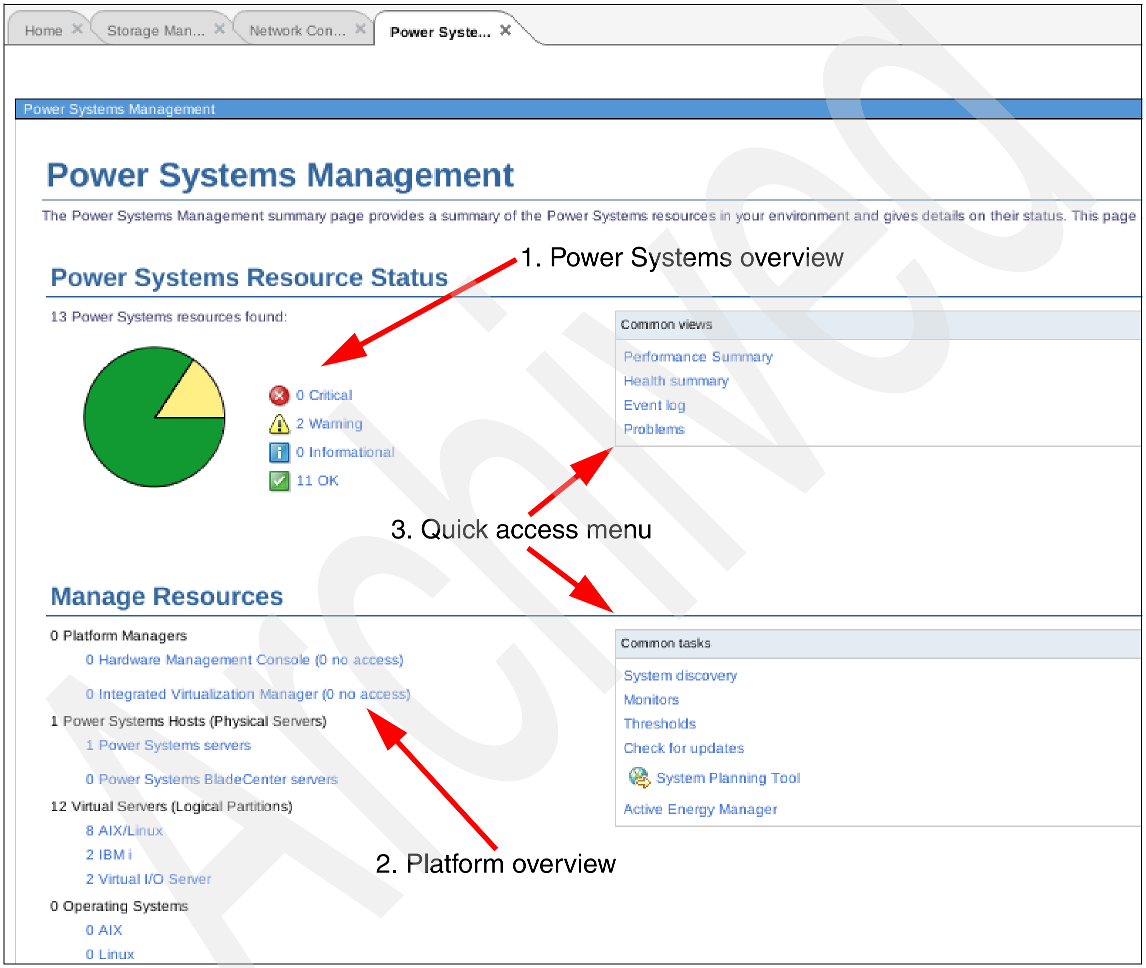

Figure 6-117 Power Systems Management

The Power Systems Management options are:

- Power Systems Overview: This option shows a tactical overview of the various states of the managed Power Systems servers and partitions.

- Manage Resources: This option shows a menu with options ordered by hardware, virtualized environment, and operating systems. Select one (in our example, the IBM Flex System p460 Compute Node, as shown in Figure 6-118). The actions menu provides options for managing your power server, such as create virtual servers, manage virtual servers, manage systems plans, and power on.

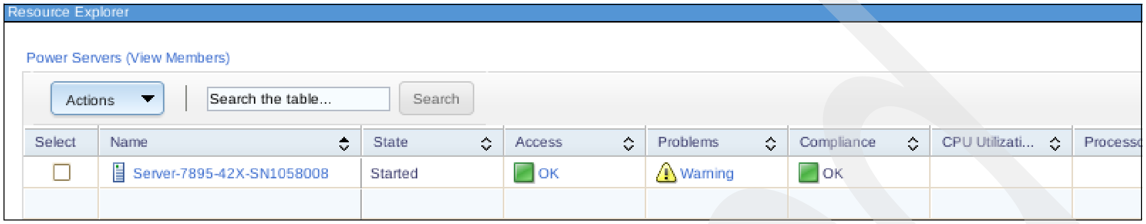

Figure 6-118 Power Server Management

By selecting a power server, in this example, the IBM Flex System p460 Compute Node, you have access to all virtual servers that run on that system. (Figure 6-119).

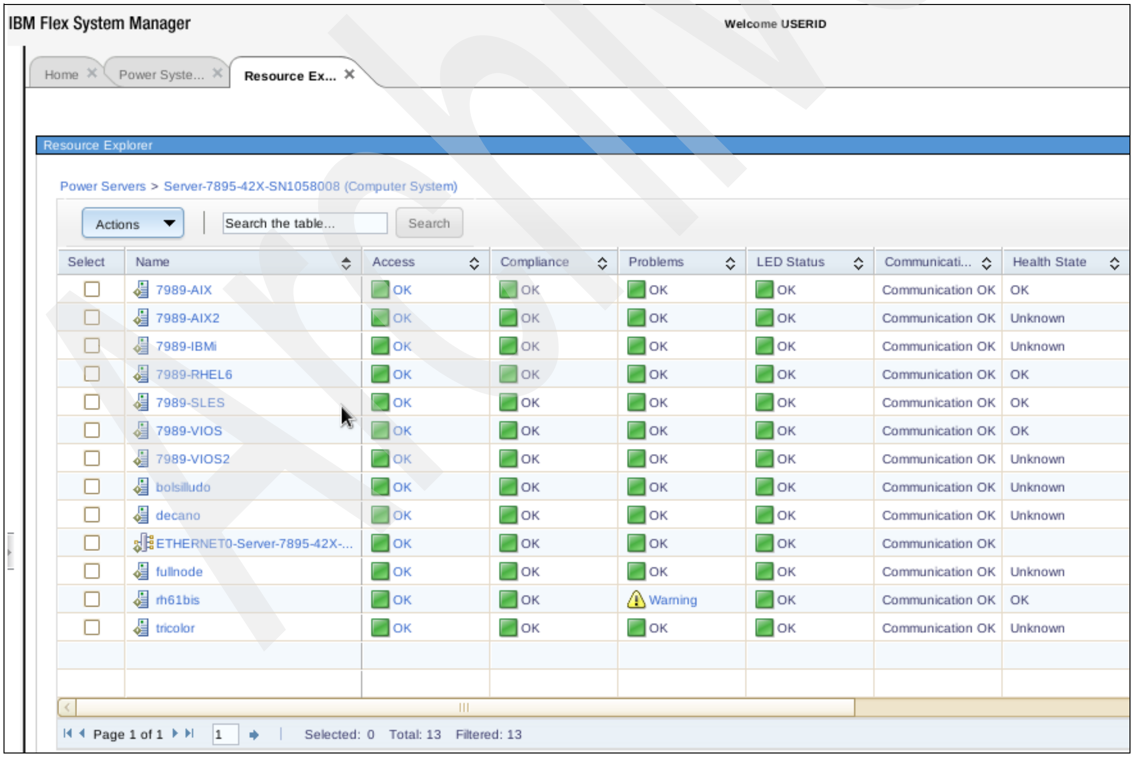

Figure 6-119 Virtual Server Management

System z Management

You can use this feature to manage the System z systems in your data center. It is similar to the Power Systems Management feature.

VMControl Enterprise Edition

With VMControl, you can manage all of your virtualized environment, from KVM to VMware servers. All of these items can be managed from this centralized interface in IBM Flex System Manager (Figure 6-120).

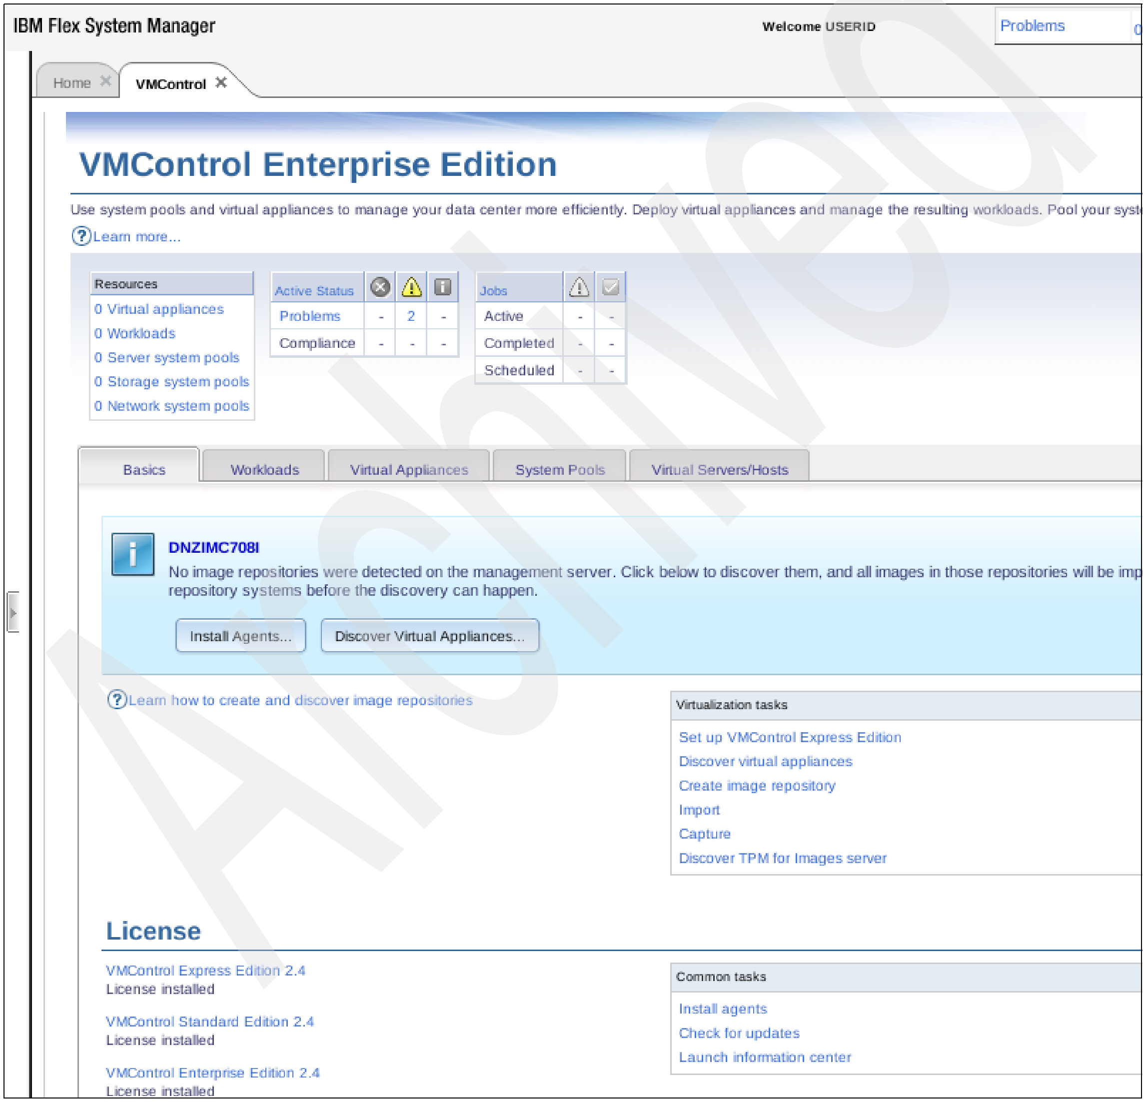

Figure 6-120 VMControl Management

Chassis Management

You can use this feature to show a tactical overview of the chassis and chassis components with problems. (For information about chassis that are not in compliance, see “Update Manager” on page 267.) Shortcuts are available for systems discovery, view, collect inventory, and so on.

Systems x Management

You can use this feature to manage the System x servers in your data center. It is similar to the Power Systems Management feature.

Service and Support Management

You can use this feature to open a tactical overview of the systems with serviceable events and service requests. There is also a quick menu for opening a new service request and managing your serviceable events.

6.7.4 Administrator tab

From the Administration tab, you can access all IBM Flex System Manager configurations for tasks, such as shut down, restart, power off, upgrade firmware, set up network, set up users, and backup and restore. See Figure 6-121.

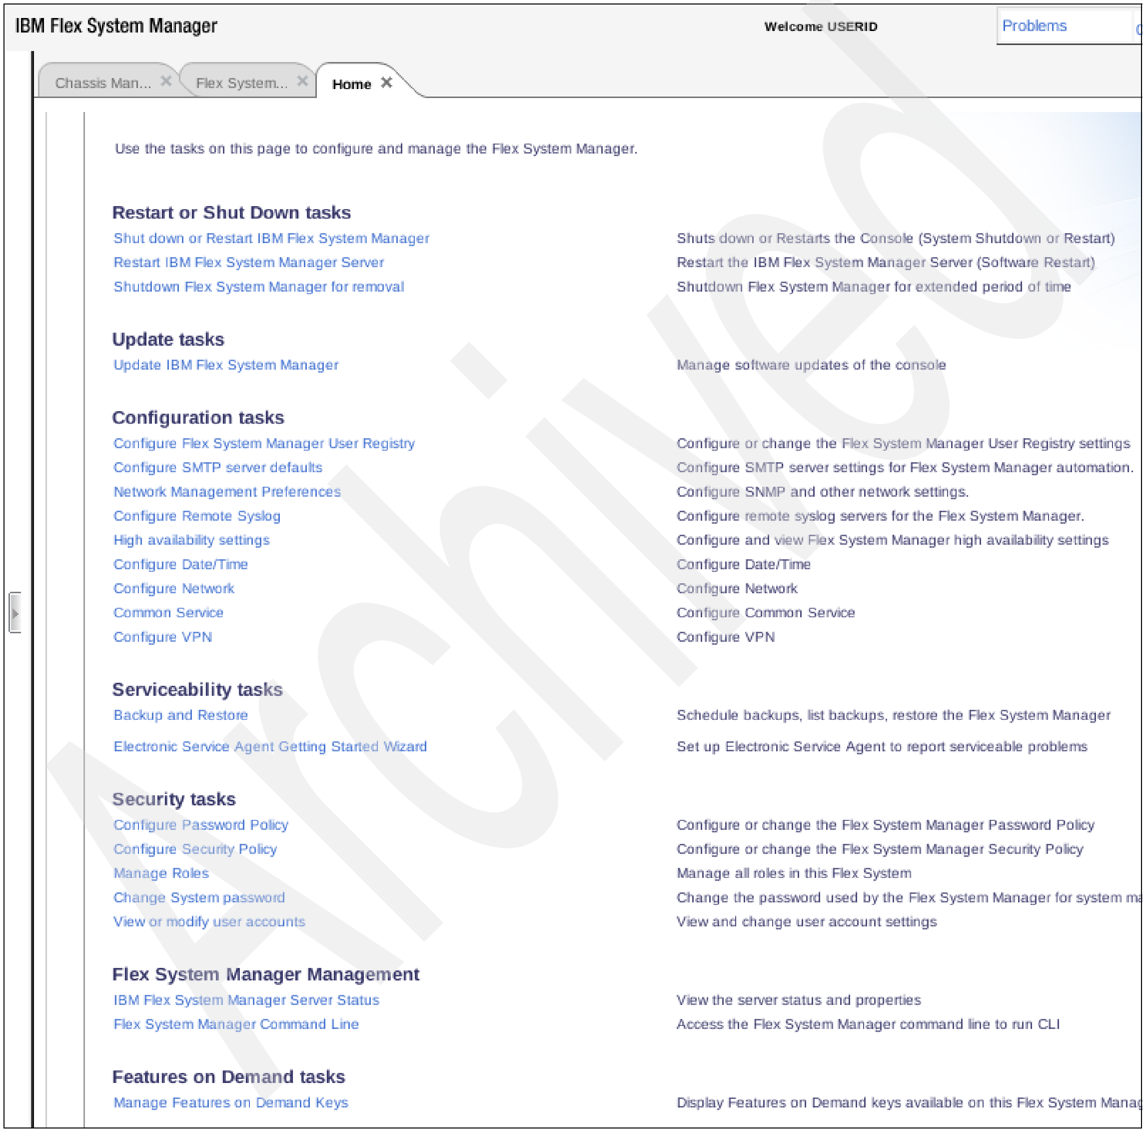

Figure 6-121 Administration

6.7.5 Learn tab

From the Learn tab, you can access IBM Flex System Manager online manuals.

7

Virtualization

If you create virtual servers (also known as *logical partitions* (LPARs)) on your Power Systems compute node, you can consolidate your workload to deliver cost savings and improve infrastructure responsiveness. As you look for ways to maximize the return on your IT infrastructure investments, consolidating workloads and increasing server use becomes an attractive proposition.

IBM Power Systems, combined with PowerVM technology, are designed to help you consolidate and simplify your IT environment. The following list details several key capabilities:

- Improve server use by consolidating diverse sets of applications.
- Share processor, memory, and I/O resources to reduce the total cost of ownership (TCO).
- Improve business responsiveness and operational speed by dynamically reallocating resources to applications as needed, to better anticipate changing business needs.
- Simplify IT infrastructure management by making workloads independent of hardware resources, so that you can make business-driven policies to deliver resources based on time, cost, and service-level requirements.
- Move running workloads between servers to maximize availability and avoid planned downtime.

7.1 PowerVM

PowerVM delivers industrial-strength virtualization for AIX, IBM i, and Linux environments on IBM POWER processor-based systems. Power Systems servers, coupled with PowerVM technology, are designed to help clients build a dynamic infrastructure, which reduces costs, manages risk, and improves service levels.

7.1.1 Features

The latest version of PowerVM contains the following features:

- Support for the following number of maximum virtual servers (or logical partitions, LPARs):
  - p260: Up to 160 virtual servers
  - p460: Up to 320 virtual servers
  - p24L: Up to 120 virtual servers
- Role Based Access Control (RBAC)

  RBAC brings an added level of security and flexibility in the administration of the Virtual I/O Server (VIOS). With RBAC, you can create a set of authorizations for the user management commands. You can assign these authorizations to a role named UserManagement, and this role can be given to any other user. So one user with the role, UserManagement, can manage the users on the system, but does not have any further access.

  With RBAC, the VIOS can split management functions that presently can be done only by the *padmin* user, providing better security by giving only the necessary access to users, and easy management and auditing of system functions.
- Suspend/Resume

  Using Suspend/Resume, you can provide long-term suspension (greater than 5 - 10 seconds) of partitions, saving partition state (memory, NVRAM, and VSP state) on persistent storage. This action makes server resources available that were in use by that partition, restoring partition state to server resources, and resuming operation of that partition and its applications, either on the same server or on another server.

  The requirements for Suspend/Resume dictate that all resources must be virtualized before suspending a partition. If the partition is resumed on another server, then the shared external I/O (disk and local area network (LAN)) needs to remain identical. Suspend/Resume works with AIX and Linux workloads when managed by the Hardware Management Console (HMC).

- Shared storage pools

  You can use VIOS 2.2 to create storage pools that can be accessed by VIOS partitions that are deployed across multiple Power Systems servers. Therefore, an assigned allocation of storage capacity can be efficiently managed and shared.

  The December 2011 Service Pack enhances capabilities by enabling four systems to participate in a Shared Storage Pool configuration. This configuration can improve efficiency, agility, scalability, flexibility, and availability. Specifically, the Service Pack enables:

  – Storage Mobility: A function that allows data to be moved to new storage devices within Shared Storage Pools, while the virtual servers remain active and available.

  – VM Storage Snapshots/Rollback: A new function that allows multiple point-in-time snapshots of individual virtual server storage. These point-in-time copies can be used to quickly roll back a virtual server to a particular snapshot image. This functionality can be used to capture a VM image for cloning purposes or before applying maintenance.

- Thin provisioning

  VIOS 2.2 supports highly efficient storage provisioning, where virtualized workloads in VMs can have storage resources from a shared storage pool dynamically added or released, as required.

- VIOS grouping

  Multiple VIOS 2.2 partitions can use a common shared storage pool to more efficiently use limited storage resources and simplify the management and integration of storage subsystems.

- Network node balancing for redundant Shared Ethernet Adapters (SEAs) (with the December 2011 Service Pack).

  This feature is useful when multiple VLANs are being supported in a dual VIOS environment. The implementation is based on a more granular treatment of trunking, where there are different trunks defined for the SEAs on each VIOS. Each trunk serves different VLANs, and each VIOS can be the primary for a different trunk. This situation occurs with just one SEA definition on each VIOS.

IBM PowerVM Workload Partitions Manager™ for AIX Version 2.2 has the following enhancements:

- When used with AIX V6.1 Technology Level 6, the following support applies:
  - Support for exporting a VIOS SCSI disk into a Workload Partition (WPAR). There is compatibility analysis and mobility of WPARs with VIOS SCSI disk. In addition to Fibre Channel devices, VIOS SCSI disks can be exported into a WPAR.
  - WPAR Manager Command-Line Interface (CLI). The WPAR Manager CLI allows federated management of WPARs across multiple systems through the command line.
  - Support for workload partition definitions. The WPAR definitions can be preserved after WPARs are deleted. These definitions can be deployed at a later time to any WPAR-capable system.
- In addition to the features supported on AIX V6.1 Technology Level 6, the following features apply to AIX V7.1:
  - Support for AIX 5L V5.2 Workload Partitions for AIX V7.1. Lifecycle management and mobility enablement for AIX 5L V5.2 Technology Level 10 SP8 Version WPARs.
  - Support for AIX 5L V5.3 Workload Partitions for AIX V7.1. Lifecycle management and mobility enablement for AIX 5L V5.3 Technology Level 12 SP4 Version WPARs.
  - Support for trusted kernel extension loading and configuration from WPARs. Enables exporting a list of kernel extensions that can then be loaded inside a WPAR, while maintaining isolation.

7.1.2 POWER Hypervisor

Combined with features designed for POWER7 processors, the POWER Hypervisor delivers functions that enable capabilities. These functions include dedicated processor partitioning, Micro-Partitioning, virtual processors, IEEE VLAN-compatible virtual switches, virtual SCSI adapters, virtual Fibre Channel adapters, and virtual consoles.

The user interface for the POWER Hypervisor on POWER based blades is traditionally based on the Integrated Virtualization Manager. With the PS700 family of blades, a second method of systems management is available: the Systems Director Management Console. A new user interface is introduced with the introduction of IBM Flex System Manager. This chapter focuses on using the IBM Flex System Manager for most configuration tasks performed on the Power Systems compute nodes.

POWER Hypervisor technology is integrated with all IBM POWER servers, including the Power Systems compute nodes. The hypervisor orchestrates and manages system virtualization, including creating logical partitions and dynamically moving resources across multiple operating environments. POWER Hypervisor is a basic component of the system firmware that is layered between the hardware and the operating system.

The POWER Hypervisor has the following features:

- Provides an abstraction layer between the physical hardware resources and the logical partitions that use them
- Enforces partition integrity by providing a security layer between logical partitions
- Controls the dispatch of virtual processors to physical processors and saves and restores all processor state information during a logical processor context switch
- Controls hardware I/O interrupt management facilities for logical partitions
- Provides virtual Ethernet switches between logical partitions that help to reduce the need for physical Ethernet adapters for interpartition communication
- Monitors the service processor and performs a reset or reload if it detects the loss of the service processor, notifying the operating system if the problem is not corrected
- Uses Micro-Partitioning to allow multiple instances of the operating system

The POWER Hypervisor is always installed and activated, regardless of system configuration. It does not own any physical I/O devices; all physical I/O devices in the system are owned by logical partitions or by the Virtual I/O Server.

Memory is required to support the resource assignment to the logical partitions on the server. The amount of memory required by the POWER Hypervisor firmware varies according to several factors. The following factors influence POWER Hypervisor memory requirements:

- Number of logical partitions
- Number of physical and virtual I/O devices used by the logical partitions
- Maximum memory values specified in the logical partition profiles

The minimum amount of physical memory to create a partition is the size of the system logical memory block (LMB). The default LMB size varies according to the amount of memory configured in the system, as shown in Table 7-1.

Table 7-1 Configured memory-to-default LMB size

Configurable memory in the system	Default LMB
Less than 4 GB	16 MB
Greater than 4 GB up to 8 GB	32 MB
Greater than 8 GB up to 16 GB	64 MB
Greater than 16 GB up to 32 GB	128 MB
Greater than 32 GB	256 MB

Physical memory assigned to partitions is in increments of LMB.

The POWER Hypervisor provides the following types of virtual I/O adapters:

- Virtual SCSI
- Virtual Ethernet
- Virtual Fibre Channel
- Virtual (TTY) console

Virtual I/O adapters are defined by system administrators during logical partition definition. Configuration information for the adapters is presented to the partition operating system.

Virtual SCSI

The POWER Hypervisor provides a virtual SCSI mechanism for virtualization of storage devices. Virtual SCSI allows secure communications between a logical partition and the I/O Server (VIOS). The storage virtualization is accomplished by pairing two adapters: a virtual SCSI server adapter on the VIOS, and a virtual SCSI client adapter on IBM i, Linux, or AIX partitions. The combination of Virtual SCSI and VIOS provides the opportunity to share physical disk adapters in a flexible and reliable manner.

Virtual Ethernet

The POWER Hypervisor provides an IEEE 802.1Q, VLAN-style virtual Ethernet switch that allows partitions on the same server to use fast and secure communication without any need for physical connection.

Virtual Ethernet support starts with AIX 5L V5.3, or the appropriate level of Linux supporting virtual Ethernet devices. The virtual Ethernet is part of the base system configuration.

Virtual Ethernet has the following major features:

- Virtual Ethernet adapters can be used for both IPv4 and IPv6 communication and can transmit packets up to 65,408 bytes in size. Therefore, the maximum transmission unit (MTU) for the corresponding interface can be up to 65,394 (65,408 minus 14 for the header) in the non-VLAN case, and up to 65,390 (65,408 minus 14, minus 4) if VLAN tagging is used.
- The POWER Hypervisor presents itself to partitions as a virtual 802.1Q-compliant switch. The maximum number of VLANs is 4096. Virtual Ethernet adapters can be configured as either untagged or tagged (following the IEEE 802.1Q VLAN standard).
- An AIX partition supports 256 virtual Ethernet adapters for each logical partition. Aside from a default port VLAN ID, the number of additional VLAN ID values that can be assigned per virtual Ethernet adapter is 20, which implies that each virtual Ethernet adapter can be used to access 21 virtual networks.
- Each operating system partition detects the virtual local area network (VLAN) switch as an Ethernet adapter without the physical link properties and asynchronous data transmit operations.

Any virtual Ethernet can also have connectivity outside of the server if a Layer 2 bridge to a physical Ethernet adapter is configured in a VIOS partition. The device configured in this fashion is the SEA.

> **Important:** Virtual Ethernet is based on the IEEE 802.1Q VLAN standard. No physical I/O adapter is required when creating a VLAN connection between partitions, and no access to an outside network is required for inter-partition communication.

Virtual Fibre Channel

A virtual Fibre Channel adapter is a virtual adapter that provides client logical partitions with a Fibre Channel connection to a storage area network (SAN) through the VIOS logical partition. The VIOS logical partition provides the connection between the virtual Fibre Channel adapters on the VIOS logical partition and the physical Fibre Channel adapters on the managed system.

N_Port ID virtualization (NPIV) is a standard technology for Fibre Channel networks. You can use NPIV to connect multiple logical partitions to one physical port of a physical Fibre Channel adapter. Each logical partition is identified by a unique worldwide port name (WWPN), which means that you can connect each logical partition to independent physical storage on a SAN.

Enabling NPIV: To enable NPIV on a managed system, you must have VIOS V2.1 or later. NPIV is only supported on 8 Gb Fibre Channel and Converged Network (Fibre Channel over Ethernet, FCoE)) adapters on a Power Systems compute node.

You can configure only virtual Fibre Channel adapters on client logical partitions that run the following operating systems:

- AIX V6.1 Technology Level 2, or later
- AIX 5L V5.3 Technology Level 9, or later
- IBM i V6.1.1, V7.1, or later
- SUSE Linux Enterprise Server 11, or later
- RHEL 5.5, 6, or later

Systems that are managed by the Integrated Virtualization Manager, a Systems Director Management Console, or IBM Flex System Manager can dynamically add and remove virtual Fibre Channel adapters from logical partitions.

Figure 7-1 shows the connections between the client partition virtual Fibre Channel adapters and external storage.

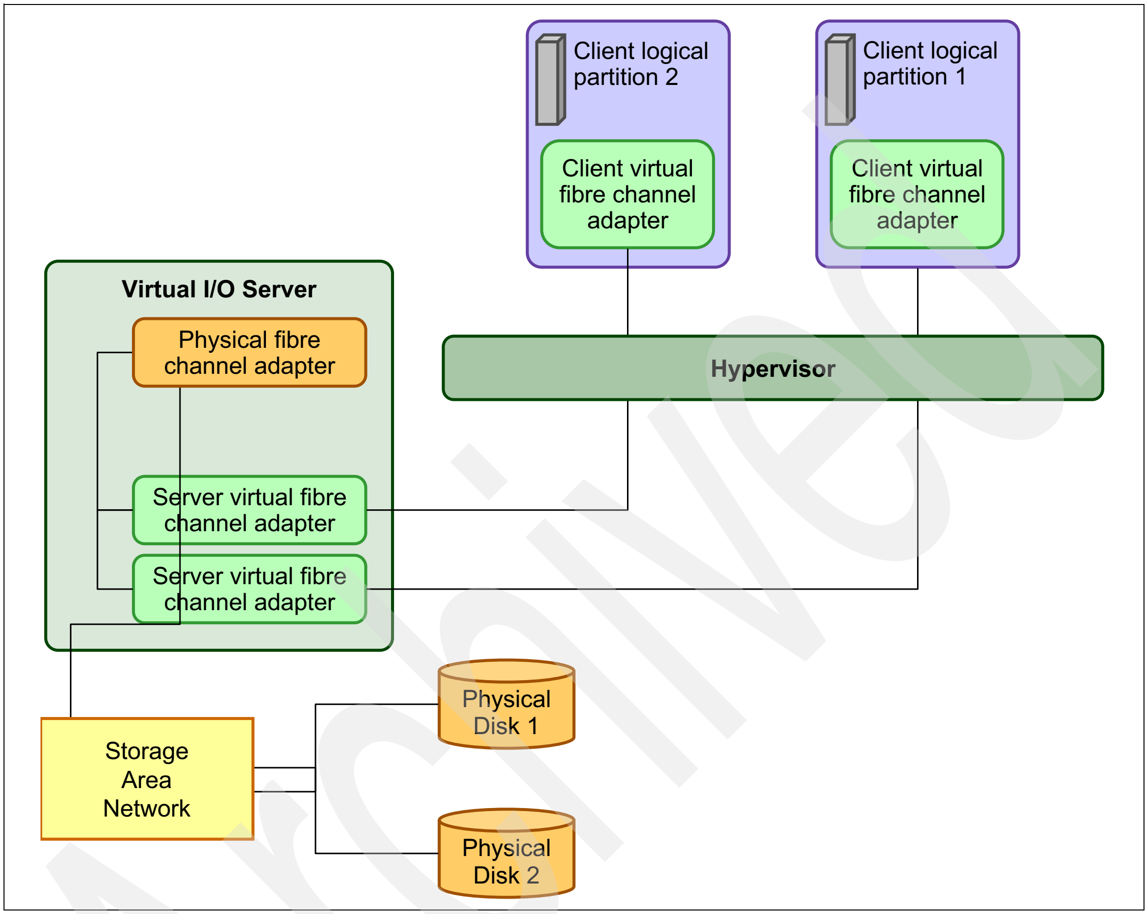

Figure 7-1 Connectivity between virtual Fibre Channel adapters and external SAN devices

Virtual serial adapters (TTY) console

Virtual serial adapters provide a point-to-point connection from one logical partition to another, or the IBM Flex System Manager to each logical partition on the managed system. Virtual serial adapters are used primarily to establish terminal or console connections to logical partitions.

Each partition must have access to a system console. Tasks such as operating system installation, network setup, and certain problem analysis activities require a dedicated system console. The POWER Hypervisor provides the virtual console using a virtual TTY or serial adapter and a set of Hypervisor calls to operate on it. Virtual TTY does not require the purchase of any additional features or software, such as the PowerVM Edition features.

For Power Systems compute nodes, the operating system console can be accessed from IBM Flex System Manager.

7.1.3 Preparing to use the IBM Flex System Manager for partitioning

FSM is used to create virtual servers on Power Systems compute nodes. This function is one of many provided by FSM.

It is presumed that FSM is set up so that it is managing the Enterprise Chassis that contains the Power Systems compute node on which the virtual servers is created. For more details, see Chapter 6, “Management setup” on page 157.

If you have experience using the Integrated Virtualization Manager, HMC, or the Systems Director Management Console to create logical partitions or virtual servers on any POWER7 system, the process is similar.

Removing an existing configuration

All Power Systems compute nodes are delivered preinstalled with an operating system. In addition, if a IBM PureFlex System configuration was ordered, there is a configuration on the Power Systems compute nodes in the chassis. The virtual servers that are already configured must be removed if you want to create a custom configuration.

Virtual server name, processor, and memory settings

Before beginning the virtual server creation tasks in FSM, document the basic information about the virtual servers to be defined (Table 7-2).

Table 7-2 Virtual server name, processor, and memory planning information

Virtual server name	Processor/UnCap/Weight	Memory
node1		
vios1	1/Y/200	2 GB
vios2	1/Y/200	2 GB
lpar1	3/Y/100	4 GB

Virtual server name	Processor/UnCap/Weight	Memory
lpar2	.5/N/-	1 GB
lpar3	.5/N/-	1 GB
node2		
vios3	1/Y/200	2 GB
vios4	1/Y/200	2 GB
lpar1	3/Y/100	4 GB
lpar2	2/Y/50	2 GB
lpar3	2/Y/50	2 GB
lpar4	1.5/N/-	1 GB
lpar5	1.5/N/-	1 GB
lpar6	1.5/N/-	1 GB

Physical adapters

For the VIOS partitions, planning for physical adapter allocation is important, because the VIOS provides virtualized access through the physical adapters to network or disk resources. If availability is a concern for the virtualized environment, use redundant physical adapters in the VIOS. For network adapters, you most likely use Etherchannel. For storage adapters, a multipathing package (for example, an MPIO-PCM or EMC PowerPath) is installed and configured in the VIOS after the operating system is installed. To further enhance availability in a virtualized configuration, implement two VIOS servers, both capable of providing the same network and storage access to the virtual servers on the Power Systems compute node. The ideal availability configuration involves redundant physical adapters in each of the two VIOS servers. Because of hardware requirements in a dual VIOS configuration, a p460 might be the better choice.

Create a similar document that shows the physical adapter assignments to each VIOS. With only two or four adapters to be assigned to the virtual servers, the document is fairly simple.

Virtual adapters

Assigning and configuring virtual adapters requires more planning and design. For virtual Ethernet adapters, the VLANs that the virtual servers require access to must be considered. The VIOS provides bridging from the virtual Ethernet adapter to the physical. Therefore, the virtual Ethernet adapter in the VIOS must be configured with all of the VLANs that are required for the virtual servers in the node.

For virtual storage access, either virtual SCSI or NPIV can be used. Virtual SCSI adapters are configured in a client-server relationship, with the client adapter in the client virtual server configured to refer to the server adapter configured in the VIOS. The server adapter in the VIOS can be configured to refer to one client adapter or allow any client to connect. NPIV configuration differs, in that the VIOS serves as a pass-through module for a virtual Fibre Channel adapter in the client virtual server. The SAN administrator assigns LUNs to the virtual Fibre Channel adapters in the virtual servers, just as they would for a real Fibre Channel adapter. The WWPNs are generated when the virtual Fibre Channel adapter is defined for the client. This configuration can be provided to the SAN administrator to ensure the LUNs are correctly mapped in the SAN.

Documenting the relationships between the VIOS and the client virtual servers leads to correctly defined virtual adapters when you created the virtual servers in FSM.

For more details about planning and configuring virtualized environments, including configuring for availability, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

For more information about planning and configuring a highly available virtual environment, see *IBM System p Advanced POWER Virtualization (PowerVM) Best Practices*, REDP-4194.

7.2 Creating the VIOS virtual server

Creating any virtual server requires basically the same steps. It is possible that you create only AIX or Linux virtual servers on your Power Systems compute node, but your options are limited by hardware, especially the expansion cards. To use the virtualization capabilities of the compute node, use a VIOS. For this reason, we focus on the creation of the VIOS.

7.2.1 Using the CLI

Many integrators and system administrators make extensive and efficient use of the CLI, rather than use a graphical interface for their virtual server creation and administration tasks. Tasks can be scripted, and often the tasks are completed faster using the command line.

> **Scripts:** In many cases, existing scripts that were written for usage on a Systems Director Management Console can run unchanged on FSM. Similarly, scripts written to run on an HMC might run with the **smcli** command, added to each command in the script.

Accessing the IBM Flex System Manager

To access the FSM, you must know the IP address or host name of the FSM node and have a valid user ID and password. You must start a Secure Shell (ssh) session with FSM and log in. This process is similar to the process of accessing the SDMC or HMC command line.

Creating the VIOS virtual server using the CLI

Creating the VIO Server can be done by using the FSM CLI. It is a good idea to consider creating the virtual server without the I/O definitions using one command, and then update the virtual server with the I/O definitions later, possibly by using the graphical interface, to minimize the complexity of the command.

To ensure that the correct I/O devices are specified in the command, understand and document the intended I/O adapters using the information described in "Assigning physical I/O" on page 301. An example of modifying a virtual server definition using the CLI is in 7.3, "Modifying the VIOS definition" on page 304.

To create a VIO Server using a single command, run the following command:

```
smcli mksyscfg -r lpar -m p4601 -i
"name=vios1,profile_name=vios1_default,lpar_env=vioserver,lpar_id=1,msp
=1,min_mem=1024,desired_mem=2048,max_mem=4096,proc_mode=shared,min_proc
_units=1.0,desired_proc_units=1.0,max_proc_units=2.0,min_procs=2,desire
d_procs=4,max_procs=8,sharing_mode=uncap,uncap_weight=200,max_virtual_s
lots=100,virtual_eth_adapters=\"11/0/1//1/0,12/0/99///0\",virtual_scsi_
adapters=13/server/10//13/0,boot_mode=norm,io_slots=\"21010201//1,21010
210//1,21010220//1\""
```

> **VIOS command:** This command creates a VIOS server that matches the one created and modified in 7.2.2, "Using the IBM Flex System Manager" on page 288, which shows the usage of the graphical interface.

Verification of success

A successful command produces a prompt with no message displayed.

To verify that the VIO Server was created, run **smcli lssyscfg**, scanning the results for the name of your virtual server:

```
sysadmin@sys1234567: ~> smcli lssyscfg -r lpar -m p4601 -F name
7989-SLES
7989-AIX
7989-RHEL6
7989-VIOS
vios1
```

To verify the content of the profile created as a result, run **smcli lssyscfg** with different parameters:

```
sysadmin@sys1234567: ~> smcli lsysscfg -r prof -m p4601 --filter
lpar_names=vios1
```

Recognition of failure

There are many reasons that your CLI command might fail. A syntax error is the most likely, producing something like the following output. Use the information in the message to correct the problem.

```
An error occurred while creating the partition named vios1.

The format of the configuration data is invalid. The correct format is
"<attribute name 1>=<value>,<attribute name 2>=<value>,..." or
""<attribute name 1>=<value 1>,<value 2>,...>",..." .  Please correct
the configuration data and retry the command.  Note that depending on
the shell being used, any nested double quote characters may need to be
preceded by an escape character, which is usually a '\' character.
```

7.2.2 Using the IBM Flex System Manager

This section describes the sequence to perform the same steps described in 7.2.1, “Using the CLI” on page 287, but with the FSM GUI instead.

Accessing the IBM Flex System Manager

IBM Flex System Manager can be accessed in one of two ways:

- Locally with a keyboard, mouse, and screen attached directly to port at the front panel of the FSM through the Console Breakout Cable.
- Through a web browser to the FSM web interface.

We access the FSM remotely using a browser. Complete the following steps:

1. Open a browser and point the browser to the following URL (where `system_name` is the host name or IP address of the FSM node):

   `https://system_name:8422`

   > **Port number:** The port you use may be different than the port we use in our examples.

   A login window opens, as shown in Figure 7-2.

Figure 7-2 IBM Flex System Manager login window

2. Enter a valid FSM user ID and password, and click **Log in**. The Welcome window opens.

3. Click **Home** opens the main window, as shown in Figure 7-3.

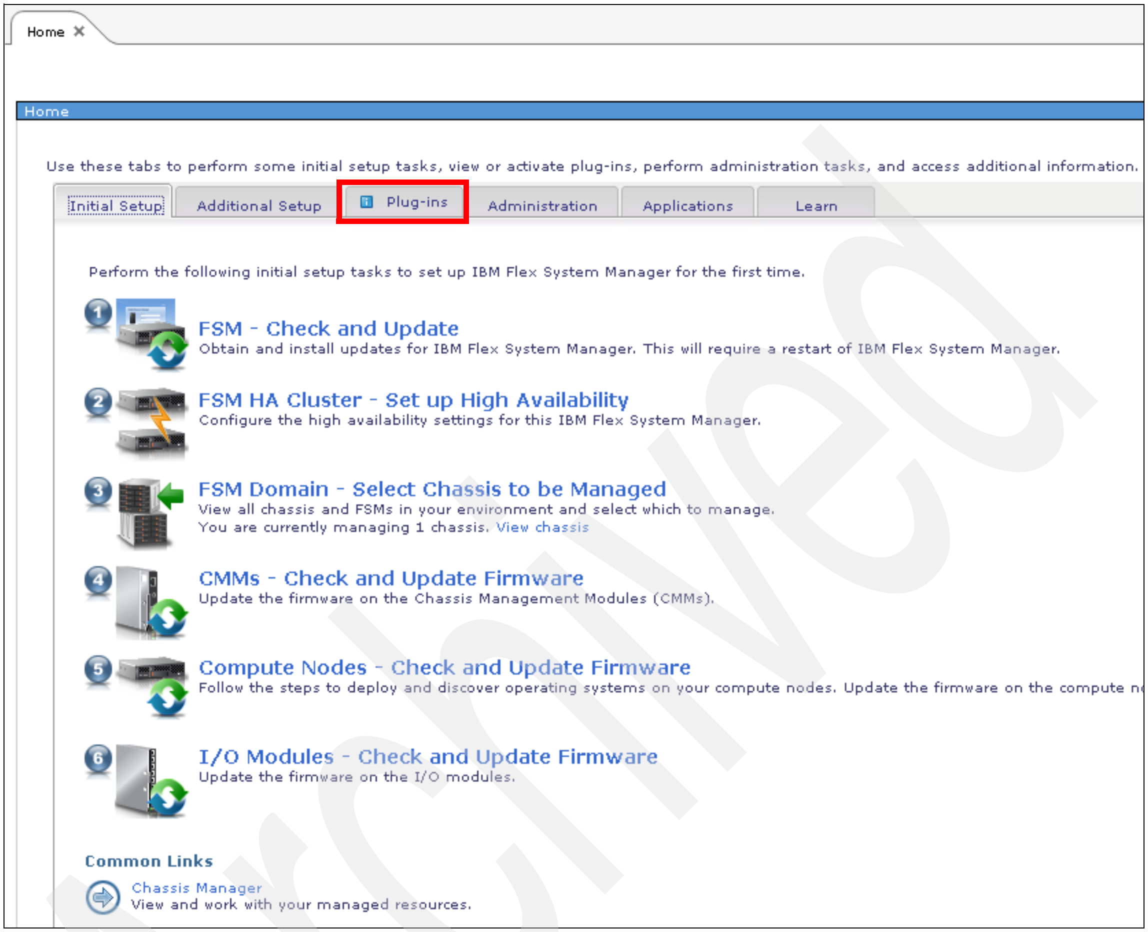

Figure 7-3 IBM Flex System Manager home window

4. Click the **Plug-ins** tab to display the list of installed plug-ins. The list of installed plug-ins opens, as shown in Figure 7-4.

Figure 7-4 IBM Flex System Manager Plug-ins tab - highlighting the Power Systems Management plug-in

5. Click the **Power Systems Management** plug-in to display the Power Systems Management main window, shown in Figure 7-5.

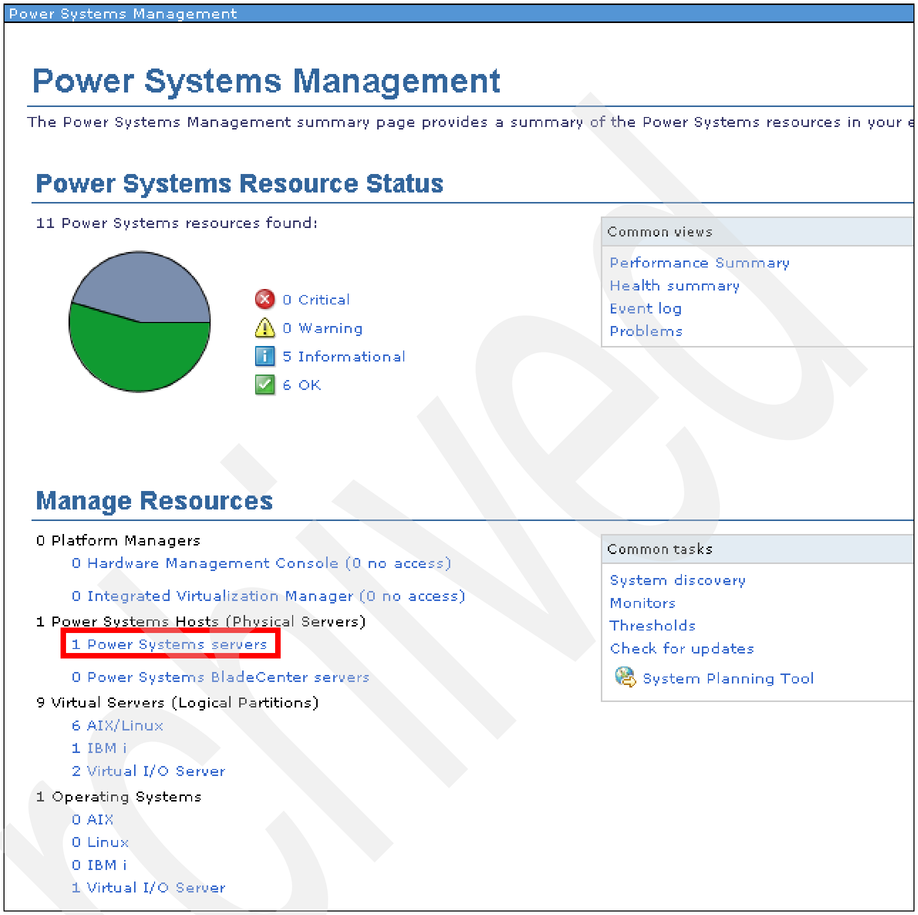

Figure 7-5 Power Systems Management main window

Creating the VIOS virtual server using the GUI

When you open the Power Systems Management main window shown in Figure 7-5, you see choices to manage hosts and virtual servers. In this section, we describe how to create the VIOS virtual server.

Creating the virtual server

To create the virtual server, complete the following steps:

1. In the Power Systems Hosts (Physical Servers) section, click **Power Systems servers**, as shown in Figure 7-5. A list of hosts that are managed by this Flex System Manager opens.

2. Select the compute node.

   If more hosts are managed by this Flex System Manager, select the one on which the VIOS virtual server is created.

3. Click **Actions** → **System Configuration** → **Create Virtual Server** to start the wizard (Figure 7-6).

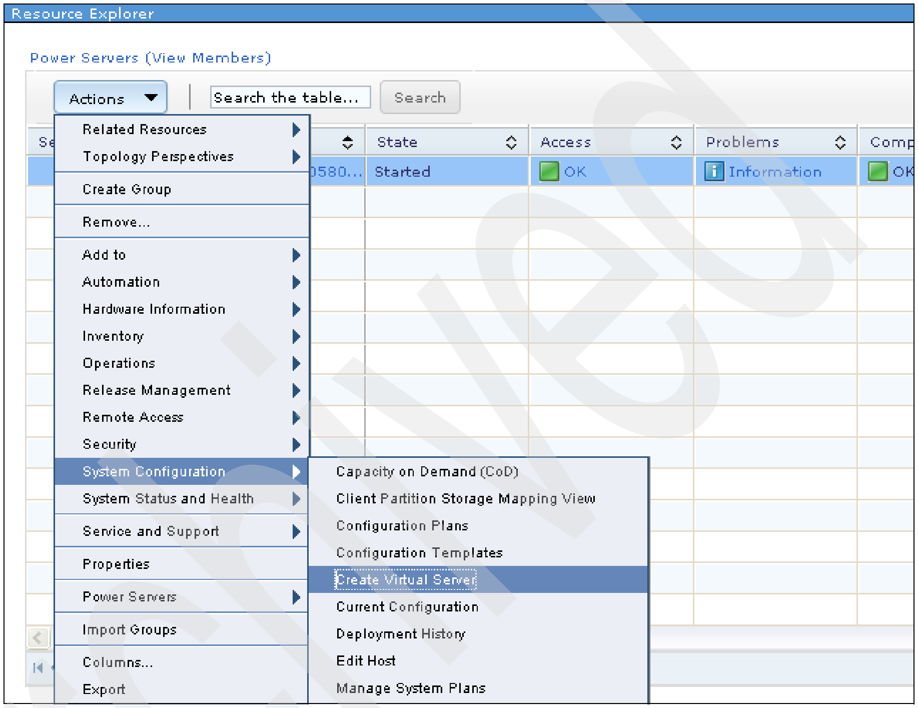

Figure 7-6 Create a virtual server menu option

The window shown in Figure 7-7 opens.

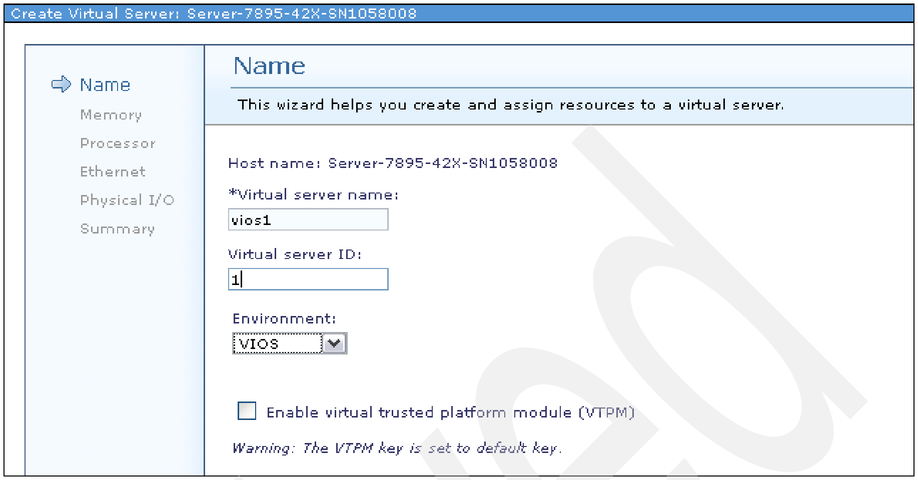

Figure 7-7 Setting the VIOS virtual server name and ID

– Enter the virtual server name. We use the name vios1.

– Enter the server ID. We give our VIOS an ID of 1.

– Specify the Environment option to identify this environment as a VIOS.

4. Click **Next**.

Memory and processor settings

The next task is to choose the amount of memory for the VIOS virtual server. Starting with Figure 7-8 (which you reach by performing the steps in “Creating the virtual server” on page 292), complete the following steps.

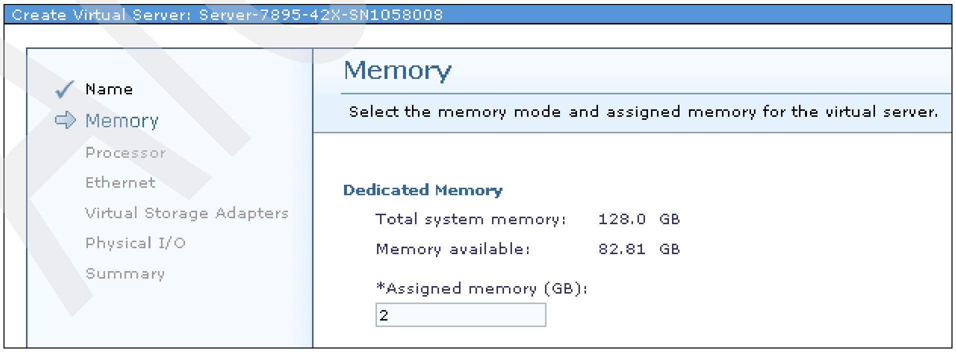

Figure 7-8 Specify the memory information for the VIOS virtual server

1. Change the value to reflect the amount of wanted memory in gigabytes.

   Decimal fractions can be specified to assign memory in megabyte increments. This memory is the amount of memory the hypervisor attempts to assign when the VIOS is activated. We assign the VIOS 2 GB of memory.

   > **Minimum and maximum values:** You cannot specify minimum or maximum settings. The value specified here is the wanted value. Minimum and maximum values can be edited after the virtual servers are created, as described in 7.3, "Modifying the VIOS definition" on page 304.

2. Click **Next** to proceed to the processor settings. Figure 7-9 opens.

   We choose to make vios1 part of the shared processor pool. We want 1.0 processing units (one processor available for threads every timer tick), so we must specify 10 in the Assigned processors field. This setting gives us 1.0 processing units and 10 virtual processors.

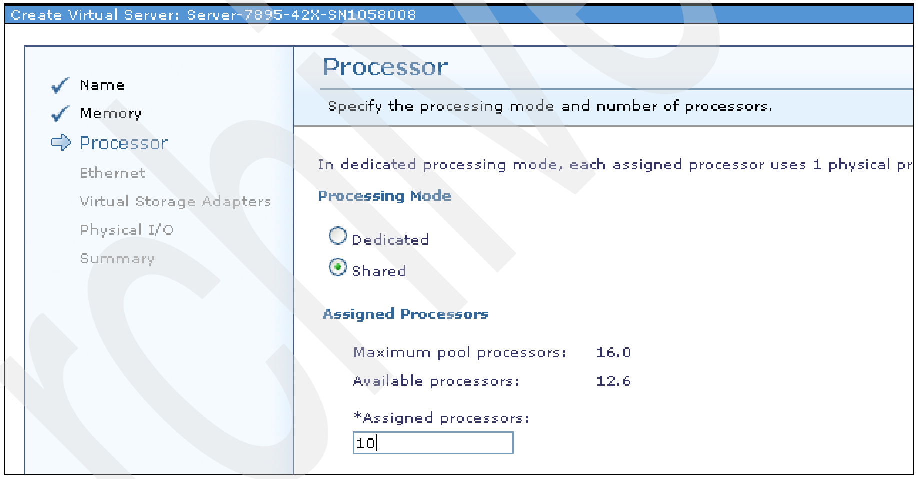

Figure 7-9 Setting the processor characteristics for the VIOS virtual server

> **Specifying processor units:** You cannot specify processing units, either uncapped or capped, or weight. These values can be edited after the virtual servers are created, as described in 7.3, "Modifying the VIOS definition" on page 304.

No memory or processing resources are committed. In this step, and in the rest of the steps for defining the virtual server, we are defining only the resources that are allocated to this virtual server after it is activated.

3. Click **Next** to move to the virtual adapter definitions.

Virtual Ethernet

In this task, the process is repeated for each virtual adapter to be defined on the VIOS, but the characteristics differ with each adapter type. The order in which the adapters are created does not matter.

Be sure to double-check your planning documentation to ensure that you are specifying the correct VLAN IDs for the virtual Ethernet adapters, that the virtual SCSI client and server adapters are correctly linked, and that the WWPN of the virtual Fibre Channel adapters are noted and provided to the SAN administrators.

If you performed the steps in “Memory and processor settings” on page 294, you should see the window shown in Figure 7-10.

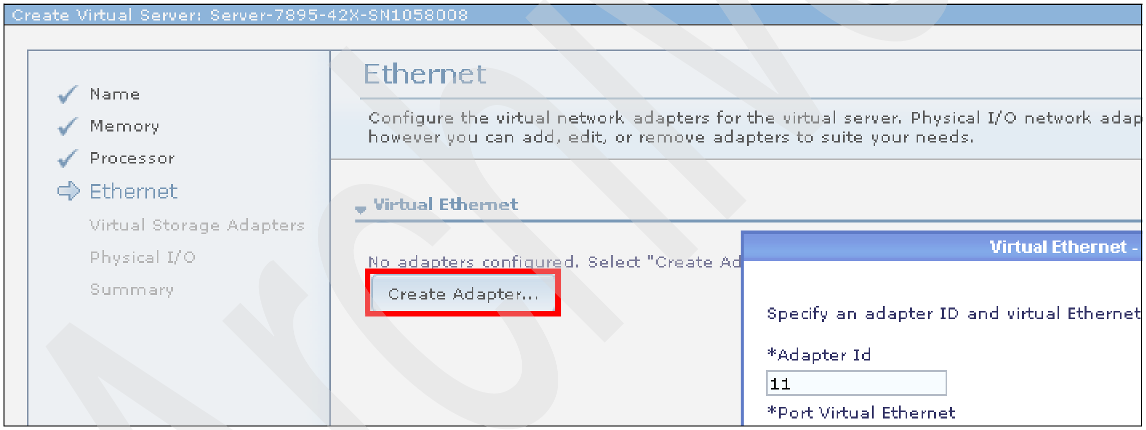

Figure 7-10 Create the bridging virtual Ethernet adapter in the VIOS

Complete the following steps:

1. Define the bridging virtual Ethernet adapter. Click **Create Adapter**, which opens the window where you create the bridging virtual Ethernet adapter, as shown in Figure 7-11.

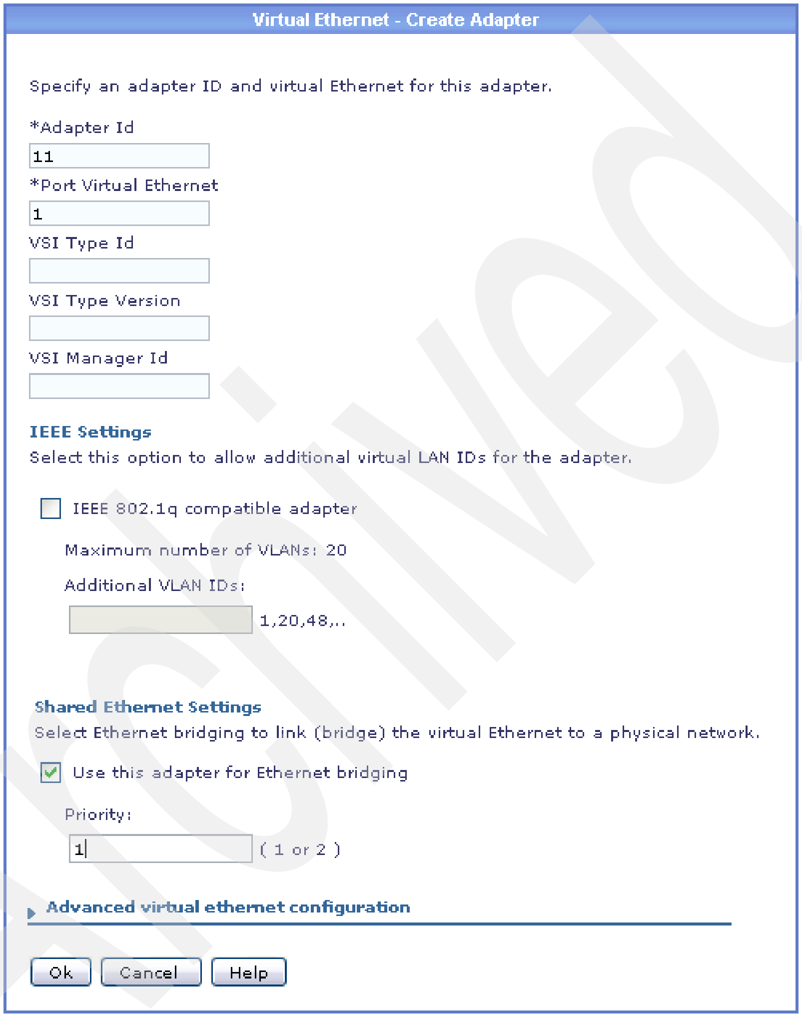

Figure 7-11 Create Adapter window

2. Enter the characteristics for the bridging virtual Ethernet adapter as follows:

   – It is standard practice to skip the first 10 adapter IDs. Start by defining the bridging virtual Ethernet adapter with an ID of 11.

   – Assume that all the packets are untagged, so leave the **Port Virtual Ethernet** option set to 1, and leave the **IEEE 802.1Q capable adapter** option unset.

   – This adapter is used in a Shared Ethernet Adapter definition, so update that section. Select the **Use this adapter for Ethernet bridging** check box, and set the Priority value. Each adapter in each VLAN must have a unique priority. No two adapters in the same VLAN should have the same priority.

3. Click **OK** when the values are specified. You return to the main virtual Ethernet window, where the newly created adapter is displayed (partially obscured in Figure 7-12).

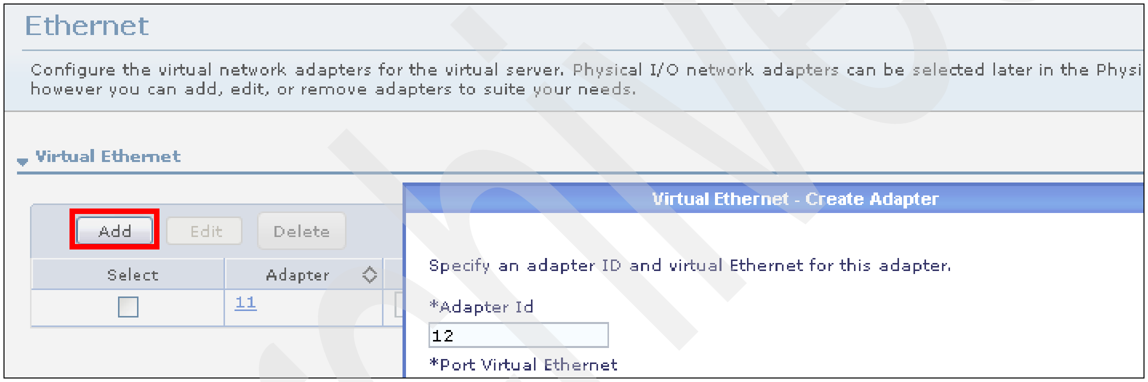

Figure 7-12 Create control channel virtual Ethernet adapter for SEA failover

4. Click **Add** to add more virtual Ethernet adapters, and a new virtual Ethernet adapter window opens.

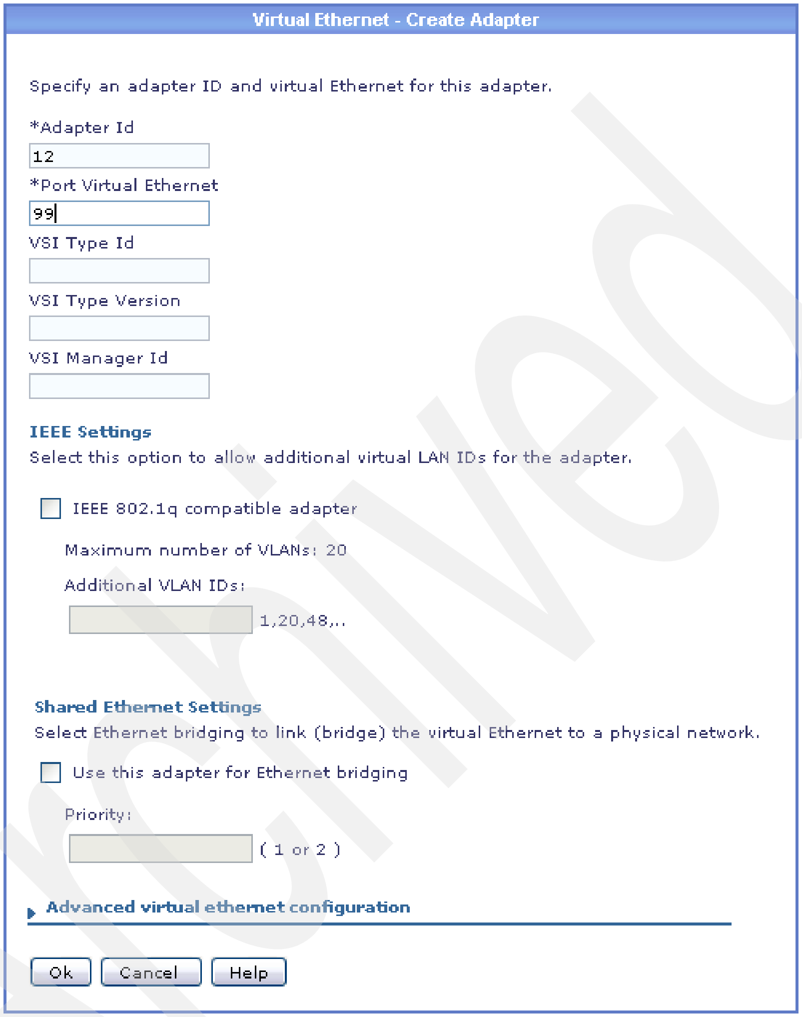

Figure 7-13 Create Adapter window

5. Create an additional virtual Ethernet adapter to use as the control channel for shared Ethernet adapter failover:

   a. Make the adapter ID 12 and the VLAN 99, leaving all other fields as they are, to create the control channel virtual Ethernet adapter.

   b. Click **OK** to return to the virtual Ethernet adapter main window.

6. Review the virtual Ethernet adapters that are defined, and click **Next** to save the settings and move on to the Virtual Storage Adapters window.

Virtual storage

Here we show an example of creating a virtual SCSI adapter for the VIOS virtual server. When creating a virtual Fibre Channel adapter, the same windows shown in "Virtual Ethernet" on page 296 are shown. However, change the Adapter type field to **Fibre Channel**.

Complete the following steps:

1. Click **Create adapter...** to start the wizard (Figure 7-14).

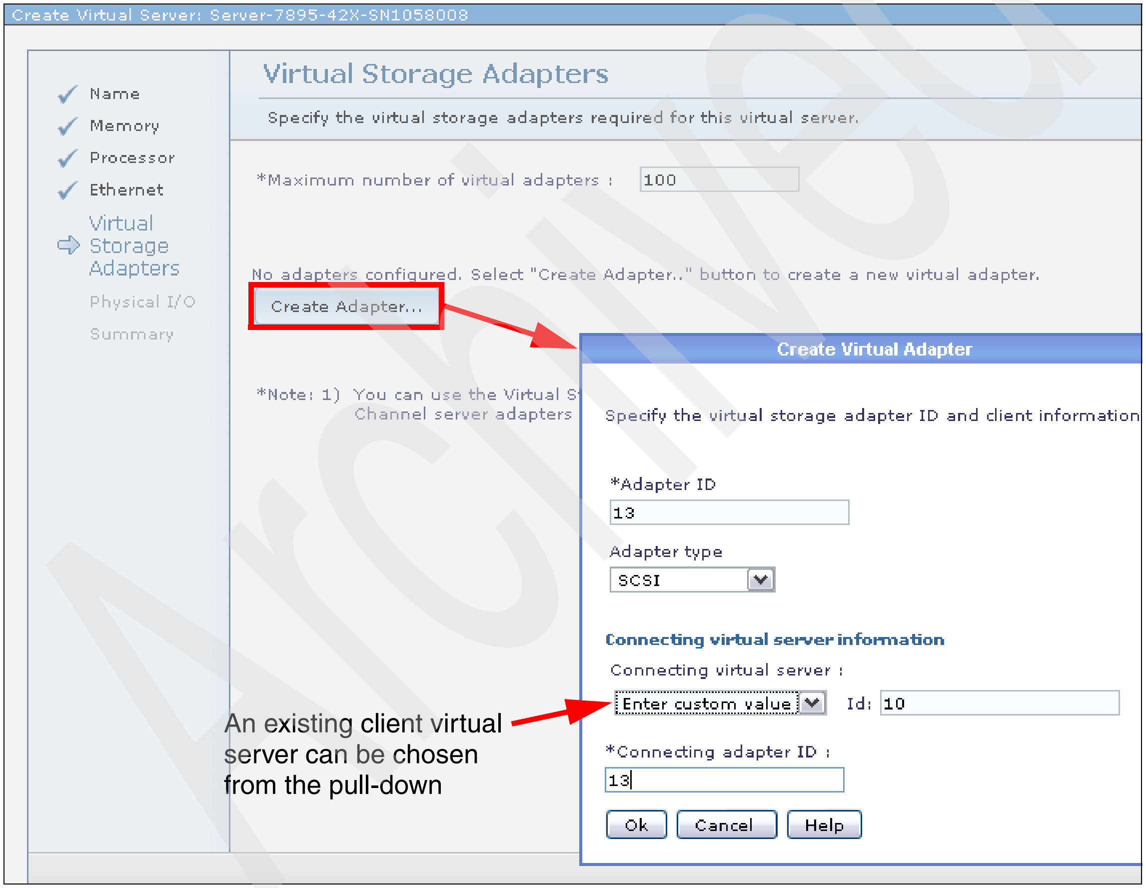

Figure 7-14 Create a virtual SCSI adapter on VIOS

2. Complete the fields in Figure 7-14 on page 300 as follows:
   – Specify 13 as the **Adapter ID**.
   – To create a virtual SCSI relationship between this VIOS and a client virtual server, specify SCSI as the **Adapter type**. Either choose an existing virtual server and supply an ID in the **Connecting adapter ID** field, or enter a new ID and connecting adapter ID for a virtual server that is not defined.

     Figure 7-14 on page 300 shows the window for creating a virtual SCSI adapter between this VIOS and a client virtual server with an ID of 10 and a connection adapter ID of 13.

   > **Note:** The number of virtual adapters allowed on the virtual server can be set in this window. Set it to one more than the highest ID number that you plan to assign. If you do not set it correctly, it automatically increases, if necessary, when assigning ID numbers to virtual adapters that exceed the current setting. This value cannot be changed dynamically after a virtual server is activated.

3. Click **OK** to save the settings for this virtual storage adapter, and return to the main virtual storage adapter window.
4. When all virtual storage adapters are defined, click **Next** in that window to save the settings and proceed to the physical adapters window (Figure 7-17 on page 303).

Assigning physical I/O

Any virtual server can be assigned physical I/O adapters from one of three sources:

- Expansion cards
  – p260 and p24L: Maximum of two
  – p460: Maximum of four
- Storage controllers
  – p260 and p24L: Maximum of one
  – p460: Maximum of one
- USB

Identifying the I/O resource in the FSM configuration menus is necessary to assigning the correct physical resources to the intended virtual servers.

Figure 7-15 shows the physical location codes on a p460. The locations codes shown in the configuration menus contain a prefix as follows:

```
Utttt.mmm.ssssss, where tttt:Machine Type, mmm:Model, ssssss: 7-digit
Serial Number
```

For example, an EN4054 4-port 10Gb Ethernet Adapter in a p460 is represented as:

```
U78AF.001.ssssss-P1-C34
```

An FC3172 2-port 8Gb FC Adapter is represented as:

```
U78AF.001-ssssss-P1-C36
```

Ports: Keep in mind that the four ports on the EN4054 4-port 10Gb Ethernet Adapter are on two busses, so you can assign two ports to one partition independent of the other two ports. The location code has a suffix of L1 or L2 to distinguish between the two pairs of ports.

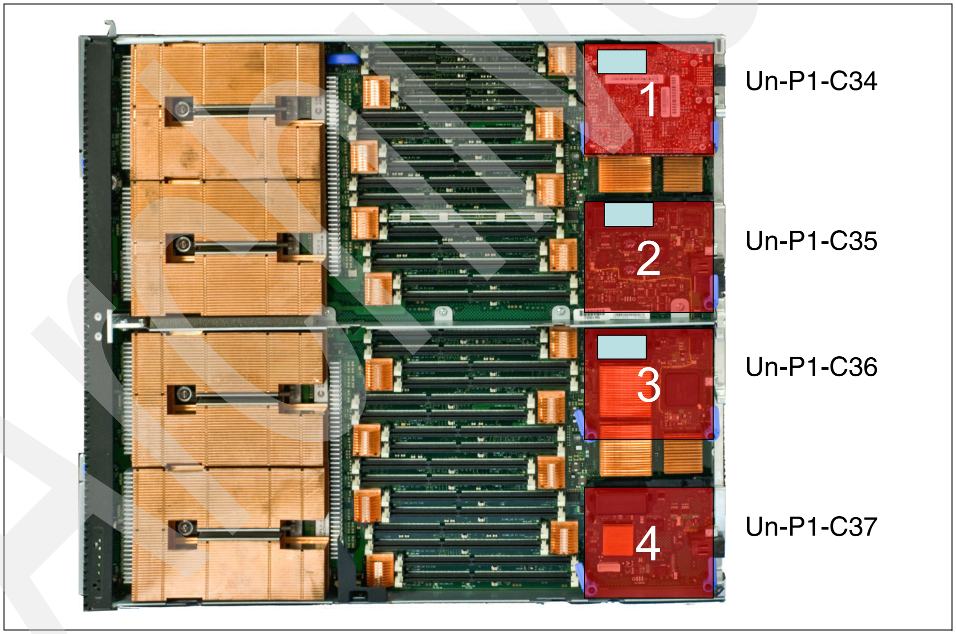

Figure 7-15 p460 adapter location codes

Figure 7-16 shows the expansion card location codes for the p260.

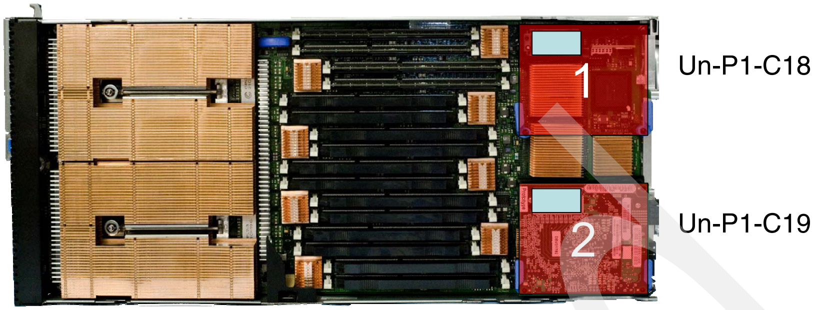

Figure 7-16   p260 adapter location codes

The storage controller, if disks were ordered, has a location code of P1-T2 on both models. The USB controller has a location code of P1-T1 on both models.

For our VIOS, we assign all four ports on an Ethernet expansion card and the storage controller.

Complete the following steps:

1. Choose the expansion card and storage controller from the list in Figure 7-17.

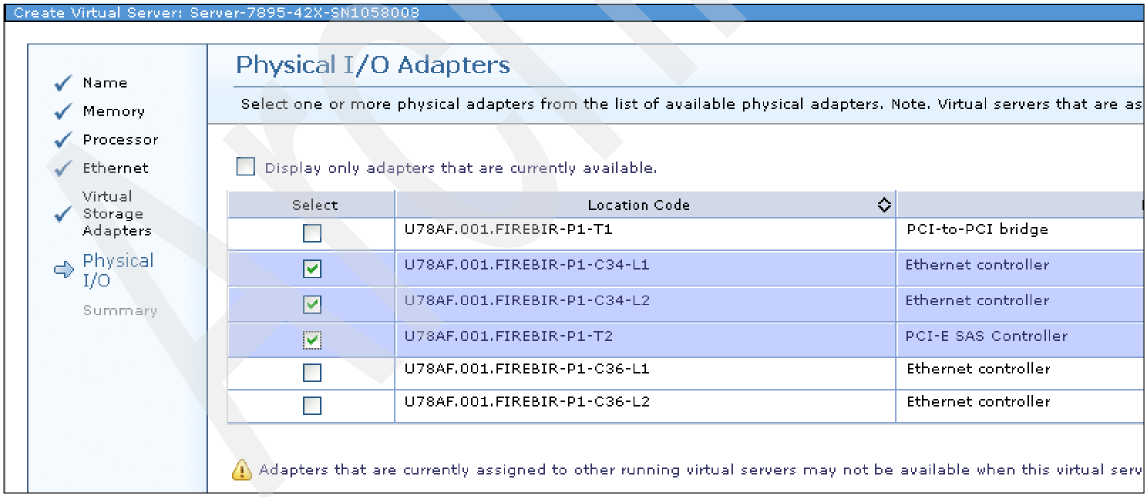

Figure 7-17   Physical adapter selections on VIOS virtual server

2. Click **Next** to proceed to the Summary window.

   Review the summary to ensure that the VIOS virtual server is created as you expect. If you need to make corrections, go back to the section where the correction must be made and change the option.

3. Click **Finish** to complete the definition of the VIOS virtual server.

4. To verify that the virtual server was defined, return to the Power Systems Management tab and click the **Virtual I/O Server** link under **Virtual Servers** in the **Manage Resources** section. You see vios1 (or whatever you named your virtual server) in the list.

7.3 Modifying the VIOS definition

Some of the values chosen by the virtual server creation wizard might not be what you want for the VIOS. In this section, we describe changing the VIOS using the FSM web interface and then making the same changes using the CLI.

7.3.1 Using the IBM Flex System Manager

To change the values using the web interface, complete the following steps:

1. Select the newly created VIOS and click **Actions** → **System Configuration** → **Manage Profiles**, as shown in Figure 7-18.

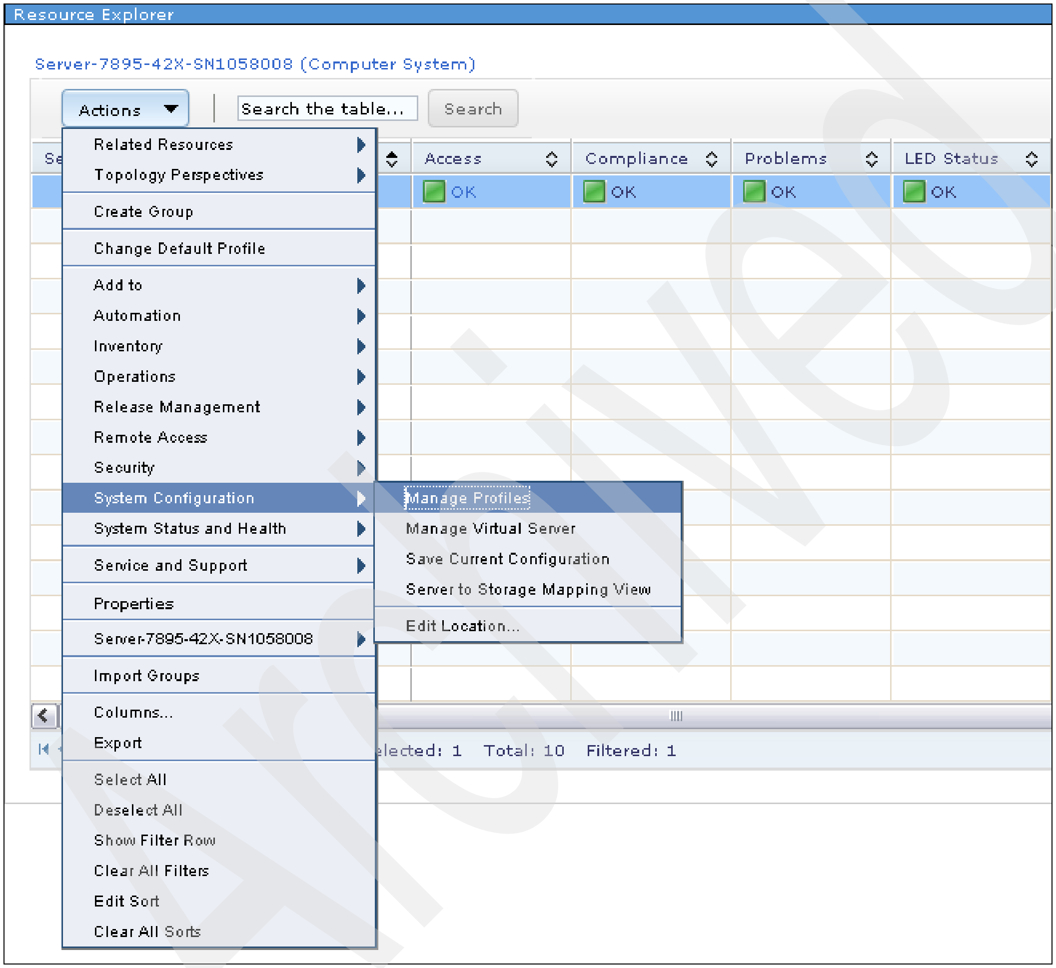

Figure 7-18 Manage VIOS profiles to change settings

A window opens and shows all of the profiles for the selected virtual server.

2. Select the profile to edit and click **Actions** → **Edit**.

3. Click the **Processors** tab to access the processor settings that were made by the wizard. The window shown in Figure 7-19 opens. Options can be changed in this window to the values planned for the VIOS virtual server.

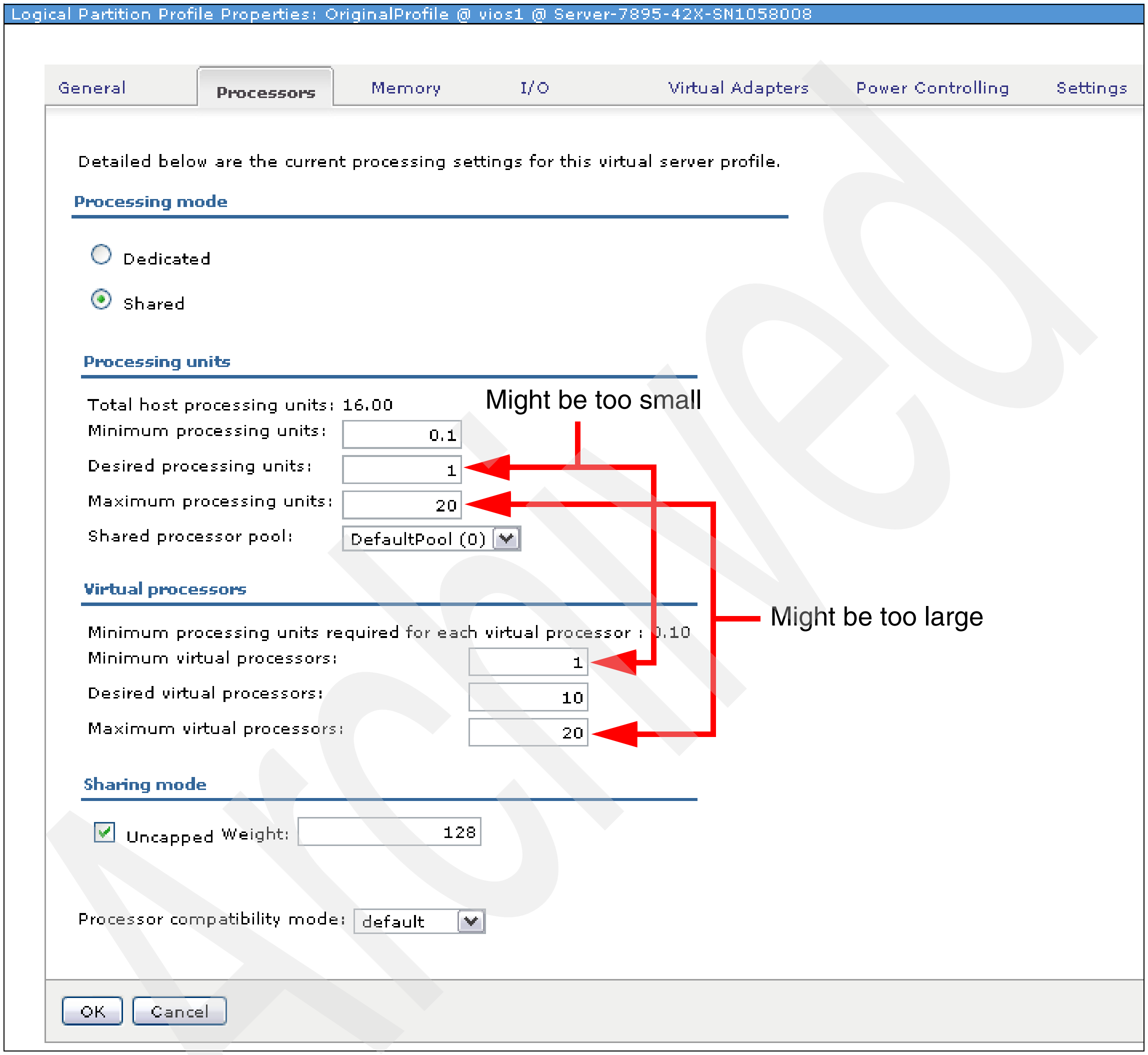

Figure 7-19 VIOS profile - changing processor settings

Note the values that were made by the wizard:

– The desired virtual processor count is 10 (as specified when creating the virtual server). This count translates to a desired processing unit setting of 1.0.

– The maximum virtual processor count is 20. The maximum count is always the desired count plus 10. The maximum processing units setting is also set to 20.

– The minimum virtual processors setting is set to 1, with the processing units set to 1.

– The sharing mode is set to uncapped with a weight of 128. Verify that this setting is what is acceptable for the virtual server.

> **Maximum virtual processors:** Based on the observed default values from the virtual server creation wizard, the maximum virtual processors probably need to be reduced in line with the number of processors in the pool if uncapped, or the processing unit value (rounded up) if capped. In addition, the minimum values are probably too small and need to be raised to a more realistic value for the use of the virtual server.

4. Similar observations and modifications can be made regarding the memory settings by using the **Memory** tab in the profile window. The default minimum memory is 256 MB. Increase this memory for an AIX virtual server.

7.3.2 Using the CLI

To change the values for the VIOS using the CLI, run the following command:

```
smcli chsyscfg -r prof -m p4601 -i
“name=OriginalProfile,lpar_name=vios1,min_proc_units=1.0,max_proc_units
=2.0,desired_procs=4,max_procs=16,uncap_weight=200,min_mem=512”
```

Success is indicated when the prompt returns and no error message is displayed.

Run **smcli lssyscfg** to confirm that the changes occurred.

7.4 Creating an AIX or Linux virtual server

Creating an AIX or Linux virtual server is similar to creating a VIOS virtual server.

Use the same process shown in 7.2, “Creating the VIOS virtual server” on page 286, but with some differences. The differences between creating a VIOS and an AIX or Linux virtual server are:

- The Environment option in the initial window is set to **AIX/Linux**.
- No physical I/O adapters must be defined, if the virtual server is virtualized. In this case, a VIOS must be defined to provide virtualized access to network and storage.
- The virtual server might use all physical resources, running as a full system partition.
- The virtual server can be defined as suspend capable.

7.5 Creating an IBM i virtual server

You can install the IBM i operating system in a client virtual server of a VIOS. Begin by completing the steps in 7.2, “Creating the VIOS virtual server” on page 286 to create the VIOS.

For more details about installing IBM i in a virtual server, see *IBM i on a POWER Blade Read-me First*, found at:

http://www.ibm.com/systems/i/advantages/v6r1/blades/pdf/i_on_blade_readme.pdf

Creating the virtual server for an IBM i installation is similar to the process for creating a VIOS. Complete the following steps:

1. Set the Environment option to **IBM i**, as shown in Figure 7-20.

Figure 7-20 Create an IBM i virtual server

2. Click **Next** to go to the Memory settings. The window shown in Figure 7-21 opens.

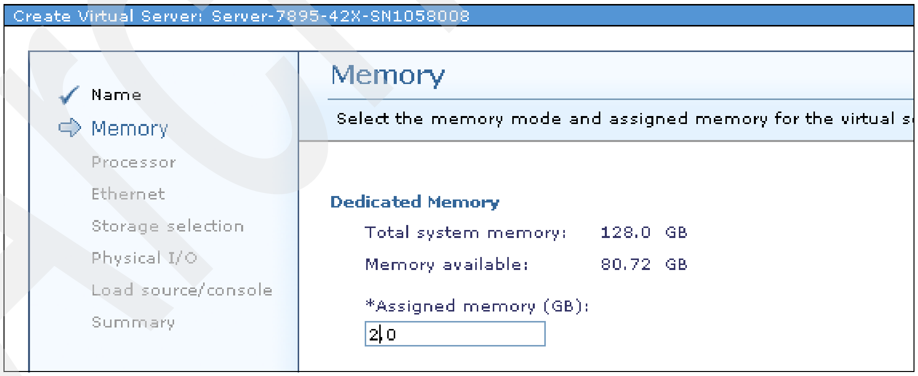

Figure 7-21 IBM i virtual server memory

3. Specify the wanted quantity of memory. Click **Next** to go to the processor settings. The window shown in Figure 7-22 opens.

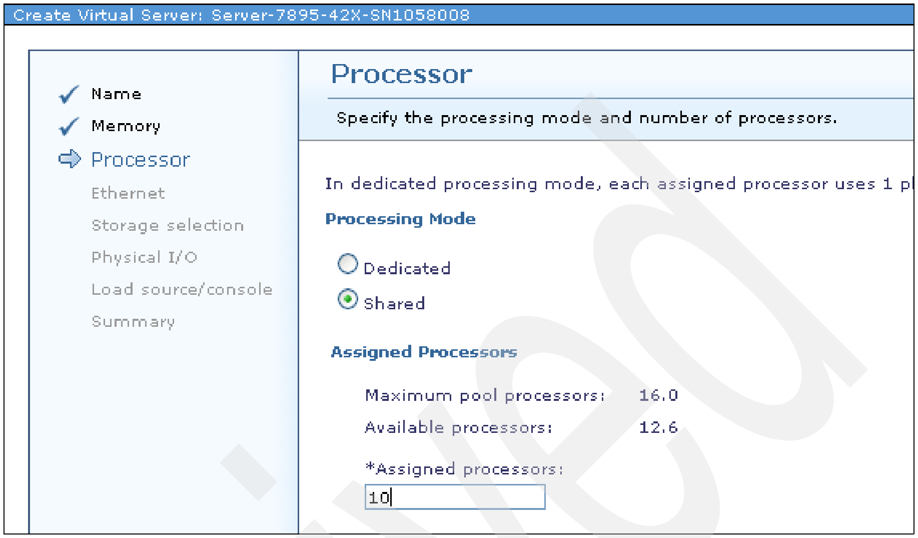

Figure 7-22 IBM i virtual server processor settings

4. Choose a quantity of processors for the virtual server and click **Next** to create the virtual Ethernet adapters. The window shown in Figure 7-23 opens.

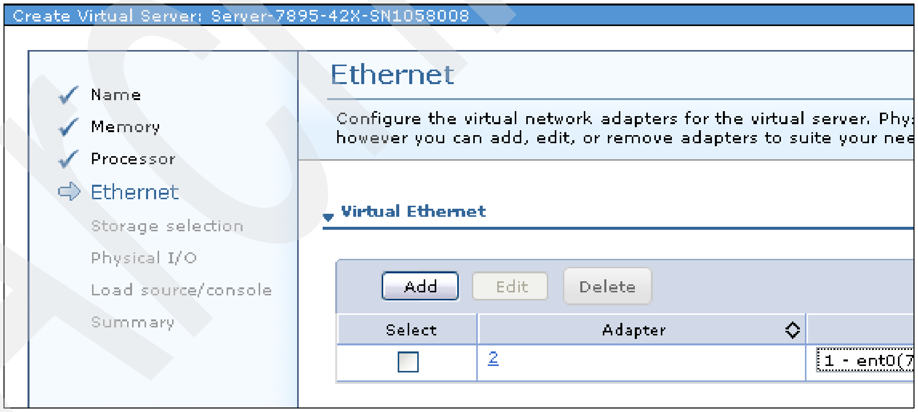

Figure 7-23 IBM i virtual server settings for virtual Ethernet

With the VIOS already defined, the FSM defines a virtual Ethernet on the same VLAN as the SEA on the VIOS. We keep that definition, as shown in Figure 7-23.

Important: These steps are critical, because the IBM i virtual server must be defined to use only virtual resources through a VIOS. At the least, a virtual Ethernet and a virtual SCSI adapter must be defined in the IBM i virtual server.

5. Click **Next** to proceed to the Virtual Storage definitions, as shown in Figure 7-24.

Figure 7-24 IBM i virtual server manual virtual storage definition

6. Indicate that you do not want automatic virtual storage definition (configure the adapters manually), and click **Next** to proceed to the main Virtual Storage window.

   Because no virtual storage adapters exist, the **Create Adapter** option is displayed, as shown in Figure 7-25. If virtual storage adapters are already created, they are shown.

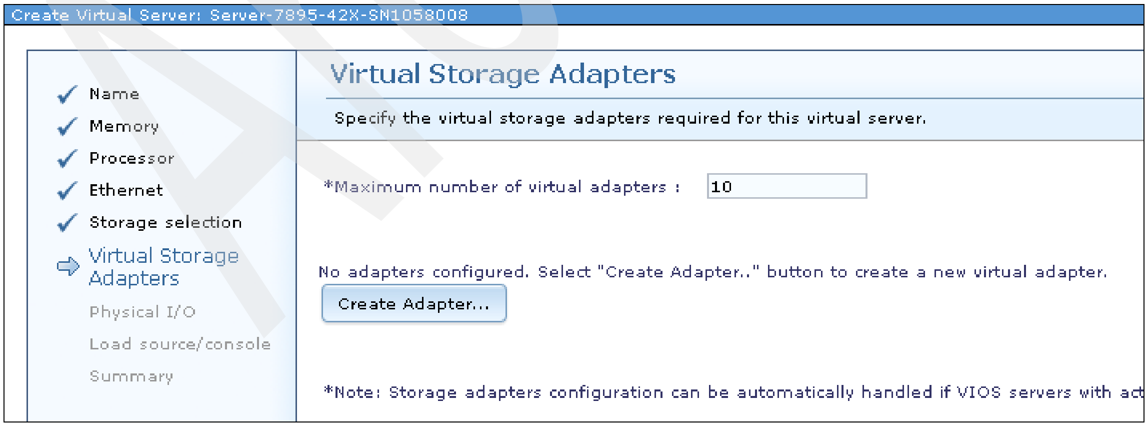

Figure 7-25 IBM I virtual server create virtual storage adapter

7. Click **Create Adapter**. The window shown in Figure 7-26 opens.

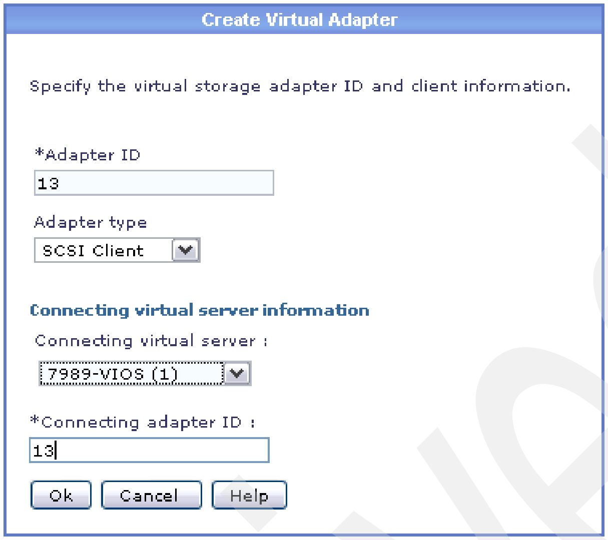

Figure 7-26 IBM i virtual server - create virtual SCSI adapter

8. Complete the fields in this window as follows:
   - Choose an adapter ID.
   - Specify **SCSI Client** for the adapter type.
   - Specify a virtual SCSI adapter on the VIOS as the Connecting virtual server.

9. Click **OK** to create this virtual SCSI adapter and return to the main Virtual Storage adapter window, as shown in Figure 7-27.

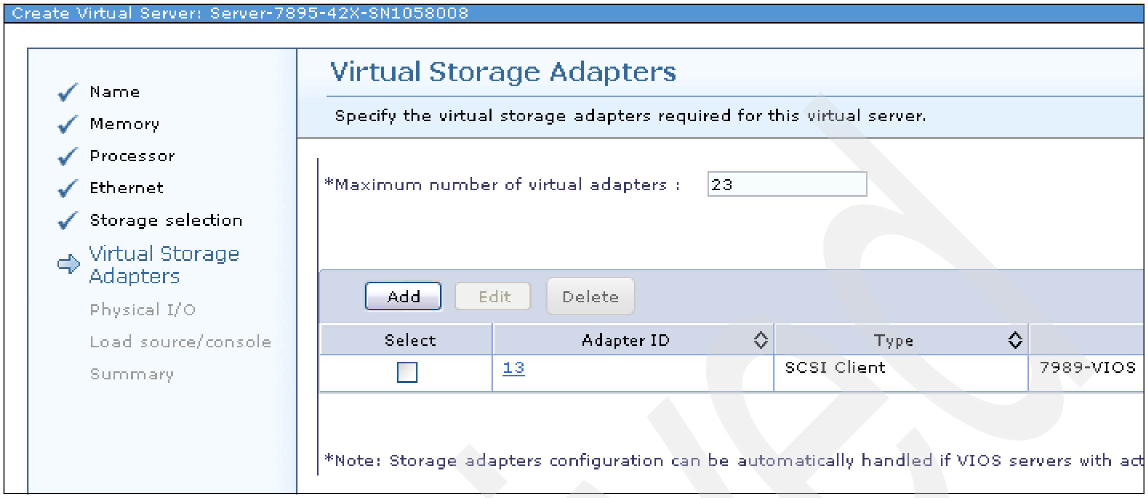

Figure 7-27 IBM i virtual server settings for virtual SCSI adapter

10. This adapter is the only virtual SCSI adapter we create, so click **Next** to proceed to the physical adapter settings, as shown in Figure 7-28.

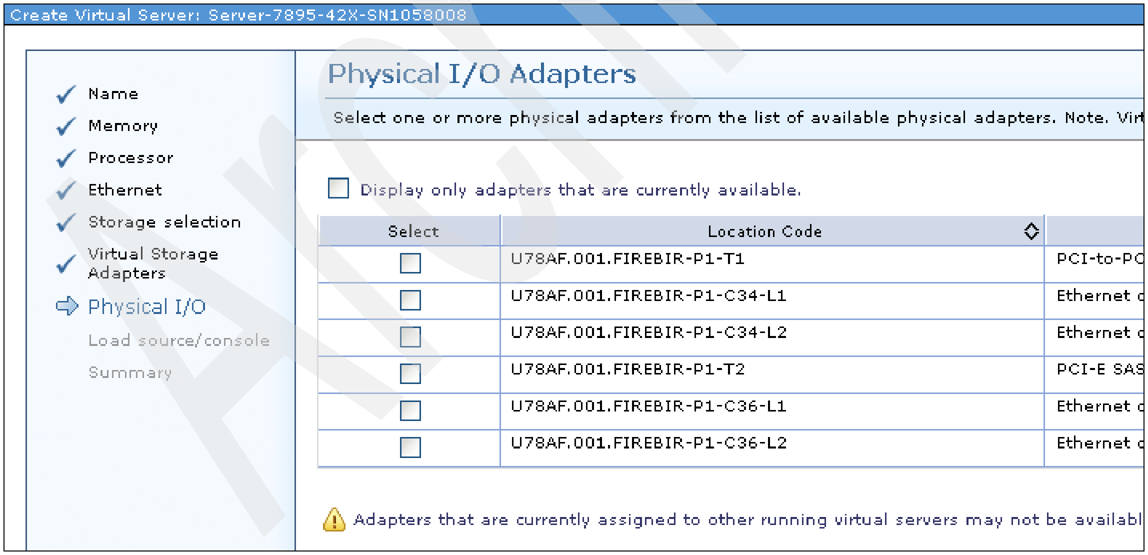

Figure 7-28 IBM i virtual server physical adapter settings

Important: Do not forget to configure the virtual SCSI server adapter on the VIOS that this virtual SCSI client adapter refers to. In addition, disks must be provisioned to the virtual SCSI server adapter in the VIOS to be used by the IBM i virtual server (operating system and data).

To use a virtual optical drive from the VIOS for the IBM i operating system installation, the installation media ISO files must be copied to the VIOS, and the virtual optical devices must be created.

11. Do not select physical adapters for IBM i virtual servers, as shown in Figure 7-28 on page 313. Click **Next** in this window to proceed to the Load Source and Console settings, as shown in Figure 7-29.

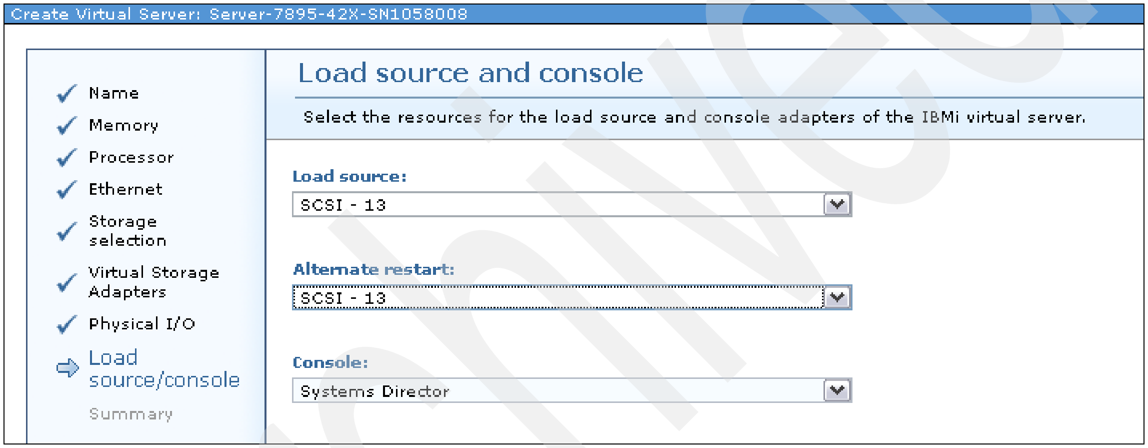

Figure 7-29 IBM i virtual server load source and console settings

12. Choose the virtual SCSI as the Load Source. Click **Next** to proceed to the Summary.
13. Review the settings on the Summary page, and click **Finish** to complete the definition.

The IBM i virtual server is now ready to be activated for load.

7.6 Preparing for a native operating system installation

If you need the entire capacity of the Power Systems compute node, an operating system can be installed natively on the node. The configuration is similar to the setup for a partitioned node, but all of the resources are assigned to a single virtual server.

The operating system can then be installed to that single virtual server, using the methods described in Chapter 8, “Operating system installation” on page 317.

7.6.1 Creating a full node server

The process to create a full node server (also known as a full system partition) is similar to the process described in “Creating the VIOS virtual server using the GUI” on page 292. Complete the steps in that section to reach the point shown in Figure 7-6 on page 293, and the window shown in Figure 7-30 opens.

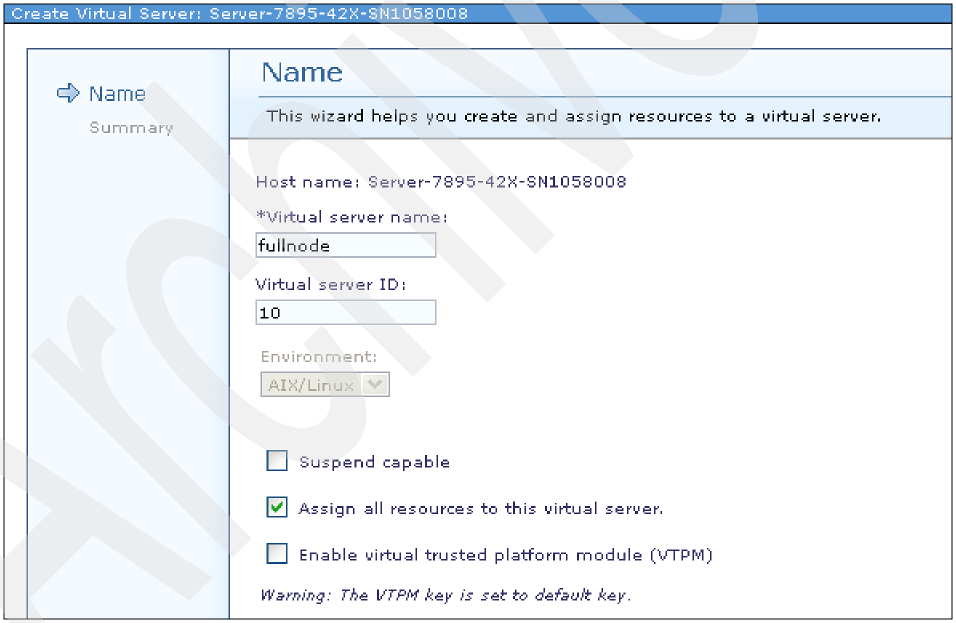

Figure 7-30 Assigning all resources to a full node server

Complete the following steps:

1. Complete the fields shown in Figure 7-30 as follows:
   – For **Virtual server name**, assign a node a name, such as fullnode.
   – Provide a **Virtual server ID**, for example, 10.

– Set the **Environment** to **AIX/Linux**.

– Select **Assign all resources to this virtual server**. This is the key selection.

2. Click **Next**. All the resources are assigned to this virtual server. The Summary window opens, as shown in Figure 7-31.

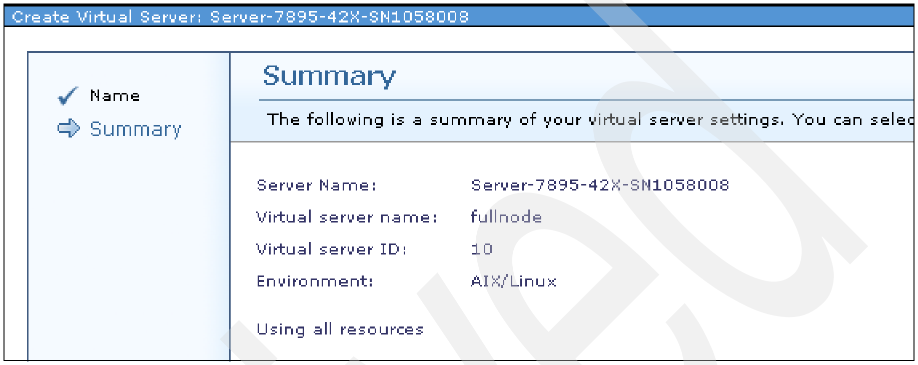

Figure 7-31 Summary window when creating full node server

3. Click **Finish** to complete the creation of the single partition.

8

Operating system installation

In this chapter, we describe how to update firmware and install various operating systems on the compute node.

We cover the following topics in this chapter:

- Firmware updates
- Methods to install operating systems
- Installation procedures
- Installing AIX
- Installing Red Hat Enterprise Linux
- Installing SUSE Linux Enterprise Server
- Installing IBM i

8.1 Firmware updates

IBM periodically makes firmware updates available for the compute node, the management module, or expansion cards in the compute node. In a compute node or chassis environment, there are multiple points to consider when considering firmware updates. In some cases, the chassis and infrastructure components can be updated concurrently, without disrupting virtual server operations.

In this chapter, we describe methods for updating the Power Systems compute nodes.

The simpler way to do a firmware update of a POWER processor-based compute node requires that a supported operating system be running on the compute node. Live Partition Mobility (LPM) can be used to avoid the disruptive nature of firmware updates. If the node is managed by IBM Flex System Manager, the firmware update can also be concurrent.

Firmware updates can provide fixes to previous versions and can enable new functions. Compute node firmware typically has a prerequisite CMM firmware level. It is best to have a program in place for reviewing the current firmware levels of the chassis components and compute nodes to ensure the best availability.

Firmware updates are available at IBM Fix Central web page at the following website:

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

8.1.1 Firmware update methods

Important: To avoid problems and to maintain proper system performance, always verify that the compute node, service processor, and diagnostic firmware levels are consistent for all compute nodes within the IBM Flex System Enterprise Chassis. For more information, see 8.1.4, “Verifying the system firmware levels” on page 325.

The firmware of the compute node can be updated in various ways:

- The IBM Flex System Manager acquires, installs, and manages firmware and device driver updates, and monitors your compute nodes to ensure that they remain current.

Figure 8-1 shows the update firmware menu in IBM Flex System Manager. Firmware updates done using IBM Flex System Manager can be non-destructive or concurrent with respect to server operations, so that a server reboot is not required. Only updates within a release can be, but are not guaranteed to be, concurrent.

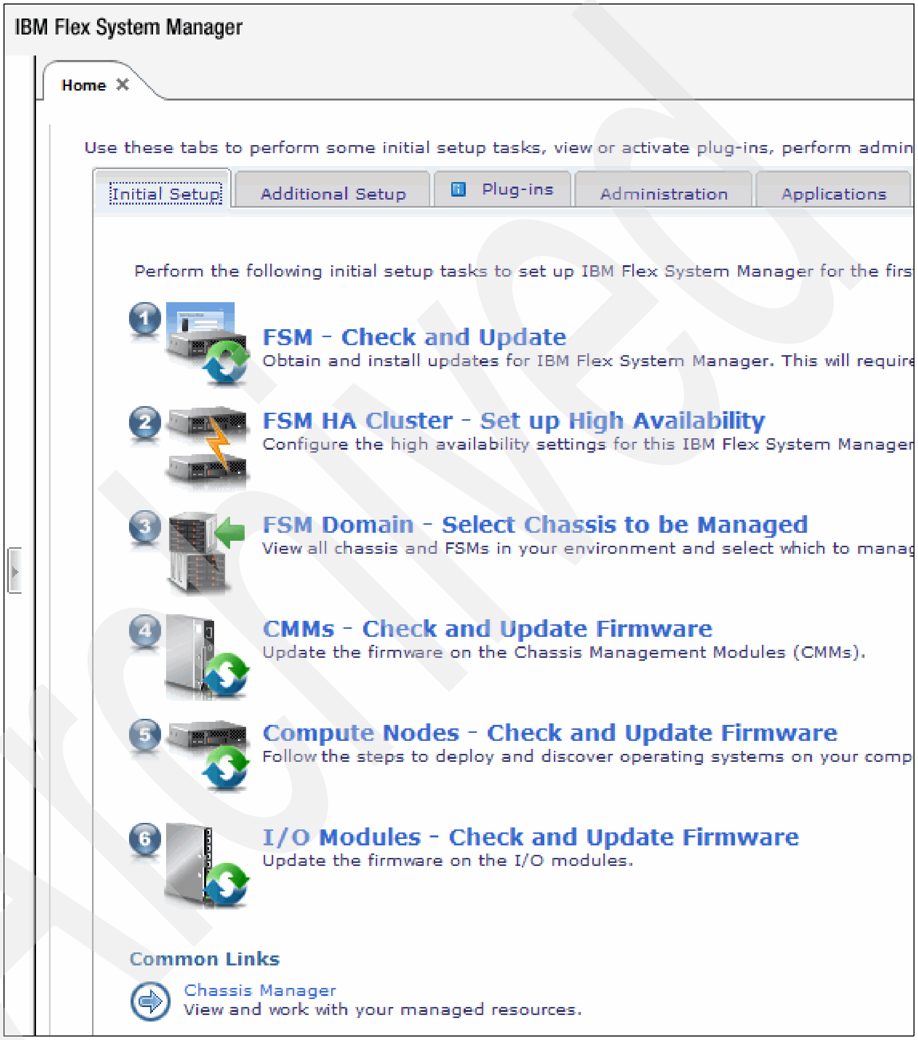

Figure 8-1 The IBM Flex System Manager compute nodes check and update firmware

- By using In-band operating system capabilities, such as the **`update_flash`** command for Linux and AIX or the **`ldfware`** command for Virtual I/O Server.

- By using the firmware update function of AIX diagnostic tests.
- The firmware update function of the stand-alone diagnostics boot image.

> **Installation of firmware:** Before the installation of the new firmware to the temporary side (firmware backup area) begins, the contents of the temporary side are copied to the permanent side. After firmware installation begins, the previous level of firmware on the permanent side is no longer available.
>
> Firmware updates can take time to load. To expedite the initial setup process, you can install your operating system while you wait for firmware updates.

In the following sections, we explain several methods.

8.1.2 Firmware in-band installation

When the compute node operating system is already installed and it is not controlled by IBM Flex System Manager, the most common way to update firmware is through the in-band method. This method is easier than the others, but it is always disruptive and always requires a reboot of the compute node.

To use this method, complete the following steps:

1. Go to the IBM Support website at:

   http://ibm.com/systems/support/

2. Select your system by specifying the machine type and model and click **Go**.
3. Click the **Download** tab, if necessary, for device driver and firmware updates.
4. Download the appropriate firmware update files to the `/tmp/fwupdate` directory of the compute node you want update.
5. Log in to the AIX or Linux system as root, or log in to the Virtual I/O Server partition as `padmin`.
6. Run the following command to identify the name of the firmware:

   ```
   ls /tmp/fwupdate
   ```

   The result of this command is a list of any firmware updates that you downloaded to the directory, for example:

   ```
   01AA7xx_yyy_zzz
   ```

7. Install the firmware update using one of the following methods:
   - Install the firmware with the in-band diagnostic tests of your AIX system by running the following command:

     ```
     diag
     ```

- Install the firmware by running **update_flash** (on AIX):

```
cd /tmp/fwupdate
/usr/lpp/diagnostics/bin/update_flash -f 01EA3xx_yyy_zzz
```

- Install the firmware by running **update_flash** (on Linux):

```
cd /tmp/fwupdate
/usr/sbin/update_flash -f 01EA3xx_yyy_zzz
```

- Install the firmware by running **ldfware** (on VIOS):

```
cd /tmp/fwupdate
ldfware -file 01EA3xx_yyy_zzz
```

8. Verify that the update installed correctly, as described in 8.1.4, “Verifying the system firmware levels” on page 325.

8.1.3 IBM system diagnostics installation

When the compute node operating system is not installed, you can update the firmware using a system diagnostics CD. The IBM stand-alone diagnostics CD for Power Systems servers is at:

http://www14.software.ibm.com/webapp/set2/sas/f/diags/download/home.html

To update the firmware, complete the following steps:

1. Run the diagnostics CD. You can:
   - Burn the ISO image to a CD and insert the CD in a CD/DVD drive
   - Use the ISO image connected to a virtual optical device
   - Boot the CD from a Network Installation Manager (NIM) server

Figure 8-2 shows the diagnostic post-boot system console definition.

```
 ******* 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.
Escriviu 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.
```

Figure 8-2 Diagnostic console definition

2. Accept the copyright notice, and then choose the task selection menu entry shown in Figure 8-3.

```
FUNCTION SELECTION

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

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



       To make a selection, type the number and press Enter.  [3 ]
```

Figure 8-3 Function selection

3. Select Microcode Tasks, as shown in Figure 8-4.

```
TASKS SELECTION LIST                                                  801004


From the list below, select a task by moving the cursor to
the task and pressing 'Enter'.
To list the resources for the task highlighted, press 'List'.

[MORE...7]
  Delete Resource from Resource List
  Display Configuration and Resource List
  Display Firmware Device Node Information
  Display Hardware Error Report
  Display Hardware Vital Product Data
  Display Multipath I/O (MPIO) Device Configuration
  Display Resource Attributes
  Display Service Hints
  Display or Change Bootlist
  Hot Plug Task
  Microcode Tasks
  Process Supplemental Media
[MORE...1]

F1=Help            F4=List              Esc+0=Exit          Enter
F3=Previous Menu
```

Figure 8-4 Microcode task selection

4. Select Download Latest Available Microcode, as shown in Figure 8-5.

```
Microcode Tasks                                                          801004

Move cursor to desired item and press Enter.

  Display Microcode Level
  Download Latest Available Microcode
  Generic Microcode Download















F1=Help             F4=List             Esc+0=Exit          Enter
F3=Previous Menu
```

Figure 8-5 Download (install) the latest available microcode

5. Insert the CD-ROM with the microcode image, or select the virtual optical device that points to the microcode image. If the system is booted from a NIM server, the microcode must be in `usr/lib/microcode` of the Shared Product Object Tree (SPOT) the client is booted from.

8.1.4 Verifying the system firmware levels

The diagnostics program displays the current system firmware levels for the TEMP and PERM images. This function also displays which image the compute node used to start.

To verify the system firmware levels, complete the following steps:

1. Start the in-band diagnostics program by running the following command:

   ```
   diag
   ```

2. From the Function Selection menu, select Task Selection and press Enter, as shown in Figure 8-3 on page 323.
3. From the Tasks Selection List menu, select Microcode Tasks → Display Microcode Level and press Enter, as shown in Figure 8-6.

```
Microcode Tasks                                                    801004

Move cursor to desired item and press Enter.

  Display Microcode Level
  Download Latest Available Microcode
  Generic Microcode Download








F1=Help             F4=List             Esc+0=Exit          Enter
F3=Previous Menu
```

Figure 8-6 Display microcode level

4. Select the system object `sys0` and press F7 to commit, as shown in Figure 8-7.

```
RESOURCE SELECTION LIST                                                  801006

From the list below, select any number of resources by moving
the cursor to the resource and pressing 'Enter'.
To cancel the selection, press 'Enter' again.
To list the supported tasks for the resource highlighted, press 'List'.

Once all selections have been made, press 'Commit'.
To avoid selecting a resource, press 'Previous Menu'.


  All Resources
      This selection will select all the resources currently displayed.
+ sys0                                      System Object





F1=Help             F4=List             F7=Commit           Esc+0=Exit
F3=Previous Menu
```

Figure 8-7 sys0 selection and commit

The Display Microcode Level menu opens. The top of the window shows the system firmware level for the permanent and temporary images and the image that the compute node used to start (Figure 8-8).

```
DISPLAY MICROCODE LEVEL                                                        802811
IBM,7895-42X

The current permanent system firmware image is AF740_051
The current temporary system firmware image is AF740_051
The system is currently booted from the temporary firmware image.

Use Enter to continue.







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

Figure 8-8 Showing the microcode level

8.1.5 Firmware update using IBM Flex System Manager

We can acquire, install, and manage firmware updates directly from the IBM support web page if the compute node is managed by IBM Flex System Manager and the FSM is connected to the Internet. In many cases, the firmware update is concurrent.

Figure 8-9 shows window where you select the type of update to search for, download, and apply. For this procedure, we use Power System Firmware.

Complete the following steps:

1. Select **PowerIOFW** from the list of available update types, as shown in Figure 8-9.

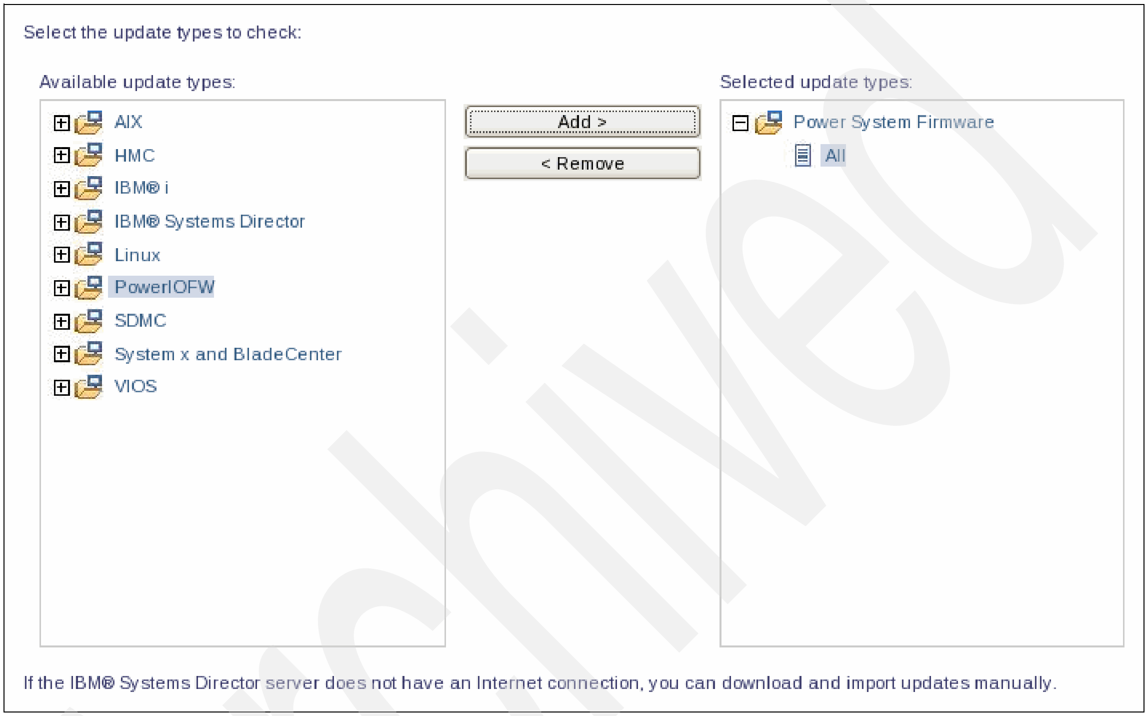

Figure 8-9 Select the type of update

2. Click **Add** to add your selection to the list of selected update types.

3. Select **Power System Firmware** from the list of selected update types. Figure 8-10 shows the firmware update that is ready to install.

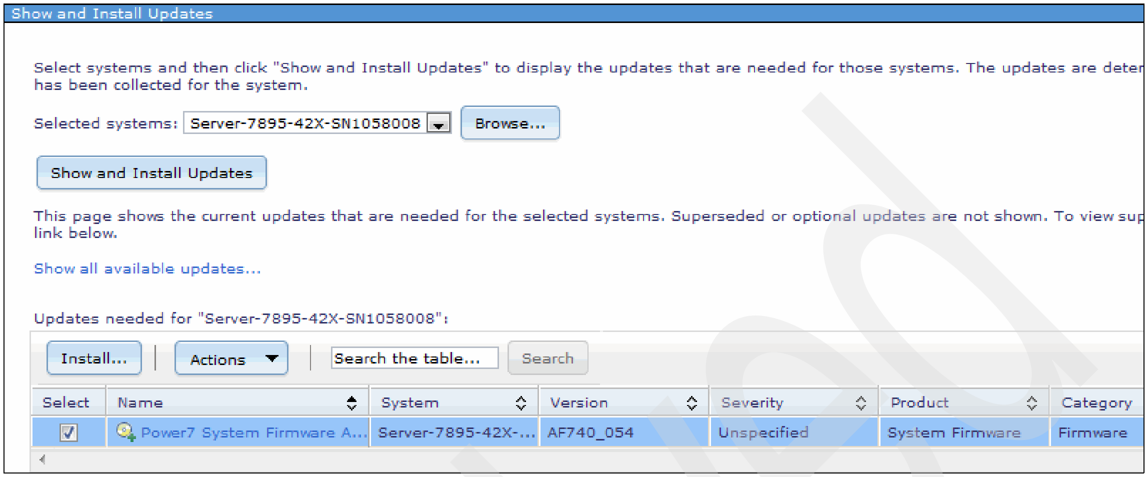

Figure 8-10 List of selected firmware

4. Review and confirm the list of updates that will be installed on the selected systems, as shown in Figure 8-11. After you confirm this list, the update begins and is concurrent (the system does not require a restart to activate the new firmware).

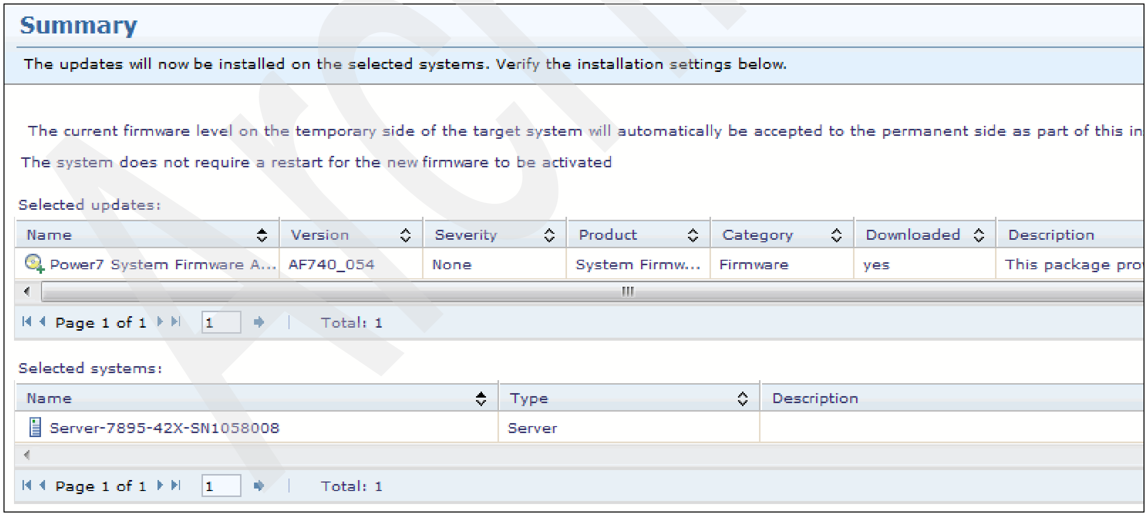

Figure 8-11 Firmware update wizard

5. If necessary, review the installation log, as shown in Figure 8-12, to determine the status of the installation.

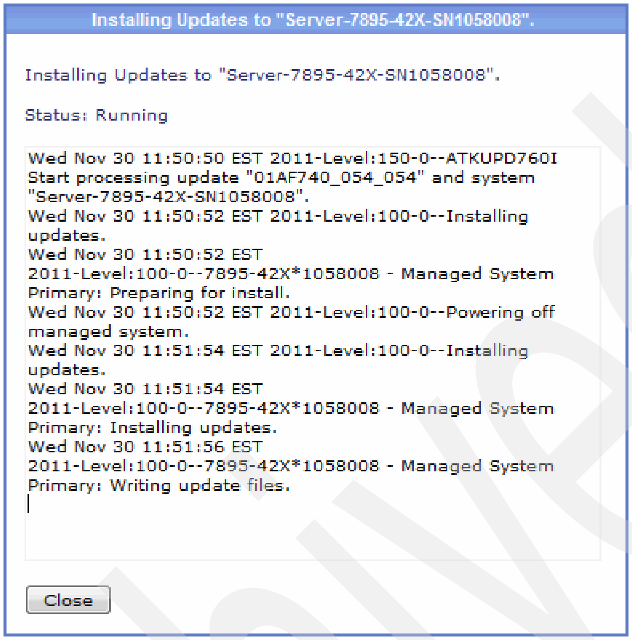

Figure 8-12 Log for installation update

8.2 Methods to install operating systems

The Power Systems compute node provides several methods for installing and deploying your operating system images.

We cover the following methods in this section:

- NIM installation
- Optical media installation
- TFTP network installation
- Cloning methods

Installation method compatibility among operating systems is shown in Table 8-1.

Table 8-1 Installation methods - compatibility among operating systems

Installation method	AIX	VIOS	Red Hat Enterprise Linux	SUSE Linux Enterprise Server	IBM i
Optical	Yes	Yes	Yes	Yes	Yes
NIM	Yes	Yes	Yes	Yes	No
TFTP or BOOTP	No	No	Yes	Yes	No
Tape Backup	Yes	Yes	Yes	Yes	Yes

8.2.1 NIM installation

The NIM installation method is the one that is used the most in a Power Systems environment. You can use NIM to install your servers and back up, restore, upgrade software, and to perform maintenance tasks on them.

More information about the NIM installation is *NIM from A to Z in AIX 5L*, SG24-7296.

To perform a NIM installation, complete the following steps:

1. Set up a Domain Name Server (DNS) or include the machine you are about to install in the `/etc/hosts` file of your AIX NIM server.
2. Now that the NIM server has your machine IP address, create the machine in the NIM environment by running the following command:

   ```
   smit nim_mkmac
   ```

3. In the next window, respond to the prompt for a machine name and the type of network connectivity you are using. The system populates the remaining fields and displays the screen shown in Figure 8-13.

```
 Define a Machine

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

[TOP]                                                 [Entry Fields]
* NIM Machine Name                                  [7989AIXtest]
* Machine Type                                      [standalone]            +
* Hardware Platform Type                            [chrp]                  +
  Kernel to use for Network Boot                    [64]                    +
  Communication Protocol used by client             []                      +
  Primary Network Install Interface
*   Cable Type                                       bnc                    +
    Network Speed Setting                           []                      +
    Network Duplex Setting                          []                      +
*   NIM Network                                     [ent-Network1]
*     Network Type                                   ent
*     Ethernet Type                                  Standard               +
*     Subnetmask                                    []
*     Default Gateway Used by Machine               [9.27.20.1]
*     Default Gateway Used by Master                [9.27.20.241.1]
*   Host Name                                        7989AIXtest
    Network Adapter Hardware Address                [0]
    Network Adapter Logical Device Name             []
  IPL ROM Emulation Device                          []                      +/
  CPU Id                                            []
  Machine Group                                     []                      +

  Managing System Information
  WPAR Options
    Managing System                                 []
    -OR-
    LPAR Options
      Identity                                      []
      Management Source                             []                      +

[MORE...1]

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

Figure 8-13 Adding a machine to the NIM environment

4. In the screen shown in Figure 8-13 on page 333, enter the remainder of the information required for the node.

   There are many options in this window, but you do not need to set them all to set up the installation. Most importantly, set the correct gateway for the machine.

   With your machine created in your NIM server, assign it the resources for the installation. When installing a system from NIM, you need to have additional resources defined, that is, at least one `spot` and one `lpp_source`, or one `spot` and one `mksysb`. These items are defined as follows:

   – `mksysb`: This item is a system image backup that can be recovered on the same or another machine.
   – `spot`: A `spot` is what your system uses from the NIM at boot time. It contains all boot elements for the NIM client machine. Spots can be created from a `mksysb` or from installation media.
   – `lpp_source`: An `lpp_source` is the place where the NIM has the packages for installation. They can be created from installation media and fix packs.

   > **Creating installation resources:** The steps for creating the installation resources are not covered here. Detailed information is in *NIM from A to Z in AIX 5L*, SG24-7296.
   >
   > The **`smit`** fast path for creating resources is **`nim_mkres`**.

5. Assign the installation resources to the machine. For this example, we are doing an RTE installation, so we use `spot` and `lpp_source` for the installation. Run the following command:

```
smit nim_mac_res
```

6. Select Allocate Network Install Resources, as shown in Figure 8-14. A list of available machines opens.

```
 Manage Network Install Resource Allocation

Move cursor to desired item and press Enter.

  List Allocated Network Install Resources
  Allocate Network Install Resources
  Deallocate Network Install Resources

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

Figure 8-14 Select Allocate Network Install Resources

7. Choose the machine you want to install (in this example, we use `7989AIXtest`). A list of the available resources to assign to that machine opens, as shown in Figure 8-15.

```
 Manage Network Install Resource Allocation

Mo+--------------------------------------------------------------------------+
  |                               Target Name                                |
  |                                                                          |
  | Move cursor to desired item and press Enter.                             |
  |                                                                          |
  |   CURSO     groups       mac_group                                       |
  |   master          machines       master                                  |
  |   STUDENT1        machines       standalone                              |
  |   STUDENT2        machines       standalone                              |
  |   STUDENT3        machines       standalone                              |
  |   STUDENT4        machines       standalone                              |
  |   STUDENT5        machines       standalone                              |
  |   STUDENT6        machines       standalone                              |
  |   tws01           machines       standalone                              |
  |   7989nimtest     machines       standalone                              |
  |   7989AIXtest     machines       standalone                              |
  |   bolsilludo      machines       standalone                              |
  |   tricolor        machines       standalone                              |
  |   decano          machines       standalone                              |
  |                                                                          |
  | F1=Help                 F2=Refresh              F3=Cancel                |
  | F8=Image                F10=Exit                Enter=Do                 |
F1| /=Find                  n=Find Next                                      |
F9+--------------------------------------------------------------------------+
```

Figure 8-15 Machine selection for resource allocation

8. Assign both `lpp_source` and `spot`. Press F7 to make multiple selections.

9. Confirm your resource selections by running **smit nim_mac_res** and selecting Select List Allocated Network Install Resources, as shown in Figure 8-16.

```
 Manage Network Install Resource Allocation

Move cursor to desired item and press Enter.

  List Allocated Network Install Resources
  Allocate Network Install Resources
  Deallocate Network Install Resources




  +--------------------------------------------------------------------------+
  |                   Available Network Install Resources                    |
  |                                                                          |
  | Move cursor to desired item and press F7.                                |
  |     ONE OR MORE items can be selected.                                   |
  | Press Enter AFTER making all selections.                                 |
  |                                                                          |
  | > LPP_AIX61_TL04_SP01_REL0944_BOS     lpp_source                         |
  | > SPOT_AIX61_TL04_SP01_REL0944        spot                               |
  |   AIX61_LAST_TL                       lpp_source                         |
  |                                                                          |
  | F1=Help                 F2=Refresh              F3=Cancel                |
  | F7=Select               F8=Image                F10=Exit                 |
F1| Enter=Do                /=Find                  n=Find Next              |
F9+--------------------------------------------------------------------------+
```

Figure 8-16 Resource selection

10. Confirm your resource selections by running **smit nim_mac_res** and selecting List Allocated Network Install Resources.

    Your machine is now created and your resources are assigned.

11. The last step on the NIM side is to launch the installation from the NIM by running **smit nim_mac_op**.

12. Select your machine as you did previously, as shown in Figure 8-15 on page 336.

13. Select the option to perform a Base Operating System (BOS) installation by selecting bos_inst - perform a BOS installation, as shown in Figure 8-17.

```
+-----------------------------------------------------------------------+
 |                           Operation to Perform                        |
 |                                                                        |
 | Move cursor to desired item and press Enter. Use arrow keys to scroll. |
 |                                                                        |
 | [TOP]                                                                  |
 |      diag              = enable a machine to boot a diagnostic image   |
 |      cust              = perform software customization                |
 |      bos_inst          = perform a BOS installation                    |
 |      maint             = perform software maintenance                  |
 |      reset             = reset an object's NIM state                   |
 |      fix_query         = perform queries on installed fixes            |
 |      check             = check the status of a NIM object              |
 |      reboot            = reboot specified machines                     |
 |      maint_boot        = enable a machine to boot in maintenance mode  |
 |      showlog           = display a log in the NIM environment          |
 |      lppchk            = verify installed filesets                     |
 |      restvg            = perform a restvg operation                    |
 | [MORE...6]                                                             |
 |                                                                        |
 | F1=Help                 F2=Refresh               F3=Cancel             |
 | F8=Image                F10=Exit                 Enter=Do              |
 | /=Find                  n=Find Next                                    |
 +-----------------------------------------------------------------------+
```

Figure 8-17 Operation on machine selection

14. Confirm your machine selection and option selection in the next window, and select additional options to further customize your installation, as shown in Figure 8-18.

```
 Perform a Network Install

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

                                                      [Entry Fields]
  Target Name                                       7989AIXtest
  Source for BOS Runtime Files                      rte                    +
  installp Flags                                   [-agX]
  Fileset Names                                    []
  Remain NIM client after install?                  yes                    +
  Initiate Boot Operation on Client?                yes                    +
  Set Boot List if Boot not Initiated on Client?    no                     +
  Force Unattended Installation Enablement?         no                     +
    ACCEPT new license agreements?                 [yes]                   +









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

Figure 8-18 Base Operating System (BOS) installation options

The selection of options on the NIM machine is complete. Next, continue the installation from the Systems Management Services (SMS) menu on the POWER7 based compute node.

15. Reboot the server and, during reboot, press the 1 key to access SMS mode, as shown in Figure 8-19.

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

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


     Memory      Keyboard     Network     SCSI     Speaker
```

Figure 8-19 SMS boot options

16. Select option 1 (SMS Menu) to open the SMS Main Menu, as shown in Figure 8-20.

```
Version AF740_051
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









-------------------------------------------------------------------------------
Navigation Keys:

                                             X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure 8-20 SMS menu options

17. Select option 2 (Setup Remote IPL (Initial Program Load) from the SMS main menu.

18.Select the adapter to use for the installation, as shown in Figure 8-21.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
NIC Adapters
     Device                          Location Code              Hardware
                                                                Address
1.  Interpartition Logical LAN      U7895.42X.1058008-V5-C4-T1  42dbfe361604






-------------------------------------------------------------------------------
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:
```

Figure 8-21 NIC adapter selection

19.Select the IP protocol version (either ipv4 or ipv6), as shown in Figure 8-22. For our example, we select ipv4.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Internet Protocol Version.

1.   IPv4 - Address Format 123.231.111.222
2.   IPv6 - Address Format 1234:5678:90ab:cdef:1234:5678:90ab:cdef








-------------------------------------------------------------------------------
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:
```

Figure 8-22 Internet protocol version selection

20.Select option 1 (BOOTP) as the network service to use for the installation, as shown in Figure 8-23.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Network Service.
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:
```

Figure 8-23 Select a network service

21.Set up your IP address and the IP address of the NIM server for the installation. To do so, select option 1 (IP Parameters), as shown in Figure 8-24.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Network Parameters
Interpartition Logical LAN: U7895.42X.1058008-V5-C4-T1
1.   IP Parameters
2.   Adapter Configuration
3.   Ping Test
4.   Advanced Setup: BOOTP




-------------------------------------------------------------------------------
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
```

Figure 8-24 Network parameters configuration

22. Perform system checks, for example, **ping** or adapter speed, to verify your selections, as shown in Figure 8-25.

```
 Version AF740_051
 SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
 IP Parameters
Interpartition Logical LAN: U7895.42X.1058008-V5-C4-T1
 1.   Client IP Address                    [9.27.20.216]
 2.   Server IP Address                    [9.42.241.191]
 3.   Gateway IP Address                   [9.27.20.1]
 4.   Subnet Mask                          [255.255.252.0]









 -------------------------------------------------------------------------------
 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:
```

Figure 8-25 IP configuration sample

23. Press M to return to the SMS main menu (see Figure 8-20 on page 341).

24. Select option 5 (Select boot options) to display the Multiboot screen, as shown in Figure 8-26, and select option 1 (Select Install/Boot Device).

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Multiboot
1.   Select Install/Boot Device
2.   Configure Boot Device Order
3.   Multiboot Startup <OFF>
4.   SAN Zoning Support
5.   Management Module Boot List Synchronization




-------------------------------------------------------------------------------
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:
```

Figure 8-26 Select boot options

25. Select option 6 (Network), as shown in Figure 8-27.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Device Type
1.   Diskette
2.   Tape
3.   CD/DVD
4.   IDE
5.   Hard Drive
6.   Network
7.   List all Devices



-------------------------------------------------------------------------------
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:
```

Figure 8-27 Select device type

26. After selecting this option, you are prompted again for the network service as you were in Figure 8-23 on page 343. Make the same selection here (option 1, (BOOTP)).

27. Select the same network adapter that you selected for Figure 8-21 on page 342), as shown in Figure 8-28.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Device
Device  Current  Device
Number  Position  Name
1.        3      Interpartition Logical LAN
       ( loc=U7895.42X.1058008-V5-C4-T1 )













-------------------------------------------------------------------------------
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:
```

Figure 8-28   Network adapter selection

28.On the Select Task screen, select option 2 (Normal Mode Boot), as shown in Figure 8-29.

```
 SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
 Select Task

Interpartition Logical LAN
    ( loc=U7895.42X.1058008-V5-C4-T1 )

 1.   Information
 2.   Normal Mode Boot
 3.   Service Mode Boot








 -------------------------------------------------------------------------------
 Navigation keys:
 M = return to Main Menu
 ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
 Type menu item number and press Enter or select Navigation key:
```

Figure 8-29 Select boot mode

29.Click X to exit SMS.

30. Respond to the prompt to confirm the exit. In the next screen, select Yes. Your installation displays a screen similar to the one shown in Figure 8-30.

```
chosen-network-type = ethernet,auto,none,auto
server IP           = 9.42.241.191
client IP           = 9.27.20.216
gateway IP          = 9.27.20.1
device              = /vdevice/l-lan@30000004
MAC address         = 42 db fe 36 16 4
loc-code            = U7895.42X.1058008-V5-C4-T1

BOOTP request retry attempt: 1







TFTP BOOT ---------------------------------------------------
Server IP.....................9.42.241.191
Client IP.....................9.27.20.216
Gateway IP....................9.27.20.1
Subnet Mask...................255.255.252.0
( 1  ) Filename................./tftpboot/vios2-7989.stglabs.ibm.com
TFTP Retries..................5
Block Size....................512
```

Figure 8-30 Machine booting from NIM

31. To proceed with the OS installation, see 8.3, “Installation procedures” on page 364.

8.2.2 Optical media installation

The optical media installation is the most commonly used method for deploying system images. All of the supported systems listed in 5.1.2, “Software planning” on page 119 are available through DVD or CD media installation.

Note: IBM i installation can be performed from optical media. The IBM i process is different from what is described here for AIX and Linux. For more information, see Section 5 of *Getting Started with IBM i on an IBM Flex System compute node*, available at:

http://www.ibm.com/developerworks/

To perform Optical media installation, you need an external USB drive (not provided with either the chassis nor the Power Systems compute node) attached to your Power Systems compute node.

To perform an optical media installation, complete the following steps:

1. Attach the external optical drive to the USB port of your Power Systems compute node.
2. Insert the installation media into the optical drive.
3. Reboot or power on the server and press the 1 key when prompted to access SMS mode, as shown in Figure 8-31.

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

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


     Memory      Keyboard      Network      SCSI      Speaker
```

Figure 8-31 SMS menu

The window shown in Figure 8-32 opens.

```
Version AF740_051
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





-------------------------------------------------------------------------------
Navigation Keys:

                                             X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure 8-32 SMS main menu options

4. Select option 5 (Select Boot Options) to display the multiboot options. The window shown in Figure 8-33 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Multiboot
1.   Select Install/Boot Device
2.   Configure Boot Device Order
3.   Multiboot Startup <OFF>
4.   SAN Zoning Support
5.   Management Module Boot List Synchronization




-------------------------------------------------------------------------------
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:
```

Figure 8-33 Multiboot options menu

5. Select option 1 (Select Install/Boot Device). The window shown in Figure 8-34 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Device Type
1.   Diskette
2.   Tape
3.   CD/DVD
4.   IDE
5.   Hard Drive
6.   Network
7.   List all Devices

-------------------------------------------------------------------------------
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:
```

Figure 8-34 Boot device options

6. Select the device type, in this case, option 3 (CD/DVD). The window shown in Figure 8-35 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Media Type
1.   SCSI
2.   SSA
3.   SAN
4.   SAS
5.   SATA
6.   USB
7.   IDE
8.   ISA
9.   List All Devices

-------------------------------------------------------------------------------
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:
```

Figure 8-35 Device type selection

7. Select option 6 (USB) media type. The window shown in Figure 8-36 opens and shows the list of available USB optical drives. In our example, a virtual optical drive is shown as item 1.What you see depends on the drive you have connected.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Media Adapter
1.         U7895.42X.1058008-V6-C2-T1   /vdevice/v-scsi@30000002
2.   List all devices






-------------------------------------------------------------------------------
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:
```

Figure 8-36 Select media adapter

8. Select your optical drive. The window shown in Figure 8-37 opens.

```
 SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
 Select Task

Interpartition Logical LAN
    ( loc=U7895.42X.1058008-V6-C4-T1 )

 1.   Information
 2.   Normal Mode Boot
 3.   Service Mode Boot




 -------------------------------------------------------------------------------
 Navigation keys:
 M = return to Main Menu
 ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
 Type menu item number and press Enter or select Navigation key:
```

Figure 8-37 Media selection

9. When you select your optical drive, you have three options. Select option 2 (Normal Mode boot), then select option 1 (Yes) on the next screen. The boot process for your CD displays, and you can continue with the installation process shown in "Installation procedures" on page 364.

8.2.3 TFTP network installation

We can use the standard tools of any Linux distribution to manage a network installation. This method is useful when an optical drive is not available or if a NIM server is not installed and configured. Any Linux x86-based computer can be used as the TFTP server and virtually any Linux distribution can be easily configured to perform this task. In this section, we describe how to implement this function.

First, you must set up three standard Linux services on the installation server:

- `tftpd`
- `dhcpd` (used only to allow `netboot` using `bootpd` to a specific MAC address)
- NFS server

SUSE Linux Enterprise Server 11

The following steps pertain to SLES 11:

1. Obtain the distribution ISO file, and copy it to a work directory of the installation server. We configure a Network File System (NFS) server (this server can be the installation server itself or another server) and mount this shared directory from the target virtual server to unload the software.
2. On the installation server, install the `tftp` and the `dhcpd` server packages (we use `dhcpd` only to run `bootp` for a specific MAC address).
3. Copy in the `tftpboot` directory (the default for SUSE Linux Enterprise Server 11 is `/tftpboot`), the `netboot` image, and the `yaboot` executable from the DVD directory, `sles11/suseboot`.
   - The `netboot` image is named `inst64`.
   - The `yaboot` executable is named `yaboot.ibm`.
4. Boot the target virtual server and access SMS (see Figure 8-38) to retrieve the MAC address of the Ethernet interface to use for the installation.

```
Version AF740_051
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




-------------------------------------------------------------------------------
Navigation Keys:

                                             X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:2
```

Figure 8-38 Setup remote IPL selection

The MAC address shown in Figure 8-39 is the Hardware Address.

```
 Version AF740_051
 SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
 NIC Adapters
      Device                          Location Code              Hardware
                                                                  Address
 1.  Interpartition Logical LAN      U8406.71Y.06ACE4A-V4-C4-T1  XXXXXXXXXXXX










 -------------------------------------------------------------------------------
 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:
```

Figure 8-39 MAC address

5. On the installation server, configure the `dhcpd.conf` file and, assuming it is the NFS server too, the `/etc/exports` file. The `dhcpd.conf` file is shown in Figure 8-40, where we must replace `XX.XX.XX.XX.XX.XX` and the network parameters with our MAC and IP addresses.

```
always-reply-rfc1048 true;
allow bootp;
deny unknown-clients;
not authoritative;
default-lease-time 600;
max-lease-time 7200;
ddns-update-style none;

subnet 10.1.0.0 netmask 255.255.0.0 {
host sles11 {
fixed-address 10.1.2.90;
hardware ethernet XX:XX:XX:XX:XX:XX;
next-server 10.1.2.56;
filename "yaboot.ibm";
}
}
```

Figure 8-40 The dhcpd.conf file for SUSE Linux Enterprise Server 11

6. Create a file in `/tftpboot` named `yaboot.conf-xx.xx.xx.xx.xx.xx` (where `xx.xx.xx.xx.xx.xx` is our MAC address), as shown in Figure 8-41. Figure 8-41 shows an example of this file that is configured to start the installer and access the DVD ISO image using NFS.

```
default=sles11
  timeout=100
  image[64bit]=inst64.sles11
    label=sles11
    append="quiet usevnc=1 vncpassword=passw0rd
install=nfs://10.1.2.51/temp/sles11"
```

Figure 8-41 yaboot.conf-xx.xx.xx.xx.xx.xx

7. Figure 8-42 shows an example of the `/etc/exports` file with the exported directory that contains the image of the SUSE Linux Enterprise Server 11 DVD.

```
/dati1/sles11/                  *(rw,insecure,no_root_squash)
```

Figure 8-42 Exports NFS server configuration sample

8. On the installation server or virtual server, start the `dhcpd` and `nfsd` services.
9. On the target virtual server, start `netboot`, as shown in the Figure 8-43.

```
Version AF740_051
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




-------------------------------------------------------------------------------
Navigation Keys:

                                            X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:5
```

Figure 8-43 Select boot options

10. Select option 5 (Select Boot Options). The window shown in Figure 8-44 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Multiboot
1.   Select Install/Boot Device
2.   Configure Boot Device Order
3.   Multiboot Startup <OFF>
4.   SAN Zoning Support
5.   Management Module Boot List Synchronization




-------------------------------------------------------------------------------
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:1
```

Figure 8-44 Select Install/Boot Device

11. Select option 1 (Select Install/Boot Device). The window shown in Figure 8-45 opens.

```
Version AF740_051
 SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
 Select Device Type
 1.   Diskette
 2.   Tape
 3.   CD/DVD
 4.   IDE
 5.   Hard Drive
 6.   Network
 7.   List all Devices



 -------------------------------------------------------------------------------
 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
```

Figure 8-45 Select a network as the installation device

12. Select option 6 (Network) as the boot device. The window shown in Figure 8-46 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Network Service.
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:1
```

Figure 8-46 Select BOOTP as the boot protocol

13. Select option 1 (BOOTP), as shown in Figure 8-46.

14.Select the network adapter and the normal mode boot, and the installation starts loading the `yaboot.ibm` boot loader through the network, as shown in Figure 8-47.

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








TFTP BOOT ---------------------------------------------------
Server IP.....................192.168.20.11
Client IP.....................192.168.20.12
Subnet Mask...................255.255.255.0
( 1  ) Filename.................yaboot.ibm
TFTP Retries..................5
Block Size....................512
FINAL PACKET COUNT = 407
FINAL FILE SIZE = 208348  BYTES
```

Figure 8-47 Netbooting the boot loader

For a description of the installation, see 8.3.3, “Installing SUSE Linux Enterprise Server” on page 381.

Red Hat Enterprise Linux 6.1

For Red Hat Enterprise Linux 6.1, we follow a procedure similar to the one shown in “SUSE Linux Enterprise Server 11” on page 355. The description that follows shows the differences between the two procedures.

Complete the following steps:

1. Obtain the ISO file of Red Hat Enterprise Linux 6.1, and copy it to a work directory of the installation server.
2. On the installation server, install the `tftp` and the `dhcpd` server packages (we use `dhcpd` to run `bootp` on a specific MAC address).
3. Copy the `yaboot` executable from the DVD directory `ppc/chrp` to the `tftpboot` directory on the installation server (`/var/lib/tftpboot/`).

Tip: The yaboot executable is named simply yaboot. We can rename it, for example, to yaboot.rh61, to avoid conflicts in the tftpboot directory.

4. The netboot image is larger than 65,500 512 bytes blocks and cannot be used due to a limitation of tftpd. We must boot the vmlinuz kernel and use the ramdisk image. Copy the two files from the ppc/ppc64 directory of the DVD to the tftpboot directory of the installation server.
5. On the installation server, create a directory named tftpboot/etc, and create a file named 00-XX-XX-XX-XX-XX-XX, replacing all characters except the 00 with the target virtual server MAC address, as shown in Figure 8-48.

```
default=rh61
  timeout=100
  image=vmlinuz
   initrd=ramdisk.image.gz
    label=rh61
```

Figure 8-48 00-XX-XX-XX-XX-XX-XX file

6. The dhcpd.conf file is displayed in Figure 8-49. It is similar to the SLES version. Again, change the network addresses and the MAC address and the IP configuration to your environment settings.

```
allow bootp;
deny unknown-clients;
not authoritative;
default-lease-time 600;
max-lease-time 7200;
ddns-update-style none;

subnet 192.168.20.0 netmask 255.255.255.0 {
host rh61-vs1 {
fixed-address 192.168.20.12;
hardware ethernet XX:XX:XX:XX:XX:XX;
next-server 192.168.20.11;
filename "yaboot.rh6";
}
}
```

Figure 8-49 The dhcpd.conf file for Red Hat Enterprise Linux 6.1

8.2.4 Cloning methods

There are two cloning methods available for an AIX installation. The most common method of cloning is to create `a mksysb` image on one machine and restore it in the cloned machine. This method clones all of your OS (`rootvg`) but no non-`rootvg vg` OSes or file systems. This method is a fast way of cloning your AIX installation, and it can be performed using tape devices, DVD media, or a NIM installation.

Ensure that the IP address is not cloned in this process. If you are using NIM to restore the `mksysb`, the IP address given to the client during the network boot overrides the IP address on the interface used by NIM.

It is also important to determine if all device drivers that are needed to support the hardware on the target system are in the `mksysb`. This task can be accomplished by installing the necessary device drivers in the image before creating the `mksysb`, or, when using NIM to restore the `mksysb`, ensure that an `lpp_source` is specified that contains the needed drivers.

You can also use the ALT_DISK_INSTALL method, but this method work only if you have SAN disks attached or removable disks that can be attached to the new server. You can use the ALT_DISK_INSTALL method to create a full copy of your system `rootvg`, and then you can remove that disk from the server and assign it to another server. When you start your system, your system is cloned.

8.3 Installation procedures

For the rest of this chapter, we describe the installation methods for the Power Systems compute node-supported operating systems. We describe the following topics:

- Installing AIX
- Installing Red Hat Enterprise Linux
- Installing SUSE Linux Enterprise Server
- Installing IBM i

8.3.1 Installing AIX

There are three possible methods to install AIX on your Power Systems compute node:

- NIM installation with lpp_source installation
- NIM installation with mksysb
- Optical media installation

To install AIX using the NIM lpp_source method, complete the following steps:

1. The first part of the process, setting up the environment for installation, is covered in 8.2.1, “NIM installation” on page 332, and we follow up after exiting to the normal boot part of the process.
2. After you exit to normal boot, a screen opens that shows the network parameters for BOOTP, as shown in Figure 8-24 on page 343.
3. Next, a screen opens that shows the AIX kernel loading. You are prompted to select the installation language (English, by default), as shown in Figure 8-50.

```
>>>  1 Type 1 and press Enter to have English during install.







    88  Help ?


>>> Choice [1]:
```

Figure 8-50 Installation language selection

4. After selecting the language, the installation options are displayed, as shown in Figure 8-51.

```
 Welcome to Base Operating System
                     Installation and Maintenance

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

>>> 1 Start Install Now with Default Settings

    2 Change/Show Installation Settings and Install

    3 Start Maintenance Mode for System Recovery

    4 Configure Network Disks (iSCSI)

    5 Select Storage Adapters




    88  Help ?
    99  Previous Menu

>>> Choice [1]:
```

Figure 8-51 Installation options

You can install the OS using option 1 or 2:

– Option 1 (Start Install Now with Default Settings) begins the installation using the default options.

– Option 2 (Change/Show Installation Settings and Install) displays several options, as shown in Figure 8-52.

```
 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)
    5  Select Edition......................express
>>> 0  Install with the settings listed above.


+-----------------------------------------------------
    88  Help ?         |  WARNING: Base Operating System Installation
will
    99  Previous Menu  |       destroy or impair recovery of ALL data
on the
                       |            destination disk hdisk0.
>>> Choice [0]:
```

Figure 8-52 Installation settings

In this screen, the following settings are available. After you change and confirm your selections, type 0 and press Enter to begin the installation. The settings are:

– Option 1 (Systems Settings) refers to the installation method and destination disk. Supported methods for AIX installation are:

- New and Complete Overwrite: Use this method when you are installing a new system or reinstalling one that needs to be erased.

- Migration installation: Use this method when you are upgrading an older version of AIX (AIX 5L V5.3 or AIX V6.1) to a newer version, such as AIX V7.1. This option retains all of your configuration settings. The `tmp` directory is erased during installation.
- Preservation installation: This method is similar to the New and Complete Overwrite option, except that it retains only the `/home` directory and other user files. This option overwrites the file systems.

– Option 2 (Primary Language Environment Settings (AFTER Install)): After you select the correct type of installation, choose the language for the installation, a keyboard, and cultural convention.

– Option 3 (Security model): You can use this option to enable the trusted computer database and other security options, as shown in Figure 8-53.

```
 Security Models

  Type the number of your choice and press Enter.


 1. Trusted AIX............................................ no

 2. Other Security Options  (Trusted AIX and Standard)
     Security options vary based on choices.
     LSPP, SbD, CAP/CCEVAL, TCB

>>> 0  Continue to more software options.

    88  Help ?
    99  Previous Menu

>>> Choice [0]:
```

Figure 8-53 Security options selection

– Option 4 (More Options (Software Install options)): You can use this option to choose whether to install graphics software, such as X Window System, to select the file system type jfs or jfs2, and to enable system backups at any time, as shown in Figure 8-54 on page 369.

```
 Install Options

 1.  Graphics Software............................................... yes
 2.  System Management Client Software............................... yes
 3.  Create JFS2 File Systems........................................ yes
 4.  Enable System Backups to install any system...................... yes
     (Installs all devices)


>>> 5.  Install More Software

    0  Install with the settings listed above.

    88  Help ?
    99  Previous Menu

>>> Choice [5]:
```

Figure 8-54 Install Options screen

5. After you complete your options selection, you are prompted to confirm your choices, as shown in Figure 8-55.

```
 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
Selected Edition:  express

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]:
```

Figure 8-55 Installation summary

6. To proceed, click option 1 (Continue with Install). The packages display as they install.

8.3.2 Installing Red Hat Enterprise Linux

This section describes the installation of Red Hat Enterprise Linux (RHEL). Detailed information about supported operating systems is listed in 5.1.2, “Software planning” on page 119.

We install the virtual servers using a virtual optical media and the ISO image of the RHEL distribution as the boot device. Figure 8-56 shows the Virtual Optical Media window in IBM Flex System Manager.

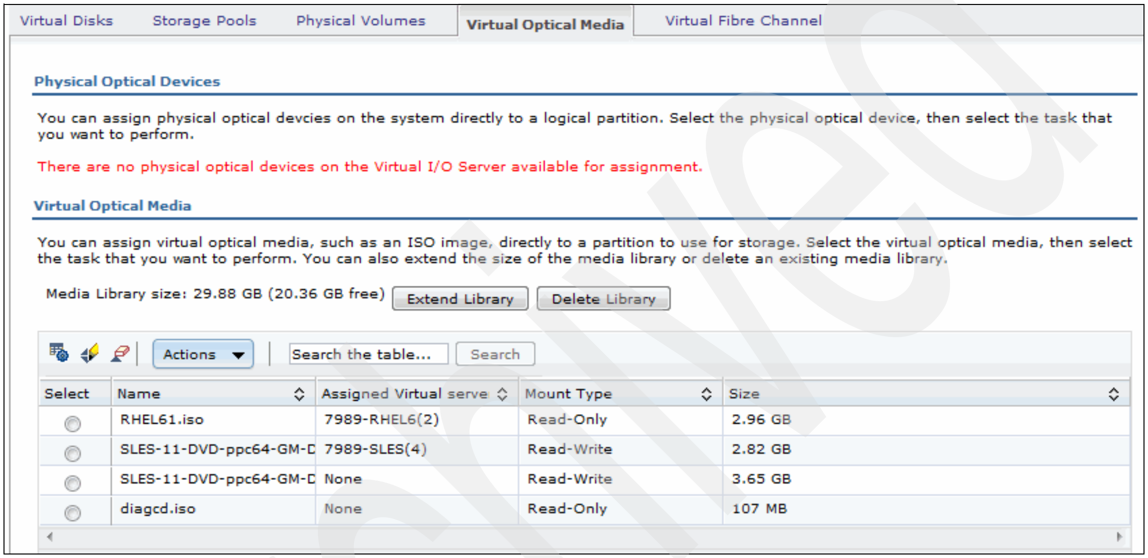

Figure 8-56 Virtual optical media management

To install RHEL, complete the following steps:

1. After the virtual media is set up, boot the server and enter SMS. The screen shown in Figure 8-57 opens.

```
Version AF740_051
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




-------------------------------------------------------------------------------
Navigation Keys:

                                            X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:5
```

Figure 8-57 Virtual server SMS menu

2. Select option 5 (Select Boot Options). The screen shown in Figure 8-58 opens.

```
 Version AF740_051
 SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
 Multiboot
 1.   Select Install/Boot Device
 2.   Configure Boot Device Order
 3.   Multiboot Startup <OFF>
 4.   SAN Zoning Support
 5.   Management Module Boot List Synchronization





 -------------------------------------------------------------------------------
 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:1
```

Figure 8-58 Select Install/Boot Device

3. Select option 1 (Select Install/Boot Device). The window shown in Figure 8-59 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Device Type
1.   Diskette
2.   Tape
3.   CD/DVD
4.   IDE
5.   Hard Drive
6.   Network
7.   List all Devices



-------------------------------------------------------------------------------
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:3
```

Figure 8-59 Select Install/Boot Device

4. We want to boot from a virtual optical drive, so we select option 3 (CD/DVD). The window shown in Figure 8-60 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Media Type
1.   SCSI
2.   SSA
3.   SAN
4.   SAS
5.   SATA
6.   USB
7.   IDE
8.   ISA
9.   List All Devices






-------------------------------------------------------------------------------
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:1
```

Figure 8-60 Selection of the SCSI DVD reader

5. For the virtual optical media, select option 1 (SCSI). The window shown in Figure 8-61 opens.

```
Version AF740_051
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-------------------------------------------------------------------------------
Select Device
Device  Current  Device
Number  Position  Name
1.        -      SCSI CD-ROM
       ( loc=U7895.42X.1058008-V2-C2-T1-L8200000000000000 )






-------------------------------------------------------------------------------
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:1
```

Figure 8-61 SCSI CD-ROM in position one

6. Select the drive you want to boot from. In Figure 8-61, there is only one drive to select, which is the virtual optical media linked to the Red Hat Enterprise Linux DVD ISO image.

   The system now boots from the ISO image. Figure 8-62 shows the boot of the virtual media and the VNC parameters.

```
Welcome to the 64-bit Red Hat Enterprise Linux 6.1 installer!
Hit <TAB> for boot options.


Welcome to yaboot version 1.3.14 (Red Hat 1.3.14-35.el6_0.1)
Enter "help" to get some basic usage information
boot:
* linux
boot: linux vnc vncpassword=mypassword
```

Figure 8-62 Installation prompt with VNC parameters

It is possible to stop the boot process by pressing the Tab key, allowing you to enter optional parameters on the command line:

- To use VNC and perform an installation in a graphic environment, run **`linux vnc vncpassword=yourpwd`**. The password must be at least six characters long.
- To install Red Hat Enterprise Linux 6.1 on a multipath external disk, run the following command:

  `linux mpath`

More details about these actions are in the Red Hat Enterprise Linux 6 *Installation Guide* and the *DM Multipath* guide, found at:

http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6

For VNC information, go to the following website:

http://www.realvnc.com/

Figure 8-63 shows the network TCP/IP configuration required to use VNC.

```
Welcome to Red Hat Enterprise Linux for ppc64


        ················· Manual TCP/IP Configuration ···················
        ·                                                               ·
        · Enter the IPv4 and/or the IPv6 address and prefix (address /  ·
        · prefix).  For IPv4, the dotted-quad netmask or the CIDR-style ·
        · prefix are acceptable. The gateway and name server fields must ·
        · be valid IPv4 or IPv6 addresses.                              ·
        ·                                                               ·
        · IPv4 address: X.XX.20.114_____ / 255.255.XXX.0___             ·
        · Gateway:      X.XX.20.1________________________________       ·
        · Name Server:  X.XX.242.28______________________________       ·
        ·                                                               ·
        ·             ······                          ········          ·
        ·             · OK ·                          · Back ·          ·
        ·             ······                          ········          ·
        ·                                                               ·
        ·                                                               ·
        ·················································································



  <Tab>/<Alt-Tab> between elements  | <Space> selects | <F12> next screen
```

Figure 8-63 Manual TCP/IP configuration for VNC installation

Figure 8-64 shows the VNC graphical console start.

```
Running anaconda 13.21.117, the Red Hat Enterprise Linux system installer - please wait.
21:08:52 Starting VNC...
21:08:53 The VNC server is now running.
21:08:53

You chose to execute vnc with a password.


21:08:53 Please manually connect your vnc client to ite-bt-061.stglabs.ibm.com:1
(9.27.20.114) to begin the install.
21:08:53 Starting graphical installation.
```

Figure 8-64 VNC server running

7. Connect to the IP address listed in Figure 8-64 with a VNC client to perform the installation. You see the graphic RHEL installer welcome window.
8. Select a preferred language for the installation process.
9. Select the keyboard language.
10. Select the storage devices to use for the installation, as shown in Figure 8-65. For virtual disks, hdisks, or SAN disks, select **Basic Storage Devices**.

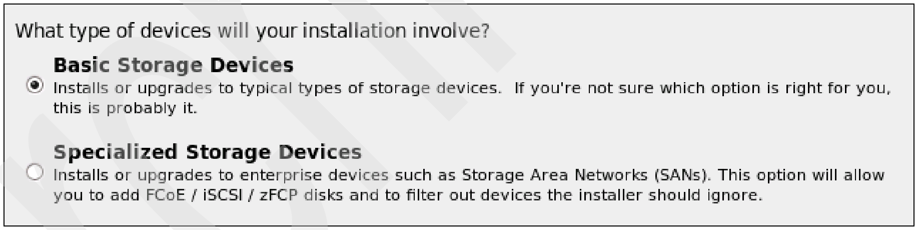

Figure 8-65 Select storage devices

11. Select either **Fresh Installation** (a new and complete overwrite) or **Upgrade an Existing Installation**, as shown in Figure 8-66.

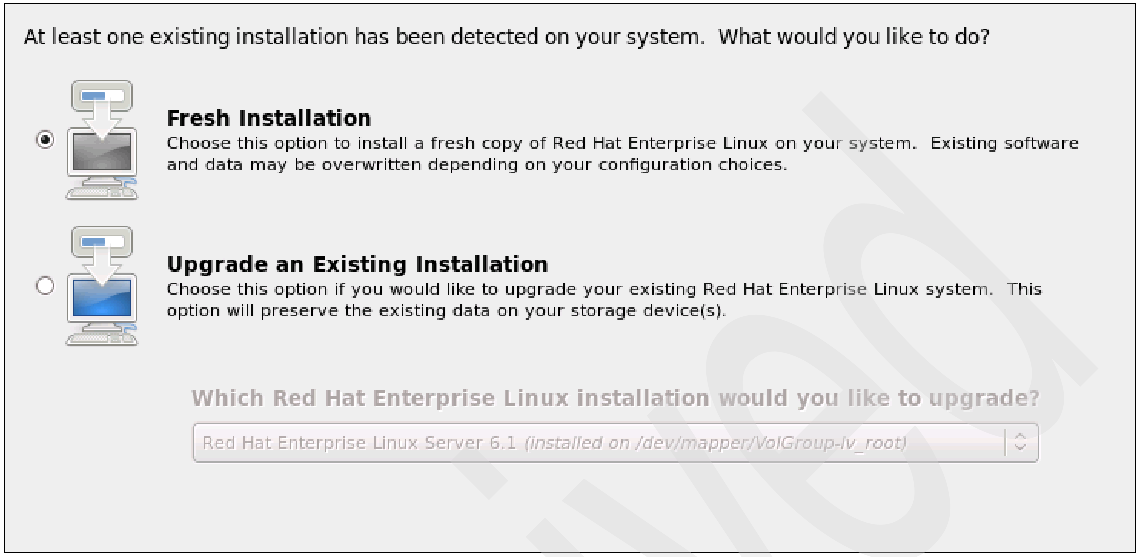

Figure 8-66 Select a fresh installation or an upgrade to an existing installation

12. Select a disk layout, as shown in Figure 8-67. You can choose from a number of installations or create a custom layout (for example, you can create a software mirror between two disks). You can also manage older RHEL installations if they are detected.

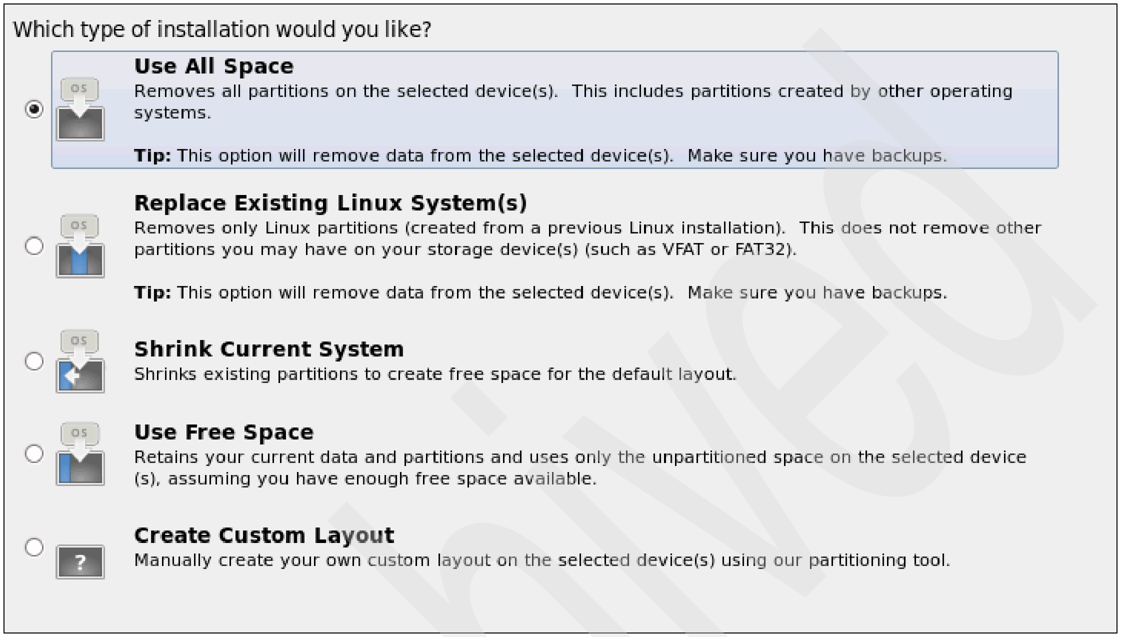

Figure 8-67 Disk space allocation selections

13. Select the software packages to install, as shown in Figure 8-68.

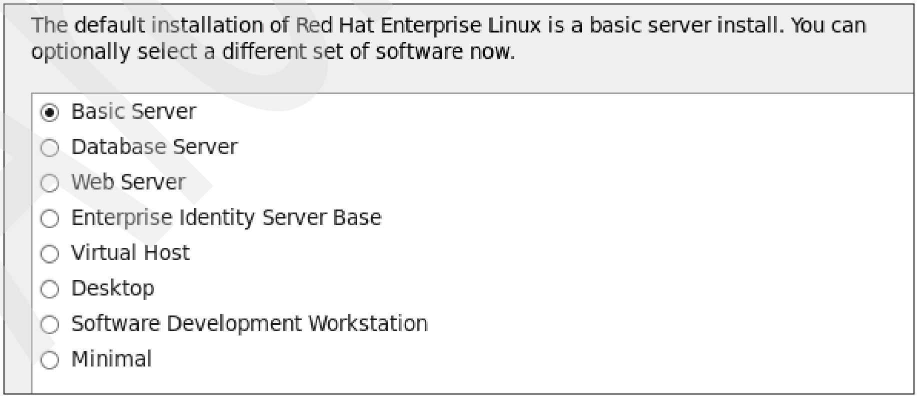

Figure 8-68 RPM packages selection

The software installation process starts.

When the VNC installation is complete, the window shown in Figure 8-69 opens. The virtual server reboots, the console returns to alphanumeric mode, and you can connect to the server using SSH or Telnet.

Figure 8-69 End of VNC installation

As the system boots, the operating system loads, as shown in Figure 8-70.

```
Starting cups: [  OK  ]
Mounting other filesystems:  [  OK  ]
Starting HAL daemon: [  OK  ]
Starting iprinit: [  OK  ]
Starting iprupdate: [  OK  ]
Retrigger failed udev events[  OK  ]
Adding udev persistent rules[  OK  ]
Starting iprdump: [  OK  ]
Loading autofs4: [  OK  ]
Starting automount: [  OK  ]
Generating SSH1 RSA host key: [  OK  ]
Generating SSH2 RSA host key: [  OK  ]
Generating SSH2 DSA host key: [  OK  ]
Starting sshd: [  OK  ]
Starting postfix: [  OK  ]
Starting abrt daemon: [  OK  ]
Starting crond: [  OK  ]
Starting atd: [  OK  ]
Starting rhsmcertd 240[  OK  ]

Red Hat Enterprise Linux Server release 6.1 (Santiago)
Kernel 2.6.32-131.0.15.el6.ppc64 on an ppc64

ite-bt-061.stglabs.ibm.com login:
```

Figure 8-70 First time login screen

The basic installation is complete. You might choose to install additional RPMs from the IBM Service and Productivity Tools website found at:

http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html

8.3.3 Installing SUSE Linux Enterprise Server

In this section, we describe the installation of SUSE Linux Enterprise Server 11 (SLES 11). We prefer to do the installation using VNC (in graphic mode) because many of the panels are complex, and it is easier to accomplish this task in graphic mode.

We do not show the initial SMS steps here, as they are described in 8.3.2, “Installing Red Hat Enterprise Linux” on page 370. Follow step 1 on page 371 to step 7 on page 377 before completing the following steps:

1. The first window is the installation mode window, shown in Figure 8-71.

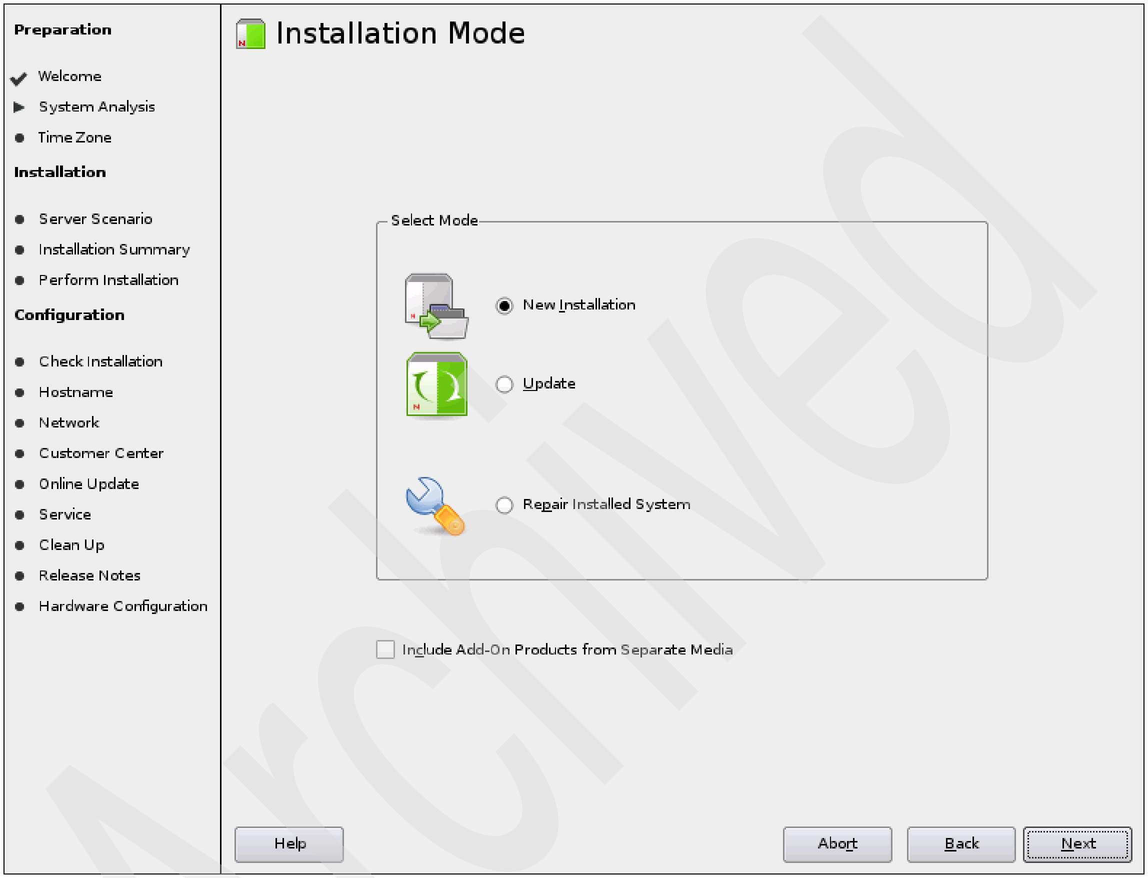

Figure 8-71 First step in graphic mode

2. Select **New installation** and click **Next**. The Installation Settings window opens (Figure 8-72).

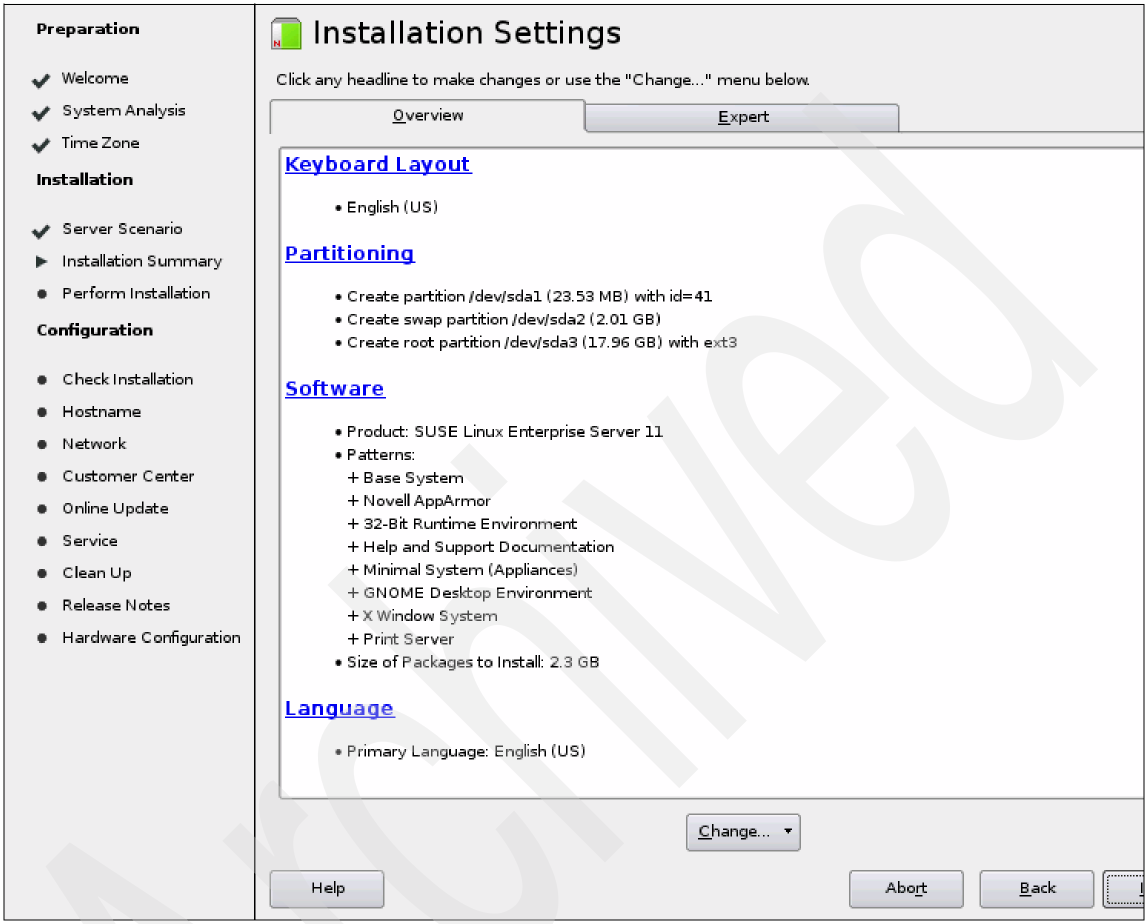

Figure 8-72 Installation settings

3. Either accept the default values or click **Change** to change the values for:
   - Keyboard layout
   - Partitioning
   - Software
   - Language

Click **Next** to continue. The Perform Installation window opens (Figure 8-73) and shows the progress of the installation.

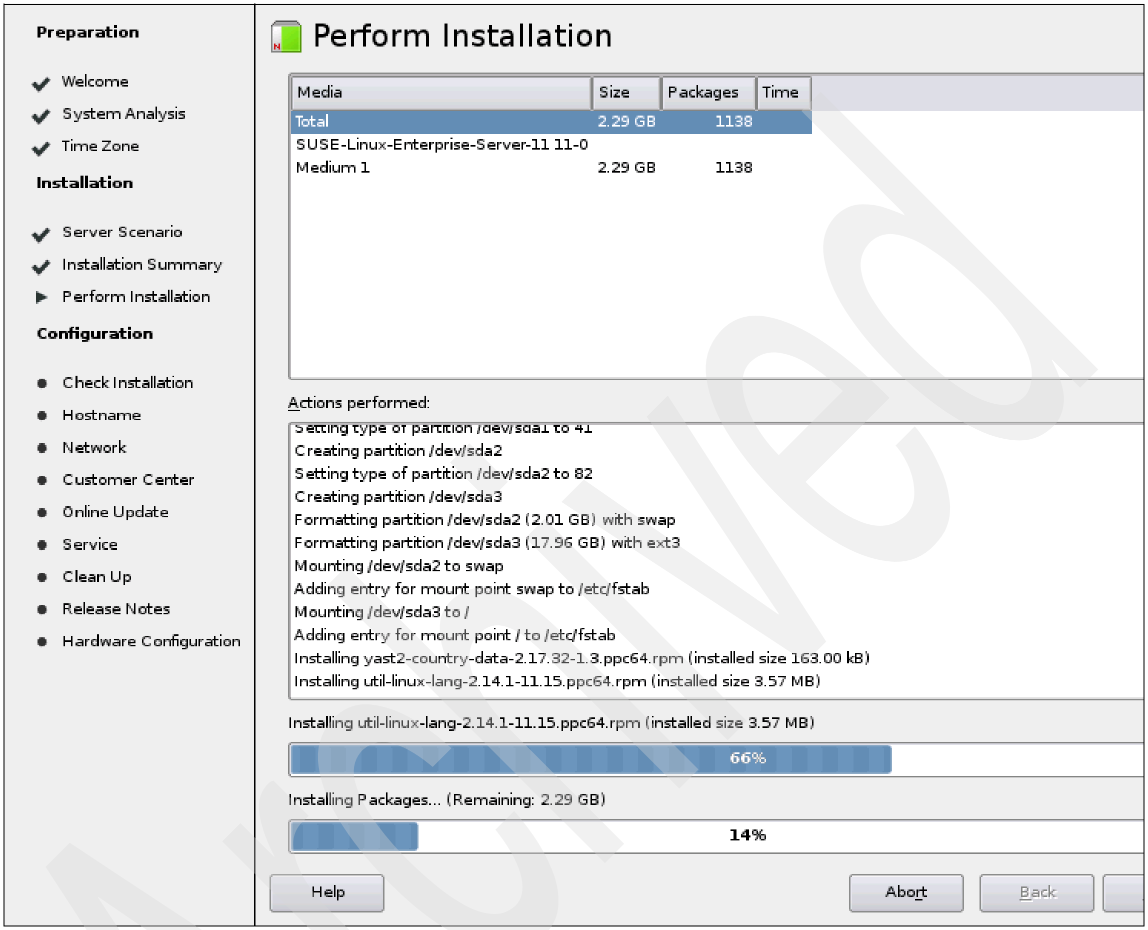

Figure 8-73 Perform Installation window

The final phase of the basic installation process is shown in Figure 8-74.

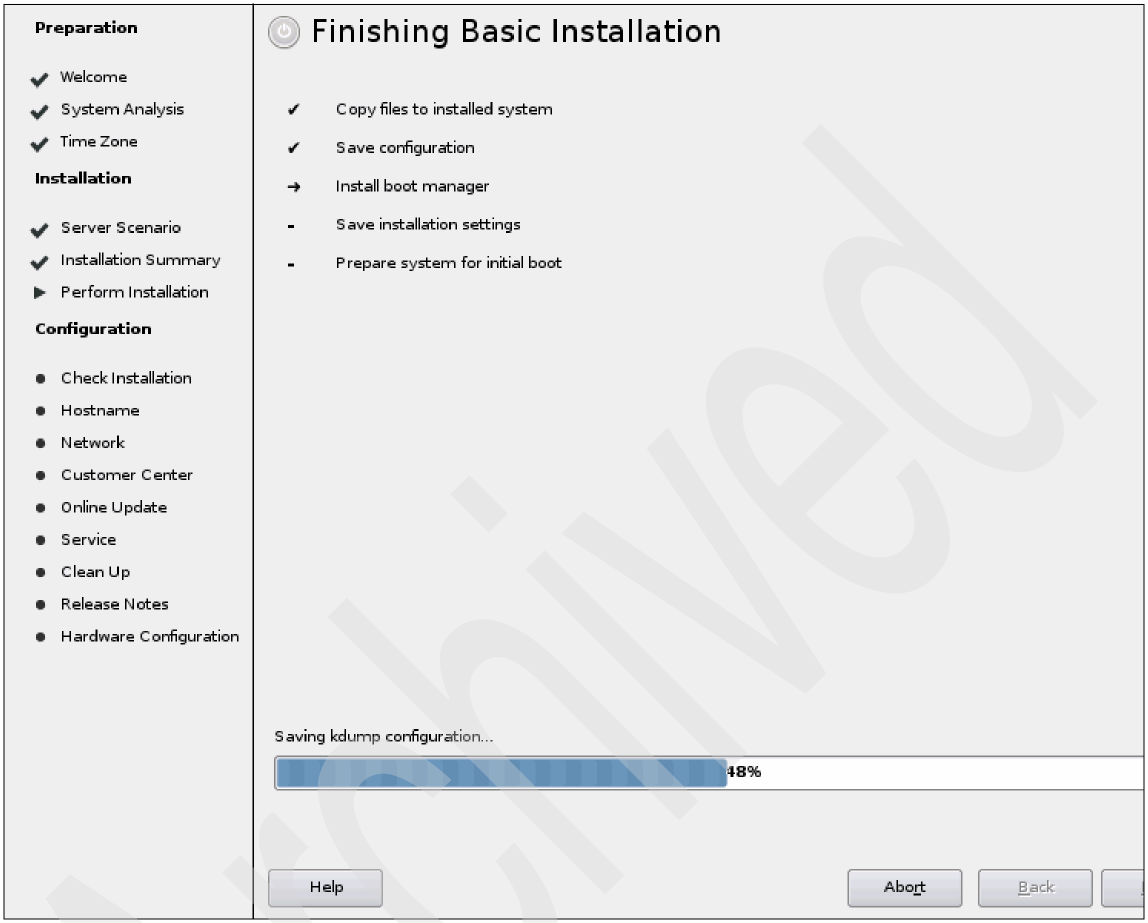

Figure 8-74 Finishing Basic Installation window

At the end of the installation, the system reboots and the VNC connection is lost.

4. Figure 8-75 shows the console while rebooting. After reboot, VNC restarts with the same configuration, and we must reconnect the VNC client.

```
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM                              IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM     STARTING SOFTWARE       IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM        PLEASE WAIT...       IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM                              IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
-
Elapsed time since release of system processors: 202 mins 30 secs

yaboot starting: loaded at 00040000 00064028 (0/0/00c3ba70; sp: 01a3ffd0)
```

Figure 8-75 Reboot and VNC automatic restart

5. The installation and configuration continues with a prompt to enter the root password. Enter the root password.

6. Other installation screens open. Enter values as needed for your environment. After the installation is complete, you see the window shown in Figure 8-76.

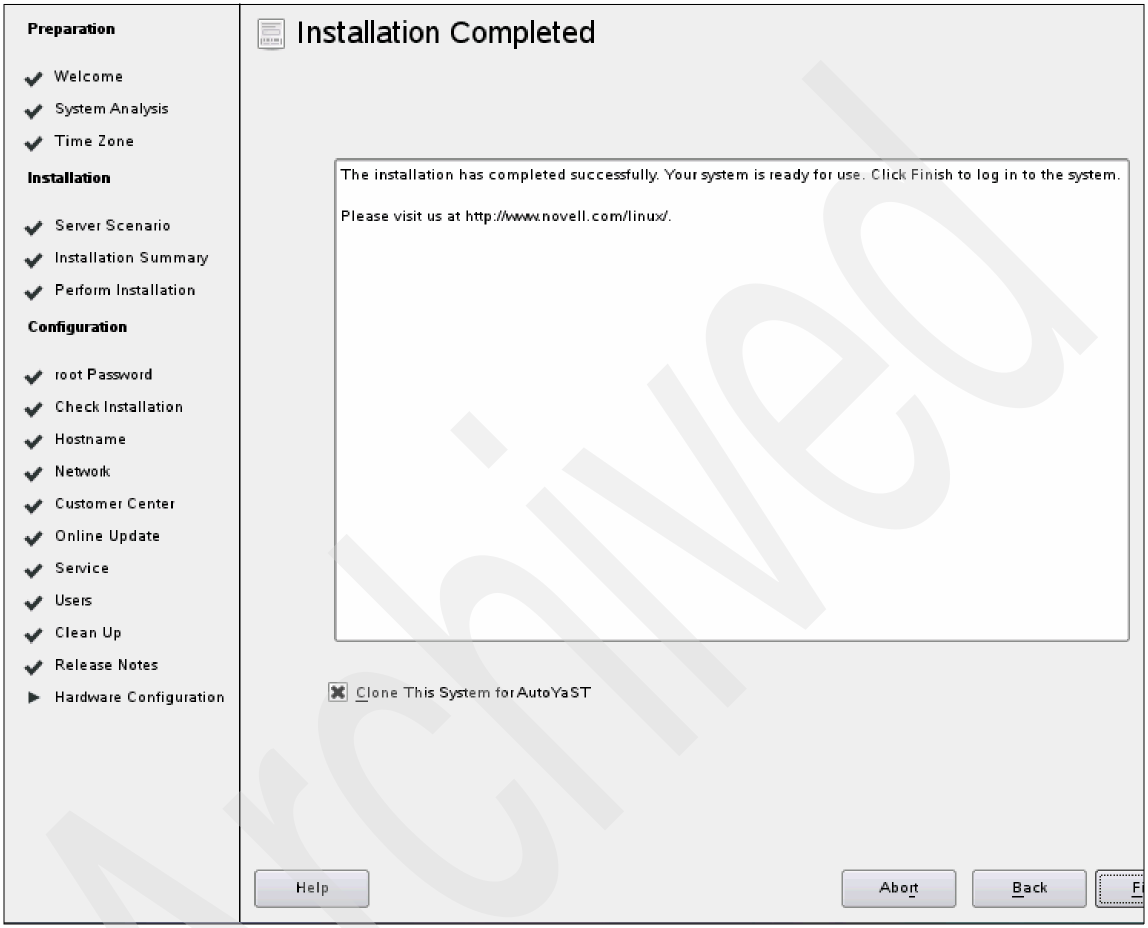

Figure 8-76 Installation Completed window

7. The virtual server reboots, the VNC server is shut down, and we can connect to the text console, through a virtual terminal, using Secure Shell (SSH) or Telnet, as shown in Figure 8-77.

```
Starting Name Service Cache Daemon                                    done
Checking ipr microcode levels
Completed ipr microcode updates                                       done
Starting ipr initialization daemon                                    done
Starting irqbalance                                                   done
Starting cupsd                                                        done
Starting rtas_errd (platform error handling) daemon:                  done
Starting ipr dump daemon                                              done
Starting SSH daemon                                                   done
Starting smartd                                                       unused
Setting up (remotefs) network interfaces:
Setting up service (remotefs) network  .  .  .  .  .  .  .  .  .  .   done
Starting mail service (Postfix)                                       done
Starting CRON daemon                                                  done
Starting INET services. (xinetd)                                      done
Master Resource Control: runlevel 3 has been                          reached
Skipped services in runlevel 3:                   smbfs nfs smartd splash


Welcome to SUSE Linux Enterprise Server 11 (ppc64) - Kernel
2.6.27.19-5-ppc64 (console).


sles11-e4kc login:
```

Figure 8-77 Login screen

The basic SLES installation is complete. You may choose to install additional RPMs from the IBM Service and Productivity Tool website at:

http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html

8.3.4 Installing IBM i

For details about installing IBM i on the Power Systems compute nodes, see *Getting Started with IBM i on an IBM Flex System compute node*, available at:

http://www.ibm.com/developerworks/

Abbreviations and acronyms

AAS	Advanced Administrative System
AC	alternating current
ACL	access control list
AME	Active Memory Expansion
AMM	Advanced Management Module
AMS	access method services
AS	Australian Standards
ASIC	application-specific integrated circuit
BIOS	basic input/output system
BOOTP	boot protocol
BOS	base operating system
BRD	board
CD	compact disk
CD-ROM	compact disc read-only memory
CFM	cubic feet per minute.
CLI	command-line interface
CMM	common management model
CPU	central processing unit
CRTM	Core Root of Trusted Measurements
CSS	cascading style sheets
CTS	clear to send
DASD	direct access storage device
DC	domain controller
DHCP	Dynamic Host Configuration Protocol
DIMM	dual inline memory module
DM	disconnected mode
DNS	Domain Name System
DPM	distributed power management
DRTM	Dynamic Root of Trust Measurement
DSA	Digital Signature Algorithm
DVD	digital video disc
EMC	electromagnetic compatibility
ESA	Electronic Service Agent
ESB	error status block
ETE	everything to everything
FC	Fibre Channel
FCP	Fibre Channel Protocol
FDR	fourteen data rate
FDX	full duplex
FSM	File System Migrator
FSP	flexible service processor
FTP	File Transfer Protocol
GB	gigabyte
GIF	graphic interchange format
GUI	graphical user interface
HA	high availability
HACMP	High-Availability Cluster Multi-Processing
HAL	hardware abstraction layer
HBA	host bus adapter
HDD	hard disk drive
HMC	Hardware Management Console
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
HTTPS	HTTP over SSL
I/O	input/output

IBM	International Business Machines
ID	identifier
IDE	integrated drive electronics
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IMM	integrated management module
IP	Internet Protocol
IPL	initial program load
ISA	industry standard architecture
ISO	International Organization for Standards
ISV	independent software vendor
IT	information technology
ITE	Internal Terminal Emulator
ITSO	International Technical Support Organization
IVM	Integrated Virtualization Manager
KB	kilobyte
KVM	keyboard video mouse
LAN	local area network
LDAP	Lightweight Directory Access Protocol
LED	light emitting diode
LMB	logical memory block
LOM	LAN on motherboard
LP	low profile
LPAR	logical partitions
LPM	lines per minute
LSPP	Labelled Security Protection Profile
LUN	logical unit number
MAC	media access control
MB	megabyte
MLC	multi-level cell
MPIO	multi-path I/O
MSI	message signaled interrupt
MTM	machine-type-model
MTS	Microsoft Transaction Server
MTU	maximum transmission unit
NASA	National Aeronautics and Space Administration
NFS	network file system
NIC	network interface card
NIM	Network Installation Management
NPIV	N_Port ID Virtualization
NVRAM	non-volatile random access memory
OS	operating system
OSPF	Open Shortest Path First
PC	personal computer
PCI	Peripheral Component Interconnect
PDU	power distribution unit
PF	power factor
PID	product ID
PSU	power supply unit
PXE	Preboot eXecution Environment
QDR	quad data rate
RAID	redundant array of independent disks
RAM	random access memory
RAS	remote access services; row address strobe
RBAC	Role Based Access Control
RDIMM	registered DIMM
RDMA	Remote Direct Memory Access
RHEL	Red Hat Enterprise Linux
RIP	Routing Information Protocol

ROM	read-only memory
RPM	Red Hat Package Manager
RSA	Remote Supervisor Adapter
RSS	receive-side scaling
RTE	Remote Terminal Emulator
SAN	storage area network
SAS	Serial Attached SCSI
SATA	Serial ATA
SCP	System Control Process
SCSI	Small Computer System Interface
SDD	Subsystem Device Driver
SDMC	Systems Director Management Console
SEA	Share Ethernet Adapter
SFF	small form factor
SFP	small form-factor pluggable
SLES	SUSE Linux Enterprise Server
SLI	Service Level Interface
SMP	symmetric multiprocessing
SMS	System Management Services
SMT	station management
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SOL	Serial over LAN
SP	service processor
SR	short range
SRAM	static RAM
SRM	system resources manager
SSA	serial storage architecture
SSD	solid-state drive
SSH	Secure Shell
SSL	Secure Sockets Layer
STP	Spanning Tree Protocol
SWMA	Software Maintenance Agreement
TB	terabyte
TCB	Transport Control Block
TCG	Trusted Computing Group
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TFTP	Trivial File Transfer Protocol
TL	technology level
TPM	Trusted Platform Module
TPMD	thermal and power management device
TSM	Tivoli Storage Manager
TTY	teletypewriter
TXT	text
UEFI	Unified Extensible Firmware Interface
UL	Underwriters Laboratories
UPS	uninterruptible power supply
URL	Uniform Resource Locator
USB	universal serial bus
VAC	volts alternating current
VGA	video graphics array
VIOS	Virtual I/O Server
VLAG	Virtual link aggregation group
VLAN	virtual LAN
VLP	very low profile
VM	virtual machine
VNC	Virtual Network Computing
VPD	vital product data
VPI	virtual path identifier
VSP	Virtual Service Providers
WPAR	workload partition
WW	world wide
WWPN	World Wide Port Name
XML	Extensible Markup Language

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

The following publications from IBM Redbooks provide additional information about IBM Flex System. They are available at:

http://www.redbooks.ibm.com/portals/puresystems

- *IBM PureFlex System and IBM Flex System Products & Technology*, SG24-7984
- *IBM Flex System Networking in an Enterprise Data Center*, REDP-4834

Chassis and Compute Nodes:

- *IBM Flex System Enterprise Chassis*, TIPS0863
- *IBM Flex System Manager*, TIPS0862
- *IBM Flex System p260 and p460 Compute Node*, TIPS0880
- *IBM Flex System x240 Compute Node*, TIPS0860

Switches:

- *IBM Flex System EN2092 1Gb Ethernet Scalable Switch*, TIPS0861
- *IBM Flex System EN4091 10Gb Ethernet Pass-thru Module*, TIPS0865
- *IBM Flex System Fabric EN4093 10Gb Scalable Switch*, TIPS0864
- *IBM Flex System FC3171 8Gb SAN Switch and Pass-thru*, TIPS0866
- *IBM Flex System FC5022 16Gb SAN Scalable Switch and FC5022 24-port 16Gb ESB SAN Scalable Switch*, TIPS0870
- *IBM Flex System IB6131 InfiniBand Switch*, TIPS0871

Adapters:

- *IBM Flex System CN4054 10Gb Virtual Fabric Adapter and EN4054 4-port 10Gb Ethernet Adapter*, TIPS0868
- *IBM Flex System EN2024 4-port 1Gb Ethernet Adapter*, TIPS0845
- *IBM Flex System EN4132 2-port 10Gb Ethernet Adapter*, TIPS0873

- *IBM Flex System FC3052 2-port 8Gb FC Adapter*, TIPS0869
- *IBM Flex System FC3172 2-port 8Gb FC Adapter*, TIPS0867
- *IBM Flex System FC5022 2-port 16Gb FC Adapter*, TIPS0891
- *IBM Flex System IB6132 2-port FDR InfiniBand Adapter*, TIPS0872
- *IBM Flex System IB6132 2-port QDR InfiniBand Adapter*, TIPS0890
- *ServeRAID M5115 SAS/SATA Controller for IBM Flex System*, TIPS0884

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

IBM education

The following are IBM educational offerings for IBM Flex System. Note that some course numbers and titles might have changed slightly after publication.

> **Note:** IBM courses prefixed with NGTxx are traditional, face-to-face classroom offerings. Courses prefixed with NGVxx are Instructor Led Online (ILO) offerings. Courses prefixed with NGPxx are Self-paced Virtual Class (SPVC) offerings.

- NGT10/NGV10/NGP10, IBM Flex System - Introduction
- NGT30/NGV30/NGP30, IBM Flex System p260 and p460 Compute Nodes
- NGT20/NGV20/NGP20, IBM Flex System x240 Compute Node
- NGT40/NGV40/NGP40, IBM Flex System Manager Node
- NGT50/NGV50/NGP50, IBM Flex System Scalable Networking

For more information about these education offerings, and many other IBM System x educational offerings, visit the global IBM Training website located at:

http://www.ibm.com/training

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

IBM Flex System p260 and p460 Planning and Implementation Guide

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

IBM Flex System p260 and p460 Planning and Implementation Guide

Describes the new POWER7 compute nodes for IBM PureFlex System

Provides detailed product and planning information

Set up partitioning and OS installation

To meet today's complex and ever-changing business demands, you need a solid foundation of compute, storage, networking, and software resources that is simple to deploy and can quickly and automatically adapt to changing conditions. You also need to be able to take advantage of broad expertise and proven preferred practices in systems management, applications, hardware maintenance, and more.

The IBM Flex System p260 and p460 Compute Nodes are IBM Power Systems servers optimized for virtualization, performance, and efficiency. The nodes support IBM AIX, IBM i, or Linux operating environments, and are designed to run various workloads in IBM PureFlex System.

This IBM Redbooks publication is a comprehensive guide to IBM PureFlex System and the Power Systems compute nodes. We introduce the offerings and describe the compute nodes in detail. We then describe planning and implementation steps and go through some of the key the management features of the IBM Flex System Manager management node.

This book is for customers, IBM Business Partners, and IBM technical specialists that want to understand the new offerings and to plan and implement an IBM Flex System installation that involves the Power Systems compute nodes.

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

SG24-7989-00 ISBN 0738436976