Redbooks
ibm.com/redbooks

IBM PurePower Technical Overview and Introduction

Patrick Lindsey
Andrea Nardone
Hansjörg Schneider
Martin Schulz

Cloud

Power Systems

IBM

Redpaper

International Technical Support Organization

IBM PurePower Technical Overview and Introduction

December 2015

REDP-5277-00

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

First Edition (December 2015)

This edition applies to the second release of the IBM PurePower System solution and the PurePower Manager (8374-01M).

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 grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com 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:

AIX®
DB2®
Electronic Service Agent™
FICON®
Global Technology Services®
IBM®
POWER®
POWER Hypervisor™
Power Systems™
Power Systems Software™
POWER7+™
POWER8®
PowerVM®
PureFlex®
PurePower System™
PureSystems®
Real-time Compression™
Redbooks®
Redpaper™
Redbooks (logo) ®
RS/6000®
Storwize®
System i®
System Storage®
System z®
WebSphere®

The following terms are trademarks of other companies:

C3, and Phyteland device are trademarks or registered trademarks of Phytel, Inc., an IBM Company.

Celeron, 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.

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

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

Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get up-to-the-minute Redbooks news and announcements
- Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App

iOS **Download Now** Android

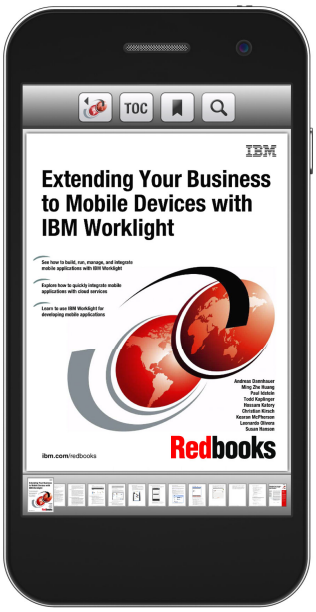

Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

This IBM® Redpaper™ publication introduces and provides a technical overview of the IBM PurePower System™ that helps support management of big data, social media, mobile, analytics, and the flow of critical information. A PurePower System can be configured in an affordable entry-level configuration in a single rack, and it is agile enough to be expanded for scalable cloud deployments. It has built-in redundancy for highly reliable and resilient operation to support demanding applications and cloud services, as required by many enterprises. A PurePower System also provides the scalability, flexibility, and versatility that you demand for business-critical workloads.

The following enhancements were announced in October 2015:

- IBM i operating system on top of a Virtual I/O Server (VIOS) now supported on the IBM Power System S822 server
- Improvements to PurePower Integrated Manager
- Integration of HMC code (virtual HMC) into the PurePower Integrated Manager
- Ability to order translated PurePower documentation that is geography-specific
- Configuration support for IBM Power System S822 and S822L server in a single rack
- PowerVC 1.2.3 Standard Edition
- Power compute node firmware SV840

Authors

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

Patrick Lindsey is an IBM WebSphere® Information Developer in Austin, Texas, in the United States. He has 15 years of experience supporting customers using IBM AIX® and WebSphere on IBM Power Systems™. Patrick has written several technical documents and product documentation for AIX and WebSphere, and he has several AIX and WebSphere administration certifications.

Andrew Nardone is a Systems Engineer at IBM Switzerland, which he joined after earning his Master's Degree in Telecommunications Engineering from the Roma Tre University in Rome, Italy. His area of expertise is mainly in IBM POWER® systems, focusing on Linux. He delivers presentations and supports customers and business partners in the German- and Italian-speaking parts of Switzerland.

Hansjörg Schneider is an IBM Power Systems Hardware Support Specialist at the Technical Support Center in Ehningen, Germany. He has 20 years of experience in Power Systems and AIX systems. He provides IBM Power Systems support to clients, especially their engineers. His areas of expertise include IBM PureFlex® and SAN, and he is also an AIX Certified Specialist.

Martin Schulz is a Power Systems Hardware Support specialist at IBM Global Technology Services® DACH office in Ehningen, Germany. He has worked in the RT PC, IBM RS/6000®, and Power Systems since 1987 and has worked at IBM Germany since 1982. His areas of expertise include Power Systems hardware. He is an AIX Advanced Technical Expert and Novell certified Linux administrator.

The project that produced this publication was managed by **Scott Vetter**, PMP.

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

George Ahrens, Kerry Anders, Kaena Freitas, Daniel Gezovich, Ray Harrington, Deborah McLemore, Jose E. Morales, Marcelo Perazolo, Marcus Richardson, Madeline Vega, and Steven A. Winch
IBM

Now you can become a published author, too

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

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

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

  ibm.com/redbooks

- Send your comments by email:

  redbooks@us.ibm.com

- Mail your comments:

  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

PurePower System description

IBM PurePower System is a complete, flexible cloud infrastructure system with integrated expertise. The system integrates and optimizes all compute, storage, and networking resources to deliver an application-ready infrastructure. These fully integrated, optimized solutions can be configured for your specific business need and can be complemented by additional support and deployment services.

1.1 Introduction

The PurePower System helps with management of big data, social media, mobile, analytics, and the flow of critical information. A PurePower System can be configured in an affordable entry-level configuration in a single rack, and it is agile enough to be expanded for scalable cloud deployments. It has built-in redundancy for highly reliable and resilient operation of demanding applications and cloud services as required by many enterprises. It also provides the scalability, flexibility, and versatility that you demand for business-critical workloads. The Ethernet and storage area networks are fully redundant.

The PurePower System solution is a complete infrastructure stack that can be deployed in hours rather than days.

PurePower System consists of the following components:

- Enterprise rack
- IBM POWER8® compute nodes
- Management nodes
- Networking and I/O switches
- IBM Storwize® V7000 storage enclosures with Version 7.4.0.2
- Software:
  - IBM AIX Version 7.1 with the 7100-03 Technology Level and Service Pack 3 and APAR IV56367 or later (5765-G98)
  - IBM i 7.1 TR11 or later (IO via VIOS only)
  - IBM i 7.2 TR3 or later (IO via VIOS only)
  - Red Hat Enterprise Linux 7.1 (BE) (5639-RH7)
  - Red Hat Enterprise Linux 7.1 (LE) (5639-RLE)
  - IBM Power Virtualization Center (PowerVC) Standard Version 1.2.3.2
  - PurePower Integrated Manager software portal plus Nagios Core open source monitoring
- Integrated software and hardware
- Three-year service and support agreement
- IBM System Lab Services initial implementation and configuration services options

1.1.1 What's new in the October 2015 release

The following enhancements were announced in October 2015:

- IBM i operating system on top of a Virtual I/O Server (VIOS) now supported on the IBM Power System S822 server

  IBM i support requires VIOS (no native I/O) and FW840.

  With firmware code level 840, the 2U Power S822 (8284-22A) supports IBM i 7.2 or IBM i 7.1 with special terms and conditions. Technology Refresh 3 or later for IBM i 7.2 or Technology Refresh 11 or later for IBM i 7.1 is required. Multiple IBM i partitions, each up to a maximum of two cores, are supported. The Power S822 software tier is P10.

The Virtual I/O Server (VIOS) is required for all of the I/Os that IBM i accesses. There is no native I/O support. Thus IBM i applications that require use of these adapters are not a good fit for the Power S822. IBM i 7.2 clients can connect to a LAN-attached OEM device that has downstream async connections.

IBM i partitions that access directly attached disk or SSD through VIOS must use 4 k byte sector drives, not 5 xx byte sector drives. The 4 k drives are required for performance reasons.

This IBM i support fulfills the Statement of Direction previously issued for PurePower systems.

Note: Async or bisync adapters or crypto-cards are not supported under the VIOS.

- Improvements to PurePower Integrated Manager

  PurePower Integrated Manager gathers and displays the current level of firmware on all components in the solution and compares those levels to the existing tested stack. It notifies the customer of any levels that are out of compliance and need manual updates to be scheduled. For more information about the PPIM, see 2.1.1, “PurePower Integrated Manager” on page 66.

- Integration of HMC code (virtual HMC) into the PurePower Management Nodes

  The vHMC enables the customer to manage the PurePower IBM POWER8 servers through a guest partition on the PurePower Integrated Manager. The HMC version changed from 8.3 to 8.4 with the October 2015 release. For more information about the vHMC, see 2.1.2, “Virtual HMC” on page 71.

- Ability to order translated PurePower documentation that is geography-specific

  Before this release, PurePower documentation was supplied only in English. With this release, clients have the ability to order publications in additional languages.

- Configuration support for IBM Power System S822 and S822L server in a single rack.

  The October 2015 release allows customers to order a mixture of compute nodes 8284-22A and 8247-22L in a single rack.

- PowerVC 1.2.3.2 Standard Edition

  With the October 2015 release, PowerVC 1.2.3.1 Standard Edition was updated to PowerVC 1.2.3.2.

- Power compute node firmware SV840

  The system firmware for the power compute node change from SV830 to SV840 with the October 2015 release.

Table 1-1 lists the changes to the software stack between the April and October 2015 releases.

Table 1-1 Changes in the software stack

April 2015	October 2015
PowerVC 1.2.3.1 Standard Edition	PowerVC 1.2.3.2 Standard Edition
Physical HMC with code 8.3	Virtual HMC with code 8.4
Power system firmware SV830	Power system firmware SV840
No IBM System i® support	IBM i version 7.1 TR11 or later support IBM i version 7.2 TR3 or later support

1.1.2 Security considerations

IBM PurePower provides prebuilt and preintegrated security at several levels, from the hardware architecture to the software stack.

With a cloud infrastructure serving several users, where many different customers share the same physical medium, it is crucial to ensure that all virtual machines are clearly separated from one another and that security exists at all levels, as shown in Figure 1-1.

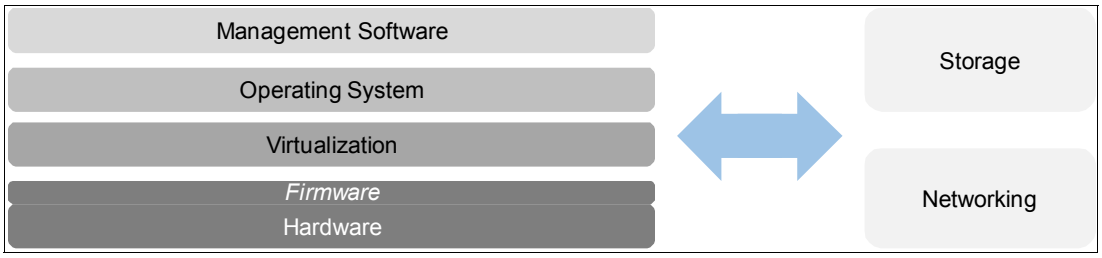

Figure 1-1 Layers of IBM PurePower that are related to security features described in this paper

Hardware

The POWER microarchitecture is designed with this security in mind. By introducing randomness and separation in those areas of the silicon that are shared by two virtual machines, Power Systems is designed so that one logical partition (LPAR) cannot steal information from another LPAR by using the same hardware resources.

IBM digitally signs all of its firmware images with a unique signature so that customers have the assurance that the updates are truly coming from IBM.

Virtualization

IBM PowerVM®, the PurePower virtualization software, has a proven track record in security because it is nested in the hardware so that security can be driven from the foundation of the system.

The capability to support more LPARs per physical server is also a security benefit. Because communication between two LPARs is more secure when they are on the same host, as compared to different hosts (in which the information needs to flow through a network cable), partitions that exchange sensitive data can be kept isolated from the physical network more efficiently.

Operating system

LPARs running IBM AIX operating system use several additional security features, including encrypted file system (EFS) technology. This ensures that data on a disk in a SAN environment, such as an IBM PurePower V7000 storage subsystem, is encrypted and remains encrypted over the network through the virtualization layer. Data is decrypted only after it arrives at the destination guest OS. This feature eases backup of the data because the data is encrypted, so there are fewer concerns about how the backed-up data is exposed.

IBM i also supports several security features, especially on the database side. Row & Column Access Control (RCAC) enables you to limit the amount of data exposed to specific users, and other options allow data to be encrypted when written to the disks. LPARs running Linux also use similar encryption features, such as the standard, platform-independent open source disk encryption specification, Linux Unified Key Setup (LUKS).

Management and monitoring environment

Security is also enforced in the main “control panel” of the system, the PurePower Integrated Manager (PPIM), which lets users choose between certificates, based on industry authentication standards. The management tools, such as the vHMC, Nagios, or PowerVC, can be all customized so that only selected users have the access rights. The vHMC administrator can also hide specific functions from the view of specific users (for role-based access control, or RBAC).

The operating system that hosts the management software, Red Hat Enterprise Linux, is a focal point for security concerns. In fact, with the amount of trust given from the managed compute nodes, it can be an attractive back door to the system. Therefore, those VMs implement several additional security measures, such as encryption of important files, including the one that contains the rack configuration, or the use of specific SELinux policy modules.

PurePower also includes Nagios, a tool that is widely used in the open source world for resource monitoring, including complete monitoring of security logs and security data. It alerts users via SMS or email when a predefined log pattern is detected.

Storage and networking

Encryption is supported by the V7000, which is the gateway to the storage area network (SAN). The storage controller encrypts the data stored in all of its attached enclosures.

On the network side, the policy used by PurePower is to keep the management network and the internal data network logically and physically separated by using different cables and switches. The vHMC communicates with the compute nodes over the Secure Sockets Layer (SSL) protocol, which provides server authentication, data encryption, and data integrity.

For a more detailed description of these security features, see 1.2.5, “Security” on page 40.

1.2 Architecture and technical overview

This section describes the physical system architecture and technical overview for the components of the PurePower System 8374-01M. For information about physical site requirements of the PurePower System see Chapter 3, “Physical planning and installation” on page 81.

Figure 1-2 shows the PurePower system.

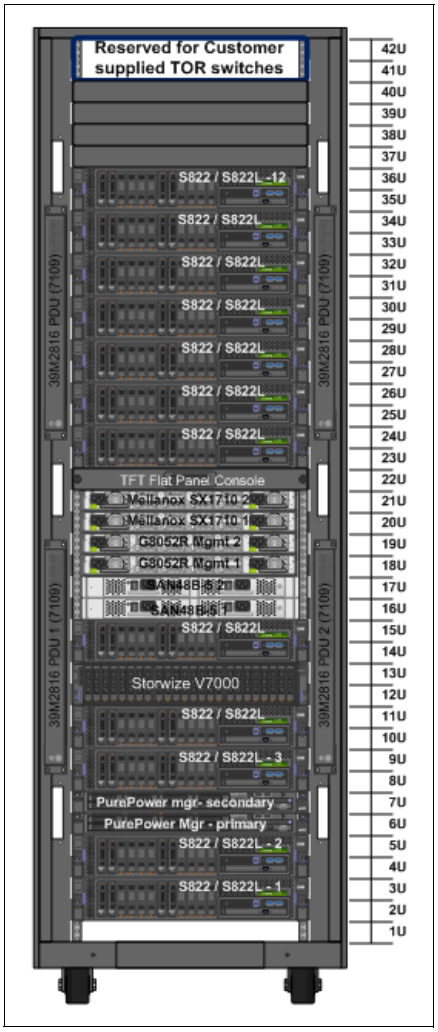

Figure 1-2 The IBM PurePower system

The PurePower solution consists of the following components:

- IBM Compute Nodes 8284-22A or 8247-22L
- IBM Storwize V7000 Gen2 storage enclosure 2076-524
- IBM System Storage® SAN48B-5 switch 2498-F48
- Mellanox SX1710 Data Network switch 8831-NF2
- Systems Management IBM 7120 Model G8052 RackSwitch 7120-48E
- PurePower System Management Node designed to provide the virtualization capabilities for the IBM PurePower 8374-01M

- IBM Service console 7316-TF4
- IBM Enterprise Rack 7014-B42

Customer supplied TOR switches may be used.

Table 1-2 lists the PurePower solution components.

Table 1-2 PurePower solution components

Component	Description	Machine type model
Compute Nodes	Power S822 or Power S822L	8284-22A 8247-22L
Storage Enclosure	Storwize V7000 Gen2 Storwize expansion enclosure (optional)	2076-524 2076-24F
Storage Switch	IBM System Storage SAN48B-5 switch	2498-F48
Internal Data Network Switch	Mellanox SX1710 40 Gb, RoCE capable Ethernet	8331-NF2
Systems Management Switch	IBM 7120 RackSwitch Model G8052 1/10 Gb	7120-48E
PurePower System Management Node	x3550 M4 server	8374-01M (based on a 7042-CR8)
Service console	TFT video and keyboard tray	7316-TF4
Rack	Enterprise rack	7014-B42
Optional Data Center Integration top-of-rack switch	Customer provided, not provided with solution	

Power S822 and Power S822L compute nodes

The PurePower Systems offers a choice of either the IBM Power System S822 8284-22A or the IBM Power System S822L 8247-22L 2U compute nodes and, with the October 2015 release, also a mix of them.

- Customers are able to add hardware components as they desire, but just not on the original order which is prebuilt and tested in manufacturing
- The POWER8 technology-based servers running Virtual I/O Server (VIOS) version 2.2.3.52 in a dual VIOS configuration (VIOS OS `rootvg` on local disk). Each POWER8 server will have four HDDs installed in a split storage backplane (6+6), two pairs of two same-size HDDs, software mirrored by the OS with one hdisk for rootvg.
- At least the first two POWER8 technology-based servers in the rack will have a DVD drive. Default memory order is 256 GB for the two processor sockets (8 DIMMs per socket x 16 GB each DIMM).
- POWER8 technology-based servers (S822) running VIOS 2.2.3.52 can deploy client guest OSes of AIX 7.1, RHEL 7.1 LE, and RHEL 7.1 BE. IBM i on top of a VIOS is supported. IBM i when running on top of a VIOS is supported.
- POWER8 technology-based servers (S822L) running PowerVM 2.2.3.5 can deploy client guest OSes of RHEL 7.1 LE and BE.
- Ubuntu 14.10 LE and SUSE Linux Enterprise Server 11 SP3 BE and SUSE Linux Enterprise Server 12 LE support will come in future releases.

System architecture for the S822 and S822L

This section describes the system architecture for the Power S822 and 822L. The bandwidths that are provided throughout the section are theoretical maximums that are used for reference.

POWER8 processor overview

The POWER8 processor is manufactured by using the IBM 22 nm Silicon-On-Insulator (SOI) technology. Each chip is 649 mm^2 and contains 4.2 billion transistors. As shown in Figure 1-3, the chip contains 12 cores, two memory controllers, PCIe Gen3 I/O controllers, and an interconnection system that connects all components within the chip. Each core has 512 KB of L2 cache, and all cores share 96 MB of L3 embedded DRAM (eDRAM). The interconnect also extends through module and system board technology to other POWER8 processors in addition to DDR3 memory and various I/O devices.

POWER8 systems use memory buffer chips to interface between the POWER8 processor and DDR3 or DDR4 memory.[1] Each buffer chip also includes an L4 cache to reduce the latency of local memory accesses.

Figure 1-3 shows the POWER8 chip.

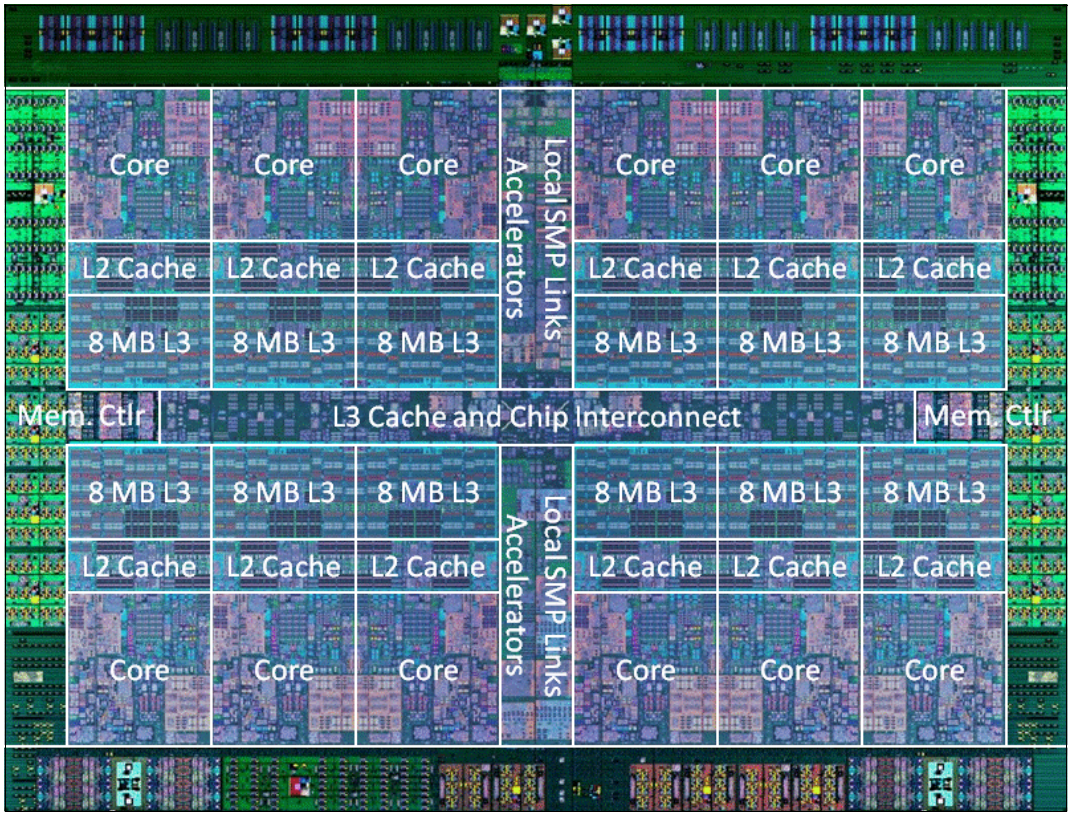

Figure 1-3 The POWER8 processor chip

[1] At the time of writing, the available POWER8 processor-based systems use DDR3 memory.

The POWER8 processor is designed for system offerings from single-socket servers to multisocket Enterprise servers. It incorporates a triple-scope broadcast coherence protocol over local and global SMP links to provide superior scaling attributes. Multiple-scope coherence protocols reduce the amount of SMP link bandwidth that is required by attempting operations on a limited scope (single chip or multi-chip group) when possible. If the operation cannot complete coherently, the operation is reissued by using a larger scope to complete the operation.

The following additional features can augment the performance of the POWER8 processor:

- Support for DDR3 and DDR4 memory through memory buffer chips that offload the memory support from the POWER8 memory controller.
- L4 cache within the memory buffer chip that reduces the memory latency for local access to memory behind the buffer chip; the operation of the L4 cache is transparent to applications running on the POWER8 processor. Up to 128 MB of L4 cache can be available for each POWER8 processor.
- Hardware transactional memory.
- On-chip accelerators, including on-chip encryption, compression, and random number generation accelerators.
- Coherent Accelerator Processor Interface (CAPI), which allow accelerators plugged into a PCIe slot to access the processor bus using a low latency, high-speed protocol interface.
- Adaptive power management.

There are two versions of the POWER8 processor chip. Both chips use the same building blocks. The scale-out systems use a 6-core version of POWER8. The 6-core chip is installed in pairs in a dual-chip module (DCM) that plugs into a socket in the system board of the systems. Functionally, it works as a single chip.

Table 1-3 summarizes the technology characteristics of the POWER8 processor.

Table 1-3 Summary of POWER8 processor technology

Technology	POWER8 processor
Die size	649 mm^2
Fabrication technology	► 22 nm lithography ► Copper interconnect ► SOI ► eDRAM
Maximum processor cores	6 or 12
Maximum execution threads core/chip	8/96
Maximum L2 cache core/chip	512 KB/6 MB
Maximum On-chip L3 cache core/chip	8 MB/96 MB
Maximum L4 cache per chip	128 MB
Maximum memory controllers	2
SMP design-point	16 sockets with IBM POWER8 processors
Compatibility	With prior generation of POWER processor

Memory subsystem

Converged Infrastructure customers can order memory following the S822 default memory configuration rules. The default memory in IBM Configurator for e-business (e-config) for a Converged Infrastructure order should be 16 x 16 GB (8 DIMMs on each processor socket). An e-config tip should advise customers that they should order at least 8 DIMMs for best performance.

The Power S822 is a two socket system that supports up to two POWER8 processor modules. The server supports a maximum of 16 DDR3 CDIMM slots, with eight DIMM slots per installed processor. Memory features that are supported are 16 GB, 32 GB, and 64 GB, which run at speeds of 1600 MHz, allowing for a maximum system memory of 1024 GB.

System bus

This section provides information about the internal buses.

The Power S822 server has internal I/O connectivity through Peripheral Component Interconnect Express (PCI Express or PCIe) Gen3 (PCI Express Gen3 or PCIe Gen3) slots and also external connectivity through SAS adapters.

The internal I/O subsystem on the Power S822 is connected to the PCIe Controllers on a POWER8 processor module in the system. Each POWER8 processor module has a bus that has 48 PCIe lanes running at 8 Gbps full-duplex and provides 96 GBps of I/O connectivity to the PCIe slots, SAS internal adapters, and USB ports.

Some PCIe devices are connected directly to the PCIe Gen3 buses on the processors, and other devices are connected to these buses through PCIe Gen3 Switches. The PCIe Gen3 Switches are high-speed devices (512 - 768 GBps each) that allow for the optimal use of the processors PCIe Gen3 x16 buses by grouping slower x8 or x4 devices that plug into a x16 slot and not use its full bandwidth.

Internal I/O subsystem

The internal I/O subsystem is on the system board, which supports PCIe slots. PCIe adapters on the Power S822 can be hot-plugged.

All PCIe slots support Enhanced Error Handling (EEH). PCI EEH-enabled adapters respond to a special data packets that are generated from the affected PCIe slot hardware by calling system firmware, which examines the affected bus, allows the device driver to reset it, and continues without a system reboot. For Linux, EEH support extends to the most frequently used devices, although certain third-party PCI devices might not provide native EEH support.

PCI Express

PCI Express (PCIe) uses a serial interface and allows for point-to-point interconnections between devices (using a directly wired interface between these connection points). A single PCIe serial link is a dual-simplex connection that uses two pairs of wires, one pair for transmit and one pair for receive, and can transmit only one bit per cycle. These two pairs of wires are called a *lane*. A PCIe link can consist of multiple lanes. In such configurations, the connection is labeled as x1, x2, x8, x12, x16, or x32, where the number is effectively the number of lanes.

The PCIe interfaces that are supported on this server are PCIe Gen3, and they are capable of 16 GBps simplex (32 GBps duplex) on a single x16 interface. PCIe Gen3 slots also support previous generations (Gen2 and Gen1) adapters, which operate at lower speeds, according to the following rules:

- Place x1, x4, x8, and x16 speed adapters in the same connector size slots first, before mixing adapter speed with connector slot size.

- Adapters with smaller speeds are allowed in larger sized PCIe connectors but larger speed adapters are not compatible in smaller connector sizes (that is, a x16 adapter cannot go in an x8 PCIe slot connector).

All adapters support EEH. PCIe adapters use a different type of slot than PCI adapters. If you attempt to force an adapter into the wrong type of slot, you might damage the adapter or slot.

I/O adapters S822

Table 1-4 provide the standard I/O adapter configuration for the S822 Power compute nodes.

Table 1-4 Standard I/O adapter configuration for the S822 Power compute nodes

FC	Description	Adapter slots	Quantity
EC3A	PCIe3 LP 2-Port 40 GbE NIC RoCE QSFP+ Adapter	C3®, C6	2
EN0B	PCIe2 LP 16 Gb 2-port Fibre Channel Adapter	C5, C7	2
EN0V	PCIe2 LP 4-port (10 Gb + 1 GbE) Copper SFP+RJ45 adapter	C2, C10	2

I/O adapters S822L

Table 1-5 provide the standard I/O adapter configuration for the S822L Power compute nodes.

Table 1-5 Standard I/O adapter configuration for the S822L Power compute nodes

FC	Description	Adapter slots	Quantity
EC3A	PCIe3 LP 2-Port 40 GbE NIC RoCE QSFP+ Adapter	C3,C6	2
EL43	PCIe3 LP 16 Gb 2-port Fibre Channel Adapter	C5,C7	2
EN0V	PCIe2 LP 4-port (10 Gb + 1GbE) Copper SFP+RJ45 Adapter	C2, C10	2

Internal storage

The internal storage on the Power S822 server depends on the DASD/Media backplane that is used. The server supports two DASD/Media backplanes: #EJ0T and #EJ0V.

The Power nodes in the PurePower use the base storage backplane (FC EJ0T) with the split 6+6 option (FC EJ0V). Four disks (2 pairs) of the same size are required to be on the order. Two disks are to be placed in each half of the split backplane. Up to 12 disks max can be ordered with each S822 node.

For configuration details and the possible minimum maximum configuration, see 1.2.4, “Capacity (min/max)” on page 36.

For more information about the Power S822 or Power S822L, see the IBM Redpaper publication titled *IBM Power System S822 Technical Overview and Introduction*, REDP-5102:

http://www.redbooks.ibm.com/abstracts/redp5102.html

IBM Storwize V7000 Gen2 SAN characteristics and system architecture

This section describes the characteristics and system architecture of the IBM Storwize V7000 Gen2 and the SAN of the PurePower System.

Storwize V7000 Gen2 and SAN characteristics

The Storwize V7000 Gen2 and SAN have the following characteristics:

- The V7000 Gen2 has an 8 x 1.2 TB HDD configured with one RAID5 array, with a spare with 8 Gb Fibre Channel (FC) connections to the SAN48B5 (IBM MTM 2498F48) a 48-port 16 Gb switch.
- A single pool configured from the MDisk is configured on the V7000 Gen2 for initial factory shipment.
- Brocade Fibre Channel switch zoning is supported through the Brocade OpenStack Cinder zone manager driver with PowerVC.
- Initial deployments using PowerVC create NPIV-based client guest OS (LPAR) on the POWER8 technology-based servers by using virtual Fibre Channel (vFC) on the client guest OS.
- If a customer wants to deploy a vSCSI-based client guest OS (LPAR) using PowerVC, the preferred SAN client attachment type can be appropriately set when the client guest OS is deployed.
- The host zones on the V7000 Gen2 for each VIOS are created by WWPN, not by port.
- The V7000 Gen2 has a dedicated iSCSI port set with an interface for the management nodes to access the recovery volume (`/data/purekvm/recovery_volume.properties`):
  - Controller (Canister) 1 – 192.168.93.4 (attached to G8052 left switch 1)
  - Controller (Canister) 2 – 192.168.93.7 (attached to G8052 right switch 2)

IBM Storwize V7000 Gen2 system architecture

The Storwize V7000 Gen2 has been improved substantially to support more advanced processors, more memory, and faster interconnects. This is also the first time that the IBM SAN Volume Controller (SVC) and Storwize share the same processors.

The following items provide details that can help you understand the changes made across both to meet the goals:

- Processors

  IBM SAN Volume Controller DH8 and Storwize V7000 Gen2. Both platforms use the Ivy Bridge processors from Intel, which has eight cores.

- Memory

  32 GB or cache and compression with an option to have another 32 GB is added for IBM Real-time Compression™ workloads.

- Peripheral Component Interconnect (PCI) Express (PCIe) technology

  Both the platforms have multiple PCIe Gen3 slots, as compared to dual PCIe Gen 2 slots in previous versions. This shift to PCIe Gen3 enables each PCIe lane to get a maximum speed of 1000 megabytes per second (MBps).

- Optional adapters

  In previous models, the only option that customers had was to add a dual-port 10 Gbps converged network adapter. The Storwize V7000 Gen1 base model has dual-port 1 Gbps adapters, plus quad-port 8 Gbps FC adapters. In both of the new platforms, the base models come with three 1 Gbps Ethernet ports. However, customers have an option to select multiple add-on adapters for driving host input/output (I/O) and offloading compression workloads.

Expansion enclosure

The expansion enclosure contains the following major components:

- SAS 3.0 expander:
  - PMC-Sierra SXP 12G chip
  - Speed-matching buffers for optimum performance with 6G devices
  - 1x links to the midplane for the internal drive slots
  - 2-off 4x Mini-SAS HD connectors on the back-panel:
    - One port is connected to the upstream enclosure.
    - The remaining port is available for attaching a downstream enclosure.
  - TWI buses 1 and 2 to the midplane to access the VPD EPROMs and the VPD and status of the power supplies
  - TWI bus 3, which is used to access canister CPLD and temperature sensors
  - Enclosure management processor and SES firmware with external memories:
    - 16 MB SPI Flash with cyclic redundancy check (CRC) protection
    - 8 MB static random access memory (SRAM)
    - At least 64 KB nonvolatile SRAM (NVSRAM)
- RS232 port on the back panel (3.5 mm stereo jack) used for configuration during manufacturing.

For more information, see *Implementing the IBM Storwize V7000 Gen2*, SG24-8244:

http://www.redbooks.ibm.com/abstracts/sg248244.html

IBM System Storage SAN48B-5 switch 2498-F48

The IBM PurePower system provides redundant SAN data network. This section describes the overall system architecture for the storage area network IBM System Storage SAN48B-5.

The SAN48B-5 switch meetss the demands of hyperscale private or hybrid cloud storage environments by delivering 16 Gbps Fibre Channel technology and capabilities that support highly virtualized environments. To enable greater flexibility and investment protection, the SAN48B-5 switch can be configured for 24, 36, or 48 ports, and it supports 2, 4, 8, 10, or 16 Gbps speeds in an efficiently designed 1U package.

This switch is now enhanced with enterprise connectivity options that add support for IBM FICON® connectivity. It can provide a highly reliable infrastructure when used with fast and scalable IBM System z® servers.

The SAN48B-5 switch supports multitenancy in cloud environments through virtual fabrics, quality of service (QoS), and fabric-based zoning features. It also enables secure metro extension to virtual private or hybrid clouds with 10 Gbps Dense Wavelength Division Multiplexing (DWDM) link support and in-flight encryption and data compression over inter-switch links (ISLs).

Organizations can have up to four ports at 8 Gbps and up to two ports at 16 Gbps of in-flight encryption and data compression per SAN48B-5 switch. The switch also features onboard data security and acceleration, which minimizes the need for separate acceleration appliances to support distance extension. Internal fault-tolerant and enterprise-class RAS features help minimize downtime to support mission-critical cloud environments.

Key features and technical capabilities

The SAN48B-5 switch includes the following features:

- 16 Gbps performance with up to 48 ports in an energy-efficient, 1U form factor
- Ports on Demand (PoD) capabilities for scaling 24 - 48 ports in 12-port increments
- 2, 4, 8, 10, or 16 Gbps speed on all ports, producing an aggregate 768 Gbps full-duplex throughput
- 16 Gbps optimized ISLs
- 128 Gbps high-performance and resilient frame-based trunking
- Simplified deployment process and point-and-click user interface

The SAN48B-5 switch provides a critical building block for today's highly virtualized private or hybrid cloud storage environments. This switch can simplify server virtualization and virtual desktop infrastructure (VDI) management while meeting the high-throughput demands of enterprise data centers with the following capabilities:

- Native 10 Gbps Fibre Channel support, which provides integrated dense wavelength division multiplexing metro connectivity (DWDM)
- Multitenancy in cloud environments through virtual fabrics, integrated routing, QoS, and fabric-based zoning features
- In-flight data compression and encryption, which provides efficient link use and security (up to four ports at 8 Gbps or up to two ports at 16 Gbps)

Technical capabilities

The Gen 5 Fibre Channel SAN48B-5 cloud storage networking-optimized switch offers the following capabilities:

- 48 x 16 Gbps Fibre Channel ports supporting 2, 4, 8, 10, and 16 Gbps speeds or FICON connections
- Scalable PoD configuration of 24 -48 ports to accommodate various sizes of SAN fabric environments
- Configurable ports for 10 Gbps metro optical connectivity
- Small footprint and low energy consumption in 1U (fewer than 18 inches wide and 18 inches deep)
- Easy to use three-step deployment with EZSwitchSetup architecture

Architecture and key components

In addition to scalability, the SAN48B-5 switch can address demanding reliability, availability, and serviceability (RAS) requirements to help minimize downtime to support mission-critical cloud environments through the following features:

- Fabric Vision technology simplifies management, reduces cost of operation, and improves application performance.
- ClearLink diagnostic technology (D-port) identifies and isolates optic and cable problems, which simplifies deployment and support of high-performance fabrics.

- Redundant, hot-pluggable components and nondisruptive software upgrades enable easier administration and increase availability.
- Real-time power monitoring enables users to monitor real-time power use of the fabric at a node level.

For more information about the storage area network 2498-F48 Switch, see the Redbooks product guide titled *IBM System Storage SAN48B-5*:

http://www.redbooks.ibm.com/abstracts/tips1125.html

Internal 40 GbE Mellanox SX1710 Data Network Switch

The IBM PurePower system will provide redundant data networks. This section describes the Internal 40 GbE Mellanox SX1710 Data Network Switch 8831-NF2

The Mellanox switch provides high-speed 40 Gb interconnection between all the compute nodes

The switch comes standard with Mellanox operating system (MLNX-OS) software. The MLNX-OS software includes CLI, GUI, SNMP, and XML gateway interfaces. The XML Gateway provides support to retrieve and set management information over HTTP/HTTPS or SSH. The MLNX-OS enables the user to define and manage logging, email alerts, and security capabilities including RADIUS, TACACS+, AAA, and LDAP.

Key features

The Mellanox 40 GbE TOR switch features the following characteristics:

- 36 ports with 40/56 GbE in a 1U form factor
- Up to 64 ports with 10 GbE
- 220 nsec for 40 GbE
- 270 nsec for 10 GbE
- CPU Intel Celeron 1047UE (x86)
- Typical power consumption per port is 1.3 W
- Redundant and hot swappable power supplies and fans provide high availability
- Unified Fabric Manager (UFM) software can maintain thousands of nodes and switches
- Configure and manage the data center from a single location
- Fewer cables with 40/56 GbE high density
- Easy deployment and maintenance

Systems Management IBM 7120 Model G8052 RackSwitch 7120-48E

The IBM PurePower system will provide redundant management networks. This section describes the Systems Management Switch Lenovo G8052 Type 7120-48E.

Features and specifications

Note: Features and specifications listed in this section are based on Networking OS 7.9.

Ports

- Forty-eight auto-sensing 10/100/1000 Mb Ethernet ports with RJ-45 connectors
- Four ports for 1 Gb or 10 Gb Ethernet SFP/SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or 10GBASE-ER) or SFP+ direct-attach copper (DAC) cables. SFP+ modules or DAC cables are not included and must be purchased separately.

- One RS-232 serial port (Mini-USB connector) that provides an additional means to configure the switch module
- One USB port for mass storage devices

Scalability and performance

- 1 Gb and 10 Gb Ethernet ports for bandwidth optimization and performance
- Non-blocking architecture with wire-speed forwarding of traffic; up to 176 Gbps of full duplex switching throughput
- Up to 132 Million packets per second (Mpps) with switching latency of 1.8 microseconds
- Media access control (MAC) address learning: Automatic update, support of up to 32,000 MAC addresses
- Static and LACP (IEEE 802.3ad) link aggregation
- Support for jumbo frames (up to 12,288 bytes)
- Broadcast/multicast storm control
- IGMP snooping for limit flooding of IP multicast traffic
- IGMP filtering to control multicast traffic for hosts participating in multicast groups
- Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses or both
- Fast port forwarding for rapid STP convergence

For detailed information about the Systems Management Switch Lenovo G8052, see the *Lenovo RackSwitch G8052 Product Guide*:

https://lenovopress.com/tips0813

PurePower Management Node

This section describes the Management Nodes 8374-01M. The PurePower Management Nodes is designed to provide the virtualization capabilities for the IBM PurePower Rack.

Since the October 2015 release, there are two IBM 8374-01Ms in the PurePower rack configuration as the management nodes.

The 8374-01M at EIA 6 is called the *Primary Management Node*, and the 8374-01M at EIA 7 is called the *Secondary Management Node*.

PurePower includes intelligently integrated network, compute and storage in a single rack configuration.

Virtualization capabilities are provided by PowerVC which leverages OpenStack and PowerVM. Monitoring of all system hardware and optionally VM will be provided by Nagios open source monitoring.

For more information about Nagios see 2.1.4, “Nagios” on page 75.

The PurePower Management Nodes include a management network to provide management access to all the components within the rack from a portal. Firmware updates are the responsibility of the customer but can be accomplished from a single point of access via the PurePower Integrated Manager, which runs on the management node.

For more information about the PurePower Integrated Manager see 2.1.1, “PurePower Integrated Manager” on page 66.

As delivered, the PurePower rack two management nodes are preconfigured identically. The POWER8 technology-based servers will be managed by both Side A and Side B, with the new virtual HMC (vHMC). Each side will manage the FSP on a unique subnet. For more information about the vHMC, see 2.1.2, “Virtual HMC” on page 71.

The management node recovery volume provides manual synchronization of the active databases, data stores, and other relevant configuration data. Data for each VM can be replicated to the preferred management node. (See “Primary and Secondary Management Node Data Considerations” on page 50.) This allows up-to-date configurations to provide redundant management capabilities of the PurePower rack.

The Primary Management node and the Secondary Management node are mirror images of the four VMs which provide the capabilities to manage the PurePower rack.

VM specifications for the Pure Power Management Nodes

On the primary management node (Side A) the VM have these attributes:

- 192.168.93.46 RHEL 7.1 PurePower Integrated Manager and Nagios VM
- 192.168.93.45 RHEL 7.1 PowerVC Standard Edition VM
- 192.168.93.47 RHEL 7.1 Service VM
- 192.168.93.61 vHMC 840 VM

On the secondary management node (Side B) the VM have these attributes:

- 192.168.93.146 RHEL 7.1 PurePower Integrated Manager and Nagios VM
- 192.168.93.145 RHEL 7.1 PowerVC Standard Edition VM
- 192.168.93.147 RHEL 7.1 Service VM
- 192.168.93.161 vHMC 840 VM

Figure 1-4 shows the PurePower Management Node network connections.

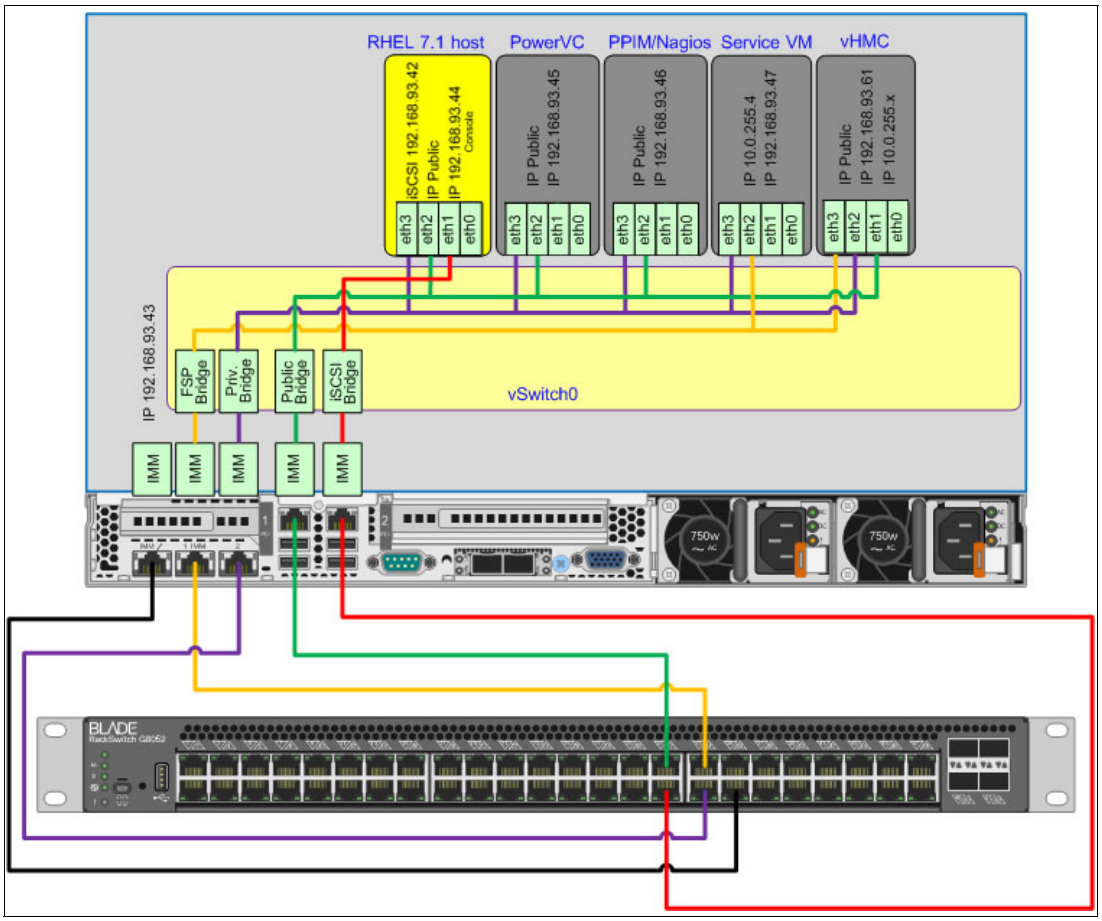

Figure 1-4 The PurePower Management Node network connections

IMM service processor cabling specifications

The primary management node IMM service processor will be cabled to the G8052 (1G) Management Switch 1:

- VLAN 3000 – Mgmt vLAN – 192.168.93.43 KVM IMM
- VLAN 3000 – Mgmt vLAN – 192.168.93.44 KVM Host

The secondary management node IMM service processor will be cabled to the G8052 (1G) management Switch 2:

- VLAN 3000 – Mgmt vLAN – 192.168.93.143 KVM IMM
- VLAN 3000 – Mgmt vLAN – 192.168.93.144 KVM Host

Hardware specifications of the 8374-01M PurePower Management Node

The hardware of the Management Node based on HMC 7042-CR8 IBM, x3550 M4 equivalent to x server 7914-F3C.

The Management Node requires the following hardware specifications:

- Intel-based, two socket (8 cores per socket, 2 threads per core, total of 32 procs), 2.0 GHz
- Processor with IVB CPU/memory support
- 12 DIMM slots
- 192GB memory installed
- Ivy bridge memory stretch
- 2 x 1TB hard disk drives with RAID1 (1 TB available storage 7.2K 6Gbps NL SATA 2.5” SFF HS HDD)
- Four integrated Ethernet RJ45 ports and an RJ45 port for the IMM2 Ethernet cabling for management capabilities
- Updated I2C topology
- SATA CD/RW DVD-RAM
- 550 W AC redundant, hot-swap power supplies
- Redundant, hot-swap fans (up to 40 fans that are 56 mm, with 6 standard, supports redundancy)
- SSD/JBOD-focused RAID controller
- Four USB ports (two front, two rear)
- One PCIe slot
- One PCI-X slot

Note: The 8374-01M Pure Power Management Node does not offer an internal or external modem.

Figure 1-5 shows the rear side with connectors of a PurePower Management Node.

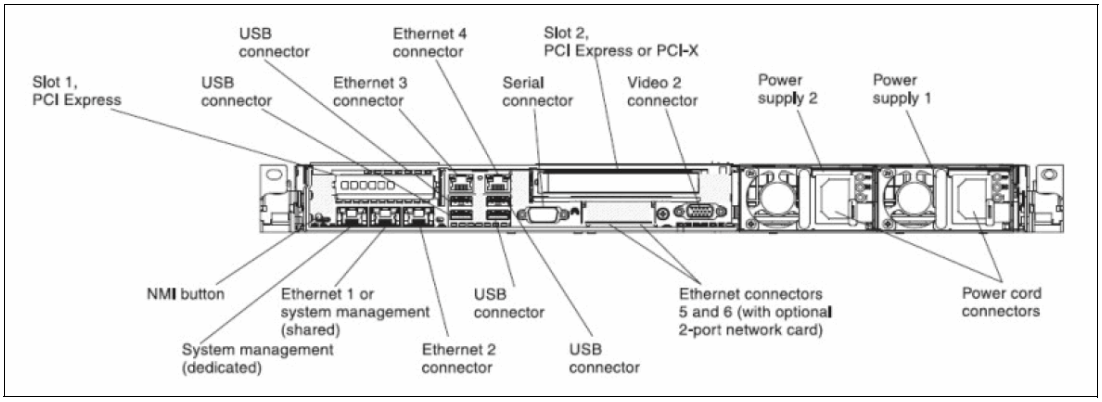

Figure 1-5 The rear side with connectors of a PurePower Management Node

Service console 7316-TF4

This section describes the 7316-TF4 18.5-Inch Flat Panel Rack-Mounted Monitor and Keyboard console.

Console specifications and features

Here are the Service console 7316-TF4 specifications and features:

- Mounts on slide-rails in the rack to enable easy movement and storage of the monitor
- Tool less rack installation in the 1U space of IBM rack configurations
- 18.5-inch, 16:9 ratio LCD panel with a VGA connection to the server or KVM switch
- Support for wide screen and previous (through scaling) resolutions
- Compatible with worldwide power and regulatory requirements
- Cable-management arm comes preinstalled on the rear of the console
- Can be shipped installed in an IBM rack

Enterprise Rack 7014-B42

This section describes the Enterprise Rack 7014-B42.

Model highlights 7014-B42

The IBM 7014 Model B42 2.0-meter rack for the IBM PurePower system enables you to maximize the space and better manage the cables in your data center. Model B42:

- Delivers 42U (EIA units*) capacity in an industry-standard 19-inch rack
- Is designed to enable cables to be routed on either side of the rack drawer mounting area
- Has a ruggedized rack option for added earthquake protection
- Includes numerous anchor points formed into the frame where cables can be tied down for strain relief
- Meets the Electronics Industries Association EIA-310-D standard

Note: An EIA unit is an industry standard for indicating vertical mounting space in a rack, where 1U stands for 44.45 mm (1.75 in).

Description

Model B42 Racks can be configured with or without side panels. Without side panels, multiple racks can be bolted together into a contiguous suite, using optional side-to-side rack-connecting hardware. When multiple racks are joined in this way, cables can be run between the racks without having to exit the continuous rack enclosure. A small gap is maintained between adjacent racks, which is filled by three matching steel trim pieces that snap into place on the front, top, and rear. No tools are required for installation.

Optional side panels can be quickly installed or removed from the outside of the rack. The removable side panels are needed only for the two end racks of the suite. On all Model B42 Racks, small side panels are included to cover the portion of the rack that extends above 1.8 meters.

The rack can be reinforced by using an optional Ruggedized Rack Feature, which includes hardware for anchoring the rack to a concrete floor. A large steel brace that bolts into the rear of the rack is hinged so it can swing out of the way for access to the rack drawers when necessary. Steel filler panels are included for any unoccupied spaces in the rack.

The standard equipment shipped with each rack includes the following components:

- A rear door
- Front and rear anti-tilt plates
- Four leveling feet
- Snap-in filler panels for any unoccupied rack space

Physical specifications of the IBM 7014 Model B42 rack

These are the physical specifications of the IBM 7014 Model B42 rack:

- Width:

  644 mm (25.4 in) with side panels

  623 mm (24.5 in) without side panels

- Depth:

  1098 mm (43.3 in) with front and rear doors

  1042 mm (41.0 in) with rear door only

  1409 mm (55.5 in) with acoustic doors

- Height:

  2015 mm (79.3 in) with standard AC power

- Weight (empty rack with doors and side panels):

  261 kg (575 lb)

Operating environment

The Operating environment of the IBM 7014 Model B42 Rack:

- Temperature: 10 - 40 degrees C
- Relative humidity: 8 - 80%
- Maximum wet bulb: 27 degrees C
- Sound power: None

Power Requirements

Model B42 rack (when ordered with either optional feature number 7189 or feature number 7196):

- Operating voltage: 200 to 240 V ac 50/60 Hz
- Power source loading: 4.8 KVA per Power Distribution Unit (PDU)
  Rack may have up to 8 PDU.

The rack power distribution system provides 200 to 240 V ac (50Hz/60Hz) to the system components (drawers). Each PDU is capable of providing up to 4.8 KVA.

For more information, see the 7014-B42 IBM Rack Model B42 web page:

http://ibm.co/1NgR3uy

Top of Rack switch

Data Center Integration Switch TOR switch will be customer placed.

Two EIA units are reserved at the top of the rack for placement of customer provided switches to ease integration with customers data center network.

1.2.1 Network components

IBM PurePower is designed to provide the maximum availability and performance for cloud workloads. On the network level, this is achieved through the use of separate high-speed networks:

Data network For the overall data traffic within the nodes (internal network) and the customer's data center network (uplink network)

Management network Dedicated to management operations and monitoring data traffic. It also connects the flexible service processors (FSP) to the vHMC

Storage network For the storage area network (SAN) between the compute nodes and the storage subsystem where most of the data is located

To provide resiliency, each network cable, switch, and adapter is redundant.

Figure 1-6 provides a high-level illustration of the networks.

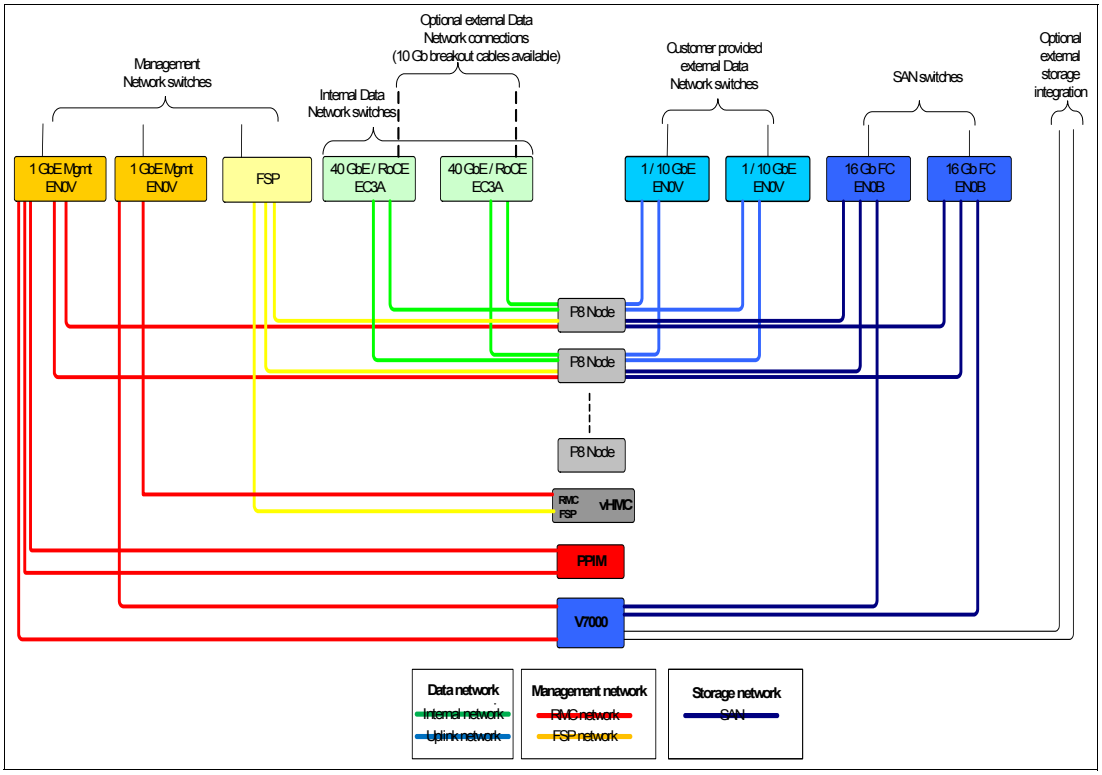

Figure 1-6 High-level illustration of the data, management and storage networks

For a more detailed one, see Figure 1-7 on page 25 and Figure 1-8 on page 27.

The sections that follow describe the three networks and their devices more in detail.

Data network

The IBM PurePower data network is actually made up of two networks:

40 Gb Internal network For east-west (between compute nodes) data traffic

10 Gb Uplink network The cables connecting to the customer's data center network

The following adapters are on the nodes servicing those networks:

- PCIe3 LP 2-Port 40 GbE NIC RoCE QSFP+ adapter (FC EC3A; CCIN 57BD)
- PCIe2 LP 4-port (10 Gb + 1 GbE) Copper SFP+RJ45 adapter (FC EN0V; CCIN 2CC3)

Internal network

IBM PurePower features a 40 Gb RoCE[2] network for the internal data traffic. Each compute node connects to this network through its two 40 GbE / RoCE adapters (EC3A). This network makes use of a Mellanox SX1710 switch, part of the IBM PurePower offering.

High availability is achieved through the redundancy of the network components and the use of *Multi-Chassis Link Aggregation* (MLAG), providing both link-level and node-level redundancy.

Multi-Chassis Link Aggregation (MLAG) on IBM PurePower

MLAG is implemented on IBM PurePower in the following manner.

Each port pair on the switch connects to two ports on each compute node, located on two different EC3A adapters; these connections are aggregated using the Etherchannel protocol, so that the overall bandwidth of the channel doubles that of the single connections. In addition to that, if one connection fails, the other one takes over automatically without disruption. This is called *link-level redundancy*, typical of a *Link Level Aggregation* (LAG).

MLAG however extends the availability provided by LAG by adding another degree of redundancy, with the use of two switches instead of one. The two communicate between themselves with a special protocol called *inter-peer link* (IPL) that enables two separate switches to act like one towards other servers or networking devices; if one switch fails, the other one will take over without disruption. This is called *node-level redundancy*.

Uplink network

For this network, customers can use either of the following options:

- The Mellanox switches that are part of the IBM PurePower offering, connected to the compute nodes' 40 GbE / RoCE adapters (EC3A)
- The top of rack (TOR) switches of their own, which are not included in the IBM PurePower offering; they will connect to the compute nodes' 10 GbE ports of the 1/10 GbE adapters (EN0V)

For the latter option, two rack units (2U) are reserved at the top of the rack for two redundant devices.

High availability is attained through the use of standard LAG. Trunk encapsulation can be implemented between the switches, in order to allow information from different VLANs, flowing on the same link, to be maintained in Ethernet frames; Cisco System's proprietary *Inter-Switch Link* protocol (ISL) can be used.

[2] RDMA over Converged Ethernet, a network protocol allowing remote direct memory access (RDMA) over an Ethernet network. This technology offloads the processor from networking tasks, improving the performance and the scalability of the processor.

Figure 1-7 on page 25 illustrates the data network in detail. For the supported network configurations of the customer data network, see 1.2.6, “Use cases for data center integration” on page 43.

Management network

The management network provides access to the hardware elements of IBM PurePower, and is kept separated from the networks dedicated to data traffic. It is used for accessing PowerVC, the vHMC, PPIM and Nagios, and the Resource, Monitoring and Control (RMC) network.[3] Part of the management network also connects the compute nodes FSP to the management nodes hosting the redundant vHMCs.

These parts of the management network are kept separated via VLAN tagging:

- VLAN 3050 and 3055: FSP network
- VLAN 3000: RMC network

On the compute nodes, the standard ports labeled “HMC1” and “HMC2” are used for the FSP network, and receive an address via DHCP from the Service VM; for the RMC network, one 1GbE port on each redundant 1/10 GbE (EN0V) adapter is used. Both of these ports are connected to the two redundant management switches.

On the management nodes, the four ports on each redundant Ethernet adapter are used this way:

Port 1	For the FSP connections. Connects to the management switch.
Port 2 and Port 3	For the RMC network. They are redundant and connect to the management switches via LAG.
Port 4	For remote access to the management node from the uplink network, does not connect to the management switches, but to the TOR switches.

The two redundant management switches, part of the IBM PurePower offering, are the Lenovo RackSwitch G8052.

[3] This network facilitates the HMC to perform its operations on the LPARs, like Live Partition Mobility (LPM), dynamic LPAR operations, VIOS management operations, and so on.

Figure 1-7 shows a detailed picture of the data network and the management network, from the point of view of one of the compute nodes.

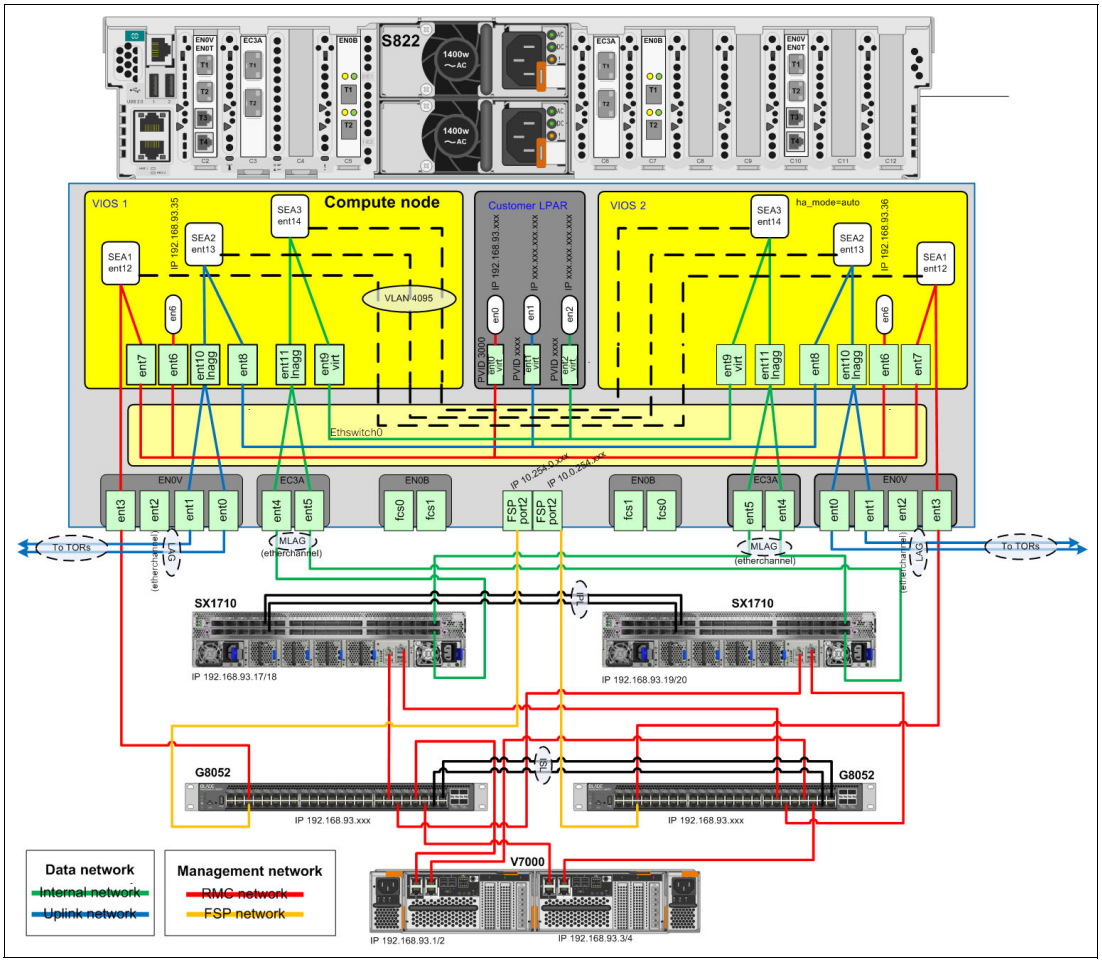

Figure 1-7 The data network and the management network from a compute node

Storage network

The connection between the compute nodes and the Storwize V7000 Gen2 enclosures is what constitutes the storage network.

The two redundant Fibre Channel (FC) adapters used by each compute node for this network are PCIe2 LP 16 Gb 2-port Fibre Channel Adapter (FC EN0B; CCIN 577F).

They connect to the two redundant 48-port, 16 Gb-capable SAN switches, the IBM System Storage SAN48B-5 (IBM MTM 2498-F48). The SAN switches connect to the V7000 but they also contain extra ports for integration into the customer's existing SAN environment or to accommodate additional V7000 controllers (not factory configured).

Each V7000 drawer can contain up to 24 hard disk drives (HDDs), each one with a capacity of 1.2 TB, configured with RAID5. IBM PurePower includes a minimum of eight drives. Additional enclosures can be attached, up to a maximum of 10, bringing the total capacity from a minimum of 9.6 TB to a maximum of 316.8 TB. For more information on the minimum and maximum capacity of the system, please refer to section 1.2.4, "Capacity (min/max)" on page 36.

All of the data will be located on the SAN. Only the VIOS LPARs of each compute node will use the internal disks located on the 6+6 split backplane (FC EJ0T with EJ0V option). Please refer to section “Software installed on the compute nodes” on page 35 for information on how the internal storage of the compute nodes is handled.

SAN zoning

Host connections are defined for every VIOS and Client OS LPAR. The VIOS will be mapped to the WWPN of the four ports of the two *physical* FC adapters owned by the node, while the OS LPARs will be mapped to WWPNs of *virtual* FC (vFC) adapters. In other words, the FC adapters will be all virtualized by the VIOS and made available to the client LPARs, no physical adapter will be assigned to the client LPARs directly.

The vFC are virtualized by the VIOS with NPIV; if customer desires to use vSCSI for newly deployed LPARs, the preferred SAN client attachment type can be appropriately set during LPAR deployment with PowerVC.

Each OS LPAR will have four vFC adapters; each vFC gets two WWPN assigned, both to be mapped to one of the four WWPNs on the V7000 Fibre Channel (FC) ports. In the end, the two WWPNs of each vFC should be mapped to the same port on the V7000.

In all, 8 zone paths must be mapped for each client OS. One vFC will follow one path, the other vFC will follow another path; these are marked with the colors green and blue in Figure 1-8 on page 27.

There should also be the following two zones:

- Port 1 of the physical Fibre Channel (FC) adapter in slot 3 of *canister 1* → port 1 of the FC adapter in slot 3 of *canister 2*
- Port 2 of the FC adapter in slot 3 of *canister 1* → port 2 of FC adapter in slot 3 of *canister 2*

After PowerVC, installed on the management node, has added the V7000 to its storage configuration and captured the OS on the compute node, it will handle all additional zoning between host vFC adapters and the V7000.

PowerVC also handles creation of the virtual Fibre Channel adapters in the host LPARs.

Figure 1-8 is a detailed picture of the storage network from the point of view of one of the compute nodes.

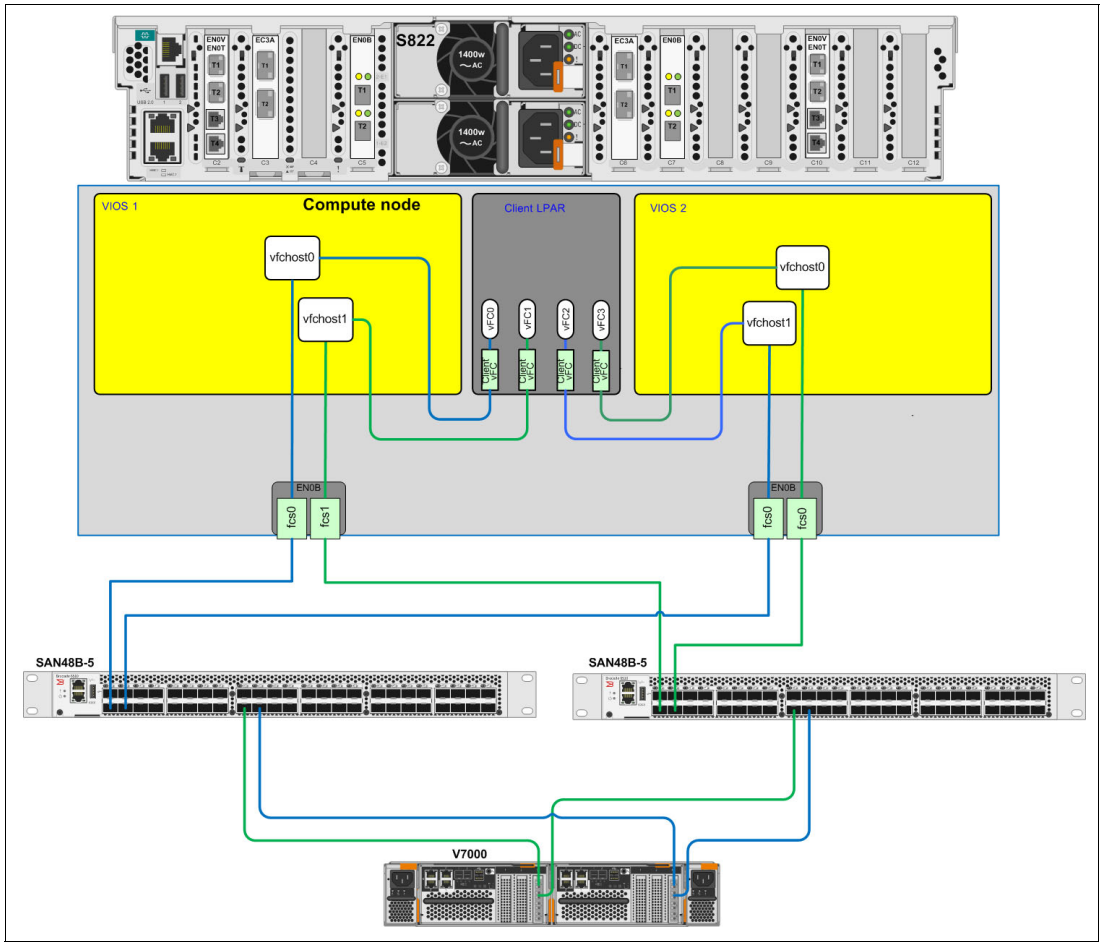

Figure 1-8 The storage network

The colors green and blue in Figure 1-8 highlight the two redundant paths to the V7000 controller.

Overall default IP addresses and user credentials

This map should help the user identify the default private IP addresses used in IBM PurePower to manage the hardware components, and the default user credentials of the VMs and LPARs. This information is also provided in the Home section of the PPIM. For more information, check “Home page” on page 68 in this paper.

Note: The average user would need to know only the IP addresses of the PPIM, which provides access the other management tools; LPARs operations performed by PowerVC and PPIM usually take care of the hardware settings themselves, without requiring the user to do anything on the storage or network GUIs.

Management nodes

This section displays the IP addresses to access the operating systems of the VMs located on the management nodes, which are placed per default at the location provided in Table 1-6.

Table 1-6 Rack location of management nodes

Management node	Rack location
Primary	Rack U6
Secondary	Rack U7

Table 1-7 shows the IP addresses and credentials to access the operating systems hosting the management tools via SSH. Their GUIs are reachable by connecting over HTTPS with a web browser, and the exact URL has also been included for reference in Table 1-8.

Table 1-7 Default IP addresses and credentials to access operating system of management nodes VMs

VM	Primary management node IP address	Secondary management node IP address	Default credentials
IMM	192.168.93.43	192.168.93.143	USERID/PASSW0RD
RHEL KVM Host	192.168.93.44	192.168.93.144	root/PASSW0RD
PowerVC VM	192.168.93.45	192.168.93.145	root/PASSW0RD
PPIM / Nagios	192.168.93.46	192.168.93.146	root/PASSW0RD or admin/PASSW0RD
Service VM	192.168.93.47	192.168.93.147	root/PASSW0RD
vHMC	192.168.93.61	192.168.93.161	hscroot/abc1234

Table 1-8 Default URLs to access the management applications from a web browser (primary management node only)

Applications	Function	URL	Default credentials
IMM	Service processor	https://192.168.93.43	USERID/PASSW0RD[a]
PowerVC	Virtualization management	https://192.168.93.45/powervc/	root/PASSW0RD[a]
PPIM	Hardware management	https://192.168.93.46/puremgr/	root/PASSW0RDa or admin/PASSW0RD[a]
Nagios	HW monitoring	https://192.168.93.46/nagios/	nagiosadmin/PASSW0RD[a]
vHMC	Hardware management	https://192.168.93.61/hmc	hscroot/abc1234

a. All PASSW0RD credentials use zero in place of the uppercase letter O.

The secondary management node has been omitted for brevity. Use the same IP addresses as in Table 1-7 on page 28.

Compute nodes

The following section outlines the default IP addresses and user credentials used in the shipped pre-configured compute nodes of IBM PurePower. The number of the compute nodes ordered by the customer may vary; in the tables below we have included up to the maximum of twelve nodes.

Every pre-configured LPAR is shipped with an IP address in the range 192.168.93.128-191, starting from the first POWER compute node. The default credentials are always **root/PASSWORD**.

Every node is pre-configured with dual Virtual I/O Servers. Their IP addresses are distributed in the range 192.168.93.89 → 112. The default credentials for all of them is **padmin/padmin**, which has been omitted for brevity from Table 1-9. VIOS can be accessed via SSH. The exact location of the compute nodes is included.

Table 1-9 IP addresses of preinstalled Virtual IO Servers

POWER compute node	Rack location	VIOS 1 IP address	VIOS 2 IP address
1	Rack U2/U3	192.168.93.89	192.168.93.90
2	Rack U4/U5	192.168.93.91	192.168.93.92
3	Rack U8/U9	192.168.93.93	192.168.93.94
4	Rack U10/U11	192.168.93.95	192.168.93.96
5	Rack U14/U15	192.168.93.97	192.168.93.98
6	Rack U23/U24	192.168.93.99	192.168.93.100
7	Rack U25/U26	192.168.93.101	192.168.93.102
8	Rack U27/U28	192.168.93.103	192.168.93.104
9	Rack U29/U30	192.168.93.105	192.168.93.106
10	Rack U31/U32	192.168.93.107	192.168.93.108
11	Rack U33/U34	192.168.93.109	192.168.93.110
12	Rack U35/U36	192.168.93.111	192.168.93.112

Internal network switches

Beginning with GA2, the Mellanox switches each have only a single external 192.168.93 address. The management port 0 on each switch is directly connected to the corresponding port on the other switch. These ports use IP addresses 10.10.20.3 for rack U20 and 10.10.20.1 for rack U21 (see Table 1-10). These ports are used as part of the MLAG configuration.

The credentials are always `admin/admin`. Their GUIs can be accessed via HTTPS from a web browser.

Table 1-10 IP addresses of switches for the internal networks

Internal switch	Rack location	Management port 0	Management port 1
1	Rack U21	10.10.20.1	192.168.93.36
2	Rack U20	10.10.20.3	192.168.93.38

TOR switches

The customer-provided TOR switches are reserved the IP addresses in Table 1-11 (two for redundancy) for the IBM PurePower internal network. For access instructions, see the TOR switches' documentation.

Table 1-11 IP addresses of TOR switches

TOR switch	Rack location	IP address 1	IP address 2
1	Rack U42	192.168.93.17	192.168.93.18
2	Rack U41	192.168.93.19	192.168.93.20

Management network switches

The IP addresses provided in Table 1-12 are assigned to the management switches (two for redundancy); the credentials are per default **`admin/admin`**. Their GUI can be accessed via HTTPS.

Table 1-12 IP addresses of management network switches

Management switch	Rack location	IP address 1	IP address 2
1	Rack U19	192.168.93.81	192.168.93.82
2	Rack U18	192.168.93.83	192.168.93.84

SAN switches

Table 1-13 provides the assignment of the IP Address of the SAN switches (two for redundancy); the default credentials are in this case `admin/password`. Their GUI can be accessed via HTTPS from a Java enabled web browser.

Table 1-13 IP addresses of SAN switches

SAN switch	Rack location	IP address 1
1	Rack U17	192.168.93.9
2	Rack U16	192.168.93.11

Storage controller

The IP addresses provided in Table 1-14 are assigned to the Storwize V7000 controller. The physical location is Rack U12/13, and the default credential is always `superuser/passw0rd`. Where available, the GUI URL has been included.

Table 1-14 IP addresses of Storwize V7000 controller

V7000 port	IP address or URL
Cluster management	`https://192.168.93.8`
Canister service port1	`https://192.168.93.2/service`
iSCSI port 1	`192.168.93.4`
Canister service port 2	`https://192.168.93.5/service`
iSCSI port 2	`192.168.93.7`

Power distribution units (PDUs)

Table 1-15 provides the IP addresses of the smart PDUs (FC 7109) part of IBM PurePower offering. Credentials to access them are always `USERID/passw0rd`. Their GUIs can be accessed via HTTPS from a web browser.

Table 1-15 IP addresses of Smart PDUs

PDU number	IP address
1 (lower left)	192.168.93.25
2 (lower right)	192.168.93.26
3 (middle left)	192.168.93.27
4 (middle right)	192.168.93.28
5 (upper left)	192.168.93.29
6 (upper right)	192.168.93.30

1.2.2 Software stack

IBM PurePower System includes an integrated, preinstalled software stack, orchestrating to provide the user with the most advanced virtualization capabilities through an intuitive interface, in order to keep the promise of quick time to value.

Many of them based on open source tools, the software components take care of the monitoring, provisioning and overall management of the cloud infrastructure; like the hardware stack, they are intelligently preconfigured and preintegrated, so that their arrangement is transparent to the user.

From a high-level capability standpoint, the IBM PurePower Software Stack provides the following functions:

Hardware Resource management Low-level management of the compute, memory, storage and network resources of the physical servers that constitute the physical nodes

Virtualization management Management of the virtual machines and the provisioning of them, which also involves the capability to capture one, store it as a backup, or move it to another physical server

System monitoring The capability to monitor system events, both hardware and software, through a centralized interface (includes alerting specific users and automating corrective measures)

These capabilities are combined across the multiple pieces of software provided with IBM PurePower, as shown in Table 1-16.

Table 1-16 Overview of the software stack on a management node

Capability	Software	Description
Hardware Resource Management	POWER hypervisor - firmware level SV840	IBM POWER long standing and most feature-rich hypervisor software, supports logical partitions (LPARs) running AIX, IBM i or Linux
	Virtual I/O Server (VIOS) 2.2.3.52	This software enables LPARs to share physical I/O resources
	PurePower Integrated Manager (PPIM) 1.1	The single point of access for all the interfaces managing the hardware components of the system; it can be used also to perform firmware updates
	Virtual Hardware Management Console (vHMC) 8.4.0	The appliance enabling administrators for planning, creating and managing resources for the LPARs; available as an LPAR in the management nodes
Virtualization Management	PowerVC 1.2.3 Standard Edition	The software allowing users to capture, store and redeploy LPARs across the physical servers
System Monitoring	Nagios Core 4.0.8	The open source software used for the monitoring of systems, networks and infrastructure; additional capabilities include alerts and automation.

Software installed on the management nodes

All of the software listed on Table 1-16 is installed on local disk drives on the management nodes, with the exception of the IBM POWER Hypervisor™ and VIOS, which reside on the compute nodes.

The primary management node (with MTM 8374-01M) of IBM PurePower provides all the management functionalities for the whole system. It is derived from the 7042-CR8 HMC hardware, with the addition of a second processor and more memory. No customer selected options are available for the management nodes.

The management node software is installed on several Red Hat Enterprise Linux (RHEL) 7.1 virtual machines (VMs) running on a Red Hat Enterprise Linux 7.1 host, using KVM virtualization technology.

A set of maintenance tools is also installed in a separate virtual machine, the Service VM. It contains tools intended to be used solely by IBM maintenance personnel; customers do not need to use this VM. Additional resources on the management nodes are reserved for future use.

Note: Since the management node is virtualized with KVM, the term used to define the virtual hosts is virtual machine, or VM, as opposed to logical partitions, or LPARs, in use on the compute nodes virtualized with PowerVM.

A second management node is also part of the solution: this is a mirror image of the primary management node, to be used in case the primary is not functional.

Table 1-17 lists the specifications of the VMs on the primary management node. For a full list of IP addresses and default credentials, including also those on the secondary management node, see “Management nodes” on page 28.

Table 1-17 Definitions of the VMs for the management tools

Virtual Machine	OS	vCPUs	Memory (GB)	Disk space (GB)
PowerVC	RHEL 7.1	4	16	204
PPIM, Nagios	RHEL 7.1	2	8	58
vHMC	Embedded OS	4	8	160
Service VM	RHEL 7.1	1	12	58
Reserved space for future use				

Figure 1-9 illustrates the overall software architecture of the management node.

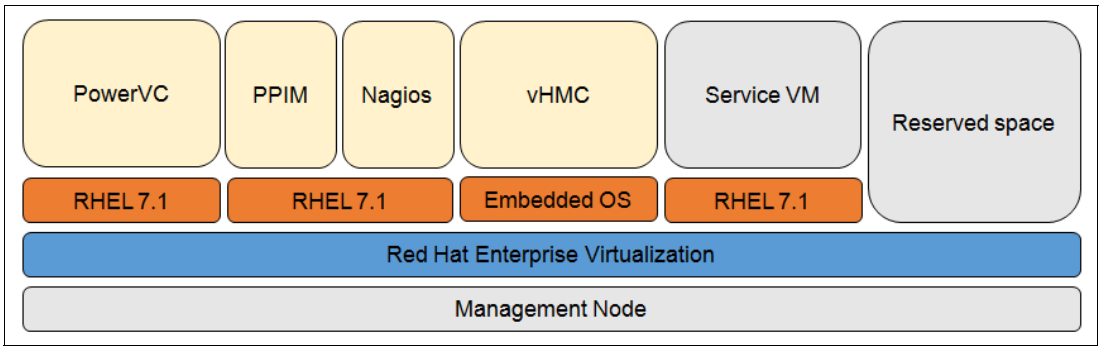

Figure 1-9 Software architecture of the management node

For a more detailed description of the components, see 2.1, “System management overview” on page 66.

You can review more detailed hardware information on the management node on section “Hardware specifications of the 8374-01M PurePower Management Node” on page 19.

Differences between management tools capabilities

This section should help you navigate through the functions offered by the IBM PurePower tools PowerVC and the vHMC (which extend to those of a usual HMC provided with other Power Systems servers), especially in those areas where they overlap.

Table 1-18, Table 1-19, and Table 1-20 address the main management areas of the IBM PurePower software stack: Cloud use, virtualization management, and hardware management.

Table 1-18 Cloud use capabilities of IBM PurePower

Capability	PowerVC	vHMC
Start/Stop VM Instance	x	x
Delete VM Instance	x	x

Table 1-19 Virtualization management capabilities of IBM PurePower

Capability	PowerVC	vHMC
VM image management	x	
VM image deployment	x	
LPAR image capture	x	
Host pools management	x	
Host and LPAR monitoring	x	
LPAR placement	x	
LPAR availability management	x	
LPAR performance management	x	
LPAR isolation management	x	

Table 1-20 Hardware management capabilities of IBM PurePower

Capability	PowerVC	vHMC
LPAR lifecycle management	x	x
Hardware configuration		x
Hardware ops controls		x
Firmware management		x
Service and support		x

Software installed on the compute nodes

VIOS and the IBM POWER Hypervisor, which are part of PowerVM Enterprise Edition, enable virtualization on the compute nodes. The POWER Hypervisor is nested in the firmware, while VIOS instances are installed on the compute nodes in a Dual VIOS configuration. Each VIOS instance has two software-mirrored drives.

The customer can add additional disks on the compute nodes, up to a total of 12, which is the total capacity of the storage backplane of the S822 and S822L in the IBM PurePower configuration. The storage backplane model is the one supporting the 6+6 split backplane option: PCIe3 x8 SAS RAID internal adapter 6Gb (FC EJ0T and FC EJ0V; CCIN 57D7).

The VIOS is the only software installed on local disks on the compute nodes. The OS and the data running on the LPARs are all located on the V7000, on IBM PurePower's SAN.

For the specifications of OSes supported on IBM PurePower, see 1.2.3, “Software support” on page 36.

1.2.3 Software support

IBM PurePower System supports the following software:

- AIX Version 7.1 with the 7100-03 Technology Level and Service Pack 3 and APAR IV56367 or later
- IBM i 7.1 TR11 or later
- IBM i 7.2 TR3 or later
- Red Hat Enterprise Linux 7.1 LE, BE
- PowerVM Enterprise Edition
- PowerVC 1.2.3 Standard Edition or later

IBM i support requires VIOS (no native I/O), along with Power system firmware 8.4.0. With firmware code level 840, the 2U Power S822 supports IBM i 7.2 or IBM i 7.1 with special terms and conditions. Technology Refresh 3 or later for IBM i 7.2 or Technology Refresh 11 or later for IBM i 7.1 is required. Multiple IBM i partitions, each up to a maximum of two cores, are supported.

For more information about the software that is available on IBM PurePower Systems, go to the IBM Power Systems Software™ web page:

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

1.2.4 Capacity (min/max)

This section describes the maximum and minimum values for the options in the system and for the single machine types. The rack content of the system can be flexibly configured when ordering, but there is no MES upgrade available for the rack content.

System configuration

The minimum number of machines that must be configured to the PurePower System is provided in the table that follows. IBM Power System S822L compute nodes and IBM Power System S822 compute nodes can be mixed. Two compute nodes are the required minimum.

Table 1-21 shows the machines types that are required for the IBM PurePower System solution.

Table 1-21 Machine types required for PurePower System

M/T-Model	Description	Default	Min	Max
8247-22L	IBM Power System S822L	2	0	12
8284-22A	IBM Power System S822	0	0	12
2498-F48	SAN48B-5 switch	2	2	2
7120-48E	Management switch	2	2	2
8831-NF2	Mellanox switch	2	2	2
8374-01M	Management nodes (primary and secondary)	2	2	2
7014-B42	Enterprise rack	1	1	1
7316-TF4	Display	1	1	1
2076-524	V7000 controller	1	1	1
2076-24F	V7000 expansion drawer	0	0	10

Compute note configuration

This tables that follow provide the default minimum configuration of the two different types of compute nodes. Table 1-22 provides the 8247-22L configuration, and Table 1-23 on page 38 provides the 8284-22A configuration. Two pairs of internal disk drives are required and are used as VIOS boot devices. Three I/O slots are available in each node for additional adapters. Eight memory DIMMs are the minimum for best performance.

Table 1-22 POWER8 S822L (8247-22L) minimum configuration

FC	Description	Default	Min	Max
4651	Rack Indicator Feature	1	1	1
EHKX	CI Base Indicator feature	1	1	1
EHKY	CI Add-On Indicator feature	0	0	12
EJTP	Front Bezel for 12-Bay BackPlane	1	1	1
EL1B	System AC Power Supply, 1400 W	2	2	2
5771	SATA Slimline DVD-RAM Drive	1	1	1
6577	Power Cord to IBM PDU	2	2	2
2147	Primary Operating System Indicator - Linux	1	1	1
5000	SW Preload Indicator	1	1	1
0266	Linux Partition Specify	2	1	250
ELPD	10 C Processor 3.425 GHz POWER8 processor	2	2	2
ELAD	One Processor Activation for Processor Feature #ELPD	20	20	20
EL3P	16 GB CDIMM (1.35 V), 1600 MHz	16	0	16
EL3Q	32 GB CDIMM (1.35 V), 1600 MHz	0	0	16

FC	Description	Default	Min	Max
EL3R	64 GB CDIMM (1.35 V), 1600 MHz	0	0	16
EL43	PCIe3 LP 16Gb 2-port Fibre Channel Adapter	2	2	5
EC3A	PCIe3 LP 2-Port 40GbE NIC RoCE QSFP+ Adapter	2	2	5
EN0V	PCIe2 LP 4-port (10Gb+1GbE) Copper SFP+RJ45 Adapter	2	2	5
EL3V	Split 6+6 storage backplane + Second SAS controller	1	1	1
ELDF	600 GB 15K HDD in Gen3 carrier	4	4	12
EC22	PowerVM Linux Edition	20	20	20
ECCJ	Yellow Ethernet Cat 5E cable (MFG Sel length)	2	2	2
ECCG	Blue Ethernet Cat 5E cable (MFG Sel length)	1	1	2
ECCH	Green Ethernet Cat 5E cable (MFG Sel length)	1	1	1
ECCK	Fiber SAN cable (MFG Sel Length)	4	4	4
ECCN	DAC QSFP+ TO QSFP+ Cable (MFG Sel Length)	4	4	4
9447	New VIOS Core counter	2	0	20
93XX 97XX	Language Specify code	1	1	1
ESC0	Free shipping	1	1	1

Table 1-23 POWER8 S822 (8284-22A) minimum configuration

FC	Description	Default	Min	Max
4651	Rack Indicator Feature	1	1	1
EHKX	CI Base Indicator feature	1	1	1
EHKY	CI Add-On Indicator feature	0	0	12
EJT4	Front Bezel for 12-Bay BackPlane	1	1	1
EB2M	System AC Power Supply, 1400 W	2	2	2
5771	SATA Slimline DVD-RAM Drive	1	1	1
6577	Power Cord to IBM PDU	2	2	2
5000	SW Preload Indicator	1	1	1
0265	AIX Partition Specify	2	0	250
0266	Linux Partition Specify	0	0	250
0267	IBM i Partition Specify	0	0	250
EPXD	10 C Processor 3.425 GHz POWER8 processor	2	2	2
EPYD	One Processor Activation for Processor Feature EPXD	20	20	20
EM83	16 GB CDIMM (1.35 V), 1600 MHz	16	0	16
EM84	32 GB CDIMM (1.35 V), 1600 MHz	0	0	16

FC	Description	Default	Min	Max
EM85	64 GB CDIMM (1.35 V), 1600 MHz	0	0	16
EN0B	PCIe2 LP 16Gb 2-port Fibre Channel Adapter	2	2	5
EC3A	PCIe3 LP 2-Port 40GbE NIC RoCE QSFP+ Adapter	2	2	5
EN0V	PCIe2 LP 4-port (10Gb+1GbE) Copper SFP+RJ45 Adapter	2	2	5
EJ0V	Split #EJ0T to 6+6 SFF-3 Bays: Add 2nd SAS controller	1	1	1
ESDF	600 GB 15K HDD in Gen3 carrier	4	4	12
5228	PowerVM Linux Edition	20	20	20
ECCJ	Yellow Ethernet Cat 5E cable (MFG Sel length)	2	2	2
ECCG	Blue Ethernet Cat 5E cable (MFG Sel length)	1	1	1
ECCH	Green Ethernet Cat 5E cable (MFG Sel length)	1	1	1
ECCK	Fiber SAN cable (MFG Sel Length)	4	4	4
ECCN	DAC QSFP+ TO QSFP+ Cable (MFG Sel Length)	4	4	4
9447	New VIOS Core counter	2	0	20
93XX 97XX	Language Specify code	1	1	1
ESC0	Free shipping	1	1	1

Storage configuration

The base storage configuration is one V7000 Controller enclosure (2076-524) with eight 1.2 TB drives and can be expanded with up to 10 additional V7000 expansion enclosures. The V7000 disks is configured with RAID5 and one hot spare drive.

Table 1-24 provides the maximum disk capacity with 1.2 TB disk drives, one V7000 Controller enclosure (2076-524) and up to 10 V7000 expansion enclosures (2076-24F) with a maximum storage capacity of up to 316.80 TB.

Table 1-24 Storage disk capacities

V7000 controller	V7000 expansion	Number of drives	Disk capacity
1	0	8	9.60 TB
1	0	24	28.80 TB
1	1	48	57.60 TB
1	2	72	86.40 TB
1	3	96	115.20 TB
1	4	120	144.00 TB
1	5	144	172.80 TB
1	6	168	201.60 TB
1	7	192	230.40 TB
1	8	216	259.20 TB

V7000 controller	V7000 expansion	Number of drives	Disk capacity
1	9	240	288.00 TB
1	10	264	316.80 TB

Virtual HMC limitations

The maximum number of supported virtual machines (VMs) is 500.

1.2.5 Security

This section describes the security features of the IBM PurePower System.

PowerVC

PowerVC is the advanced virtualization management offering, built on OpenStack, that delivers advanced virtualization management for IBM AIX, IBM i, and Linux environments on IBM PurePower Systems.

Management of PowerVC is through its default management network.

PowerVC provides security services that support a secure environment and, in particular, the following security features:

- LDAP support for authentication and authorization information (users and groups)
- The PowerVC Apache web server is configured to use secured https protocol
- Host key and certificate verification of hosts, storage, and switches
- Audit logs, recorded and available

If SELinux enforcing mode is detected while installing PowerVC, the PowerVC SELinux policy module is installed automatically. If enforcing mode was not detected while installing PowerVC and you later choose to enable SELinux enforcing mode, you must manually install the PowerVC SELinux policy module. PowerVC will not work properly in SELinux enforcing mode without this module.

See *Cloud Security Guidelines for IBM Power Systems*, SG24-8242 for a detailed description of the PowerVC security features:

http://www.redbooks.ibm.com/abstracts/sg248242.htm

PowerVM

The virtualization software suite used on IBM Power Systems is PowerVM, which includes the POWER Hypervisor and the Virtual I/O Server (VIOS) for I/O sharing between virtual machines (VMs). Security at the hypervisor level is very important, because every cloud workload relies on it. PowerVM has a proven track record in security, enforcing separation between LPARs and minimizing the possibility for malicious users to find vulnerabilities to use.

The Common Vulnerability and Exposure (CVE) database is an openly accessible dictionary of common names for publicly known information security vulnerabilities across any piece of software. PowerVM so far has zero vulnerabilities found in this database. To verify the amount of vulnerabilities discovered in PowerVM, use “PowerVM” as the keyword to search on the CVE website:

http://cve.mitre.org/find/index.html

PowerVM supports industry-standard methods for strong isolation, such as virtual local area network (VLAN) technology. Communication between partitions on the *same* physical host is highly secure. It is impossible to snoop on the type of traffic being sent, because it moves directly between partitions, as opposed to traffic between partitions located on *different* physical hosts, because that traffic is sent via a physical Ethernet connection. Given the virtualization capabilities of POWER, such as the high number of virtual machines or logical partitions running on the same physical host, this feature is a compelling advantage.

See Chapter 4 in *Cloud Security Guidelines for IBM Power Systems*, SG24-8242, cited previously, for a detailed description of the PowerVM security features.

PurePower Integrated Manager

PurePower Integrated Manager will create a *puremgradmin* group and a *puremgrviewers* group. An `admin` user will be created in the puremgradmin group with root level authority. The file and directory ownership of IBM included code will be placed under puremgradmin and user scope to admin.

The PowerVC Guest OS defines the following for backing up files:

- When a new update of the PowerVC Guest OS is installed (via download from Fix Central), then the **`setupguests.sh -u`** option will be used, which will take the archived backed up files for the PowerVC Guest OS and load the files on the newly installed PowerVC Guest OS achieving an upgrade of the PowerVC Guest OS.
- Manual steps to restore the backed up PowerVC data files and databases using the PowerVC Restore command are required.

PurePower Integrated Manager web server

The PurePower Integrated Manager web server configures HTTPS using self-signed X.509 certificates. If you prefer, the self-signed certificates can be replaced with custom certificate authority (CA) certificates.

PureKVM Host OS and Guest OS SSH Key Exchanges

During the bare metal install the RHEL 7.1 KVM Host OS will exchange SSH keys with the four guest OSes (PurePower Integrated Manager, PowerVC, Service, and the vHMC guest OSes).

After each Fix Central Guest OS update the `setupguests.sh` will also exchange SSH keys with each guest OS. If you prefer, the SSH `/root/.ssh/authorized_keys` can be removed from each of the guest OSes to remove the password-less SSH capabilities.

HMC and vHMC

The Hardware Management Console (HMC) is a Linux-based appliance used for configuration, management, and maintenance of IBM Power Systems servers.

Security certificates ensure that the HMC can operate securely in the client/server mode. The managed machines are servers and the managed users are clients. Servers and clients communicate over the Secure Sockets Layer (SSL) protocol, which provides server authentication, data encryption, and data integrity.

See Chapter 3 in *Cloud Security Guidelines for IBM Power Systems*, SG24-8242 for a detailed description of the HMC security features.

AIX Version 7.1

This section describes the security features of the AIX Version 7.1 operation system:

- Hardware acceleration support for Encrypted File Systems, IPSec and Trusted Execution
- Domain support in role-based access control
- Complaince with Common Criteria CAPP/EAL4+ security certification
- IBM Compliance Expert Express Edition (5765-G82) updated to include a new security profile designed to aid compliance with the Sarbanes Oxley (SOX) and COBIT standards
- Audit subsystem

POWER8 systems contain a cryptographic function that can be used with the Open Secure Sockets Layer (OpenSSL) version 1.0.1.510 fileset and with AIX 7 with 7100-03 with Service Pack 3.

For more information, see the Security in AIX 7.1 section of the IBM Knowledge Center:

http://ibm.co/1RJcV4T

IBM i 7.1 TR11 and IBM i 7.2 TR3

This section describes the security features of the IBM i 7.1 and IBM i 7.2 operating system:

- IBM i disk data encryption and IBM i provided tape data encryption.
- All information on the system is encapsulated within Objects, which have defined security attributes.
- Object accessibility control is implicit.
- OpenSSL provides cryptographic protocols to applications.
- IBM DB2® data encryption operates on the system layer.

See *Security Guide for IBM i V6.1*, SG24-7680 for a detailed description of the top security management practices for an IBM i system:

https://www.redbooks.ibm.com/abstracts/sg247680.html

Red Hat Enterprise Linux 7.1

Red Hat Enterprise Linux allows a secure implementation of a partition by providing the following security features:

- Vulnerability assessment tools
- Support for DNSSEC
- Virtual Private Network (VPN) by using IPSec tunneling protocol
- OpenSSL provides cryptographic protocols to applications
- Disk Encryption with Linux Unified Key Setup-on-disk-format (LUKS)
- Public-Key Cryptography by using openCryptoki, a Linux implementation of PKCS#11

Nagios

Nagios can be set up to monitor security logs and security data and can alert the administrators for certain log pattern that may occur.

1.2.6 Use cases for data center integration

IBM PurePower can be integrated into the customer's data center environment in several ways. This section addresses some of those.

Networking integration

This section describes the supported configurations by which IBM PurePower can be connected to the customer's data center network.

Use case 1. BM PurePower switch for both internal and external data access

In this case the customer makes use of the Mellanox MSX 1710 40/56 Gb RoCE capable switch provided with the IBM PurePower offering not only for the internal data access, as it is primarily designed for, but also for the external traffic, such as the connection to the customer's data center network (uplink).

This configuration is typical for users who want to benefit from the high performance of the networking device, and are not averse to integrating a new switch into their existing networking environment.

Figure 1-10 shows that the adapters of the compute nodes and the switches used for the uplink are the same as those for the internal network, such as the 40 GbE and RoCE adapters (EC3A) and the Mellanox MSX 1710 40/56 Gb RoCE capable switch, respectively.

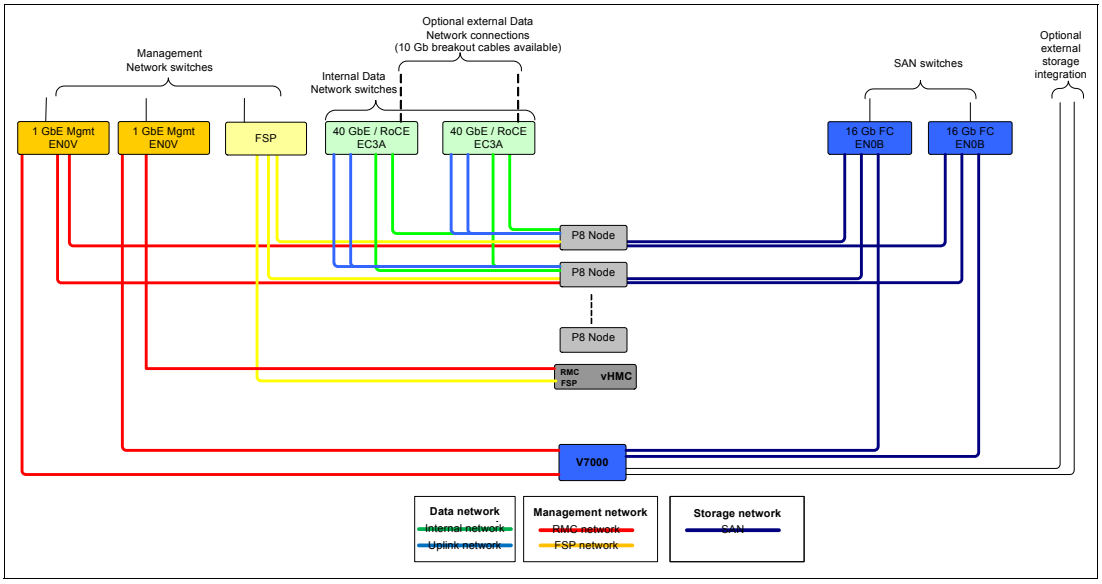

Figure 1-10 Use case 1 for network integration: Same switch for internal and uplink network

Use case 2. Customer-provided TOR switch for external data access

If customers prefer to use their own switches for the uplink, the space at the top of the rack (two rack units, 2U) can be used for one or two redundant customer provided TOR switches.

This configuration adapts to those environments where the requirements of the data center network cannot include the Mellanox 40 GbE switch in the IBM PurePower offering for technical or strategic reasons. The switch is then used only for its primary function, the servicing of the internal data network of IBM PurePower.

Figure 1-11 illustrates how, in this case, the compute nodes use the two 10 GbE ports on the 1/10GbE adapter (EN0V) to connect to the customer-provided TOR switches.

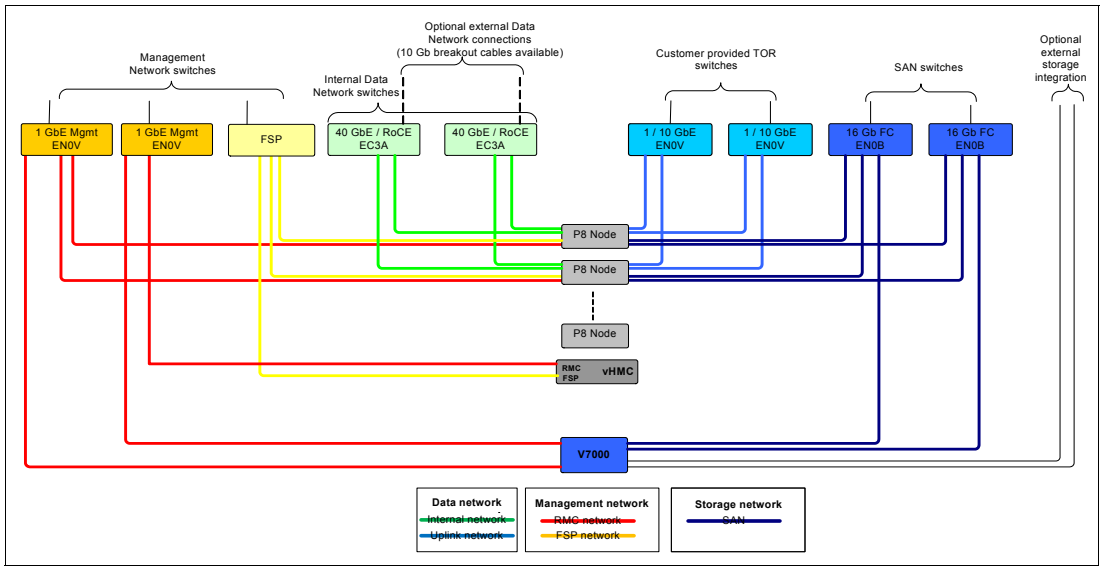

Figure 1-11 Use case 2 for network integration: Customer-provided TOR switches for the uplink

The TOR switches are not part of the IBM PurePower offering, but compatibility has been tested with three models. See “Uplink network” on page 23 and “Top of Rack switch” on page 22 for more information about the supported TOR models.

Use case 3. Customer-provided remote switch for the external data access

IBM PurePower supports connections not only to TOR switches, but also to customer-provided remote switches located on other racks. For these connections, the compute nodes’ 10 GbE ports of the 1/10GbE adapters (EN0V) are used.

Figure 1-12 illustrates this case.

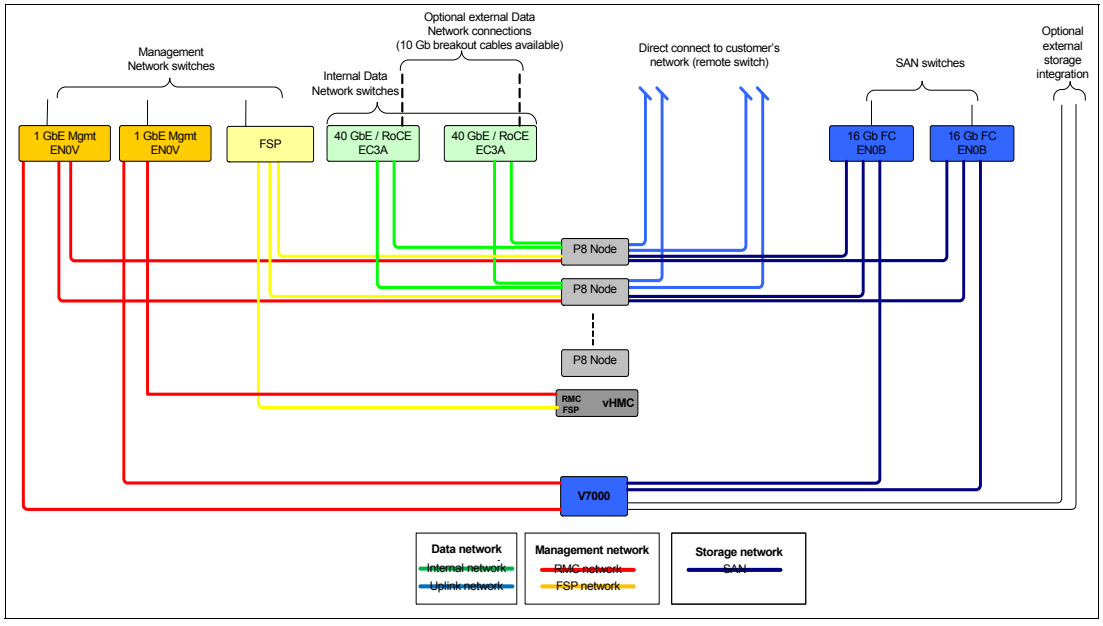

Figure 1-12 Use case 3 for network integration: Customer-provided remote switches for the uplink

Other forms of data center integration

There are several other ways to integrate IBM PurePower in the customer's data center; storage and compute devices can be added to the solution in order to increase its capacity, and thanks to OpenStack drivers in PowerVC, IBM PurePower can also be managed by other OpenStack compatible software.

Note: Adding other components to the IBM PurePower or letting the solution be managed by other software does not terminate support. However, only the original components will be supporte. Any newly added storage or compute devices, as well as the new management software, will not be supported under IBM PurePower license, but they under their own support contracts, instead.

Storage

Extra ports are reserved on the IBM System Storage SAN48B-5 switch for integration into the customer's existing SAN, while OpenStack drivers are available to facilitate direct access of various third-party storage solutions.

The V7000 is also equipped with the SAN Volume Controller functionality, which enables also other storage to be virtualized and used: extra ports on the V7000 are available to allow this without having to integrate the internal SAN with the customer's SAN.

For more information about the V7000, the SAN Volume Controller, and the SAN Switch, see the following Redbooks publications:

- *Implementing the IBM System Storage SAN Volume Controller V7.4*, SG24-7933

  http://www.redbooks.ibm.com/abstracts/sg247933.html
- *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521

  https://www.redbooks.ibm.com/redbooks.nsf/RedbookAbstracts/sg247521.html
- *IBM System Storage SAN48B-5*, TIPS1125

  http://www.redbooks.ibm.com/abstracts/tips1125.html

Compute

No special codes are included in the IBM Power Systems S822 and S822L compute nodes part of the offering. This means that other IBM Power System servers can also be added to IBM PurePower, provided that they are equipped with the necessary adapters (so they also have FC and Ethernet adapters and PowerVM).

Management software

IBM PurePower virtualization management software PowerVC is based on OpenStack, which are a set of software tools for building and managing cloud computing platforms for public and private clouds. OpenStack is managed by the OpenStack Foundation, a nonprofit organization that oversees both development and community-building related to the project.

Because it is using OpenStack drivers, which are industry standard API, PowerVC can also be integrated into a customer's management solution also using those drivers. At the time of writing, VMware vRealize Automation is one of the cloud management tools able to provision and manage LPARs on IBM Power Systems. This means that a customer using this management tool can perform cloud operations on IBM POWER, Linux on Z and x86 virtualized environments from the same pane of glass.

1.3 PurePower reliability, availability, and serviceability

This section provides information about IBM PurePower system reliability, availability, and serviceability (RAS) design and features.

There are three basic elements of RAS:

Reliability Indicates how infrequently a defect or fault in a server occurs

Availability Indicates how infrequently the functioning of a system or application is impacted by a fault or defect

Serviceability Indicates how well faults and their effects are communicated to system managers and how efficiently and non-disruptively the faults are repaired

RAS designs and features for the IBM PurePower components:

- IBM PurePower
- IIBM Compute Nodes 8284-22A or 8247-22L
- IIBM Storwize V7000 Gen2 storage enclosure 2076-524
- IBM System Storage SAN48B-5 switch 2498-F48
- Mellanox SX1710 Data Network Switch 8831-NF2
- Systems Management IBM 7120 Model G8052 RackSwitch 7120-48E
- Pure Power Management Node designed to provide the Virtualization capabilities for the IBM PurePower Rack 8374-01M
- IBM Service console 7316-TF4
- IBM Enterprise Rack 7014-B42

IBM PurePower

“S822 or S822L compute nodes” on page 46 lists RAS designs and features for IBM PurePower

Table 1-25 RAS designs and features for IBM PurePower

Feature	Implement through	Notes
Redundant hardware	Full hardware redundancy	Supported with PowerVC v1.2.3
VIOS	Redundant VIOS	Supported with PowerVC
Workload mobility	Live Partition Mobility, LPM	Supported
Dynamic re-size of workload	DLPAR	Supported

S822 or S822L compute nodes

The POWER8 technology-based servers running VIOS version 2.2.3.52 in a dual VIOS configuration (VIOS OS rootvg on local disk, each server has four HDDs installed in a split storage backplane (6+6), two pairs of two same-size HDDs, software mirrored by the OS with one hdisk for rootvg.

RAS introduction of the POWER8 processor

The POWER8 processor modules support an enterprise level of reliability and availability. The processor design has extensive error detection and fault isolation (ED/FI) capabilities to allow for a precise analysis of faults, whether they are hard or soft. They use advanced technology, including stacked latches and Silicon-on-Insulator (SOI) technology, to reduce susceptibility to soft errors, and advanced design features within the processor for correction or try again after soft error events.

The design also incorporates spare capacity that is integrated into many elements to tolerate certain faults without requiring an outage or parts replacement. Advanced availability techniques are used to mitigate the impact of other faults that are not directly correctable in the hardware.

Features within the processor and throughout the system are incorporated to support design verification. During the design and development process, subsystems go through rigorous verification and integration testing processes by using these features. During system manufacturing, systems go through a thorough testing process to help ensure high product quality levels, again taking advantage of the designed ED/FI capabilities.

Fault isolation and recovery of the POWER8 processor and memory subsystems are designed to use a dedicated service processor and are meant to be largely independent of any operating system or application deployed.

RAS enhancements of POWER8 processor-based servers

In addition to being built on advanced RAS characteristics of the POWER8 processor, the Power S822 and S822L server offers reliability and availability features. Some of these features are improvements for POWER8 or features that were found previously only in higher-end Power Systems.

The following list is a brief summary of these features:

- Processor Enhancements Integration

  POWER8 processor chips are implemented using 22 nm technology and integrated on to SOI modules.

  The processor design now supports a spare data lane on each fabric bus, which is used to communicate between processor modules. A spare data lane can be substituted for a failing one dynamically during system operation.

  A POWER8 processor module has improved performance compared to POWER7+™, including support of a maximum of 12 cores compared to a maximum of eight cores in IBM POWER7+. This is because doing more work with less hardware in a system supports greater reliability.

  The processor module integrates a new On Chip Controller (OCC). This OCC is used to handle Power Management and Thermal Monitoring without the need for a separate controller, which was required in POWER7+. In addition, the OCC can also be programmed to run other RAS-related functions independent of any host processor.

  The memory controller within the processor is redesigned. From a RAS standpoint, the ability to use a replay buffer to recover from soft errors is added.

- I/O subsystem

  The POWER8 processor now integrates PCIe controllers. PCIe slots that are directly driven by PCIe controllers can be used to support I/O adapters directly in the systems or, as a Statement of Direction, be used to attach external I/O drawers. For greater I/O capacity, the POWER8 processor-based servers also support a PCIe switch to provide additional integrated I/O capacity.

These integrated I/O adapters can be repaired in these servers concurrently, which is an improvement over comparable POWER7/7+ systems that did not allow for adapter “hot-plug.”

- Memory subsystem

  Custom DIMMs (CDIMMS) are used, which, in addition to the ability to correct a single DRAM fault within an error-correcting code (ECC) word (and then an additional bit fault) to avoid unplanned outages, also contain a spare DRAM module per port (per nine DRAMs for x8 DIMMs), which can be used to avoid replacing memory.

  You can find detailed RAS information about the POWER8 compute in PurePower in *IBM Power System S822 Technical Overview and Introduction*, REDP-5102:

  http://www.redbooks.ibm.com/abstracts/redp5102.html

IBM Storwize V7000 Gen2

Reliability, availability, and serviceability (RAS) are important concepts in the design of IBM Storwize V7000 Gen2 system. Hardware features, software features, design considerations, and operational guidelines all contribute to make the IBM Storwize V7000 Gen2 system reliable.

Reliability and availability

Fault tolerance and high levels of availability are achieved by these methods:

- The Redundant Array of Independent Disks (RAID) capabilities of the underlying disks
- IBM Storwize V7000 Gen2 nodes clustering using a Compass architecture
- Auto-restart of hung nodes
- Integrated UPS units to provide memory protection if there is a site power failure

The heart of IBM Storwize V7000 Gen2 system is a pair of node canisters. These two canisters share the data transmitting and receiving load between the attached hosts and the disk arrays.

IBM System Storage SAN48B-5 switch 2498-F48

The IBM PurePower system will provide redundant SAN networks with two redundant SAN48B-5 switch for storage network 2498-F48.

Nondisruptive software upgrades and redundant, hot-pluggable components help maximize availability.

SAN48B-5 Switch Hot-swap components:

These are the SAN48B-5 switch for storage network 2498-F48 hot-swap components:

- Power supplies
- Fan modules
- Small form-factor pluggable (SFP)

Mellanox SX1710 40 GbE Data Network Switch 8831-NF2

The IBM PurePower system will provide redundant data networks with two redundant Mellanox 40 GbE switches.

The Mellanox 40 GbE switch supports system reliability and availability through the use of redundant components, which is also referred to as *N+1 redundancy*. Redundant components include those in the following list:

- AC power supply, which has these characteristics:
  - If a power supply fails, the other power supply assumes the full load.
  - If a fan fails, it might cause the power supply to fail and turn off, but the other power supply will assume the load.
  - Any power supply issue, including a fan rotor failure, results in an error log entry that flags the power supply for replacement.
- Fan units, which have these characteristics:
  - If one or both of the rotors on the same fan fail, there will be adequate cooling from the remaining fan to continue to run in the worst case environment.
  - A loss of a fan (either one or two rotors) results in an error logged and the fan is flagged for replacement as a repair action.

Systems management IBM 7120 Model G8052 RackSwitch 7120-48E

Redundant management network with two redundant Systems Management G8052 Switches with access to management ports on all components.

- AC power supply, which has these characteristics:
  - The G8052 switch has two redundant 150W AC power supplies.
  - Each internal power supply has an individual IEC 320 power connector on the rear panel.
  - The power cord attaches to a universal, grounded AC power source.
  - Power supplies are hots-wappable and are replaced by the customer.
- Fan units, which has these characteristics:
  - Up to four internal fans cool the switch unit.
  - Three fans are required for normal operation.
  - If an individual fan fails, the other fans continue to run, and the switch unit continues to operate normally.
  - These hot-swap fan modules are replaced by the customer.

PurePower System Management Nodes

The two redundant PurePower System Management Nodes are equipped with two redundant 1TB disk drives (RAID1) and redundant hot-swap power supplies and fans on each node.

RAID level-1 uses data mirroring. Two physical drives are combined into an array, and data is striped across the array. The first half of a stripe is the original data; the second half of a stripe is a mirror (or a copy) of the data, but it is written to the other drive in the RAID level-1 array. RAID level-1 provides data redundancy and high levels of performance.

Primary and Secondary Management Node Data Considerations

Leveraging the management node recovery volume manual synchronization of the active databases, datastores and other relevant configuration data on each VM will be able to be replicated to the preferred management node to allow up-to-date configurations to provide continuous high availability of the management capabilities of the PurePower rack.

The `/data` partition on the primary and secondary management nodes hypervisor is used to store data that the customer can expect to have mirrored to an external storage volume named `recovery` (200G). It is used to perform a bare metal restoration of the PurePower rack in case of disaster recovery remediation. The configuration and use of an external storage device is optional and is configured in a properties file. There is an automatic cron job run nightly to copy the recovery volume backup data to the specified backup location.

IBM Service console 7316-TF4

The IBM Flat Panel Service Console Kit is not redundant, but is usually used only for the initial system installation. Later, the system management is done remote over the redundant PurePower System Management Nodes.

IBM Rack 7014-B42

Power for PurePower is fully redundant. The system includes four monitored Power Distribution Units (PDUs), FC 7109. The PDUs are mounted in four side pockets of the rack. System monitoring software should be capable of displaying the power supplied by each PDU and the total power for the system. The power cables are plugged so that the pair of PDUs on the left side of the rack are sufficient to power the entire system.

Similarly, the pair of PDUs on the right side of the rack are sufficient to power the entire rack. The rack can be powered from just one side of PDU during power outages or power service events. The system is not intended to run for extended periods of time with just one pair of PDUs powered.

Serviceability

The purpose of serviceability is to repair the system while attempting to minimize or eliminate service cost (within budget objectives) and maintain application availability and high customer satisfaction.

Serviceability includes considering the following tasks:

- System installation
- Miscellaneous equipment specification (MES) (system upgrades/downgrades)
- System maintenance and repair

Depending on the system and warranty contract, service might be performed by the customer, an IBM System Services Representative (SSR), or an authorized warranty service provider.

Serviceability of the S822 or S822L IBM POWER8 compute nodes

The serviceability features that are delivered in this system provide a highly efficient service environment by incorporating the following attributes:

- Design for customer setup (CSU), customer installed features (CIF), and customer-replaceable units (CRU)
- Error Detection and Fault Isolation (ED/FI)
- First Failure Data Capture (FFDC)
- Lightpath service indicators

- Service labels and service diagrams available on the system and in the IBM Knowledge Center
- Step-by-step service procedures documented in the IBM Knowledge Center or available through the Hardware Management Console CRU videos planned to be available on the web at general availability
- Mobile access to important customer service functions available by scanning a QR label

This section provides an overview of how these attributes contribute to efficient service in the progressive steps of error detection, analysis, reporting, notification, and repair found in all POWER processor-based systems.

Error detecting introduction

The first and most crucial component of a solid serviceability strategy is the ability to detect accurately and effectively errors when they occur.

Although not all errors are a guaranteed threat to system availability, those that go undetected can cause problems because the system has no opportunity to evaluate and act if necessary. POWER processor-based systems employ IBM System z server-inspired error detection mechanisms, extending from processor cores and memory to power supplies and hard disk drives (HDDs).

Error checkers, fault isolation registers, and First-Failure Data Capture

IBM POWER processor-based systems contain specialized hardware detection circuitry that is used to detect erroneous hardware operations. Error-checking hardware ranges from parity error detection that is coupled with Processor Instruction Retry and bus try-again to ECC correction on caches and system buses.

Within the processor/memory subsystem error-checker, error-checker signals are captured and stored in hardware FIRs. The associated logic circuitry is used to limit the domain of an error to the first checker that encounters the error. In this way, runtime error diagnostic tests can be deterministic so that for every check station, the unique error domain for that checker is defined and mapped to field-replaceable units (FRUs) that can be repaired when necessary.

Integral to the Power Systems design is the concept of FFDC. First Failure Data Capture (FFDC) is a technique that helps ensure that when a fault is detected in a system, the root cause of the fault will be captured without the need to re-create the problem or run any sort of extending tracing or diagnostics program. For the majority of faults, a good FFDC design means that the root cause can also be detected automatically without servicer intervention.

FFDC information, error data analysis, and fault isolation are necessary to implement the advanced serviceability techniques that enable efficient service of the systems and to help determine the failing items.

In the rare absence of FFDC and Error Data Analysis, diagnostics are required to re-create the failure and determine the failing items.

Service processor

The service processor provides the capability to diagnose, check the status of, and sense the operational conditions of a system. It runs on its own power boundary and does not require resources from a system processor to be operational to perform its tasks.

Under PowerVM the service processor supports surveillance of the connection to the HMC and to the system firmware (hypervisor). It also provides several remote power control options, environmental monitoring, reset, restart, remote maintenance, and diagnostic functions, including console mirroring. The service processors menus (ASMI) can be accessed concurrently with system operation, allowing nondisruptive abilities to change system default parameters.

Diagnostics

General diagnostic objectives are to detect and identify problems so that they can be resolved quickly. The IBM diagnostic strategy includes the following elements:

- Provide a common error code format that is equivalent to a system reference code, system reference number, checkpoint, or firmware error code
- Provide fault detection and problem isolation procedures. Support a remote connection ability that is used by the IBM Remote Support Center or IBM Designated Service
- Provide interactive intelligence within the diagnostic tests with detailed online failure information while connected to IBM back-end system

Using the extensive network of advanced and complementary error detection logic that is built directly into hardware, firmware, and operating systems, the IBM Power Systems servers can perform considerable self-diagnosis.

Because of the FFDC technology that is designed in to IBM servers, re-creating diagnostic tests for failures or requiring user intervention is not necessary. Solid and intermittent errors are designed to be correctly detected and isolated at the time that the failure occurs. Runtime and boot time diagnostic tests fall into this category.

Boot time

When an IBM Power Systems server powers up, the service processor initializes the system hardware. Boot-time diagnostic testing uses a multitier approach for system validation, starting with managed low-level diagnostic tests that are supplemented with system firmware initialization and configuration of I/O hardware, followed by OS-initiated software test routines.

To minimize boot time, the system determines which of the diagnostic tests are required to be started to ensure correct operation, which is based on the way that the system was powered off, or on the boot-time selection menu.

Run time

All Power Systems servers can monitor critical system components during run time, and they can take corrective actions when recoverable faults occur. The IBM hardware error-check architecture can report non-critical errors in the Central Electronics Complex in an *out-of-band* communications path to the service processor without affecting system performance.

A significant part of IBM Runtime Diagnostic capabilities originate with the service processor. Extensive diagnostic and fault analysis routines were developed and improved over many generations of POWER processor-based servers, and enable quick and accurate predefined responses to both actual and potential system problems.

The service processor correlates and processes runtime error information by using logic that is derived from IBM engineering expertise to count recoverable errors (called *thresholding*) and predict when corrective actions must be automatically initiated by the system. These actions can include the following items:

- Requests for a part to be replaced
- Dynamic invocation of built-in redundancy for automatic replacement of a failing part
- Dynamic deallocation of failing components so that system availability is maintained

Device drivers

In certain cases, diagnostic tests are best performed by operating system-specific drivers, most notably adapters or I/O devices that are owned directly by a logical partition. In these cases, the operating system device driver often works with I/O device microcode to isolate and recover from problems.

Potential problems are reported to an operating system device driver, which logs the error. In non-HMC managed servers, the OS can start the Call Home application to report the service event to IBM. The event is reported to the HMC, which can initiate the Call Home request to IBM. I/O devices can also include specific exercisers that can be started by the diagnostic facilities for problem recreation (if required by service procedures).

Reporting

In the unlikely event that a system hardware or environmentally induced failure is diagnosed, IBM Power Systems servers report the error through various mechanisms. The analysis result is stored in system NVRAM. Error log analysis (ELA) can be used to display the failure cause and the physical location of the failing hardware.

Using the Call Home infrastructure, the system automatically can send an alert through a phone line to a pager, or call for service if there is a critical system failure. A hardware fault also illuminates the amber system fault LED, which is on the system unit, to alert the user of an internal hardware problem.

On POWER8 processor-based servers, hardware and software failures are recorded in the system log of the virtual machine. An ELA routine analyzes the error, forwards the event to the Service Focal Point (SFP) application running on the management console, and can notify the system administrator that it isolated a likely cause of the system problem. The service processor event log also records unrecoverable checkstop conditions, forwards them to the SFP application, and notifies the system administrator.

After the information is logged in the SFP application, if the system is correctly configured, a Call Home service request is initiated and the pertinent failure data with service parts information and part locations is sent to the IBM service organization. This information also contains the client contact information as defined in the IBM Electronic Service Agent™ (ESA) guided setup wizard.

With the introduction of HMC V8R8.1.0 a Serviceable Event Manager is available to block problems from being automatically transferred to IBM.

Figure 1-13 shows a graphical overview of the service event flow.

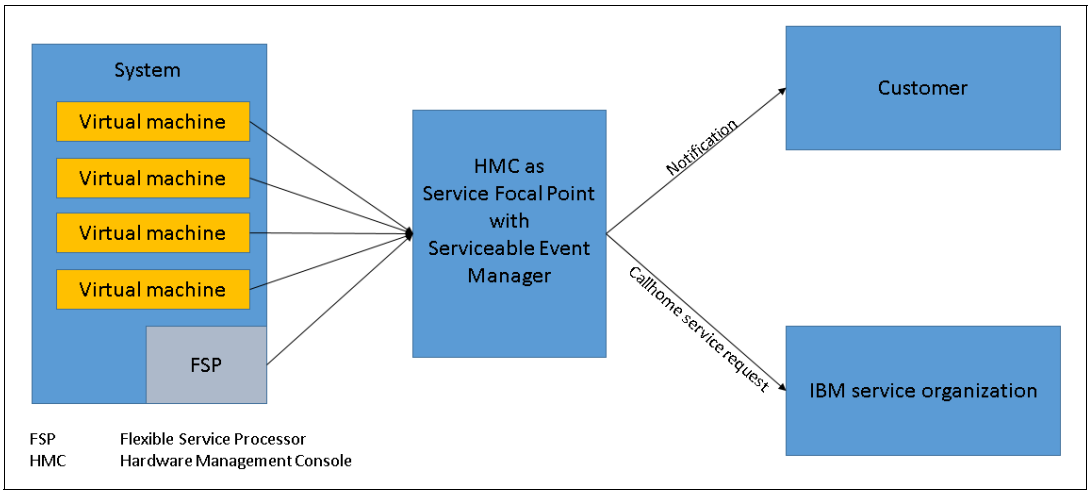

Figure 1-13 Overview of the service event flow

Error logging and analysis

When the root cause of an error is identified by a fault isolation component, an error log entry is created with basic data, such as the following examples:

- An error code that uniquely describes the error event
- The location of the failing component
- The part number of the component to be replaced, including pertinent data such as engineering and manufacturing levels
- Return codes
- Resource identifiers
- FFDC data

Data that contains information about the effect that the repair has on the system is also included. Error log routines in the operating system and FSP can then use this information and decide whether the fault is a Call Home candidate. If the fault requires support intervention, a call is placed with service and support, and a notification is sent to the contact that is defined in the ESA-guided setup wizard.

Remote support

The Remote Management and Control (RMC) subsystem is delivered as part of the base operating system, including the operating system that runs on the HMC. RMC provides a secure transport mechanism across the LAN interface between the operating system and the HMC and is used by the operating system diagnostic application for transmitting error information. It performs several other functions, but they are not used for the service infrastructure.

Service Focal Point application for partitioned systems

A critical requirement in a logically partitioned environment is to ensure that errors are not lost before being reported for service, and that an error should be reported only once, regardless of how many logical partitions experience the potential effect of the error. The SFP application on the management console is responsible for aggregating duplicate error reports, and it ensures that all errors are recorded for review and management.

The SFP application also provides other service-related functions, such as controlling service indicators, setting up Call Home, and providing guided maintenance.

When a local or globally reported service request is made to the operating system, the operating system diagnostic subsystem uses the RMC subsystem to relay error information to the HMC. For global events (platform unrecoverable errors, for example), the service processor also forwards error notification of these events to the HMC, providing a redundant error-reporting path in case there are errors in the RMC subsystem network.

The first occurrence of each failure type is recorded in the Manage Serviceable Events task on the management console. This task then filters and maintains a history of duplicate reports from other logical partitions or from the service processor. It then looks at all active service event requests within a predefined timespan, analyzes the failure to ascertain the root cause and, if enabled, initiates a Call Home for service. This method ensures that all platform errors are reported through at least one functional path, which resultis in a single notification for a single problem.

Extended error data

Extended error data (EED) is additional data that is collected either automatically at the time of a failure or manually at a later time. The data that is collected depends on the invocation method, but includes information such as firmware levels, operating system levels, additional fault isolation register values, recoverable error threshold register values, system status, and any other pertinent data.

The data is formatted and prepared for transmission back to IBM either to assist the service support organization with preparing a service action plan for the IBM SSR or for additional analysis.

System dump handling

In certain circumstances, an error might require a memory dump to be automatically or manually created. In this event, the memory dump may be offloaded to the HMC. Specific management console information is included as part of the information that optionally can be sent to IBM Support for analysis.

If additional information that relates to the memory dump is required, or if viewing the memory dump remotely becomes necessary, the management console memory dump record notifies the IBM Support center regarding on which managements console the memory dump is located. If no management console is present, the memory dump might be either on the FSP or in the operating system, depending on the type of memory dump that was initiated and whether the operating system is operational.

Notification

After a Power Systems server detects, diagnoses, and reports an error to an appropriate aggregation point, it then takes steps to notify the client and, if necessary, the IBM Support organization. Depending on the assessed severity of the error and support agreement, this client notification might range from a simple notification to having field service personnel automatically dispatched to the client site with the correct replacement part.

Client Notify

When an event is important enough to report but does not indicate the need for a repair action or the need to Call Home to IBM Support, it is classified as *Client Notify*. Clients are notified because these events might be of interest to an administrator. The event might be a symptom of an expected systemic change, such as a network reconfiguration or failover testing of redundant power or cooling systems. These events include the following examples:

- Network events, such as the loss of contact over a local area network (LAN)
- Environmental events, such as ambient temperature warnings.
- Events that need further examination by the client (although these events do not necessarily require a part replacement or repair action).

Client Notify events are serviceable events because they indicate that something happened that requires client awareness if the client wants to take further action. These events can be reported to IBM at the discretion of the client.

Call Home

Call Home refers to an automatic or manual call from a customer location to an IBM Support structure with error log data, server status, or other service-related information. The Call Home feature starts the service organization so that the appropriate service action can begin. Although configuring a Call Home function is optional, clients are encouraged to implement this feature to obtain service enhancements, such as reduced problem determination and faster and potentially more accurate transmission of error information. In general, using the Call Home feature can result in increased system availability.

Vital product data and inventory management

Power Systems store vital product data (VPD) internally, which keeps a record of how much memory is installed, how many processors are installed, the manufacturing level of the parts, and so on. These records provide valuable information that can be used by remote support and IBM SSRs, enabling the IBM SSRs to assist in keeping the firmware and software current on the server.

IBM Service and Support Problem Management database

At the IBM Support center, historical problem data is entered into the IBM Service and Support Problem Management database. All of the information that is related to the error, along with any service actions that are taken by the IBM SSR, is recorded for problem management by the support and development organizations. The problem is then tracked and monitored until the system fault is repaired.

Locating and servicing

The final component of a comprehensive design for serviceability is the ability to effectively locate and replace parts requiring service. POWER processor-based systems use a combination of visual cues and guided maintenance procedures to ensure that the identified part is replaced correctly, every time.

Packaging for service

The following service enhancements are included in the physical packaging of the systems to facilitate service:

- Color coding (touch points)

  Terracotta-colored touch points indicate that a component (FRU or CRU) can be concurrently maintained.

  Blue-colored touch points delineate components that may not be concurrently maintained (they might require that the system is turned off for removal or repair).

- Tool-less design

  Selected IBM systems support tool-less or simple tool designs. These designs require no tools, or require basic tools such as flathead screw drivers, to service the hardware components.

- Positive retention

  Positive retention mechanisms help ensure proper connections between hardware components, such as from cables to connectors, and between two cards that attach to each other. Without positive retention, hardware components risk become loose during shipping or installation, which prevents a good electrical connection. Positive retention mechanisms such as latches, levers, thumb-screws, pop plastic latches, and cables are included to help prevent loose connections and aid in installing (seating) parts correctly. These positive retention items do not require tools.

Light Path

The Light Path LED function is for scale-out systems, including Power Systems, such as model Power S822 and S822L, that can be repaired by clients. In the Light Path LED implementation, when a fault condition is detected on the POWER8 processor-based system, an amber FRU fault LED is illuminated (turned on solid), which is then rolled up to the system fault LED. The Light Path system pinpoints the exact part by lighting the amber FRU fault LED that is associated with the part that must be replaced.

The servicer can clearly identify components for replacement by using specific component level identify LEDs, and can also guide the IBM SSR directly to the component by signaling (flashing) the FRU component identify LED, and rolling up to the blue enclosure Locate LED.

After the repair, the LEDs shut off automatically when the problem is fixed. The Light Path LEDs are only visible while system is in standby power. There is no gold cap or battery implemented.

Service labels

Service providers use these labels to assist with maintenance actions. Service labels are in various formats and positions, and are intended to transmit readily available information to the IBM SSR during the repair process.

Several of these service labels and their purposes are described in the following list:

- *Location diagrams* are strategically positioned on the system hardware and relate information about the placement of hardware components. Location diagrams can include location codes, drawings of physical locations, concurrent maintenance status, or other data that is pertinent to a repair. Location diagrams are especially useful when multiple components are installed, such as DIMMs, sockets, processor cards, fans, adapter, LEDs, and power supplies.

- *Remove or replace procedure labels* contain procedures that are often found on a cover of the system or in other locations that are accessible to the IBM SSR. These labels provide systematic procedures, including diagrams, detailing how to remove and replace certain serviceable hardware components.
- *Numbered arrows* are used to indicate the order of operation and serviceability direction of components. Various serviceable parts, such as latches, levers, and touch points, must be pulled or pushed in a certain direction and order so that the mechanical mechanisms can engage or disengage. Arrows generally improve the ease of serviceability.

The operator panel on a POWER processor-based system is an LCD display (two rows by 16 elements) that is used to present boot progress codes, indicating advancement through the system power-on and initialization processes. The operator panel is also used to display error and location codes when an error occurs that prevents the system from booting. It includes several buttons, enabling an IBM SSR or client to change various boot-time options and for other limited service functions.

Concurrent maintenance

The IBM POWER8 processor-based systems are designed with the understanding that certain components have higher intrinsic failure rates than others. These components can include fans, power supplies, and physical storage devices. Other devices, such as I/O adapters, can begin to wear from repeated plugging and unplugging. For these reasons, these devices are designed to be concurrently maintainable when properly configured. Concurrent maintenance is facilitated because of the redundant design for the power supplies, fans, and physical storage.

In addition to the previously mentioned components, the operator panel can be replaced concurrently by using service functions of the ASMI menu.

Repair and verify services

Repair and verify (R&V) services are automated service procedures that are used to guide a service provider, step-by-step, through the process of repairing a system and verifying that the problem was repaired. The steps are customized in the appropriate sequence for the particular repair for the specific system being serviced. The following scenarios are covered by R&V services:

- Replacing a defective FRU or a CRU
- Reattaching a loose or disconnected component
- Correcting a configuration error
- Removing or replacing an incompatible FRU.
- Updating firmware, device drivers, operating systems, middleware components, and IBM applications after replacing a part

R&V procedures can be used by IBM SSR providers who are familiar with the task and those who are not. Education-on-demand content is placed in the procedure at the appropriate locations. Throughout the R&V procedure, repair history is collected and provided to the Service and Support Problem Management Database for storage with the serviceable event to ensure that the guided maintenance procedures are operating correctly.

Clients can subscribe through the subscription services on the IBM Support Portal to obtain notifications about the latest updates that are available for service-related documentation.

IBM Knowledge Center

The IBM Knowledge Center provides you with a single information center where you can access product documentation for IBM systems hardware, operating systems, and server software.

The latest version of the documentation is accessible through the Internet; however, a CD-ROM based version is also available.

The purpose of the IBM Knowledge Center, in addition to providing client-related product information, is to provide softcopy information to diagnose and fix any problems that might occur with the system. Because the information is electronically maintained, changes due to updates or addition of new capabilities can be used by service representatives immediately.

The IBM Knowledge Center contains sections specific to each server model, and include detailed service procedures for a number of potential repair situations. The service procedure repository for a particular server model can be found in the “Troubleshooting, service and support” sections. See the IBM Knowledge Center:

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

QR code labels for servicing information

A label containing a QR code can be found on the top service cover of the Power S822 and S822L server. This can be scanned with an appropriate app on a mobile device to link to a number of sources of information that simplify the servicing of the system.

From this quick access link, you can find information on these topics:

- Installing and configuring the system
- Troubleshooting and problem analysis
- Reference code lookup tables
- Part location guides
- Removing and replacing field replaceable units
- Video guides for removal and installation of customer replaceable units
- Warranty and maintenance contracts
- Full product documentation

Serviceability of the V7000 storage components

This section lists the methods that make the high levels of serviceability of the V7000 storage available:

- Cluster error logging
- Asynchronous error notification
- Dump capabilities to capture software detected failures
- Concurrent diagnostic procedures
- Directed maintenance procedures with simplified drive replacement process
- Concurrent log analysis and memory dump data recovery tools
- Concurrent maintenance of all IBM Storwize V7000 components
- Concurrent upgrade of IBM Storwize V7000 Software and microcode
- Concurrent addition or deletion of node canisters in a clustered system
- Automatic software version correction when replacing a node
- Detailed status and error conditions displayed on the service panel
- Error and event notification through SNMP, syslog, and email

Serviceability of the PurePower Management Node

The PurePower Management node provides servicability features, including redundant, hot-pluggable power supplies and fans.

Serviceability of the Storage SAN48B-5 Switch 2498-F48

The SAN48B-5 switch can address demanding reliability, availability, and serviceability (RAS) requirements to help minimize downtime to support mission-critical cloud environments through the following features:

- Extensive diagnostics and system monitoring capabilities to enhance high RAS
- Diagnostic Port (D-Port) feature for physical media diagnostic, troubleshooting, and verification services
- Redundant, hot-pluggable components and nondisruptive software upgrades

Serviceability of the Mellanox SX1710 Data Network Switch 8831-NF2

The Mellanox SX1710 Data Network Switch is designed for serviceability and provides following features:

- System health monitoring.
- Redundant hot-pluggable power supplies (ac) and four redundant hot-pluggable fans.
- Redundant pair of management interfaces.

Serviceability of the Systems Management G8052 Switch 7120-48E

The G8052 Switch provides redundant hot-pluggable power supplies and fans.

IBM PurePower System Firmware Maintenance

Firmware updates are the responsibility of the customer, but can be accomplished for all components from a single point of access via the PurePower Integrated Manager, which runs on the management node.

Firmware entitlement

Firmware installations are restricted to entitled servers. The client must be registered with IBM and entitled with a service contract. During the initial machine warranty period, the access key is installed in the machine by manufacturing. The key is valid for the regular warranty period, plus additional time.

Any firmware release that was made available during the entitled time frame can still be installed. For example, if the entitlement period ends 31 December 2016 and a new firmware release is released before the end of that entitlement period, it can still be installed. If that firmware is downloaded after 31 December 2016, but it was made available before the end of the entitlement period, it still can be installed. Any newer release requires a new update access key.

Firmware and Software updates

Within the PurePower Integrated Manager check the Compliance page for available updates of the single components.

IBM provides updates if needed via Fix Central to the PurePower Integrated Management Node appliance images. This includes all noncritical and critical PSIRT/Security vulnerability fixes and defect fixes.

RPM fixes for RHEL and built-in components can be downloaded directly from the Red Hat repository using traditional Red Hat Satellite Subscription services.

Fixes for add-on components, such as Nagios, Python, and Perl can be downloaded and applied by the customer from the respective open source communities and forums.

Future functions will automate compliance, acquisition, and installation of critical fixes and quarterly appliance updates.

POWER8 compute nodes firmware update steps

To update the POWER8 system firmware, complete the following steps:

1. Log on to the PurePower Integrated Manager, and select the HMC link (see Figure 1-14) to start the HMC web server interface.

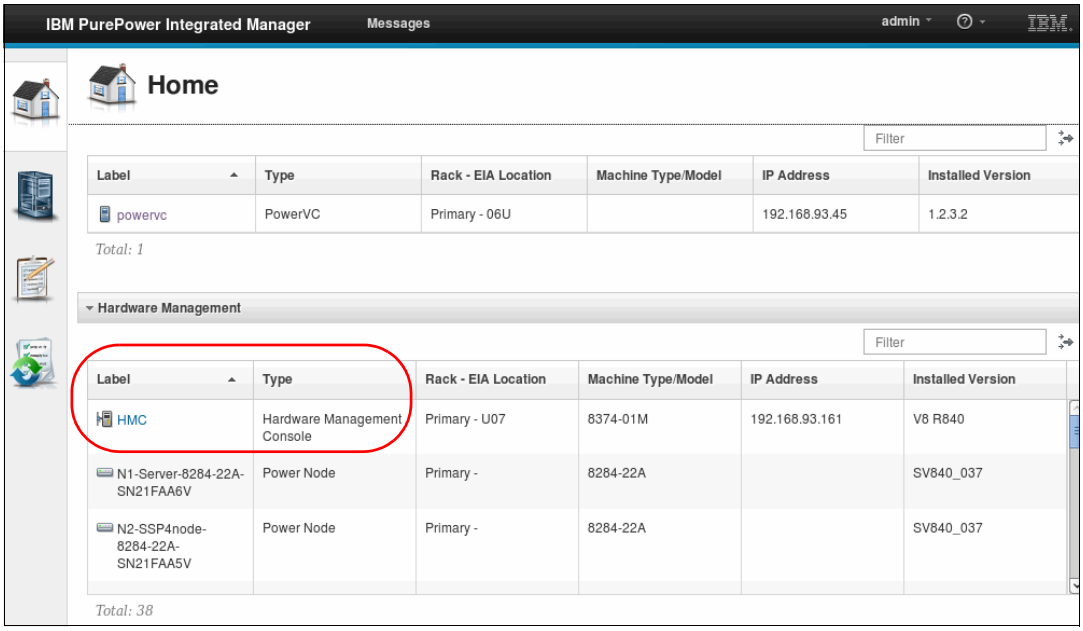

Figure 1-14 Start the HMC web interface

When the HMC is selected, the Welcome screen shown in Figure 1-15 opens.

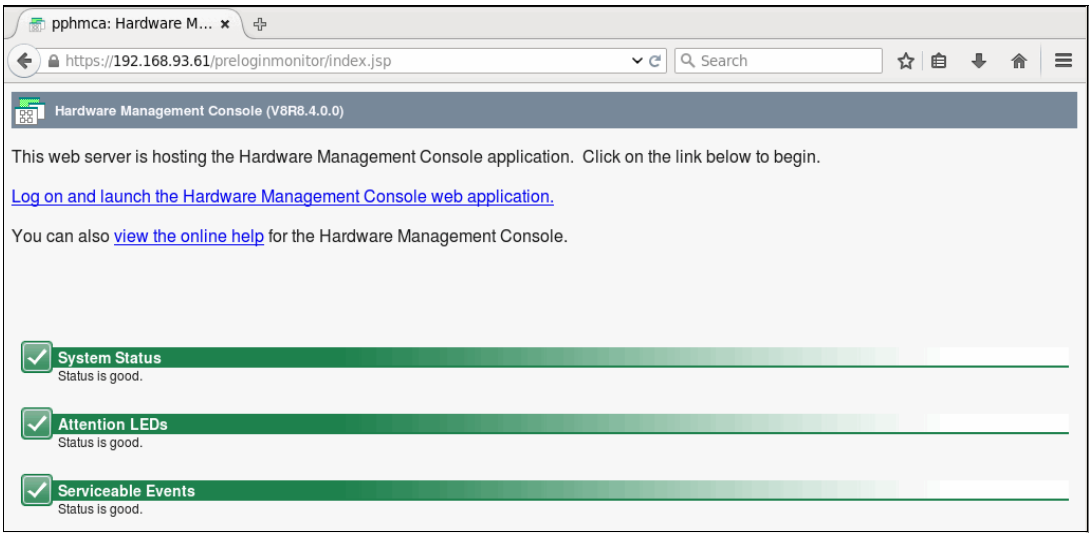

Figure 1-15 HMC welcome screen

2. To log on to the HMC, click **Log on** and launch the Hardware Management Console web application from the Welcome window showing in Figure 1-16.

   The HMC has a predefined user ID, `hscroot`, and the default password is `abc123`.

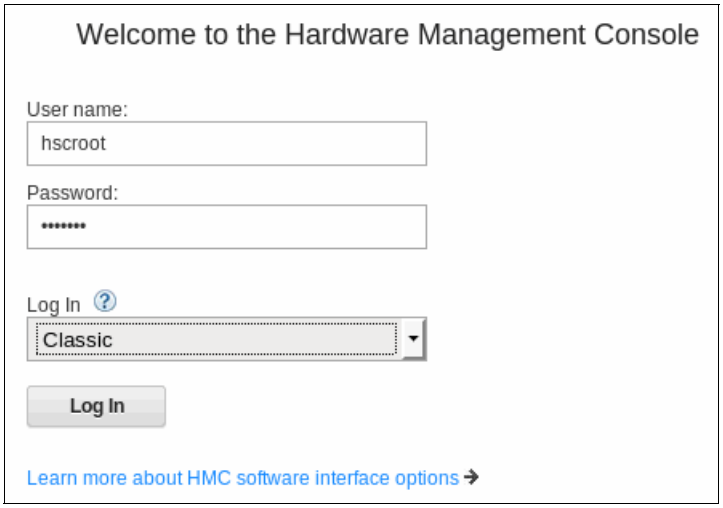

Figure 1-16 HMC logon screen

3. After logging in to the HMC, click **Systems Management** → **Server** and, in the work pane, select the server to update.

   Then, in the Task pane (Figure 1-17), expand **Updates**.

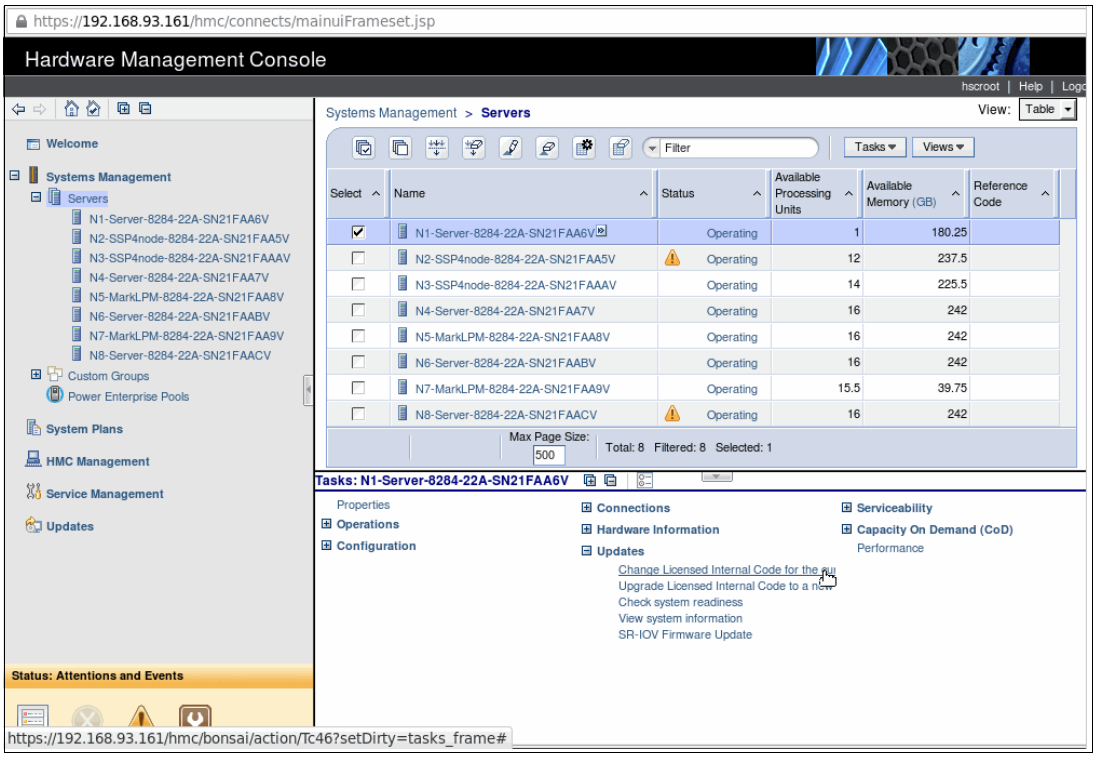

Figure 1-17 Start Change Licensed Internal Code for the current release

4. Select **Change Licensed Internal Code for the current release** for an update of the currently installed system firmware release or select **Upgrade Licensed Internal Code to a new release** for an upgrade to a new system firmware release.

   Then, select **Start Change Start Change Licensed Internal Code wizard** as shown in Figure 1-18.

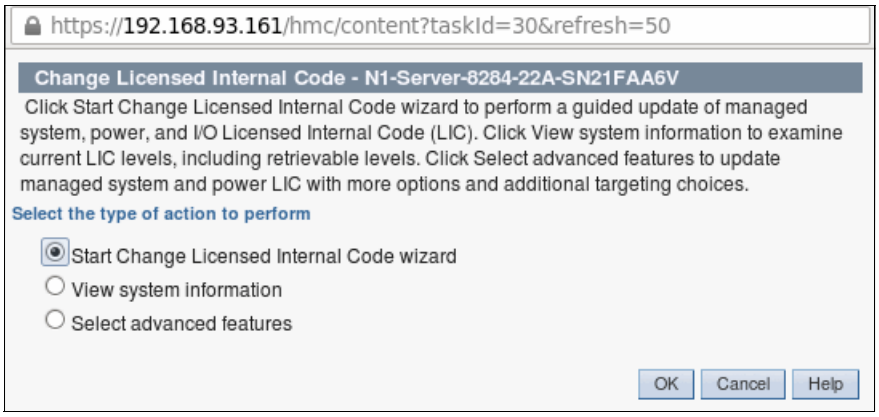

Figure 1-18 Select Start Change Licensed Internal Code wizard

5. Specify the LIC repository that you want to use, as Figure 1-19 shows.

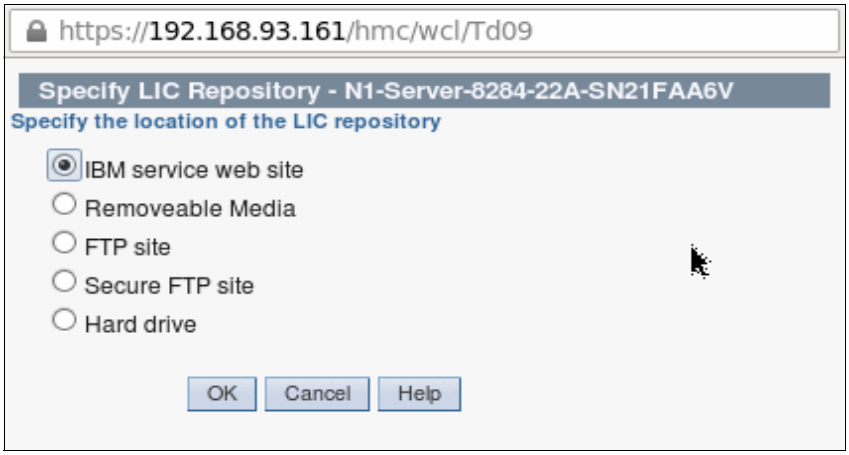

Figure 1-19 Select the LIC Repository

6. Select **Next** to start the check for available firmware updates (see Figure 1-20).

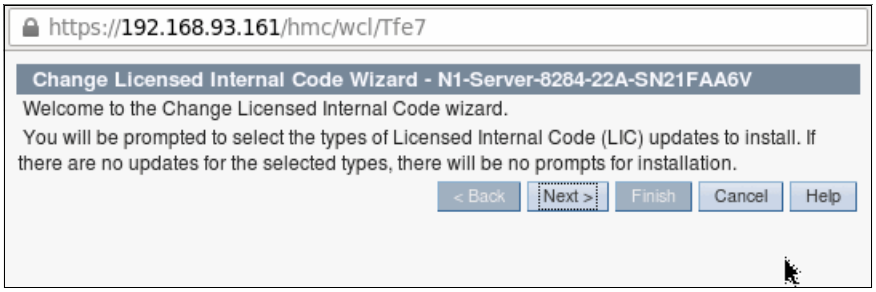

Figure 1-20 Change LIC welcome screen

For this example, no updates are available, as Figure 1-21 indicates.

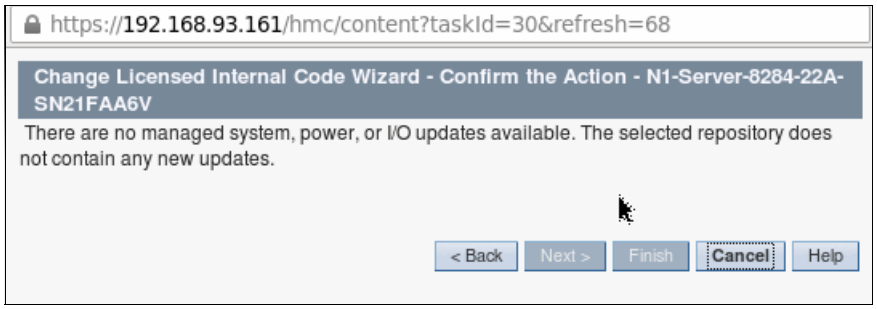

Figure 1-21 No updates are available

If any update is shown as available, make sure that you already checked the Compliance page in the PurePower Integrated Manager for compatibility with the PurePower System.

2

Managing PurePower

The IBM PurePower System offering combines various devices and capabilities into a single infrastructure. The software to manage that offering is the IBM PurePower Integrated Manager.

2.1 System management overview

By using PurePower Integrated Manager, you can manage the discrete parts of the offering from a single integrated interface. It provides the following capabilities:

- Links to the PowerVC, HMC, Nagios Core Open Source software, and the individual management interfaces provided by each managed piece of hardware.
- The ability to work with the hardware interfaces and manage your network configuration.
- Interfaces to work with virtual machines, and deploy managing agents to those machines.
- Functions to assess the software compliance of resources in your environment, and provide links to download software updates for those resources.

PurePower Integrated Manager provides the following user interfaces:

- A browser-based graphical user interface (GUI) that makes it easy for you to work with your physical and virtual resources. The GUI groups resources for you in the following ways:
  - Converged Infrastructure Management

    Links to the Nagios Core Open Source software from which you can define monitoring and alerting for elements within your environment.
  - Virtualization Management

    Links to the Power Virtualization Center (PowerVC) from which you can manage your virtual machines.
  - Hardware Management

    Lists the hardware resources that are monitored within your environment.
  - Unmonitored Resources

    Lists any other resources that have been added to PurePower Integrated Manager for inventory purposes but are not monitored.
  - Compliance page

    Provides a combined view of existing hardware components and the fix pack version installed on the components. See “Compliance page” on page 70 for details.
- A command-line interface that rounds out the functions of the GUI.

For more information about the IBM PurePower System offering, including the PurePower Integrated Manager, see the IBM PurePower System section of the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/PurePower

2.1.1 PurePower Integrated Manager

The main software to manage IBM PurePower is the PurePower Integrated Manager (PPIM), a single integrated interface for the various devices and capabilities that IBM PurePower combines in a single infrastructure.

It provides the following functions, aspects, and uses:

- Tools to manage hardware interfaces and the network configuration
- Links from a single pane to the other management and monitoring tools like PowerVC, the vHMC, Nagios, and the individual GUIs of the hardware components

- Interfaces used to work with the LPARs
- Functions to assess the software compliance of resources, providing also links to their software updates

To use the functions offered by PPIM, you can either use a graphical user interface (GUI), accessible with a web browser connected to the IBM PurePower network, or a traditional command line interface (CLI) that rounds out the functions of the GUI.

This section provides a brief description of PPIM. For a more advanced and detailed tutorial on this software and its more advanced features, like backing up the system components of IBM PurePower with PPIM, see “Configuring and managing an IBM PurePower System” in IBM Knowledge Base:

http://ibm.co/1HUOkKb

PPIM virtual machine

PPIM, like any other management tool in IBM PurePower, is preinstalled and preintegrated in the IBM PurePower offering. Therefore, it is already available when the system is delivered at the customer’s site. The software is installed on both management nodes (one is a redundant instance) in a VM, which is shared with Nagios and is virtualized with KVM on a RHEL 7.1 host. Table 2-1 shows the details of the VM.

Table 2-1 VM details

Software	OS	vCPU	Memory	Disk space
PPIM 1.1, Nagios	RHEL 7.1	2	12 GB	58 GB

For more information about the other software components installed on the management node, see 1.2.2, “Software stack” on page 31.

Accessing PPIM

First, after setting up PurePower System and connecting it to the data center network, access PPIM is good to get an overview of the system and its capabilities.

Note: This section is based on the assumption that you are using default IP addresses and default login credentials for IBM PurePower software components. If you have changed the defaults, simply replace the defaults with your custom IP addresses or credentials.

There are two ways to access PPIM:

- Using the flat panel rack-mounted monitor and keyboard
- Connecting to the PPIM remotely

Using the flat panel rack-mounted monitor and keyboard

Follow these steps to use the monitor and keyboard:

1. Power on the management nodes by pressing the buttons on the right side of the node chassis, and wait five minutes for the nodes to power on.
2. After sliding out the console display and keyboard and seeing the display powering on, press the PRTSC key to enable the KVM switch, which will prompt the management nodes. Select the primary one, and then press Enter.
3. You will now see a login prompt. The following credentials are the defaults:

   User name `admin` or `root`
   Password `PASSWORD`

4. After log in, double-click the PurePower Integrated Manager icon to launch the PPIM user interface. You will be asked for the PPIM login credentials, which are as follows by default:

   User name `admin` or `root`
   Password `PASSWORD`

Connecting to the PPIM remotely

Follow these steps to connect remotely:

1. From a computer with a desktop environment and access to the PurePower private network, open a web browser and type the PPIM address, which is as follows by default:

   `https://192.168.93.46/puremgr/`

2. Enter your login credentials for PowerVC, which are as follows by default:

   User name `admin` or `root`
   Password `PASSWORD`

PPIM pages

PPIM functions are grouped into several pages. This section describes the information and available on these pages, as well as the actions that the user can perform. For detailed tutorials, see the “Learn More” link that is available on each page in PPIM.

Home page

The home page (Figure 2-1) provides a single integrated interface to access element managers or applications. Resources are listed in tables and, where available, a link is associated with each element or application, like PowerVC or the compute note 5. Clicking the link launches the element manager or application in a new browser window.

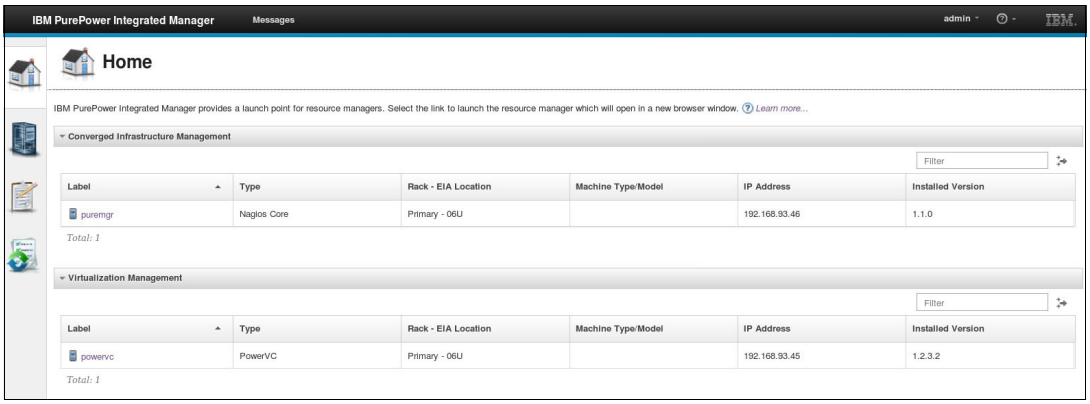

Figure 2-1 PPIM home page

The home page groups resources in the following manner:

Converged Infrastructure Management Links to Nagios, from which the user can define monitoring and alerting for elements within the environment.

Virtualization Management Links to PowerVC, which enables the handling of the LPARs across the physical infrastructure

Hardware Management Lists the hardware resources that are monitored within IBM PurePower environment; the user can access the GUIs of the specific hardware.

Unmonitored Resources	Lists any other resources that have been added to PurePower Integrated Manager for inventory purposes, but are not monitored.

The tables include the following information:

- The resource type
- The rack identifier and EIA location
- The machine type and model number
- The management interface IPv4 address
- The installed version

Virtual Machines page

The Virtual Machines page lists the LPARs identified on the IBM PurePower infrastructure. This section is fed by PowerVC, which has the primary role of defining and collecting LPAR information. If the information is outdated, the Recollect Machine Information button queries PowerVC and updates the LPAR list. From this section, the user can manage or monitor agents and change IP addresses on the target LPARs.

Monitoring agents are what enable Nagio Core and, in turn, PPIM, to monitor the status and the information on the LPARs. To do that, a user ID and a password must be defined for PPIM to access the LPARs. This action can also be performed from this page. Figure 2-2 shows this page, in a situation where no LPAR has been yet created.

The agent deployed for monitoring can be customized with open source plug-ins for Nagios, with many choices available to monitor applications or infrastructure running in the virtual machines deployed via PowerVC. The Virtual Machines page provides a way to select custom agent packages for deployment.

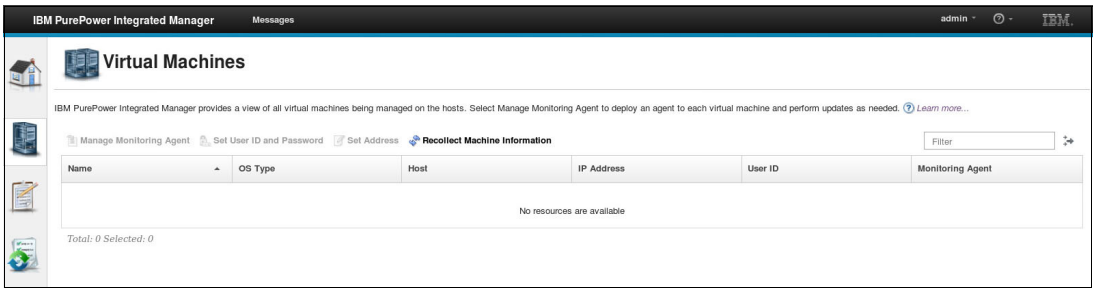

Figure 2-2 PPIM Virtual Machines page (no VM yet added)

Note: The Set User ID and Password function does not update any values on any LPAR. It only updates the information used by PurePower Integrated Manager to access the LPAR. If the user ID and password combination used on this function is not a valid combination for the selected resources, the function will fail. The user will need to first log in to the LPAR, create a new set of user name and password for PPIM to use, and then set these new parameters on this section.

Hardware Inventory page

Hardware Inventory page (Figure 2-3) lists the preconfigured hardware of IBM PurePower in the environment. From this page the user can add, edit, and remove hardware resources. The user can also make adjustments to the management network configuration.

This section also includes other information about the resource, such as the EIA location, the user ID to manage the resource, and the fix pack-installed version of each device.

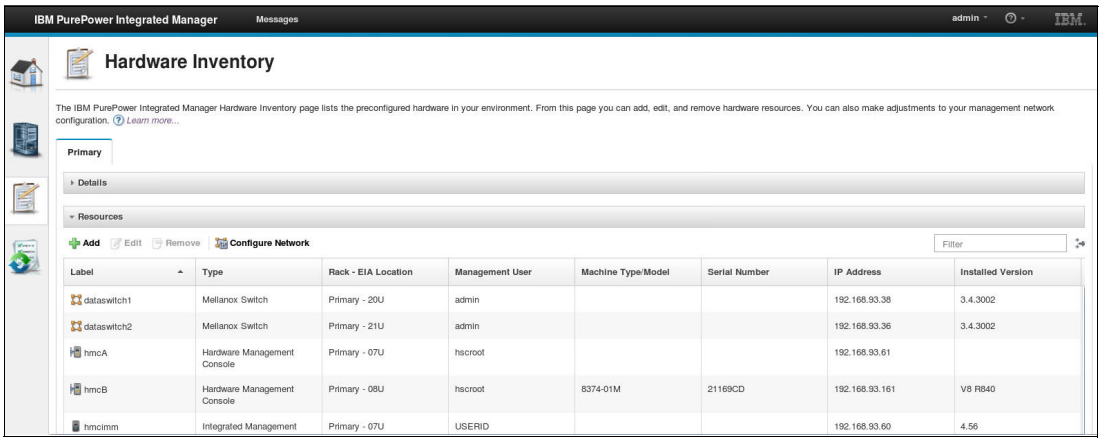

Figure 2-3 PPIM Hardware Inventory page

Note: The Add Resource function does not perform any configuration on the device, but only sets the credentials and IP address of the device in PPIM. It is important then that these parameters are configured on the target devices beforehand.

Compliance page

The Compliance page on PPIM provides a combined view of existing hardware components and the fix pack version installed on the components. It also provides a comparison between the base and the recommended fix pack version of the components.

The *base* version represents the level of the firmware or software for the device that was tested as part of the release of PurePower Integrated Manager.

The *recommended* version is the currently recommended level of the firmware or software that has been tested on a PurePower Systems rack.

A device is *compliant* when the version of the firmware or software that is installed on the device is the same as the base version or the recommended version. In case of non-compliance, the page provides a link through which the user can download and install the new version to make the device compliant.

Compliance is dependent upon the version of PPIM installed. This page, shown in Figure 2-4, offers the opportunity to change the version of PPIM by which to assess the compliance of the resources with the Assess Compliance for PurePower Version field.

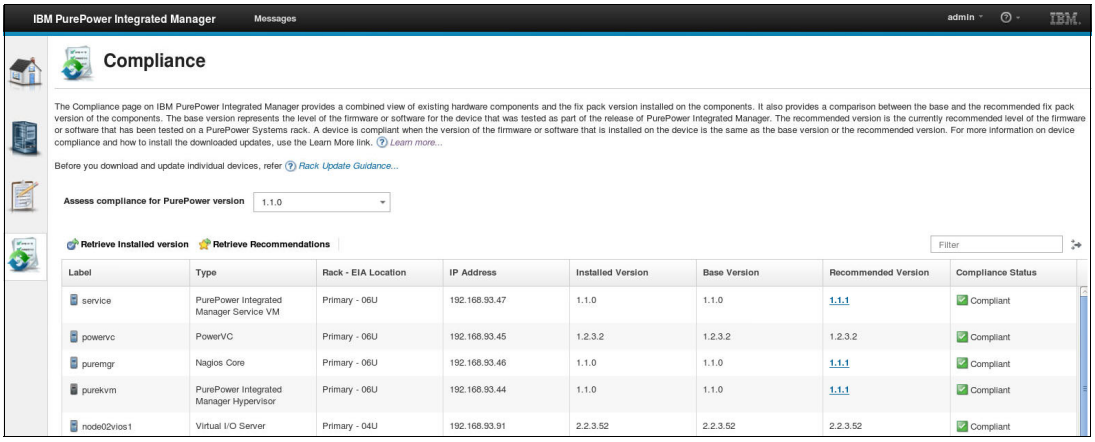

Figure 2-4 PPIM Compliance page

2.1.2 Virtual HMC

This section describes the new virtual HMC (vHMC) to provide the virtualization capabilities for the PurePower System. With the October 2015 release, the new vHMC is a VM on both PurePower Management Nodes IBM 8374-01M.

The virtual HMC (vHMC) can be started with the IBM PurePower Integrated Manager which is in another VM on the PurePower Management Nodes. For more information about the PPIM see 2.1.1, “PurePower Integrated Manager” on page 66.

The new virtual HMC (vHMC) based on the 840 code levels instead of a traditional hardware-based second HMC 7042-CR8 with 830 code level.

It is not supported to connect external IBM Power Systems to the PurePower System virtual HMC at the time of writing.

General HMC introduction

The HMC connects to one or more managed systems to perform various functions, including the following primary functions:

- Providing a console for system administrators and service providers to manage server hardware
- Creating and maintaining a multiple partitioned environment on a managed system
- Detecting, reporting, and storing changes in hardware conditions
- Acting as a service focal point for service providers to determine an appropriate service strategy
- Displaying operating system session terminals for each partition

Figure 2-5 shows the server on a virtual HMC.

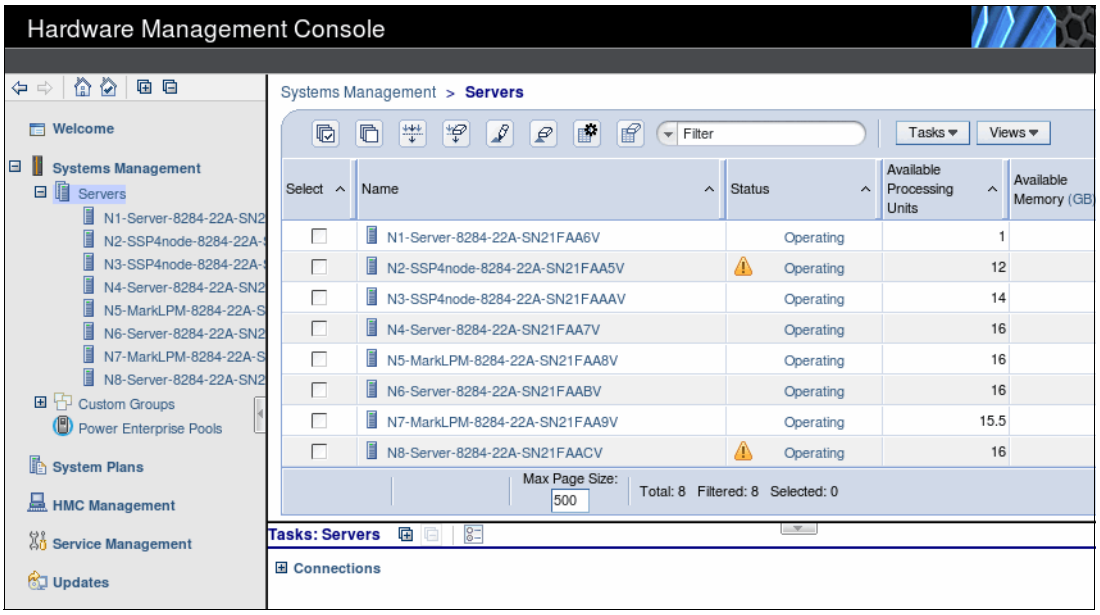

Figure 2-5 The server on a virtual HMC

These are the specifications of the vHMC VM:

- 8 GB of memory
- 4 virtual CPUs
- 4 network interfaces vNICs
- 160 GB disk

Figure 2-6 shows IBM PurePower Integrated Manager with the virtual HMC.

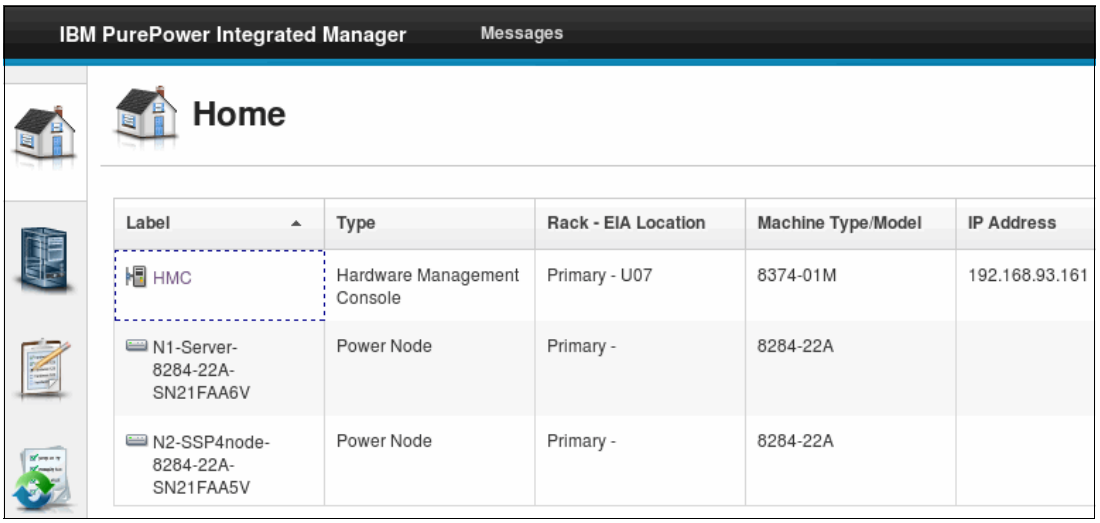

Figure 2-6 IBM PurePower Integrated Manager with the virtual HMC

Virtual HMC (vHMC) network description

Initially, the PurePower rack as delivered from manufacturing to the customer has the two management nodes preconfigured identically. The POWER8 technology-based servers are managed by both the Side A and the Side B vHMC. Each side manages the FSP on a unique subnet. For example, Side A vHMC uses 10.0.255.x subnet for Switch 1 FSP physical connections (VLAN 3050), and Side B vHMC uses 10.255.0.x subnet for Switch 2 FSP physical connections (VLAN 3055). This is a standard dual-HMC redundant configuration.

Either the vHMC manages the PurePower environment or the Power managers. Both cannot manage at the same time.

2.1.3 PowerVC

IBM Power Virtualization Center (PowerVC) is the next generation of enterprise virtualization management tools for IBM Power Systems. With a powerful yet simple and intuitive GUI and deep integration with IBM PowerVM virtualization technologies, PowerVC simplifies management of virtualization for Power Systems servers that run AIX, IBM i, and Linux operating systems.

For both PowerVM and PowerKVM systems, PowerVC includes the following primary functions:

- Simplified virtualization management for IBM Power Systems servers
- Easy replication of virtual machines for consistency and fast deployment
- Resource pooling and dynamic logical partition (LPAR) placement
- Monitoring the use of the resources

Available on any IBM Power Systems server, PowerVC 1.2.3 Standard Edition for PowerVM management has been preintegrated and preconfigured on IBM PurePower to provide its virtualization features across the whole infrastructure, without requiring the customer the initial installation or configuration.

One example in which PowerVC comes in handy is with the deployment of new LPARs. Without PowerVC, the user should perform several operations, including the creation of the LPAR, the adapters assignment, the assignment of the storage space on the SAN, and so on. This is usually done through the use of several GUIs and much manual intervention, with a high margin of error. PowerVC instead automates all these tasks, as it has control over the compute, storage and network resources, allowing users an easier implementation of complex cloud operations.

Its virtualization management functions are complementary to those of the PPIM, which in turn are designed mainly for hardware management using the vHMC APIs. PowerVC software is accessible through the PPIM interface and through a dedicated IP address.

Note: PowerVC can manage LPARs and VMs installed on IBM POWER servers both from the same or another POWER server, or from an x86 server. For IBM PurePower, PowerVC is installed on the management nodes, so the x86 version is used.

This section outlines the basics details about PowerVC on IBM PurePower. For more information about the use of PowerVC, see the IBM Redbooks publication titled *IBM PowerVC Version 1.2.3: Introduction and Configuration*, SG24-8199.

PowerVC virtual machine

PowerVC is preinstalled and preintegrated, but it is not preconfigured. Configuration must be done directly or through purchased setup services. Either way, you must have the PowerVC configuration, including adding hosts, storage, and fabrics complete after delivery.

Table 2-2 shows VM details.

Table 2-2 Details for the PowerVC VM

Software	OS	vCPU	Memory	Disk space
PowerVC 1.2.3 Standard Edition	RHEL 7.1	4	16 GB	204 GB

For more information about the other software components installed on the management node, see 1.2.2, “Software stack” on page 31.

Accessing PowerVC in IBM PurePower

PowerVC can be accessed either from the PPIM consolor or by connecting with the PowerVC IP address directly.

From the PPIM console, complete the following steps:

1. From a computer with a desktop environment and access to IBM PurePower private network, open a web browser and type the PPIM address. This is the default:

   `https://192.168.93.46`

2. Enter your login credentials, which are as follows by default:

   User name `root` or `admin`
   Password `PASSWORD`

3. You are now in the PPIM main page (Figure 2-7). Click **Home** at the top of the window, and then click **powervc**.

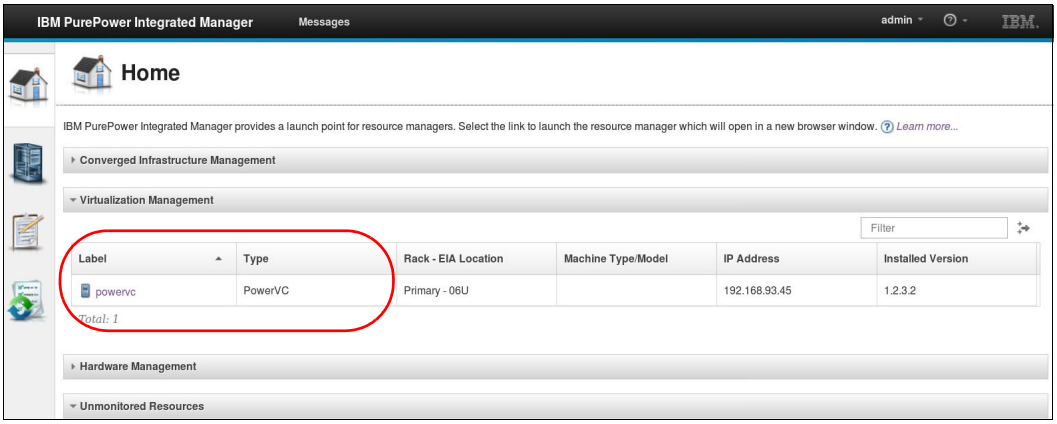

Figure 2-7 Start the PowerVC interface

4. Enter your login credentials for PowerVC, which are as follows by default:

   User name `root` or `admin`
   Password `PASSWORD`

To connect to the PowerVC IP address directly, complete these steps:

1. From a computer with a desktop environment and access to IBM PurePower private network, open a web browser and type the PowerVC address, which per default is:

   `https://192.168.93.45`

2. Enter your login credentials for PowerVC, which are as follows by default:

   User name `root` or `admin`
   Password `PASSWORD`

You are now inside PowerVC (see Figure 2-8).

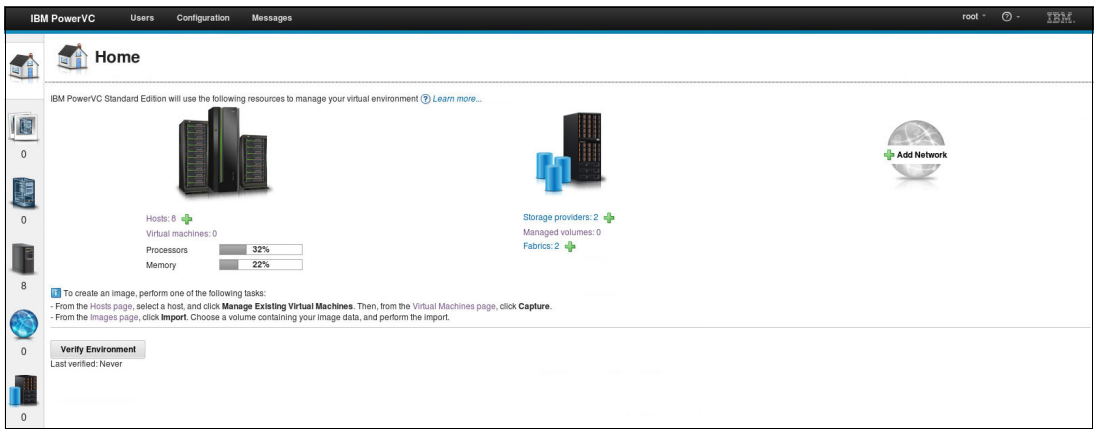

Figure 2-8 PowerVC main page

Note: This section assumes default IP addresses and default login credentials for IBM PurePower software components. If you have changed the defaults, simply replace the defaults with your custom IP addresses or credentials.

2.1.4 Nagios

This section provides an overview of the capabilities of the Nagios Core manager.

To access the Nagios Core web interface (Figure 2-9 on page 76), log on to the PurePower Integrated Manager and select the **puremgr** link with Nagios Core type to start Nagios.

The default user ID used to log in to the PurePower Integrated Manager Nagios Core web interface is `nagiosadmin` and the password is PASSW0RD (with a zero, not a capital O). More accounts can be created, or the password can be updated by using the command line.

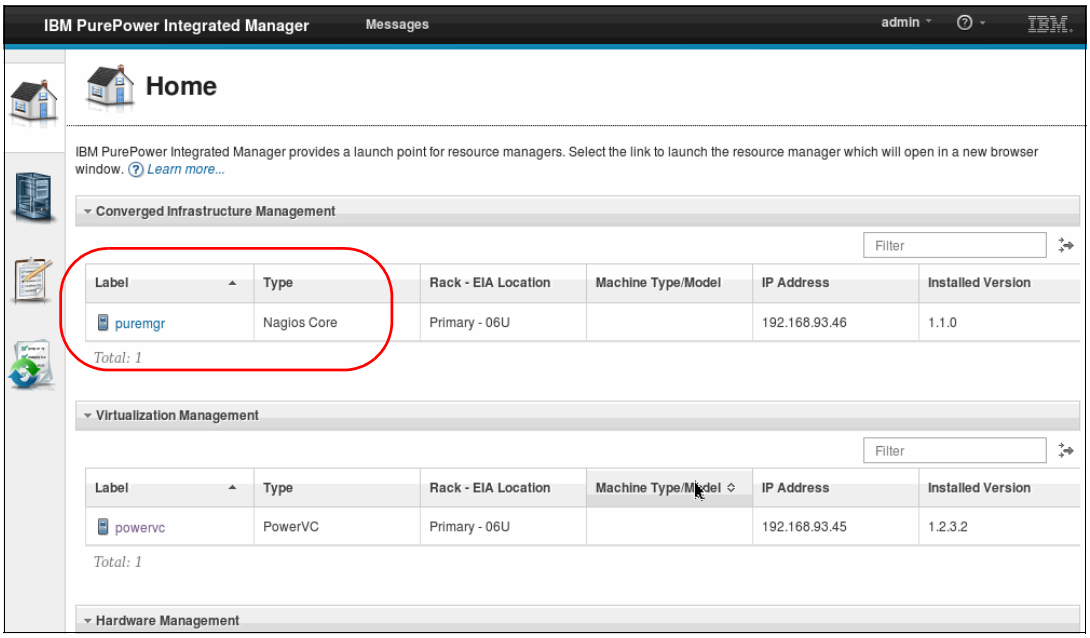

Figure 2-9 Start the Nagios Core web interface

Core functions of Nagios

The core functions of Nagios are shown in the following list:

- SNMP monitoring provides complete integration with equipment.
- SNMP trap Integration provides complete handling of alarms and alerts.
- Web-based user interface provides awareness and advanced configuration.
- Customizable dashboards allow for per-user customization.
- Monitoring plug-ins makes it easy to monitor new devices, applications, and services.
- Integrated graphs provide trending and capacity planning information.
- Advanced reports provide data export capabilities.
- Problem remediation includes alert clearing and automated restart of failed elements.
- Scheduled downtime allows for alert suppression during planned outages.

Nagios offers leading market penetration with 10 years of development and thousands of freely available custom plug-ins (commercial offerings also available).

Monitoring

Monitors all mission-critical infrastructure components, including, operating systems, network protocols, system metrics, and network infrastructure.

Figure 2-10 shows the Nagios Tactical Monitoring Overview.

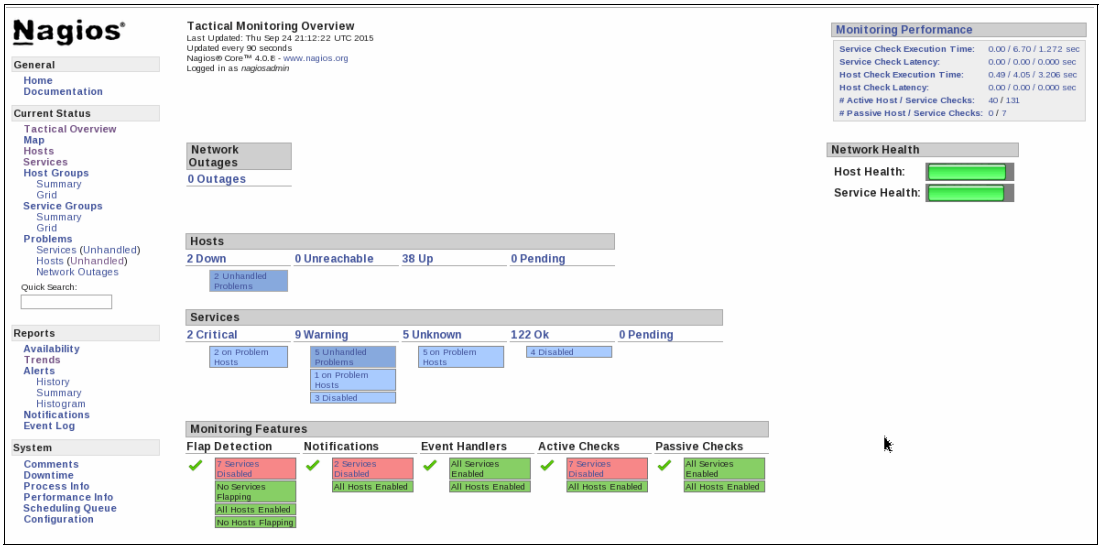

Figure 2-10 Nagios Tactical Monitoring Overview

Figure 2-11 shows the Power Distribution Unit (PDU) monitoring with Nagios.

Figure 2-11 PDU monitoring with Nagios

Centralized view

Provides a central view of your entire IT operations network and business processes. Plug-ins for all components are available.

Figure 2-12 shows the Nagios Host Status Totals.

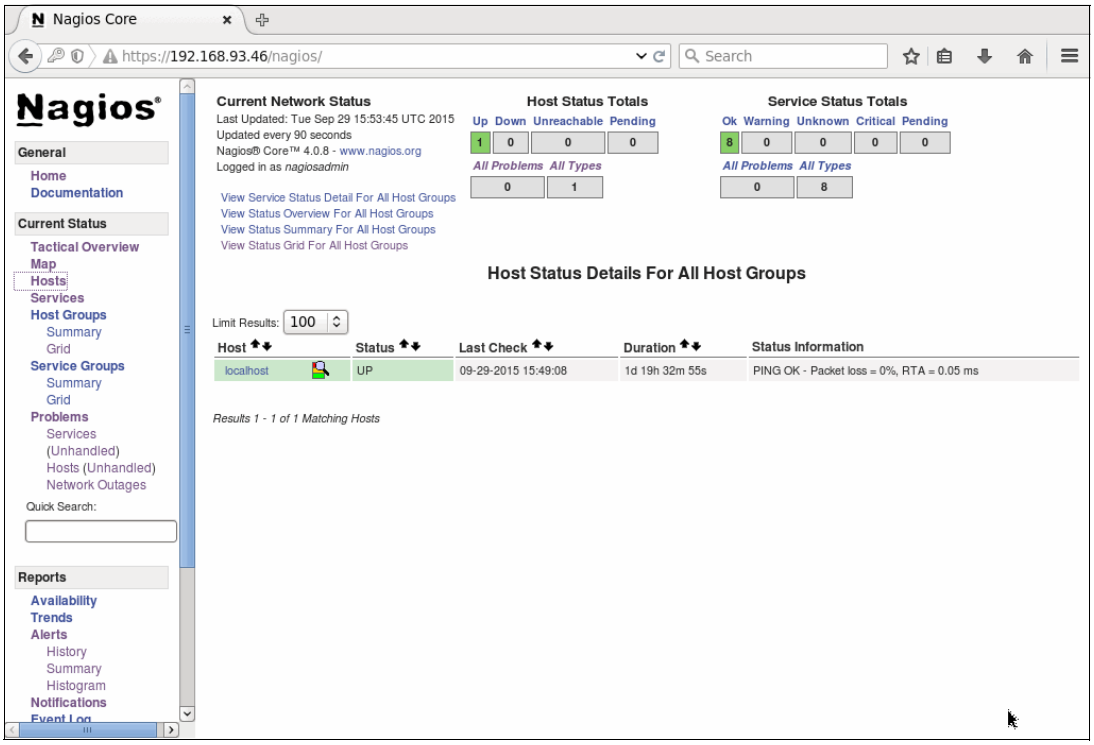

Figure 2-12 Nagios Host Status Totals

Alerts

Alerts are delivered to you via email and SMS. Multiuser notifications ensure that alerts reach the attention of the appropriate people.

Figure 2-13 shows the Nagios Alert History of all hosts and services.

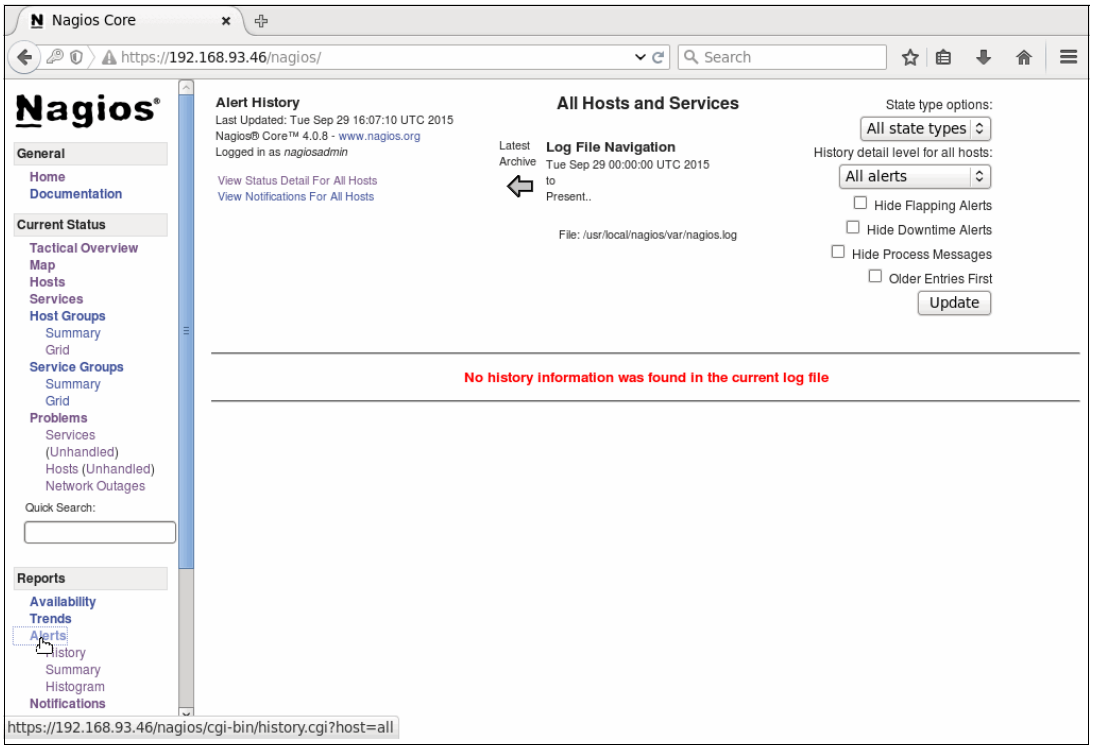

Figure 2-13 The Nagios Alert History of all hosts and services

Event handlers

Nagios event handlers can be configured with custom scripts or executable to try solve a problem before sending out an alert.

Event handlers can handle tasks that clients specify, such as restarting failed applications/devices (as an example).

For more information, see the Nagios website:

http://www.nagios.org

3

Physical planning and installation

Good planning helps you identify the unique needs and determine the best placement for your IBM PurePower System.

Planning ensures that you have everything you need and that you meet all of the PurePower System prerequisites. This minimizes errors during installation and allows for a quicker upgrade or installation. The planning information in this section helps you place the system, plan for power and environmental needs, prepare based on the size of the system ordered (rack configuration), and determine if special considerations are required for how you use the system (for example, Internet connectivity).

3.1 Planning activities

Planning scenarios that you are typically responsible for can include any or all of these items:

- A new system installation
- Hardware modifications (feature additions, removals, conversions, or relocations)
- Software modifications (additions, removals, updates, or other changes)
- Services modifications (additions, removals)
- To ensure that planning is completed successfully, you need to assign a planning project manager who provides a documented plan that includes:
- A timeline for the activities to be completed
- Each major activity phase and the wanted outcome
- A list of responsibilities and the person assigned
- A current system diagram and configuration listing, including all hardware content, software content, cabling, and other pertinent configuration items (if it is a modification to an existing system)
- An endsystem diagram that shows hardware content and configuration details, including cabling
- A key contacts list, including off-hours contact information for all key task or activity participants
- A plan for the extent of IBM Systems Lab Services startup services that are required:
  - See the IBM US Announcement letter ENUS615-014 for services options and descriptions.
  - Service options include PurePower startup with basic network integration, PurePower advanced network integration, PurePower basic SAN integration, PurePower migration of workloads, transition of FSM to PurePower Management System, as well as custom services.
  - A plan from IBM Lab Services for or communicating appropriate elements of your plan with key personnel (for example, seller, installer, management)

Planning your system requires attention to numerous details.

3.2 System placement considerations

When determining the placement of your system, consider the following factors:

- Adequate space for the devices
- Working environment of personnel who will be using the devices (their comfort, ability to access the devices, supplies, and reference materials)
- Adequate space for maintaining and servicing the devices
- Physical security requirements necessary for the devices
- Weight of the devices
- Heat output of the devices
- Operating temperature requirements of the devices
- Humidity requirements of the devices

- Air flow requirements of the devices
- Air quality of the location where the devices will be used (for example, excess dust cah damage your system)
- Altitude limitations of the devices
- Noise emission levels of the devices
- Any vibration of equipment near where the devices will be placed
- Paths of power cords

For detailed information, see “Site preparation and physical planning” in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/PurePower/p8ebe/p8ebe_kickoff.htm

3.3 Physical site requirements

System specifications provide detailed information about PurePower Systems, including dimensions, electrical data, power, temperature, and environmental requirements. The following five tables list specific data:

- Table 3-1 lists rack dimensions.
- Table 3-2 lists weight of building blocks.
- Table 3-3 lists electrical characteristics (also see footnotes).
- Table 3-4 on page 85 continues the list of electrical characteristics.
- Table 3-5 on page 86 lists environment requirements.

Table 3-1 Rack dimensions

Width	Depth	Height	EIA units	Weight
644 mm (25.4 in.)	1098 mm (43.3 in.)	2015 mm (79.3 in.)	42	289 kg (636 lb)

Table 3-2 Weight of building blocks

Building blocks	Weight
IBM Power System S822L (8247-22L) or IBM Power System S822 (8284-22A) compute node	29.5 kg (65 lb)
8374-01M management node	15.9 kg (35 lb)
IBM 7316-TF4 Flat Panel Rack-Mounted Monitor and Keyboard	10.4 kg (23 lb)
8831-NF2 Mellanox MSX1710-BS2F2	9.5 kg (21 lb)
IBM RackSwitch G8052 (7120-48E) 48-port management switch	5.4 kg (12 lb)
2498-F48 Brocade SAN48B-5 Fibre Channel switch	9.1 kg (20 lb)
IBM 2076-524 Storwize V7000 SFF Control Enclosure Model 524	36.7 kg (81 lb)
IBM 2076-24F Storwize V7000 Small Form Factor Expansion Enclosure Model 24F	24.9 kg (55 lb)

Table 3-3 Electrical characteristics

	IBM Power System S822L (8247-22L) or IBM Power System S822 (8284-22A) compute node	8374-01M management node	IBM 7316-TF4 flat panel rack-mounted monitor and keyboard	8831-NF2 Mellanox MSX1710-BS2F 2
Electrical characteristics	**Properties**	**Properties**	**Properties**	**Properties**
Maximum kVA	1.276	0.306	0.031	0.247
Amps	6.38 A	1.53 A	0.16 A	1.24 A
Voltage	200 - 240 V ac	200 - 240 V ac	200 - 240 V ac	200 - 240 V ac
Frequency	50 or 60 Hz	50 or 60 Hz	50 or 60 Hz	50 or 60 Hz
Maximum thermal output	4266 BTU/hr	1024 BTU/hr	75 BTU/hr	802 BTU/hr

	IBM Power System S822L (8247-22L) or IBM Power System S822 (8284-22A) compute node	8374-01M management node	IBM 7316-TF4 flat panel rack-mounted monitor and keyboard	8831-NF2 Mellanox MSX1710-BS2F 2
Maximum power consumption	1250 W	300 W	22 W	235 W

Notes:

- The maximum measured power consumption for the 8247-22L and 8284-22A that are used in PurePower Systems is less than the maximum measured power consumption for fully configured 8247-22L and 8284-22A stand-alone servers. Because the PurePower System is an integrated offering that consists of multiple servers, operating all of the servers at their maximum use simultaneously is unlikely. However, workload diversity with some applications might result in higher power consumption at times. Because there is no control over hardware or software content after the PurePower System is shipped, IBM recommends that the power consumption is monitored and reviewed periodically. The power consumption data can be used to manage the electrical environment, including, but not limited to, rack power distribution units (PDUs), universal power supplies (UPSes), and facility phase balancing. In some jurisdictions, the electrical code might dictate the need to measure power and adjust the installation so that switch gear and over-current protection margins are adhered to.
- The input voltage is based on the PDU and associated power cord. In countries that support line-to-line 200 - 208 V AC, the PDU input voltage matches the PDU output voltage. In countries that support line-to-line 380 - 415 V ac, the PDU output voltage is in the range of 220 - 240 V ac. For more information, see Power distribution unit and power cord specifications.
- All drawers mounted in the rack are rated 200 - 240 V ac. The power supplies automatically accept any voltage within the rated voltage range. With dual power supplies installed and operating, the power supplies draw approximately equal current from the utility (mains) and provide approximately equal current to the load.
- When planning the electrical system, it is important to use maximum values to account for internal or environmental conditions that result in power consumption that increases beyond typical values.

Table 3-4 Electrical characteristics (continued)

	IBM RackSwitch G8052 (7120-48E) 48-port management switch	2498-F48 Brocade SAN48B-5 Fibre Channel switch	IBM Storwize V7000 2076-524 SFF Control Enclosure Model 524	IBM Storwize V7000 2076-24F Small Form Factor Expansion Enclosure Model 24F
Electrical characteristics	**Properties**	**Properties**	**Properties**	**Properties**
Maximum kVA	0.211	0.132	0.541	0.304
Amps	1.05 A	0.66 A	2.71 A	1.52 A

	IBM RackSwitch G8052 (7120-48E) 48-port management switch	2498-F48 Brocade SAN48B-5 Fibre Channel switch	IBM Storwize V7000 2076-524 SFF Control Enclosure Model 524	IBM Storwize V7000 2076-24F Small Form Factor Expansion Enclosure Model 24F
Voltage	200 - 240 V ac	200 - 240 V ac	200 - 240 V ac	200 - 240 V ac
Frequency	50 or 60 Hz	50 or 60 Hz	50 or 60 Hz	50 or 60 Hz
Maximum thermal output	683 BTU/hr	427 BTU/hr	1846 BTU/hr	1038 BTU/hr
Maximum power consumption	200 W	125 W	541 W	304 W

Table 3-5 Environment requirements

Environment	Preferred operating temperature	Allowable operating	Nonoperating
ASHRAE class		A3	
Airflow direction		Front-to-back	
Temperature	18 - 27°C (64 - 80°F)	5 - 40°C (41 - 104°F)	5 - 45°C (41 - 113°F)
Humidity range	5.5°C (42°F) dew point (DP) to 60% relative humidity (RH) and 15°C (59°F) dew point	-12.0°C (10.4°F) DP and 8 - 85% RH	8% - 85% RH
Maximum dew point		24°C (75°F)	27°C (80°F)
Maximum operating altitude		3050 m (10000 ft)	
Shipping temperature			-40 - 60°C (-40 - 140°F)
Shipping relative humidity			5% - 100%

Note: Derate maximum allowable dry-bulb temperature 1°C/175 m above 950 m.

3.3.1 Ordering and configurations

The PurePower System can be ordered through normal channels such as IBM Configurator for e-business, an IBM Business Partner or calling IBM at 1-866-872-3902. For more information, see the IBM PurePower System web page:

http://www.ibm.com/systems/power/hardware/purepower/index.html

3.3.2 System setup

After you order an PurePower System, you receive the preconfigured, factory-racked system. You must then install the system into your site and perform post-installation tasks.

A typical installation scenario follows these steps:

1. You determine which configuration to order and then order the system.
2. IBM sends links to the planning information.
3. IBM builds and ships your system.
4. You accept the shipment.
5. You follow the unpacking instructions that came with the system.
6. You move the rack to the installation site and perform rack installation tasks.
7. You connect the rack power cords to the PDU.
8. You plug the rack into the main power source.
9. You wait until all LEDs on the front of the rack are green and are blinking slowly and that the physical operator panel on any compute nodes display 01 N PVM.
10. If you ordered IBM Systems Lab Services for startup services and basic network integration, stop here and let IBM Systems Lab Services complete the remaining steps while training your personnel.
11. You access the rack-mounted monitor and keyboard that is preinstalled in the rack.
12. You access the PurePower Integrated Manager by double-clicking the desktop shortcut using the Keyboard/Video/Mouse on the rack to check your hardware inventory.
13. You access the Nagios Core manager to ensure that your components are functioning properly.

Preparing the rack for installation

You must position the rack, lower the rack-leveling pads, and check the rack power connections before you install the system.

1. Position the rack where you want it to be installed.
2. Use the open-end wrench that comes with the hardware kit to lower each of the four leveling pads just enough so that they touch the floor. The rack casters support the weight of the rack cabinet. The pads prevent the rack from rolling.
3. Ensure that the power cords are securely attached to each power distribution unit (PDU).

Connecting power cords and verifying the attention LEDs are green

Before powering on the system, you must connect the power cords to the power source and verify that the attention light is green and flashing for each component.

To connect the power cords, follow these steps:

1. Connect the rack power cords to the power distribution unit (PDU).
2. Plug the rack into the alternating current (AC) power source.
3. Wait at least 10 minutes for the system to complete its power application process.

4. Ensure that the system attention LEDs are green and flashing for each component (see Figure 3-1).

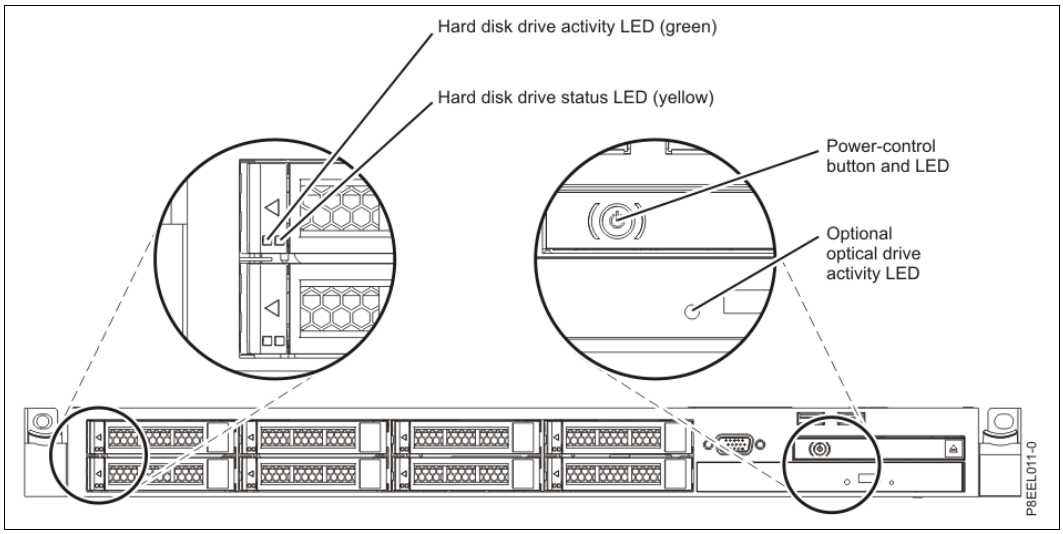

Figure 3-1 System attention LED locations

5. Ensure that the Operator panel displays 01 N PVM (see Figure 3-2).

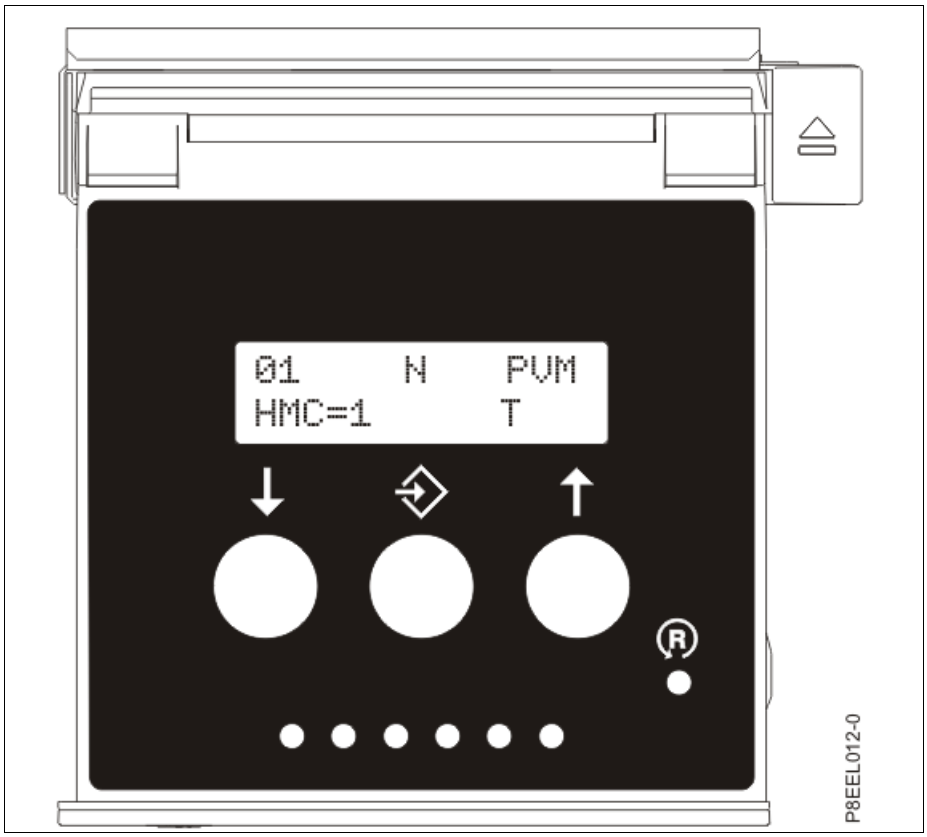

Figure 3-2 Operator panel display

Accessing the PurePower Integrated Manager user interface

After you have cabled the power cords and verified that the system LEDs are green and flashing, use the PurePower Integrated Manager user interface to ensure that your components are present and configured.

Important: If you ordered IBM Systems Lab Services startup services, do not perform these steps. IBM Systems Lab Services represenatatives perform these steps as part of the startup services, while training your personnel.

To use the PurePower Integrated Manager UI to ensure that your components are present and configured, do the following tasks:

1. Start the management nodes by pressing the buttons on the side of the node chassis.
2. Wait 5 minutes for the management nodes to start.
3. Slide out the console display and keyboard. Lift the display to access the keyboard.
4. The display powers on.
5. Press the **PRTSC** key to enable the KVM switch. The management nodes are displayed.
6. Port 1 is the PurePower Integrated Manager. Press Enter.
7. Log on to the console (Figure 3-2 on page 88). The default user ID is `admin`, and the password is `PASSWORD`.

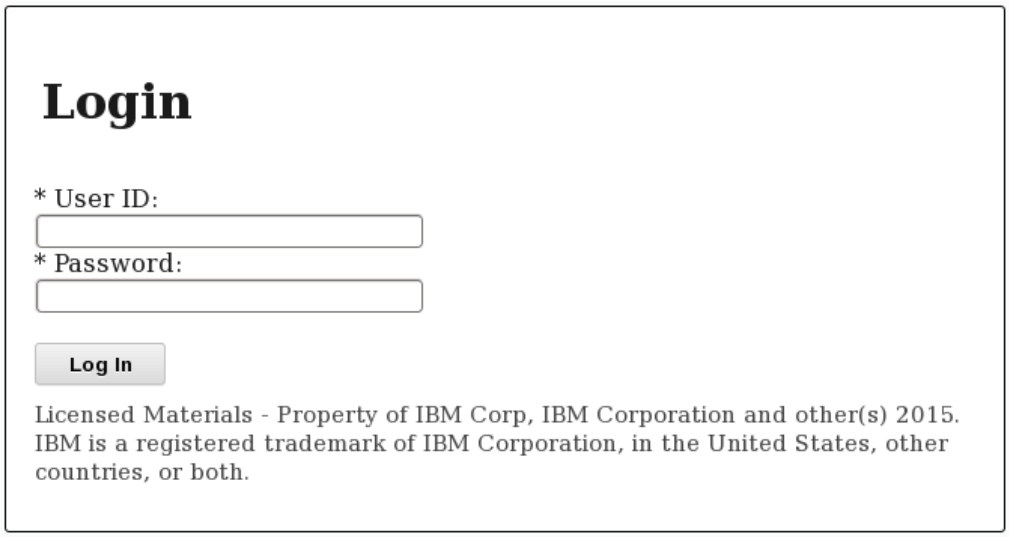

Figure 3-3 Logging in to the console

8. Double-click the **PurePower Integrated Manager** icon to launch the PurePower Integrated Manager UI.
9. Log in to the PurePower Integrated Manager UI. The default ID and passwordare `admin` and `PASSWORD`.

10. The PurePower Integrated Manager Home window opens (Figure 3-4).

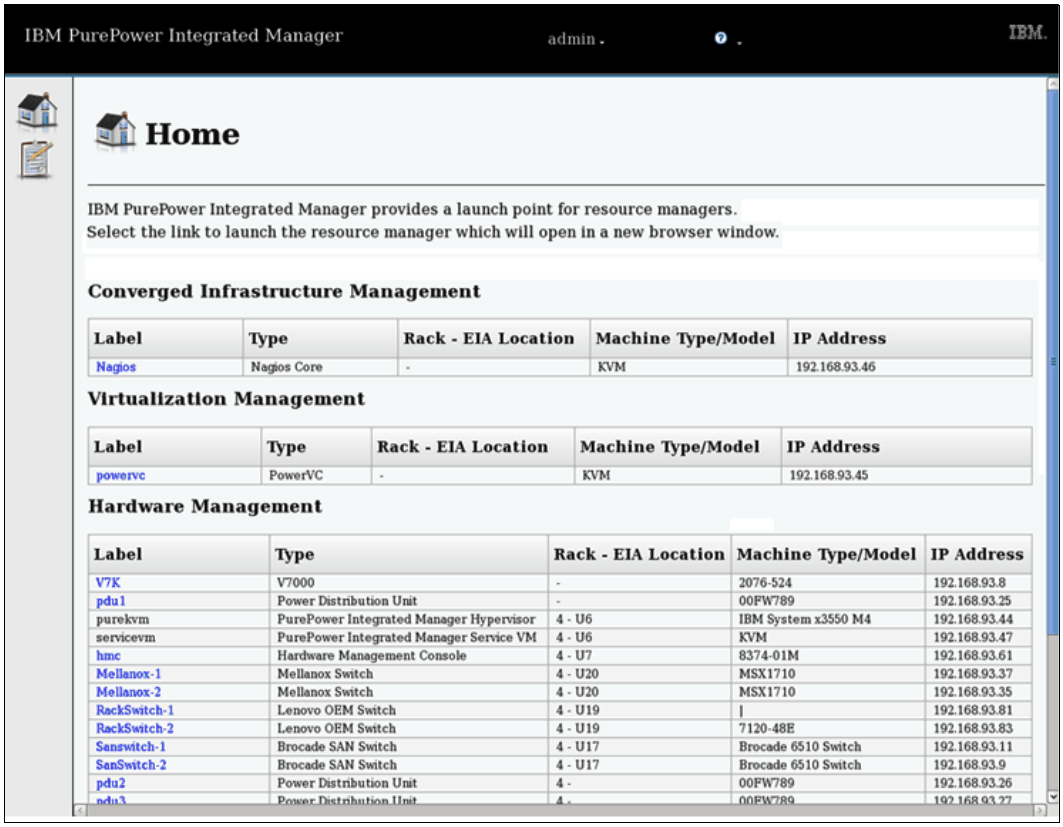

Figure 3-4 The PurePower Integrated Manager Home window

11. In the navigation area, click the **Hardware Inventory** icon. The Hardware Inventory window displays (Figure 3-5) information about hardware resources.

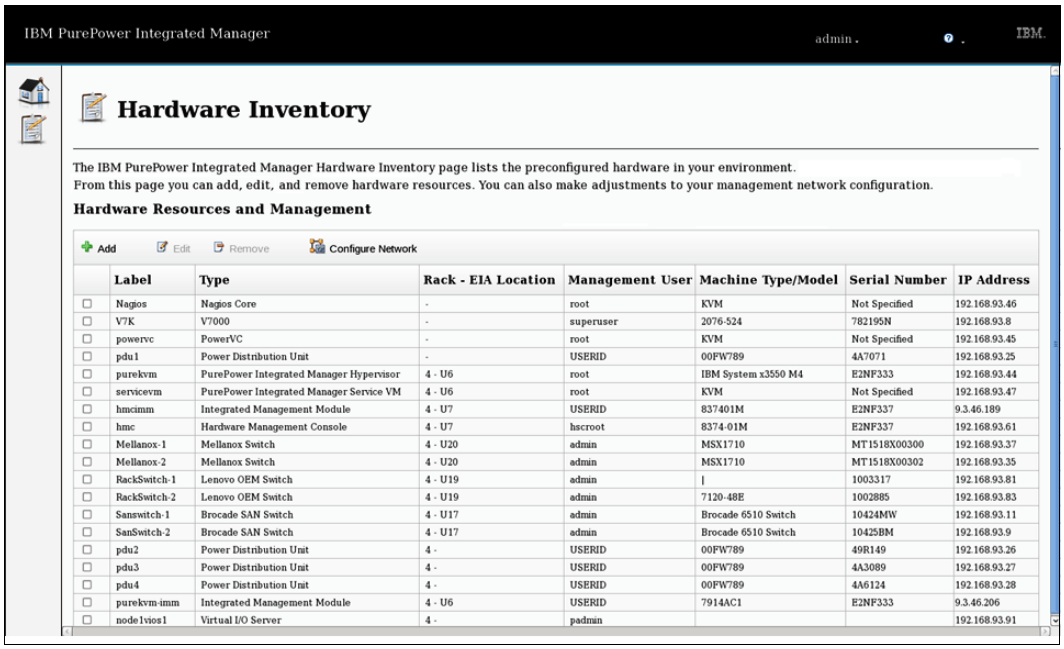

Figure 3-5 The PurePower Integrated Manager Hardware Inventory window

12. Check the Nagios interface to ensure that the components of the system are functioning properly. To access Nagios, complete these tasks:
    a. Click **Home**.
    b. From the PurePower Integrated Manager Home window, click Nagios.
    c. Log in to Nagios Core manager. The default ID and password is nagiosadmin/PASSW0RD.
    d. Ensure that each of your installed hardware components is functioning properly. To check for problems, in the navigation area, click **Host** and then click **Problems**.

Configuring the network by using the PurePower Integrated Manager

Learn how to configure the network by using the PurePower Integrated Manager.

To configure the network, complete the following steps:

1. Log in to the PurePower Integrated Manager.
2. Click the **Hardware Inventory** icon.
3. Click **Configure Network**.
4. When the Configure Network window opens, select one of the following options:
   – Make changes to the subnet mask used. Changing the subnet mask affects the addresses that can be used on the Configure Network page.
   – To apply a global network address scheme, set the subnet mask, update the network address field as needed, and select Apply address to all elements. The address of each resource is updated based on the network address.
   – If a gateway is used, click **Specify gateway on devices**, and then specify the IP address of the gateway.
5. Click **OK**. The PurePower Integrated Manager is updated to access the resources at the addresses provided.

Identifying operating system ports

To learn about the operating system ports, including the PureKVM host operating system ports, the PurePower Integrated Manager guest operating system ports, and the Service guest operating system ports, see the “Identifying operating system ports” page in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/PurePower/p8ef9/p8ef9_ports_kickoff.htm

3.3.3 Firmware and software levels

You can use the Fix Central website to get fixes for the PurePower System components.

Getting fixes for the PurePower System components

When service packs become available for the components of the PurePower System, you can use the Fix Central website to get preferred practice information and download instructions.

The April 2015 announcement of the PurePower System used firmware version 830. The October release uses firmware version 840.

To get Fix Packs for the PurePower System firmware, complete the following steps:

1. Go to the Fix Central website:

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

2. Click the **Select product** tab.
3. From the **Product Group** list, select IBM **PureSystems®**.
4. Follow the onscreen instructions to learn about preferred practices and to download and install the firmware fixes to your system.

Using the PurePower Integrated Manager to manage PurePower System

The PurePower Integrated Manager is used to manage the components of the PurePower System.

The Integrated Manager gathers and displays the current level of firmware on all components in the offering and compares those levels against the current existing tested stack. It notifies you of any levels that are out of compliance so you can schedule manual updates.

The PurePower Integrated Manager includes the following preintegrated resource managers:

- Virtualization management by using PowerVC
- Hardware management by using the PurePower Integrated Manager
- Converged infrastructure monitoring by using Nagios Core Open Source software

The PurePower Integrated Manager also provides the ability for you to complete system management tasks. You can track hardware inventory, monitor your systems, and manage your network.

The Integrated Manager offers the following capabilities:

- A single user interface to manage multiple compute, network, and storage resources
- Cloud and virtualization integration by using the PowerVC

Starting the PowerVC manager from the PurePower Integrated Manager

You can start the PowerVC manager from the PurePower Integrated Manager.

To start he PowerVC manager, complete the following steps:

1. Log in to the PurePower Integrated Manager.
2. Click the **Home** icon.
3. In the Virtualization Management area, click the label of the PowerVC resource that you want. The PowerVC Login pane appears.
4. Type your user name and password and click **OK**. The PowerVC window opens.

Using the Nagios Core manager from the PurePower Integrated Manager

After you start the Nagios Core manager from within the PurePower Integrated Manager, you can monitor your infrastructure.

To start and use the Nagios Core manager, complete the following steps:

1. Log in to the PurePower Integrated Manager, and click the **Home** icon.
2. To log in to the Nagios console, click the label of the resource that is listed in the Converged Infrastructure Management area.
3. When the Authentication Required window opens, type your user name and password and click **OK**.
4. When the Nagios Core window opens, in the **Current Status** menu, select one of the following options:
   a. Select **Host** to monitor the overall status of the devices in the rack.
   b. Select **Services** to monitor in greater detail the status of the devices in the rack.

A

IBM PurePower solutions

IBM Power Systems offer an open innovation platform for organizations of all sizes. They provide solutions to transform big data and analytics, cloud and mobile into competitive advantage, in ways never before possible.

This appendix describes the following IBM Power Systems solutions:

- Big data and analytics solutions
- Cloud computing solutions
- Linux on Power solutions
- Mobile solutions
- Industry solutions

Big data and analytics solutions

Unlock the value of data and deliver insights faster with IBM systems and solutions optimized for big data and analytics applications.

The waitless world is one in which the explosive growth of both structured and unstructured data from multiple sources requires businesses to derive insights faster than ever to keep pace. Explore how systems designed for big data provide the most scalable analytics platforms in the industry.

IBM Power Systems big data and analytics

http://www.ibm.com/systems/power/solutions/bigdata-analytics/index.html

Cloud computing solutions

Designed to help deliver on the promise of private and hybrid cloud and take advantage of superior cloud economics.

How do you choose the correct IT infrastructure for the waitless world? Outcomes matter, and the landscape is changing rapidly. To stay relevant and be successful, your business needs a systematic and integrated approach to business innovation. IBM Power Systems with POWER8 offer an answer. They use open technologies and deliver superior cloud economics with 12 times the work of x86 commodity servers, and they offer a virtualization hypervisor with zero documented vulnerabilities. A secure, open, and waitless choice.

IBM Cloud on Power Systems

http://www.ibm.com/systems/power/solutions/cloud/index.html

Linux on Power solutions

Gain performance, portability and scalability through the innovation of an open ecosystem.

Supported on every Power Systems server that we make, Linux on Power Systems is the only Linux infrastructure that offers both scale-out and scale-up choices that align to your needs and meet the pressing demands of the waitless world. Optimize your infrastructure for emerging business challenges with industry standard Red Hat, SUSE, and Ubuntu Linux distributions.

Linux on IBM Power Systems

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

Mobile solutions

Extend your enterprise to mobile devices with a range of solutions on a flexible, secure platform.

Only a small part of the effort in supporting enterprise mobility is visible on smart devices. The underlying applications and architecture dealing with data synchronization, app management, security, integration, analytics, and other enterprise mobile challenges deliver the majority of the value. That's why your infrastructure and strategy matter.

Mobile solutions on IBM Power Systems

http://www.ibm.com/systems/power/solutions/mobile/index.html

Industry solutions

Gain greater actionable insights for competitive advantage through industry solutions with leading analytics capabilities.

Healthcare and life sciences solutions

http://www.ibm.com/systems/power/solutions/industry/healthcare.html

Retail solutions

http://www.ibm.com/systems/power/solutions/industry/retail.html

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

You can search for, view, download, or order these documents and other IBM Redbooks, Redpapers, Web Docs, drafts, and additional materials on the Redbooks website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM Fix Central

  http://www.ibm.com/support/fixcentral/
- IBM Knowledge Center

  http://www.ibm.com/support/knowledgecenter/
- IBM Power Systems

  http://www.ibm.com/systems/power/
- IBM POWER8 systems information, IBM Knowledge Center

  http://www.ibm.com/support/knowledgecenter/POWER8/p8hdx/POWER8welcome.htm
- IBM System Planning Tool

  http://www.ibm.com/systems/support/tools/systemplanningtool/
- IBM Power Systems support and services

  http://www.ibm.com/systems/power/support/
- IBM PurePower System overview

  http://www.ibm.com/systems/power/hardware/purepower/index.html
- *Cloud Security Guidelines for IBM Power Systems*, SG24-8242

  http://www.redbooks.ibm.com/abstracts/sg248242.htm
- Nagios

  http://www.nagios.org

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5277-00

ISBN 0738454729

Printed in U.S.A.