


IBM Power Systems S922, S914, and S924 Technical Overview and Introduction

Young Hoon Cho

Gareth Coates

Bartlomiej Grabowski

Volker Haug



Power Systems



International Technical Support Organization

**IBM Power Systems S922, S914, and S924: Technical
Overview and Introduction**

July 2018

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

First Edition (July 2018)

This edition applies to IBM Power Systems S922, S914, and S924, machine types and model numbers 9009-22A, 9009-41A, and 9009-42A.

© Copyright International Business Machines Corporation 2018. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	x
Comments welcome	xi
Stay connected to IBM Redbooks	xi
Chapter 1. General description	1
1.1 Systems overview	2
1.1.1 Power S922 server	2
1.1.2 Power S914 server	3
1.1.3 Power S924 server	4
1.1.4 Common features	5
1.2 Operating environment	6
1.3 Physical package	8
1.3.1 Tower model	8
1.3.2 Rack-mount model	8
1.4 System features	9
1.4.1 Power S922 server features	9
1.4.2 Power S914 server features	10
1.4.3 Power S924 server features	11
1.4.4 Minimum features	13
1.4.5 Power supply features	13
1.4.6 Processor module features	14
1.4.7 Memory features	17
1.4.8 Peripheral Component Interconnect Express slots	18
1.5 Disk and media features	18
1.6 I/O drawers for Power S922, Power S914, and Power S924 servers	31
1.6.1 PCIe Gen3 I/O expansion drawer	31
1.6.2 I/O drawers and usable PCI slots	33
1.6.3 EXP24SX (#ESLS) and EXP12SX (#ESLL) SAS Storage Enclosures	35
1.7 IBM i Solution Editions for Power S914	36
1.8 IBM Capacity BackUp	37
1.8.1 Power S922 (9009-22A) IBM Capacity BackUp offering	37
1.8.2 Power S914 (9009-41A) IBM Capacity BackUp offering	38
1.8.3 Power S924 (9009-42A) IBM Capacity BackUp offering	40
1.9 System racks	41
1.9.1 IBM 7014 Model T00 rack	42
1.9.2 IBM 7014 Model T42 rack	43
1.9.3 IBM 42U Slim Rack 7965-94Y	44
1.9.4 1.8 Meter Rack (#0551)	44
1.9.5 2.0 Meter Rack (#0553)	45
1.9.6 Rack (#ER05)	45
1.9.7 The AC power distribution unit and rack content	45
1.9.8 Rack-mounting rules	49
1.9.9 Useful rack additions	49
1.9.10 Original equipment manufacturer racks	52

1.10 Hardware Management Console	53
1.10.1 New features	53
1.10.2 Hardware Management Console overview	53
1.10.3 Hardware Management Console code level	55
1.10.4 Two architectures of Hardware Management Console	55
1.10.5 Hardware Management Console connectivity to POWER9 processor-based systems	56
1.10.6 High availability Hardware Management Console configuration	57
Chapter 2. Architecture and technical overview	59
2.1 The IBM POWER9 processor	63
2.1.1 POWER9 processor overview	63
2.1.2 POWER9 processor features	64
2.1.3 POWER9 processor core	64
2.1.4 Simultaneous multithreading	65
2.1.5 Processor feature codes	66
2.1.6 Memory access	67
2.1.7 On-chip L3 cache innovation and intelligent caching	68
2.1.8 Hardware transactional memory	69
2.1.9 Coherent Accelerator Processor Interface 2.0	69
2.1.10 Power management and system performance	71
2.1.11 Comparison of the POWER9, POWER8, and POWER7+ processors	71
2.2 Memory subsystem	72
2.2.1 Memory placement rules	73
2.2.2 Memory bandwidth	77
2.3 System bus	80
2.4 Internal I/O subsystem	81
2.4.1 Slot configuration	81
2.4.2 System ports	84
2.5 Peripheral Component Interconnect adapters	85
2.5.1 Peripheral Component Interconnect Express	85
2.5.2 LAN adapters	86
2.5.3 Graphics accelerator adapters	89
2.5.4 SAS adapters	90
2.5.5 Fibre Channel adapters	91
2.5.6 InfiniBand host channel adapter	93
2.5.7 Cryptographic coprocessor	94
2.5.8 Coherent Accelerator Processor Interface adapters	95
2.5.9 USB adapters	96
2.6 Internal storage	97
2.6.1 Backplane (#EJ1C)	97
2.6.2 Split backplane option (#EJ1E)	98
2.6.3 Backplane (#EJ1D)	99
2.6.4 Expanded Function Backplane (#EJ1M)	100
2.6.5 NVMe support	101
2.6.6 Backplane (#EJ1F)	104
2.6.7 Expanded Function Storage Backplane (#EJ1G)	105
2.6.8 RAID support	105
2.6.9 Easy Tier	106
2.6.10 External SAS ports	109
2.6.11 Media drawers	110
2.6.12 External DVD drives	110
2.6.13 RDX removable disk drives	110

2.7 External IO subsystems	111
2.7.1 Peripheral Component Interconnect Express Gen3 I/O expansion drawer	111
2.7.2 PCIe Gen3 I/O expansion drawer optical cabling	112
2.7.3 PCIe Gen3 I/O expansion drawer system power control network cabling	114
2.8 External disk subsystems	114
2.8.1 EXP24SX SAS Storage Enclosure and EXP12SX SAS Storage Enclosure.	114
2.8.2 IBM Storage	116
2.9 Operating system support	117
2.9.1 AIX operating system	118
2.9.2 IBM i	119
2.9.3 Linux operating system	119
2.9.4 Virtual I/O Server	120
2.10 POWER9 reliability, availability, and serviceability capabilities	120
2.11 Manageability	122
2.11.1 Service user interfaces	122
2.11.2 IBM Power Systems Firmware maintenance	129
2.11.3 Concurrent firmware maintenance improvements	132
2.11.4 Electronic Services and Electronic Service Agent	133
Chapter 3. Virtualization	137
3.1 POWER9 Hypervisor	138
3.1.1 Virtual SCSI	140
3.1.2 Virtual Ethernet	140
3.1.3 Virtual Fibre Channel	140
3.1.4 Virtual (TTY) console	140
3.2 POWER processor modes	141
3.3 Active Memory Expansion.	142
3.4 Single Root I/O Virtualization	143
3.5 PowerVM.	143
3.5.1 Multiple shared processor pools	145
3.5.2 Virtual I/O Server	146
3.5.3 Live Partition Mobility	147
3.5.4 Active Memory Sharing.	147
3.5.5 Active Memory Deduplication	148
3.5.6 Remote Restart.	148
Related publications	149
IBM Redbooks	149
Online resources	149
Help from IBM	150

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.


Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
C3®
Easy Tier®
IBM®
IBM Spectrum®
Interconnect®
Micro-Partitioning®
POWER®

Power Architecture®
POWER6™
POWER7®
POWER8®
POWER9™
PowerHA®
PowerPC®
PowerVM®

Redbooks®
Redbooks (logo) ®
RS/6000™
Storwize®
System Storage™
XIV®

The following terms are trademarks of other companies:

Intel, Intel Xeon, 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.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

LTO, Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Fedora, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Preface

This IBM® Redpaper publication is a comprehensive guide that covers the IBM Power System S922 (9009-22A), IBM Power System S914 (9009-41A), and IBM Power System S924 (9009-42A) servers that support IBM AIX®, IBM i, and Linux operating systems. The objective of this paper is to introduce the major innovative Power S914, Power S922, and Power 924 offerings and their relevant functions:

- ▶ The new IBM POWER9™ processor, which is available at frequencies of 2.3 - 3.8 GHz, 2.8 - 3.8 GHz, 2.9 - 3.8 GHz, 3.4 - 3.9 GHz, 3.45 - 3.9 GHz, 3.5 - 3.9 GHz, and 3.8 - 4.0 GHz.
- ▶ Significantly strengthened cores and larger caches.
- ▶ Two integrated memory controllers that double the memory footprint of IBM POWER8® servers.
- ▶ Integrated I/O subsystem and hot-pluggable Peripheral Component Interconnect® Express (PCIe) Gen4 and Gen3 I/O slots.
- ▶ I/O drawer expansion options offer greater flexibility.
- ▶ Support for Coherent Accelerator Processor Interface (CAPI) 2.0.
- ▶ New IBM EnergyScale technology offers new variable processor frequency modes that provide a significant performance boost beyond the static nominal frequency.

This publication is for professionals who want to acquire a better understanding of IBM Power Systems products. The intended audience includes the following roles:

- ▶ Clients
- ▶ Sales and marketing professionals
- ▶ Technical support professionals
- ▶ IBM Business Partners
- ▶ Independent software vendors (ISVs)

This paper expands the current set of IBM Power Systems documentation by providing a desktop reference that offers a detailed technical description of the Power S914, Power S922, and Power S924 systems.

This paper does not replace the current marketing materials and configuration tools. It is intended as an extra source of information that, together with existing sources, can be used to enhance your knowledge of IBM server solutions.

Authors

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

Young Hoon Cho is a Power Systems Top Gun with the post-sales Technical Support Team for IBM in Korea. He has over 10 years of experience working on IBM RS/6000™, IBM System p, and Power Systems products. He provides second-line technical support to Field Engineers working on Power Systems and system management.

Gareth Coates has been working with AIX since 1988 and with Linux since 1991. He has trained IBMers and customers on POWER4 through POWER9 processor-based systems and AIX and Linux. He was appointed to the Worldwide Board of the IBM Learning Profession, and also spent time in the Server Services Group. In 2009, he joined the IBM EMEA Advanced Technology Services team, where he specializes in new product introduction of IBM POWER® based hardware, IBM PowerVM®, and the Hardware Management Console (HMC). He presents at technical conferences and has co-authored numerous IBM Certification Tests.

Bartłomiej Grabowski is an IBM Champion and a Principal Systems Specialist in DHL IT Services. He has over 14 years of experience in enterprise solutions. He holds a bachelor degree in computer science from the Academy of Computer Science and Management in Bielsko-Biala. His areas of expertise include IBM Power Systems, IBM i, PowerHA®, PowerVM, and storage solutions. Bartłomiej is also a developer of IBM certification exams. He is a Gold Redbooks® author.

Volker Haug is an Executive IT Specialist & Open Group Distinguished IT Specialist in Ehningen, Germany. He holds a Diploma degree in Business Management from the University of Applied Studies in Stuttgart. His career includes more than 32 years of experience with Power Systems, AIX, and PowerVM virtualization. He has written several IBM Redbooks publications about Power Systems and PowerVM and is a Platinum Redbooks author. Volker is IBM POWER9 Champion and a member of the German Technical Expert Council, which is an affiliate of the IBM Academy of Technology.

The project that produced this publication was managed by:

Scott Vetter, PMP
IBM Austin

Thanks to the following individuals for their contribution and support of this publication:

Ron Arroyo, Tonny JH Bastiaans, Matthew Butterbaugh, Jason Clegg, Nigel Griffiths, Daniel Henderson, Jeanine Hinck, Rex Kuo, Ray Laning, Chris Mann, Benjamin Mashak, Stephen Mroz, Thoi Nguyen, Kanisha Patel, William Starke, Jeff Stuecheli, Justin Thaler, Brian W Thompto, Julie Villarreal
IBM

John Banchy
SIRIUS Computer Solutions

Neil Palmer
DPS Inc.

Now you can become a published author, too!

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

Find out more about the residency program, browse the residency index, and apply online at:
ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

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

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

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

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

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



General description

The next generation of Power Systems servers with POWER9 technology is built with innovations that can help deliver security and reliability for the data-intensive workloads of today's enterprises. POWER9 technology is designed from the ground up for data-intensive workloads, such as databases and analytics. Changes in the memory subsystem and the use of industry-standard memory DIMMs take POWER9 technology to the next level by superseding a number of existing price/performance offerings. Designed to run commercial, cognitive, and database workloads, POWER9 technology provides a highly competitive server platform. Client references indicate POWER servers help provide a robust and secure backbone for their IT infrastructure. More companies are using POWER technology in their IT infrastructure, from the shop level to large data center deployments.

The Power S922 server supports two processor sockets, offering 10-core or 20-core typical 2.9 - 3.8 GHz (maximum), 8-core or 16-core typical 3.4 - 3.9 GHz (maximum), or 4-core typical 2.8 - 3.8 GHz (maximum) POWER9 configurations in a 19-inch rack-mount, 2U (EIA units) drawer configuration. All the cores are active. The Power S922 server supports a maximum of 32 DDR4 DIMM slots. Memory features that are supported are 8 GB, 16 GB, 32 GB, 64 GB, and 128 GB, and run at different speeds of 2133, 2400, and 2666 Mbps, offering a maximum system memory of 4096 GB.

The Power S914 server supports one-processor sockets, offering 4-core typical 2.3 - 3.8 GHz (maximum), 6-core typical 2.3 - 3.8 GHz (maximum), or 8-core typical 2.8 - 3.8 GHz (maximum) POWER9 processor-based configurations in a 19-inch rack-mount, 4U (EIA units) drawer or desk-side configuration. All the cores are active. The Power S914 server supports a maximum of 16 DDR4 DIMM slots. Memory features that are supported are 8 GB, 16 GB, 32 GB, and 64 GB, and run at speeds of 2133, 2400, and 2666 Mbps, offering a maximum system memory of 1024 GB.

The Power S924 server supports two processor sockets, offering 8-core or 16-core typical 3.8 - 4.0 GHz (maximum), 10-core or 20-core typical 3.5 - 3.9 GHz (maximum), 11-core or 22-core typical 3.45 - 3.9 GHz, or 24-core typical 3.4 - 3.9 GHz (maximum) configurations in a 19-inch rack-mount, 4U (EIA units) drawer configuration. All the cores are active. The Power S924 server supports a maximum of 32 DDR4 DIMM slots. Memory features that are supported are 8 GB, 16 GB, 32 GB, 64 GB, and 128 GB, and run at different speeds of 2133, 2400, and 2666 Mbps, offering a maximum system memory of 4096 GB.

1.1 Systems overview

The following sections provide detailed information about the Power S922, Power S914, and Power S924 servers.

1.1.1 Power S922 server

The Power S922 (9009-22A) server is a powerful one- or two-socket server that ships with up to 20 activated cores. If only one socket is populated at the time of the order, the second socket can be populated later. It has the I/O configuration flexibility to meet today's growth and tomorrow's processing needs. This server supports two processor sockets, offering 4-core, 8-core, or 10-core processors running 2.8 - 3.9 GHz in a 19-inch rack-mount, 2U (EIA units) drawer configuration. All the cores are active.

The Power S922 server supports a maximum of 32 DDR4 Registered DIMM (RDIMM) slots. If only one processor socket is populated, then only 16 RDIMMs can be used. The memory features that are supported are 16 GB, 32 GB, 64 GB, and 128 GB, and run at speeds of 1600, 2133, and 2666 MHz, offering a maximum system memory of 2 TB if one socket is single-chip populated and 4 TB with both sockets populated.

The IBM Active Memory Expansion feature enables memory expansion by using compression and decompression of memory content, which can effectively expand the maximum memory capacity if extra server workload capacity and performance are available.

Two features are available for the storage backplane:

- ▶ #EJ1F: Eight SFF-3 bays with an optional split card (#EJ1H)
- ▶ #EC59: Optional PCIe3 Non-Volatile Memory express (NVMe) carrier card with two M.2 module slots

Each of these backplane options uses leading-edge, integrated SAS RAID controller technology that is designed and patented by IBM.

The NVMe option offers fast boot times and is ideally suited for the rootvg of Virtual I/O Server (VIOS) partitions.

The Power S922 server is shown in Figure 1-1.




Figure 1-1 Front view of the Power S922 server

1.1.2 Power S914 server

The Power S914 (9009-41A) server is a powerful one-socket server that ships with up to eight activated cores. It has the I/O configuration flexibility to meet today's growth and tomorrow's processing needs. A one-socket system with a 4-core or 6-core POWER9 processor is available in either rack (19-inch rack-mount 4U (EIA units)) or tower configurations. The 8-core higher performance system is available only in a rack configuration.

The Power S914 server supports a maximum of 16 DDR4 error-correcting code (ECC) RDIMM slots. The memory features that are supported are 16 GB, 32 GB, 64 GB, and 128 GB, and run at speeds of 2133 - 2666 MHz, offering a maximum system memory of 1 TB.

If you use the 4-core processor #EP10, the system is limited to four memory DDR4 ECC RDIMMS supporting only the 8_GB RDIMMS (#EM60), 16 GB RDIMMs (#EM62), or the 32 GB RDIMMS (#EM63), which offers a maximum of 64 GB of RAM per system.

The IBM Active Memory Expansion feature enables memory expansion by using compression and decompression of memory content, which can effectively expand the maximum memory capacity if extra server workload capacity and performance are available.

Several different features are available for the storage backplane:

- ▶ #EJ1C: Twelve SFF-3 bays with an optional split card (#EJ1E)
- ▶ #EJ1D: Eighteen SFF-3 bays/Dual IOA with Write Cache
- ▶ #EJ1M: Twelve SFF-3 bays/RDX bays:
 - Split feature to 6+6 small form factor (SFF) bays: Add a second SAS Controller.
 - Twelve SFF-3 bays/RDX bay/2 EXT PT.
- ▶ #EC59: Optional PCIe3 NVMe carrier card with two M.2 module slots

Each of these backplane options uses leading-edge, integrated SAS RAID controller technology that is designed and patented by IBM.

The NVMe option offers fast boot times and is ideally suited for the rootvg of VIOS partitions.

For more information about the NVMe technology, see "PCIe3 NVMe carrier card with two M.2 module slots (#EC59)" on page 102.

The Power S914 tower server is shown in Figure 1-2.



Figure 1-2 The Power S914 tower

1.1.3 Power S924 server

The Power S924 (9009-42A) server is a powerful one- or two-socket server that ships with up to 24 activated cores. It has the I/O configuration flexibility to meet today's growth and tomorrow's processing needs. This server supports two processor sockets, offering 8 or 16 cores at 3.8 - 4.0 GHz, 10 or 20 cores at 3.5 - 3.9 GHz, 11 or 22 cores at 3.45 - 3.9 GHz or 24 cores at 3.4 - 3.9 GHz. The systems are 19-inch rack-mount, 4U (EIA units) drawer configuration. All the cores are active.

The Power S924 server supports a maximum of 32 DDR4 RDIMM slots per processor socket that is populated. The memory features that are supported are 16 GB, 32 GB, 64 GB, and 128 GB, offering a maximum system memory of 4 TB.

The Active Memory Expansion feature enables memory expansion by using compression and decompression of memory content, which can effectively expand the maximum memory capacity if more server workload capacity and performance are available.

Several different features are available for the storage backplane:

- ▶ #EJ1C: Twelve SFF-3 bays with an optional split card (#EJ1E)
- ▶ #EJ1D: Eighteen SFF-3 bays/Dual IOA with Write Cache
- ▶ #EJ1M: Twelve SFF-3 bays/RDX bays:
 - Split feature to 6+6 SFF bays: Add a second SAS Controller.
 - Twelve SFF-3 bays/RDX bay/2 EXT PT.
- ▶ #EC59: Optional PCIe3 NVMe carrier card with two M.2 module slots

Each of the three backplane options uses leading-edge, integrated SAS RAID controller technology that is designed and patented by IBM.

The Power S924 server is shown in Figure 1-3.



Figure 1-3 Front view of the Power S924 server

1.1.4 Common features

Many features are common to all of the servers being considered in this book, and some of them are described in this section.

There is no internal DVD

There is no internal DVD option, although an external USB DVD drive is available as #EUA5. Customers are encouraged to use USB flash drives to install operating systems and VIOS whenever possible because they are much faster than DVD.

The operator panel

The operator panel is now composed of two parts. All of the servers have the first part, which provides the power switch and LEDs, as shown in Figure 1-4.



Figure 1-4 Operator panel: Power switch and LEDs

The second part is an LCD panel with three buttons, as shown in Figure 1-5.




Figure 1-5 Operator panel: LCD and switches

The LCD panel is mandatory in the Power S914 Tower. It is also required if the server runs IBM i. In the Power S914, Power S922, and Power S924 rack-mounted servers, it is optional, but if a rack contains any of these servers, one of the them must have an LCD panel.

The LCD panel can be moved (by using the correct procedure) from one server to another server, which enables appropriate service to be carried out.

1.2 Operating environment

Table 1-1 lists the electrical characteristics for the servers.

Table 1-1 Electrical characteristics for Power S922, Power S914, and Power S924 servers

Electrical characteristics	Properties		
	Power S922 server	Power S914 server	Power S924 server
Operating voltage	1400 W power supply: 200 - 240 V AC	900 W power supply: 100 - 127 V AC or 200 - 240 V AC 1400 W power supply: 200 - 240 V AC	1400 W power supply: 200 - 240 V AC
Operating frequency	47/63 Hz	47/63 Hz	47/63 Hz
Thermal output	6,416 Btu/hour (maximum)	5,461 Btu/hour (maximum)	9,386 Btu/hour (maximum)
Power consumption	1880 watts (maximum)	1600 watts (maximum)	2750 watts (maximum)
Power-source loading	1.94 kVa (maximum configuration)	1.65 kVa (maximum configuration)	2.835 kVa (maximum configuration)
Phase	Single	Single	Single

Note: The maximum measured value is the worst-case power consumption that is expected from a fully populated server under an intensive workload. The maximum measured value also accounts for component tolerance and non-ideal operating conditions. Power consumption and heat load vary greatly by server configuration and utilization. The [IBM Systems Energy Estimator](#) should be used to obtain a heat output estimate that is based on a specific configuration.

Table 1-2 lists the environment requirements for the servers.

Table 1-2 Environment requirements for Power S922, Power S914, and Power S924 servers

Environment	Recommended operating	Allowable operating	Non-operating
Temperature	18 - 27°C (64.4 - 80.6°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) or 15°C (59°F) dew point	8% - 85% RH	8% - 80% RH
Maximum dew point	N/A	24°C (75°F)	27°C (80°F)
Maximum operating altitude	N/A	3050 m (10000 ft)	N/A

Note: IBM does not recommend operation above 27C, however, one can expect full performance up to 35C for these systems. Above 35C, the system is capable of operating, but possible reductions in performance may occur to preserve the integrity of the system components. Above 40C there may be reliability concerns for components within the system.

Table 1-3 lists the noise emissions for the systems

Table 1-3 Noise emissions for Power S922, Power S914, and Power S924 servers

Product	Declared A-weighted sound power level, L_{WAd} (B)		Declared A-weighted sound pressure level, L_{pAm} (dB)	
	Operating	Idle	Operating	Idle
Power S922 server	7.8	6.9	61	53
Power S914 (tower) server	5.8	5.3	39	34
Power S914 (rack) server	5.9	5.3	41	34
Power S924 server	6.4	5.2	46	34

Tip:

- ▶ Declared level L_{WAd} is the upper-limit A-weighted sound power level. Declared level L_{pAm} is the mean A-weighted emission sound pressure level that is measured at the 1-meter bystander positions.
- ▶ All measurements are made in conformance with ISO 7779 and declared in conformance with ISO 9296.
- ▶ 10 dB (decibel) equals 1 B (bel).

1.3 Physical package

The Power S914 server is available in both rack-mount and tower form factors. The Power S922 and Power S924 servers are available in rack-mount form factor only.

1.3.1 Tower model

Table 1-4 shows the physical dimensions of the Power S914 tower chassis.

Table 1-4 Physical dimensions of the Power S914 tower chassis

Dimension	Power S914 server (9009-41A)
Width	182.4 mm (7.18 in.)
Width with stand	328.5 mm (12.93 in.)
Depth	751.7 mm (29.59 in.)
Depth with front-rotatable door	814.7 mm (32.07 in.)
Height	486.1 mm (19.14 in.)
Height with handle	522 mm (20.55 in.)
Weight	46.94 kg (103.5 lb)

1.3.2 Rack-mount model

Table 1-5 shows the physical dimensions of the Power S922 rack-mounted chassis. The server is available only in a rack-mounted form factor and takes 2U (2 EIA units) of rack space.

Table 1-5 Physical dimensions of the Power S922 rack-mounted chassis

Dimension	Power S922 server (9009-22A)
Width	482 mm (18.97 in.)
Depth	766.5 mm (30.2 in.)
Height	86.7 mm (3.4 in.)
Weight	30.4 kg (67 lb)

Figure 1-6 show the front view of the Power S922 server.



Figure 1-6 Front view of the Power S922 server

Table 1-6 shows the physical dimensions of the rack-mounted Power S914 and Power S924 chassis. The server is available only in a rack-mounted form factor and takes 4U (4 EIA units) of rack space.

Table 1-6 Physical dimensions of the rack-mounted Power S914 and Power S924 chassis

Dimension	Power S914 server (9009-41A)	Power S924 server (9009-42A)
Width	482 mm (18.97 in.)	482 mm (18.97 in.)
Depth	769.6 mm (30.3 in.)	769.6 mm (30.3 in.)
Height	173.3 mm (6.8 in.)	173.3 mm (6.8 in.)
Weight	36.3 kg (80 lb)	39.9 kg (88 lb)

Figure 1-7 shows the front view of the Power S924 server.



Figure 1-7 Front view of the Power S924 server

1.4 System features

The system chassis contains one processor module (Power S914 server) or up to two processor modules (Power S922 and Power S924 servers). Each of the POWER9 processor chips in the server has a 64-bit architecture, up to 512 KB of L2 cache per core, and up to 10 MB of L3 cache per core. All the cores are active.

1.4.1 Power S922 server features

This summary describes the standard features of the Power S922 server:

- ▶ POWER9 processor modules:
 - 4-core typical 2.8 - 3.8 GHz (maximum) POWER9 processor.
 - 8-core typical 3.4 - 3.9 GHz (maximum) POWER9 processor.
 - 10-core typical 2.9 - 3.8 GHz (maximum) POWER9 processor.
- ▶ High-performance Mbps DDR4 ECC memory:
 - 8 GB, 16 GB, 32 GB, 64 GB, or 128 GB memory. Different sizes/configurations run at different frequencies of 2133, 2400, and 2666 Mbps.
 - Up to 4 TB of DDR4 memory with two POWER processors.
 - Up to 2 TB of DDR4 memory with one POWER processor.

- ▶ Storage feature: Eight SFF bays, one integrated SAS controller without cache, and JBOD, RAID 0, RAID 5, RAID 6, or RAID 10:
 - Optionally, split the SFF-3 bays and add a second integrated SAS controller without cache.
 - Expanded Function Storage Backplane 8 SFF-3 bays/Single IOA with Write Cache. Optionally, attach an EXP12SX/EXP24SX SAS HDD/solid-state drive (SSD) Expansion Drawer to the single IOA.
- ▶ Up to two PCIe3 NVMe carrier cards with two M.2 module slots (with up to four Mainstream 400 GB SSD NVMe M.2 modules). One PCIe3 NVMe carrier card can be ordered only with a storage backplane. If a PCIe3 NVMe carrier card is ordered with a storage backplane, then the optional split feature is not supported.
- ▶ Peripheral Component Interconnect Express (PCIe) slots with single processor:
 - One x16 Gen4 low-profile (LP), half-length slot (Coherent Accelerator Processor Interface (CAPI)).
 - One x8 Gen4 LP, half-length slot (with x16 connector) (CAPI).
 - Two x8 Gen3 LP, half-length slots (with x16 connectors).
 - Two x8 Gen3 LP, half-length slots. (One of these slots is used for the required base LAN adapter.)
- ▶ PCIe slots with two processors:
 - Three x16 Gen4 LP, half-length slots (CAPI).
 - Two x8 Gen4 LP, half-length slots (with x16 connectors) (CAPI).
 - Two x8 Gen3 LP, half-length slots (with x16 connectors).
 - Two x8 Gen3 LP, half-length slots. (One of these slots is used for the required base LAN adapter.)
- ▶ Integrated:
 - Service processor.
 - EnergyScale technology.
 - Hot-plug and redundant cooling.
 - Two front USB 3.0 ports.
 - Two rear USB 3.0 ports.
 - Two Hardware Management Console (HMC) 1 GbE RJ45 ports.
 - One system port with RJ45 connector.
 - Two hot-plug, redundant power supplies.
 - 19-inch rack-mounting hardware (2U).

1.4.2 Power S914 server features

This summary describes the standard features of the Power S914 server:

- ▶ POWER9 processor modules:
 - 4-core, typical 2.3 - 3.8 GHz (maximum) POWER9 processor.
 - 6-core, typical 2.3 - 3.8 GHz (maximum) POWER9 processor.
 - 8-core, typical 2.8 - 3.8 GHz (maximum) POWER9 processor (rack-mounted configuration only).

- ▶ High-performance Mbps DDR4 ECC memory:
 - 8 GB, 16 GB, 32 GB, or 64 GB memory.
 - Up to 1024 GB of DDR4 memory with one POWER processor.
- ▶ Storage feature:
 - Twelve SFF-3 bays/RDX bays. Optionally, split the SFF-3 bays and add a second integrated SAS controller without cache.
 - Eighteen SFF-3 bays/Dual IOA with Write Cache and External SAS port.
 - Twelve SFF-3 bays/RDX Bay/Dual IOA with Write Cache and External SAS port. Optionally, attach an EXP12SX/EXP24SX SAS HDD/SSD Expansion Drawer to the dual IOA.
- ▶ Up to two PCIe3 NVMe carrier cards with two M.2 module slots (with up to four Mainstream 400 GB SSD NVMe M.2 modules). One PCIe3 NVMe carrier card can be ordered only with a storage backplane. If a PCIe3 NVMe carrier card is ordered with a storage backplane, then the optional split feature is not supported.
- ▶ PCIe slots with single processor:
 - One x16 Gen4 full-height, half-length (CAPI).
 - One x8 Gen4 full-height, half-length (with x16 connector) (CAPI).
 - Two x8 Gen3 full-height, half-length (with x16 connectors).
 - Four x8 Gen3 full-height, half-length. (One of these slots is used for the required base LAN adapter.)
- ▶ Integrated:
 - Service processor.
 - EnergyScale technology.
 - Hot-swap and redundant cooling.
 - One front USB 3.0 port.
 - Two rear USB 3.0 ports.
 - Two HMC 1 GbE RJ45 ports.
 - One system port with RJ45 connector.
 - Four hot-plug, redundant power supplies.
 - 19-inch rack-mounting hardware (4U).

1.4.3 Power S924 server features

This summary describes the standard features of the Power S924 servers:

- ▶ POWER9 processor modules:
 - 8-core typical 3.8 - 4.0 GHz (maximum) POWER9 processor.
 - 10-core typical 3.5 - 3.9 GHz (maximum) POWER9 processor.
 - 11-core typical 3.45 - 3.9 GHz (maximum) POWER9 processor.
 - 12-core typical 3.4 - 3.9 GHz (maximum) POWER9 processor.
- ▶ High-performance Mbps DDR4 ECC memory:
 - 8 GB, 16 GB, 32 GB, 64 GB, or 128 GB memory. Different sizes/configurations run at different frequencies of 2133, 2400, and 2666 Mbps.

- Up to 4 TB of DDR4 memory with two POWER processors.
- Up to 2 TB of DDR4 memory with one POWER processor.
- ▶ Storage Backplane feature:
 - Base 12 SFF-3 bays/RDX bay. Optionally, split the SFF-3 bays and add a second integrated SAS controller without cache.
 - Expanded function 18 SFF-3 bays/Dual IOA with Write Cache and Optional External SAS port.
 - Expanded function 12 SFF-3 bays/RDX Bay/Dual IOA with Write Cache and Optional External SAS port.
 - Optionally, attach an EXP12SX/EXP24SX SAS HDD/SSD Expansion Drawer to the dual IOA.
- ▶ Up to two PCIe3 NVMe carrier cards with two M.2 module slots (with up to four Mainstream 400 GB SSD NVMe M.2 modules). One PCIe3 NVMe carrier card can be ordered only with a storage backplane. If the PCIe3 NVMe carrier card is ordered with a storage backplane, then the optional split feature is not supported.
- ▶ PCIe slots with single processor:
 - One x16 Gen4 full-height, half-length (CAPI).
 - One x8 Gen4, full-height, half-length (with x16 connector) (CAPI).
 - Two x8 Gen3, full-height, half-length (with x16 connectors).
 - Four x8 Gen3 full-height, half-length. (One of these slots is used for the required base LAN adapter.)
- ▶ PCIe slots with two processors:
 - Three x16 Gen4 full-height, half-length (CAPI).
 - Two x8 Gen4 full-height, half-length (with x16 connectors) (CAPI).
 - Two x8 Gen3 full-height, half-length (with x16 connectors).
 - Four x8 Gen3 full-height, half-length. (One of these slots is used for the required base LAN adapter.)
- ▶ Integrated:
 - Service processor.
 - EnergyScale technology.
 - Hot-swap and redundant cooling.
 - One front USB 3.0 port.
 - Two rear USB 3.0 ports.
 - Two HMC 1 GbE RJ45 ports.
 - One system port with RJ45 connector.
 - Four hot-plug, redundant power supplies.
 - 19-inch rack-mounting hardware (4U).

1.4.4 Minimum features

The minimum Power S922 or Power S914 initial order must include a processor module, two 8 GB DIMMs, two power supplies and power cords, an operating system indicator, a cover set indicator, and a Language Group Specify. Also, it must include one of the following storage options and one of the following network options:

- ▶ Storage options:
 - For boot from NVMe: One NVMe carrier and one NVMe M.2 Module.
 - For boot from a local SFF-3 HDD/SDD: One storage backplane and one SFF-3 HDD or SDD.
 - For boot from SAN: An internal HDD or SSD and RAID card are not required if Boot from SAN (#0837) is selected. A Fibre Channel adapter must be ordered if #0837 is selected.
- ▶ Network options:
 - One PCIe2 4-port 1 Gb Ethernet adapter.
 - One of the supported 10 Gb Ethernet adapters.

The minimum Power S924 initial order must include a processor module, two 8 GB DIMMs, four power supplies and power cords, an operating system indicator, a cover set indicator, and a Language Group Specify. Also, it must include one of the following storage options and one of the following network options:

- ▶ Storage options:
 - For boot from NVMe: One NVMe carrier and one NVMe M.2 Module.
 - For boot from a local SFF-3 HDD/SDD: One storage backplane and one SFF-3 HDD or SDD.
 - For boot from SAN: An internal HDD or SSD and RAID card are not required if Boot from SAN (#0837) is selected. A Fibre Channel adapter must be ordered if #0837 is selected.
- ▶ Network options:
 - One PCIe2 4-port 1 Gb Ethernet adapter.
 - One of the supported 10 Gb Ethernet adapters.

1.4.5 Power supply features

The Power S922 server supports two 1400 W 200 - 240 V AC (#EB2M) power supplies. Two power supplies are always installed. One power supply is required during the boot phase and for normal system operation, and the second is for redundancy.

The Power S914 server supports the following power supplies:

- ▶ Four 900 W 100 - 127 V AC or 200 - 240 V AC options (#EB2L) power supplies supporting a tower chassis. Two power supplies are required during the boot phase and for normal system operation, and the third and fourth are for redundancy.
- ▶ Two 1400 W 200 - 240 V AC (#EB2M) options power supplies supporting a rack chassis. One power supply is required during the boot phase and for normal system operation, and the second is for redundancy.

The Power S924 supports four 1400 W 200 - 240 V AC (#EB2M) power supplies. Four power supplies are always installed. Two power supplies are required during the boot phase and for normal system operation, and the third and fourth are for redundancy.

1.4.6 Processor module features

The following section describes the processor modules that are available for the Power S922, Power S914, and Power S924 servers.

Power S922 processor modules

A maximum of one processor with four processor cores (#EP16), or a maximum of two processors of either eight processor cores (#EP18) or 10 processor cores (#EP19) is allowed. All processor cores must be activated. The following list defines the allowed quantities of processor activation entitlements:

- ▶ One 4-core, typical 2.8 - 3.8 GHz (maximum) processor (#EP16) requires that four processor activation codes be ordered. A maximum of four processor activations (#EP46) is allowed.
- ▶ One 8-core, typical 3.4 - 3.9 GHz (maximum) processor (#EP18) requires that eight processor activation codes be ordered. A maximum of eight processor activations (#EP48) is allowed.
- ▶ Two 8-core, typical 3.4 - 3.9 GHz (maximum) processors (#EP18) require that 16 processor activation codes be ordered. A maximum of 16 processor activations (#EP48) is allowed.
- ▶ One 10-core, typical 2.9 - 3.8 GHz (maximum) processor (#EP19) requires that 10 processor activation codes be ordered. A maximum of 10 processor activation feature codes (FCs) (#EP49) is allowed.
- ▶ Two 10-core, typical 2.9 - 3.8 GHz (maximum) processors (#EP19) require that 20 processor activation codes be ordered. A maximum of 20 processor activation FCs (#EP49) is allowed.

Table 1-7 summarizes the processor features that are available for the Power S922 server.

Table 1-7 Processor features for the Power S922 server

Feature code	Processor module description
#EP16	4-core typical 2.8 - 3.8 GHz (maximum) POWER9 processor
#EP18	8-core typical 3.4 - 3.9 GHz (maximum) POWER9 processor
#EP19	10-core typical 2.9 - 3.8 GHz (maximum) POWER9 processor
#EP46	One processor entitlement for #EP16
#EP48	One processor entitlement for #EP18
#EP49	One processor entitlement for #EP19

Power S914 processor modules

A maximum of one processor with four processor cores (#EP10), one processor with six processor cores (#EP11), or one processor with eight processor cores (#EP12) is allowed. All processor cores must be activated. The following list defines the allowed quantities of processor activation entitlements:

- ▶ One 4-core, typical 2.3 - 3.8 GHz (maximum) processor (#EP10) requires that four processor activation codes be ordered. A maximum of four processor activations (#EP40) is allowed.
- ▶ One 6-core, typical 2.3 - 3.8 GHz (maximum) processor (#EP11) requires that six processor activation codes be ordered. A maximum of six processor activation FCs (#EP41) is allowed.
- ▶ One 8-core, typical 2.8 - 3.8 GHz (maximum) processor (#EP12) requires that eight processor activation codes be ordered. A maximum of eight processor activation FCs (#EP42) is allowed.

Table 1-8 summarizes the processor features that are available for the Power S914 server.

Table 1-8 Processor features for the Power S914 server

Feature code	Processor module description
#EP10	4-core typical 2.3 - 3.8 GHz (maximum) POWER9 processor
#EP11	6-core typical 2.3 - 3.8 GHz (maximum) POWER9 processor
#EP12	8-core typical 2.8 - 3.8 GHz (maximum) POWER9 processor
#EP40	One processor entitlement for #EP10
#EP41	One processor entitlement for #EP11
#EP42	One processor entitlement for #EP12

Power S924 processor modules

A maximum of two processors with eight processor cores (#EP1E), two processors with 10 processor cores (#EP1F), or two processors with 12 processor cores (#EP1G) is allowed. All processor cores must be activated. The following list defines the allowed quantities of processor activation entitlements:

- ▶ One 8-core, typical 3.8 - 4.0 GHz (maximum) processor (#EP1E) requires that eight processor activation codes be ordered. A maximum of eight processor activations (#EP4E) is allowed.
- ▶ Two 8-core, typical 3.8 - 4.0 GHz (maximum) processors (#EP1E) require that 16 processor activation codes be ordered. A maximum of 16 processor activations (#EP4E) is allowed.
- ▶ One 10-core, typical 3.5 - 3.9 GHz (maximum) processor (#EP1F) requires that 10 processor activation codes be ordered. A maximum of 10 processor activation FCs (#EP4F) is allowed.
- ▶ Two 10-core, typical 3.5 - 3.9 GHz (maximum) processors (#EP1F) require that 20 processor activation codes be ordered. A maximum of 20 processor activation FCs (#EP4F) is allowed.
- ▶ Two 11-core, typical 3.45 - 3.9 GHz (maximum) processors (#EP1H) require that 22 processor activation codes be ordered. A maximum of 22 processor activation FCs (#EP4H) is allowed.

- ▶ Two 12-core, typical 3.4 - 3.9 GHz (maximum) processors (#EP1G) require that 24 processor activation codes be ordered. A maximum of 24 processor activation FCs (#EP4G) is allowed.

Table 1-9 summarizes the processor features that are available for the Power S924 server.

Table 1-9 Processor features for the Power S924 server

Feature code	Processor module description
#EP1E	8-core typical 3.8 - 4.0 GHz (maximum) POWER9 processor
#EP1F	10-core typical 3.5 - 3.9 GHz (maximum) POWER9 processor
#EP1H	11-core typical 3.45 - 3.9 GHz (maximum) POWER9 processor
#EP1G	12-core typical 3.4 - 3.9 GHz (maximum) POWER9 processor
#EP4E	One processor entitlement for #EP1E
#EP4F	One processor entitlement for #EP1F
#EP4H	One processor entitlement for #EP1H
#EP4G	One processor entitlement for #EP1G

1.4.7 Memory features

A minimum of 32 GB of memory is required on the Power S922, Power S914, and Power S924 servers. Memory upgrades require memory pairs. The base memory is two 8 GB DDR4 memory modules (#EM60).

Table 1-10 lists the memory features that are available for the Power S922, Power S914, and Power S924 servers.

Table 1-10 Summary of memory features for Power S922, Power S914, and Power S924 servers

Feature code	DIMM capacity	Minimum quantity	Maximum quantity S922, and S924/S914
#EM60	8 GB	0	32/16 ^a
#EM62	16 GB	0	32/16
#EM63	32 GB	0	32/16
#EM64	64 GB	0	32/16
#EM65 ^b	128 GB	0	32/16

a. The maximum number of DIMMs for the S914 is 16.

b. The memory #EM65 is not available for the Power S914 server.

Note: Different sizes/configurations run at different frequencies of 2133, 2400, and 2666 Mbps.

1.4.8 Peripheral Component Interconnect Express slots

The following section describes the available PCIe slots:

- ▶ The Power S922 server has up to nine PCIe hot-plug slots:
 - With two POWER9 processor single-chip modules (SCMs), nine PCIe slots are available: Three are x16 Gen4 LP, half-length slots (CAPI), two are x8 Gen4 LP, half-length slots (with x16 connectors) (one is CAPI capable), two are x8 Gen3 LP, half-length slots (with x16 connectors), and two are x8 Gen3 LP, half-length slots.
 - With one POWER9 processor SCM, six PCIe slots are available: One is a x16 Gen4 LP half-length slot (CAPI), one is a x8 Gen4 LP, half-length slot (with x16 connector) (CAPI), two are x8 Gen3 LP, half-length slots (with x16 connectors), and two are x8 Gen3 LP, half-length slots.
- ▶ The Power S914 server has up to eight PCIe hot-plug slots. With one POWER9 processor SCM, eight PCIe slots are available: One is a x16 Gen4 full-height, half-length slot (CAPI), one is a x8 Gen4 full-height, half-length slot (with x16 connector) (CAPI), two are x8 Gen3 full-height, half-length slots (with x16 connectors), and four are x8 Gen3 full-height, half-length slots.
- ▶ The Power S924 server has up to 11 PCIe hot-plug slots:
 - With two POWER9 processor (SCMs), 11 PCIe slots are available: Three are x16 Gen4 full-height, half-length slots (CAPI), two are x8 Gen4 full-height, half-length slots (with x16 connectors) (one is CAPI capable), two are x8 Gen3 full-height, half-length slots (with x16 connectors), and four are x8 Gen3 full-height, half-length slots.
 - With one POWER9 processor SCM, eight PCIe slots are available: One is a x16 Gen4 full-height, half-length slot (CAPI), one is a x8 Gen4 full-height, half-length slot (with x16 connector) (CAPI), two are x8 Gen3 full-height, half-length slots (with x16 connectors), and four are x8 Gen3 full-height, half-length slots.

The x16 slots can provide up to twice the bandwidth of x8 slots because they offer twice as many PCIe lanes. PCIe Gen4 slots can support up to twice the bandwidth of a PCIe Gen3 slot, and PCIe Gen3 slots can support up to twice the bandwidth of a PCIe Gen2 slot, assuming an equivalent number of PCIe lanes.

At least one PCIe Ethernet adapter is required on the server by IBM to ensure proper manufacture, test, and support of the server. One of the x8 PCIe slots is used for this required adapter.

These servers are smarter about energy efficiency when cooling the PCIe adapter environment. They sense which IBM PCIe adapters are installed in their PCIe slots and, if an adapter requires higher levels of cooling, they automatically speed up fans to increase airflow across the PCIe adapters. Faster fans increase the sound level of the server.

1.5 Disk and media features

Three backplane options are available for the Power S922 servers:

- ▶ Base Storage Backplane 8 SFF-3 bays (#EJ1F)
- ▶ 4 + 4 SFF-3 bays split backplane (#EJ1H)
- ▶ Expanded function Storage Backplane 8 SFF-3 bays/Single IOA with Write Cache (#EJ1G)

The Storage Backplane option (#EJ1F) provides eight SFF-3 bays and one SAS controller with zero write cache.

By optionally adding the Split Backplane (#EJ1H), a second integrated SAS controller with no write cache is provided, and the eight SFF-3 bays are logically divided into two sets of four bays. Each SAS controller independently runs one of the four-bay sets of drives.

Four backplane options are available for the Power S914 and Power S924 servers:

- ▶ Base Storage Backplane 12 SFF-3 bays/RDX bay (#EJ1C)
- ▶ 6 +6 SFF-3 bays split backplane for #EJ1C (#EJ1E)
- ▶ Expanded function Storage Backplane 18 SFF-3 bays/Dual IOA with Write Cache and optional external SAS port (#EJ1D)
- ▶ Expanded function Storage Backplane 12 SFF-3 bays/RDX bay/Dual IOA with Write Cache and optional external SAS port (#EJ1M)

The Base Storage Backplane option (#EJ1C) provides 12 SFF-3 bays and one SAS controller with zero write cache.

By optionally adding the Split Backplane (#EJ1E), a second integrated SAS controller with no write cache is provided, and the 12 SFF-3 bays are logically divided into two sets of six bays. Each SAS controller independently runs one of the six-bay sets of drives.

The backplane options provide SFF-3 SAS bays in the system unit. These 2.5-inch or SFF SAS bays can contain SAS drives (HDD or SSD) mounted on a Gen3 tray or carrier. Thus, the drives are designated SFF-3. SFF-1 or SFF-2 drives do not fit in an SFF-3 bay. All SFF-3 bays support concurrent maintenance or hot-plug capability.

These backplane options use leading-edge, integrated SAS RAID controller technology that is designed and patented by IBM. A custom-designed PowerPC® based ASIC chip is the basis of these SAS RAID controllers, and provides RAID 5 and RAID 6 performance levels, especially for SSD. Internally, SAS ports are implemented and provide plenty of bandwidth. The integrated SAS controllers are placed in dedicated slots and do not reduce the number of available PCIe slots.

This backplane option supports HDDs or SSDs or a mixture of HDDs and SSDs in the SFF-3 bays. Mixing HDDs and SSDs applies even within a single set of six bays of the split backplane option.

Note: If mixing HDDs and SSDs, they must be in separate arrays (unless you use the IBM Easy Tier® function).

This backplane option can offer different drive protection options: RAID 0, RAID 5, RAID 6, or RAID 10. RAID 5 requires a minimum of three drives of the same capacity. RAID 6 requires a minimum of four drives of the same capacity. RAID 10 requires a minimum of two drives. Hot-spare capability is supported by RAID 5, RAID 6, or RAID 10.

This backplane option is supported by AIX and Linux, VIOS, and IBM i. It is highly recommended but not required that the drives be protected. With IBM i, all drives are required to be protected by either RAID or mirroring.

Unlike the hot-plug PCIe slots and SAS bays, concurrent maintenance is not available for the integrated SAS controllers. Scheduled downtime is required if a service action is required for these integrated resources.

In addition to supporting HDDs and SSDs in the SFF-3 SAS bays, the Expanded Function Storage Backplane (#EJ1G) supports the optional attachment of an EXP12SX/EXP24SX drawer. All bays are accessed by both of the integrated SAS controllers. The bays support concurrent maintenance (hot-plug).

Table 1-11 shows the available disk drive FCs that can be installed in the Power S922 server.

Table 1-11 Disk drive feature code description for the Power S922 server

Feature code	CCIN	Description	Maximum	OS support
ESNM	5B43	300 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	672	AIX and Linux
1953	19B1	300 GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESNK	5B41	300 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	8	AIX and Linux
ESDB	59E0	300 GB 15K RPM SAS SFF-3 Disk Drive (AIX/Linux)	8	AIX and Linux
ES94	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ES90	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESGV	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESB0	5B19	387 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESB2	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESB8	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESBA	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESGT	5B19	387 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	8	AIX and Linux
1964	19B3	600 GB 10k RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESEV	59D2	600 GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESD5	59D0	600 GB 10K RPM SAS SFF-3 Disk Drive (AIX/Linux)	8	AIX and Linux
ESF5	59D3	600 GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	8	AIX and Linux
ESNR	5B47	600 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	672	AIX and Linux
ESNP	5B45	600 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	8	AIX and Linux
ESNA	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESNC	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESGZ	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESB4	5B1A	775 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESB6	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBE	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESBG	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESGX	5B1A	775 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	8	AIX and Linux

Feature code	CCIN	Description	Maximum	OS support
ESJ0	5B29	931 GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ8	5B2B	931 GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESF3	59DA	1.2 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESF9	59DB	1.2 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	8	AIX and Linux
ESNE	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBJ	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESBL	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESNG	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESFT	59DD	1.8 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESFV	59DE	1.8 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	8	AIX and Linux
ESJ2	5B21	1.86 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJA	5B20	1.86 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ESJ4	5B2D	3.72 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJC	5B2C	3.72 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ES62	5B1D	3.86-4.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	336	AIX and Linux
ESJ6	5B2F	7.45 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJE	5B2E	7.45 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	8	AIX and Linux
ES64	5B1F	7.72-8.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	336	AIX and Linux
1929	19B1	Quantity 150 of #1953	4	
1818	19B3	Quantity 150 of #1964	4	
EQ62	5B1D	Quantity 150 of #ES62 3.86-4.0 TB 7200 rpm 4k LFF-1 Disk	2	
EQ64	5B1F	Quantity 150 of #ES64 7.72-8.0 TB 7200 rpm 4k LFF-1 Disk	2	
EQEV	59D2	Quantity 150 of #ESEV (600 GB 10k SFF-2)	4	
EQF3	59DA	Quantity 150 of #ESF3 (1.2 TB 10k SFF-2)	4	
EQFT	59DD	Quantity 150 of #ESFT (1.8 TB 10k SFF-2)	4	
ERHN	5B2F	Quantity 150 of #ESHN 7.45 TB SSD 4k SFF-2	2	
ERM8	5B2D	Quantity 150 of #ESM8 3.72 TB SSD 4k SFF-2	2	
ESPM	5B43	Quantity 150 of #ESNM (300 GB 15k SFF-2)	4	
ESPR	5B47	Quantity 150 of #ESNR (600 GB 15k SFF-2)	4	
ER94	5B10	Quantity 150 of ES94 387 GB SAS 4k	2	
ERGV	5B16	Quantity 150 of ESGV 387 GB SSD 4k	2	
ERGZ	5B17	Quantity 150 of ESGZ 775 GB SSD 4k	2	

Feature code	CCIN	Description	Maximum	OS support
ERJ0	5B29	Quantity 150 of ESJ0 931 GB SAS 4k	2	
ERJ2	5B21	Quantity 150 of ESJ2 1.86 TB SAS 4k	2	
ERJ4	5B2D	Quantity 150 of ESJ4 3.72 TB SAS 4k	2	
ERJ6	5B2F	Quantity 150 of ESJ6 7.45 TB SAS 4k	2	
ERNA	5B11	Quantity 150 of ESNA 775 GB SSD 4k	2	
ESQ2	5B16	Quantity 150 of ESB2 387 GB SAS 4k	2	
ESQ6	5B17	Quantity 150 of ESB6 775 GB SAS 4k	2	
ESQA	5B10	Quantity 150 of ESBA 387 GB SAS 4k	2	
ESQG	5B11	Quantity 150 of ESG 775 GB SAS 4k	2	
ESQL	5B12	Quantity 150 of ESBL 1.55 TB SAS 4k	2	
ERNE	5B12	Quantity 150 of ESNE 1.55 TB SSD 4k	2	

Table 1-12 shows the available disk drive FCs that can be installed in the Power S914 server.

Table 1-12 Disk drive feature code description for the Power S914 server

Feature code	CCIN	Description	Maximum	OS support
1948	19B1	283 GB 15k RPM SAS SFF-2 Disk Drive (IBM i)	672	IBM i
ESNJ	5B41	283 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (IBM i)	18	IBM i
ESDA	59E0	283 GB 15K RPM SAS SFF-3 Disk Drive (IBM i)	18	IBM i
ESNM	5B43	300 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	672	AIX and Linux
1953	19B1	300 GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESNK	5B41	300 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	18	AIX and Linux
ESDB	59E0	300 GB 15K RPM SAS SFF-3 Disk Drive (AIX/Linux)	18	AIX and Linux
ES94	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ES95	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ES90	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ES91	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESGV	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESG6	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for IBM i	336	IBM i
ESGT	5B19	387 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB0	5B19	387 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB2	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux

Feature code	CCIN	Description	Maximum	OS support
ESB8	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB9	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBA	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBB	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
1962	19B3	571 GB 10k RPM SAS SFF-2 Disk Drive (IBM i)	672	IBM i
ESEU	59D2	571 GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4224	672	IBM i
ESD4	59D0	571 GB 10K RPM SAS SFF-3 Disk Drive (IBM i)	18	IBM i
ESF4	59D3	571 GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4224	18	IBM i
ESNQ	5B47	571 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)	672	IBM i
ESDN	59CF	571 GB 15K RPM SAS SFF-2 Disk Drive - 528 Block (IBM i)	672	IBM i
ESNN	5B45	571 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (IBM i)	18	IBM i
1964	19B3	600 GB 10k RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESEV	59D2	600 GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESD5	59D0	600 GB 10K RPM SAS SFF-3 Disk Drive (AIX/Linux)	18	AIX and Linux
ESF5	59D3	600 GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	18	AIX and Linux
ESNR	5B47	600 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	672	AIX and Linux
ESDP	59CF	600 GB 15K RPM SAS SFF-2 Disk Drive - 5xx Block (AIX/Linux)	672	AIX and Linux
ESNP	5B45	600 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	18	AIX and Linux
ESNA	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESNB	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESNC	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESND	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESGZ	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESGG	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for IBM i	336	IBM i
ESGX	5B1A	775 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB4	5B1A	775 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB6	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBE	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESBF	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBG	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux

Feature code	CCIN	Description	Maximum	OS support
ESBH	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJ0	5B29	931 GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ1	5B29	931 GB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJ8	5B2B	931 GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJ9	5B2B	931 GB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESD2	59CD	1.1 TB 10K RPM SAS SFF-2 Disk Drive (IBM i)	672	IBM i
ESF2	59DA	1.1 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4224	672	IBM i
ESF8	59DB	1.1 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4224	18	IBM i
ESD3	59CD	1.2 TB 10K RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESF3	59DA	1.2 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESF9	59DB	1.2 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	18	AIX and Linux
ESNE	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESNF	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESNG	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESNH	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBJ	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESBK	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBL	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBM	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESFS	59DD	1.7 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4224	672	IBM i
ESFU	59DE	1.7 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4224	18	IBM i
ESFT	59DD	1.8 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESFV	59DE	1.8 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	18	AIX and Linux
ESJ2	5B21	1.86 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ3	5B21	1.86 TB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJA	5B20	1.86 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJB	5B20	1.86 TB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESJ4	5B2D	3.72 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ5	5B2D	3.72 TB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJC	5B2C	3.72 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJD	5B2C	3.72 TB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ES62	5B1D	3.86-4.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	336	AIX and Linux

Feature code	CCIN	Description	Maximum	OS support
ESJ6	5B2F	7.45 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ7	5B2F	7.45 TB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJE	5B2E	7.45 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJF	5B2E	7.45 TB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ES64	5B1F	7.72-8.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	336	AIX and Linux
1927	19B1	Quantity 150 of #1948	4	
1929	19B1	Quantity 150 of #1953	4	
1817	19B3	Quantity 150 of #1962	4	
1818	19B3	Quantity 150 of #1964	4	
EQ62	5B1D	Quantity 150 of #ES62 3.86-4.0 TB 7200 rpm 4k LFF-1 Disk	2	
EQ64	5B1F	Quantity 150 of #ES64 7.72-8.0 TB 7200 rpm 4k LFF-1 Disk	2	
EQD2	59CD	Quantity 150 of #ESD2 (1.1 TB 10k SFF-2)	4	
EQD3	59CD	Quantity 150 of #ESD3 (1.2 TB 10k SFF-2)	4	
EQDN	59CF	Quantity 150 of #ESDN (571 GB 15K RPM SAS SFF-2 for IBM i)	4	
EQDP	59CF	Quantity 150 of #ESDP (600 GB 15K RPM SAS SFF-2 for AIX/LINUX)	4	
EQEU	59D2	Quantity 150 of #ESEU (571 GB 10k SFF-2)	4	
EQEV	59D2	Quantity 150 of #ESEV (600 GB 10k SFF-2)	4	
EQF2	59DA	Quantity 150 of #ESF2 (1.1 TB 10k SFF-2)	4	
EQF3	59DA	Quantity 150 of #ESF3 (1.2 TB 10k SFF-2)	4	
EQFS	59DD	Quantity 150 of #ESFS (1.7 TB 10k SFF-2)	4	
EQFT	59DD	Quantity 150 of #ESFT (1.8 TB 10k SFF-2)	4	
EQG6	5B16	Quantity 150 of #ESG6 (387 GB SAS 5xx)	2	
EQGG	5B17	Quantity 150 of #ESGG (775 GB SAS 5xx)	2	
ESPL	5B43	Quantity 150 of #ESNL (283 GB 15k SFF-2)	4	
ESPM	5B43	Quantity 150 of #ESNM (300 GB 15k SFF-2)	4	
ESPQ	5B47	Quantity 150 of #ESNQ (571 GB 15k SFF-2)	4	
ESPR	5B47	Quantity 150 of #ESNR (600 GB 15k SFF-2)	4	
ER94	5B10	Quantity 150 of ES94 387 GB SAS 4k	2	
ER95	5B10	Quantity 150 of ES95 387 GB SAS 4k	2	
ERGV	5B16	Quantity 150 of ESGV 387 GB SSD 4k	2	
ERGZ	5B17	Quantity 150 of ESGZ 775 GB SSD 4k	2	

Feature code	CCIN	Description	Maximum	OS support
ERJ0	5B29	Quantity 150 of ESJ0 931 GB SAS 4k	2	
ERJ1	5B29	Quantity 150 of ESJ1 931 GB SAS 4k	2	
ERJ2	5B21	Quantity 150 of ESJ2 1.86 TB SAS 4k	2	
ERJ3	5B21	Quantity 150 of ESJ3 1.86 TB SAS 4k	2	
ERJ4	5B2D	Quantity 150 of ESJ4 3.72 TB SAS 4k	2	
ERJ5	5B2D	Quantity 150 of ESJ5 3.72 TB SAS 4k	2	
ERJ6	5B2F	Quantity 150 of ESJ6 7.45 TB SAS 4k	2	
ERJ7	5B2F	Quantity 150 of ESJ7 7.45 TB SAS 4k	2	
ERNA	5B11	Quantity 150 of ESNA 775 GB SSD 4k	2	
ERNB	5B11	Quantity 150 of ESNB 775 GB SSD 4k	2	
ERNE	5B12	Quantity 150 of ESNE 1.55 TB SSD 4k	2	
ERNF	5B12	Quantity 150 of ESNF 1.55 TB SSD 4k	2	
ESQ2	5B16	Quantity 150 of ESB2 387 GB SAS 4k	2	
ESQ6	5B17	Quantity 150 of ESB6 775 GB SAS 4k	2	
ESQA	5B10	Quantity 150 of ESBA 387 GB SAS 4k	2	
ESQB	5B10	Quantity 150 of ESBB 387 GB SAS 4k	2	
ESQG	5B11	Quantity 150 of ESBG 775 GB SAS 4k	2	
ESQH	5B11	Quantity 150 of ESBH 775 GB SAS 4k	2	
ESQL	5B12	Quantity 150 of ESBL 1.55 TB SAS 4k	2	
ESQM	5B12	Quantity 150 of ESBM 1.55 TB SAS 4k	2	

Table 1-13 shows the available disk drive FCs that can be installed in the Power S924 server.

Table 1-13 Disk drive feature code description for the Power S924 server

Feature code	CCIN	Description	Maximum	OS support
ESNL	5B43	283 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)	672	IBM i
1948	19B1	283 GB 15k RPM SAS SFF-2 Disk Drive (IBM i)	672	IBM i
ESNJ	5B41	283 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (IBM i)	18	IBM i
ESDA	59E0	283 GB 15K RPM SAS SFF-3 Disk Drive (IBM i)	18	IBM i
ESNM	5B43	300 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	672	AIX and Linux
1953	19B1	300 GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESNK	5B41	300 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	18	AIX and Linux
ESDB	59E0	300 GB 15K RPM SAS SFF-3 Disk Drive (AIX/Linux)	18	AIX and Linux
ES94	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ES95	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ES90	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ES91	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESGV	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESG6	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for IBM i	336	IBM i
ESGT	5B19	387 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB0	5B19	387 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB2	5B16	387 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESB8	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB9	5B13	387 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBA	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBB	5B10	387 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
1962	19B3	571 GB 10k RPM SAS SFF-2 Disk Drive (IBM i)	672	IBM i
ESEU	59D2	571 GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4224	672	IBM i
ESD4	59D0	571 GB 10K RPM SAS SFF-3 Disk Drive (IBM i)	18	IBM i
ESF4	59D3	571 GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4224	18	IBM i
ESNQ	5B47	571 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)	672	IBM i
ESDN	59CF	571 GB 15K RPM SAS SFF-2 Disk Drive - 528 Block (IBM i)	672	IBM i
ESNN	5B45	571 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (IBM i)	18	IBM i

Feature code	CCIN	Description	Maximum	OS support
1964	19B3	600 GB 10k RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESEV	59D2	600 GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESD5	59D0	600 GB 10K RPM SAS SFF-3 Disk Drive (AIX/Linux)	18	AIX and Linux
ESF5	59D3	600 GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	18	AIX and Linux
ESNR	5B47	600 GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	672	AIX and Linux
ESDP	59CF	600 GB 15K RPM SAS SFF-2 Disk Drive - 5xx Block (AIX/Linux)	672	AIX and Linux
ESNP	5B45	600 GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	18	AIX and Linux
ESNA	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESNB	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESNC	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESND	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESGZ	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESGG	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for IBM i	336	IBM i
ESGX	5B1A	775 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB4	5B1A	775 GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESB6	5B17	775 GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBE	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESBF	5B14	775 GB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBG	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBH	5B11	775 GB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJ0	5B29	931 GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ1	5B29	931 GB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJ8	5B2B	931 GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJ9	5B2B	931 GB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESD2	59CD	1.1 TB 10K RPM SAS SFF-2 Disk Drive (IBM i)	672	IBM i
ESF2	59DA	1.1 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4224	672	IBM i
ESF8	59DB	1.1 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4224	18	IBM i
ESD3	59CD	1.2 TB 10K RPM SAS SFF-2 Disk Drive (AIX/Linux)	672	AIX and Linux
ESF3	59DA	1.2 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESF9	59DB	1.2 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	18	AIX and Linux
ESNE	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux

Feature code	CCIN	Description	Maximum	OS support
ESNF	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESNG	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESNH	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBJ	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESBK	5B15	1.55 TB Enterprise SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESBL	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESBM	5B12	1.55 TB Enterprise SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESFS	59DD	1.7 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4224	672	IBM i
ESFU	59DE	1.7 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4224	18	IBM i
ESFT	59DD	1.8 TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	672	AIX and Linux
ESFV	59DE	1.8 TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	18	AIX and Linux
ESJ2	5B21	1.86 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ3	5B21	1.86 TB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJA	5B20	1.86 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJB	5B20	1.86 TB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ESJ4	5B2D	3.72 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ5	5B2D	3.72 TB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJC	5B2C	3.72 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJD	5B2C	3.72 TB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ES62	5B1D	3.86-4.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	336	AIX and Linux
ESJ6	5B2F	7.45 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	336	AIX and Linux
ESJ7	5B2F	7.45 TB Mainstream SAS 4k SFF-2 SSD for IBM i	336	IBM i
ESJE	5B2E	7.45 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	18	AIX and Linux
ESJF	5B2E	7.45 TB Mainstream SAS 4k SFF-3 SSD for IBM i	18	IBM i
ES64	5B1F	7.72-8.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	336	AIX and Linux
1927	19B1	Quantity 150 of #1948	4	
1929	19B1	Quantity 150 of #1953	4	
1817	19B3	Quantity 150 of #1962	4	
1818	19B3	Quantity 150 of #1964	4	
EQ62	5B1D	Quantity 150 of #ES62 3.86-4.0 TB 7200 rpm 4k LFF-1 Disk	2	
EQ64	5B1F	Quantity 150 of #ES64 7.72-8.0 TB 7200 rpm 4k LFF-1 Disk	2	
EQD2	59CD	Quantity 150 of #ESD2 (1.1 TB 10k SFF-2)	4	

Feature code	CCIN	Description	Maximum	OS support
EQD3	59CD	Quantity 150 of #ESD3 (1.2 TB 10k SFF-2)	4	
EQDN	59CF	Quantity 150 of #ESDN (571 GB 15K RPM SAS SFF-2 for IBM i)	4	
EQDP	59CF	Quantity 150 of #ESDP (600 GB 15K RPM SAS SFF-2 for AIX/LINUX)	4	
EQEU	59D2	Quantity 150 of #ESEU (571 GB 10k SFF-2)	4	
EQEV	59D2	Quantity 150 of #ESEV (600 GB 10k SFF-2)	4	
EQF2	59DA	Quantity 150 of #ESF2 (1.1 TB 10k SFF-2)	4	
EQF3	59DA	Quantity 150 of #ESF3 (1.2 TB 10k SFF-2)	4	
EQFS	59DD	Quantity 150 of #ESFS (1.7 TB 10k SFF-2)	4	
EQFT	59DD	Quantity 150 of #ESFT (1.8 TB 10k SFF-2)	4	
EQG6	5B16	Quantity 150 of #ESG6 (387 GB SAS 5xx)	2	
EQGG	5B17	Quantity 150 of #ESGG (775 GB SAS 5xx)	2	
ESPL	5B43	Quantity 150 of #ESNL (283 GB 15k SFF-2)	4	
ESPM	5B43	Quantity 150 of #ESNM (300 GB 15k SFF-2)	4	
ESPQ	5B47	Quantity 150 of #ESNQ (571 GB 15k SFF-2)	4	
ESPR	5B47	Quantity 150 of #ESNR (600 GB 15k SFF-2)	4	
ER94	5B10	Quantity 150 of ES94 387 GB SAS 4k	2	
ER95	5B10	Quantity 150 of ES95 387 GB SAS 4k	2	
ERGV	5B16	Quantity 150 of ESGV 387 GB SSD 4k	2	
ERGZ	5B17	Quantity 150 of ESGZ 775 GB SSD 4k	2	
ERJ0	5B29	Quantity 150 of ESJ0 931 GB SAS 4k	2	
ERJ1	5B29	Quantity 150 of ESJ1 931 GB SAS 4k	2	
ERJ2	5B21	Quantity 150 of ESJ2 1.86 TB SAS 4k	2	
ERJ3	5B21	Quantity 150 of ESJ3 1.86 TB SAS 4k	2	
ERJ4	5B2D	Quantity 150 of ESJ4 3.72 TB SAS 4k	2	
ERJ5	5B2D	Quantity 150 of ESJ5 3.72 TB SAS 4k	2	
ERJ6	5B2F	Quantity 150 of ESJ6 7.45 TB SAS 4k	2	
ERJ7	5B2F	Quantity 150 of ESJ7 7.45 TB SAS 4k	2	
ERNA	5B11	Quantity 150 of ESNA 775 GB SSD 4k	2	
ERNB	5B11	Quantity 150 of ESNB 775 GB SSD 4k	2	
ERNE	5B12	Quantity 150 of ESNE 1.55 TB SSD 4k	2	
ERNF	5B12	Quantity 150 of ESNF 1.55 TB SSD 4k	2	
ESQ2	5B16	Quantity 150 of ESB2 387 GB SAS 4k	2	

Feature code	CCIN	Description	Maximum	OS support
ESQ6	5B17	Quantity 150 of ESB6 775 GB SAS 4k	2	
ESQA	5B10	Quantity 150 of ESBA 387 GB SAS 4k	2	
ESQB	5B10	Quantity 150 of ESBB 387 GB SAS 4k	2	
ESQG	5B11	Quantity 150 of ESBG 775 GB SAS 4k	2	
ESQH	5B11	Quantity 150 of ESBH 775 GB SAS 4k	2	
ESQL	5B12	Quantity 150 of ESBL 1.55 TB SAS 4k	2	
ESQM	5B12	Quantity 150 of ESBM 1.55 TB SAS 4k	2	

The RDX docking station #EUA4 accommodates RDX removable disk cartridges of any capacity. The disk is in a protective rugged cartridge enclosure that plugs into the docking station. The docking station holds one removable rugged disk drive/cartridge at a time. The rugged removable disk cartridge and docking station perform saves, restores, and backups similar to a tape drive. This docking station can be an excellent entry capacity/performance option.

The Standalone USB DVD drive (#EUA5) is an optional, stand-alone external USB-DVD device. It requires high current at 5V and must use the front USB 3.0 port.

1.6 I/O drawers for Power S922, Power S914, and Power S924 servers

If more Gen3 PCIe slots beyond the system node slots are required, PCIe Gen3 I/O drawers can be attached to the Power S922, Power S914, and Power S924 servers.

EXP24SX /EXP12SX SAS Storage Enclosures (#ESLS/ELLS or #ESLL/ELLL) are also supported, providing storage capacity.

Note: The PCIe expansion drawer (#EMX0) and EXP24SX /EXP12SX SAS Storage Enclosures (#ESLS/ELLS or #ESLL/ELLL) do not apply to the 4-core configuration Power S914 server.

The 7226-1U3 offers a 1U rack-mountable dual bay enclosure with storage device options of LTO5, 6, 7, and 8 tape drives with both SAS and Fibre Channel interface. The 7226 also offers DVD-RAM SAS and USB drive features and RDX 500 GB, 1 TB, and 2 TB drive options. Up to two drives (or four DVD-RAM) can be installed in any combination in the 7226 enclosure.

Note: The tower system configuration of Power S914 does not allow attachment of any I/O expansion drawer.

1.6.1 PCIe Gen3 I/O expansion drawer

This 19-inch, 4U (4 EIA) enclosure provides PCIe Gen3 slots outside of the system unit. It has two module bays. One 6-Slot fan-out Module (#EMXH or #EMXG) can be placed in each module bay. Two 6-slot modules provide a total of 12 PCIe Gen3 slots. Each fan-out module is

connected to a PCIe3 Optical Cable adapter that is in the system unit over an active optical CXP cable (AOC) pair or CXP copper cable pair.

The PCIe Gen3 I/O Expansion Drawer has two redundant, hot-plug power supplies. Each power supply has its own separately ordered power cord. The two power cords plug into a power supply conduit that connects to the power supply. The single-phase AC power supply is rated at 1030 W and can use 100 - 120V or 200 - 240V. If using 100 - 120V, then the maximum is 950 W. It is a preferred practice that the power supply connects to a power distribution unit (PDU) in the rack. Power Systems PDUs are designed for a 200 - 240V electrical source.

A blind swap cassette (BSC) is used to house the full-height adapters that go into these slots. The BSC is the same BSC that is used with the previous generation server's 12X attached I/O drawers (#5802, #5803, #5877, and #5873). The drawer is shipped with a full set of BSCs, even if the BSCs are empty.

Concurrent repair and add/removal of PCIe adapters are done by HMC-guided menus or by operating system support utilities.

Figure 1-8 shows a PCIe Gen3 I/O expansion drawer.



Figure 1-8 PCIe Gen3 I/O expansion drawer

1.6.2 I/O drawers and usable PCI slots

Figure 1-9 shows the rear view of the PCIe Gen3 I/O expansion drawer that is equipped with two PCIe3 6-slot fan-out modules with the location codes for the PCIe adapter slots.




Figure 1-9 Rear view of a PCIe Gen3 I/O expansion drawer with PCIe slots location codes

Table 1-14 provides details about the PCI slots in the PCIe Gen3 I/O expansion drawer that is equipped with two PCIe3 6-slot fan-out modules.

Table 1-14 PCIe slot locations for the PCIe Gen3 I/O expansion drawer with two fan-out modules

Slot	Location code	Description
Slot 1	P1-C1	PCIe3, x16
Slot 2	P1-C2	PCIe3, x8
Slot 3	P1-C3	PCIe3, x8
Slot 4	P1-C4	PCIe3, x16
Slot 5	P1-C5	PCIe3, x8
Slot 6	P1-C6	PCIe3, x8
Slot 7	P2-C1	PCIe3, x16
Slot 8	P2-C2	PCIe3, x8
Slot 9	P2-C3	PCIe3, x8
Slot 10	P2-C4	PCIe3, x16
Slot 11	P2-C5	PCIe3, x8
Slot 12	P2-C6	PCIe3, x8

In the table:

- ▶ All slots support full-length, regular-height adapters or short (LP) adapters with a regular-height tailstock in single-wide, Gen3, BSC.
- ▶ Slots C1 and C4 in each PCIe3 6-slot fan-out module are x16 PCIe3 buses, and slots C2, C3, C5, and C6 are x8 PCIe buses.
- ▶ All slots support enhanced error handling (EEH).
- ▶ All PCIe slots are hot-swappable and support concurrent maintenance.

Table 1-15 summarizes the maximum number of I/O drawers that are supported and the total number of PCI slots that are available.

Table 1-15 Maximum number of I/O drawers that are supported and total number of PCI slots

Server	Maximum number of I/O Exp Drawers	Maximum number of I/O fan-out modules	Maximum PCIe slots
Power S922 (1-socket)	1	1	11
Power S922 (2-socket)	2	3	24
Power S914 (1-socket) ^a	1	1	13
Power S924 (1-socket)	1	1	13
Power S924 (2-socket)	2	3	26

a. The PCIe expansion drawer (#EMX0) do not apply to the 4-core configuration.

1.6.3 EXP24SX (#ESLS) and EXP12SX (#ESLL) SAS Storage Enclosures

If you need more disks than are available with the internal disk bays, you can attach more external disk subsystems, such as an EXP24SX SAS Storage Enclosure (#ESLS) or EXP12SX SAS Storage Enclosure (#ESLL).

The EXP24SX is a storage expansion enclosure with twenty-four 2.5-inch SFF SAS bays. It supports up to 24 hot-plug HDDs or SSDs in only 2 EIA of space in a 19-inch rack. The EXP24SX SFF bays use SFF Gen2 (SFF-2) carriers or trays.

The EXP12SX is a storage expansion enclosure with twelve 3.5-inch large form factor (LFF) SAS bays. It supports up to 12 hot-plug HDDs in only 2 EIA of space in a 19-inch rack. The EXP12SX SFF bays use LFF Gen1 (LFF-1) carriers/trays. The 4 KB sector drives (#4096 or #4224) are supported. SSDs are not supported.

With AIX, Linux, and VIOS, the EXP24SX or the EXP12SX can be ordered with four sets of six bays (mode 4), two sets of 12 bays (mode 2), or one set of 24-four bays (mode 1). With IBM i, only one set of 24 bays (mode 1) is supported. It is possible to change the mode setting in the field by using software commands along with a specifically documented procedure.

Important: When changing modes, a skilled, technically qualified person should follow the special documented procedures. Improperly changing modes can potentially destroy existing RAID sets, prevent access to existing data, or allow other partitions to access another partition's existing data.

Four mini-SAS HD ports on the EXP24SX or EXP12SX are attached to PCIe Gen3 SAS adapters or attached to an integrated SAS controller in the Power S914, Power S922, or Power S924 servers.

The attachment between the EXP24SX or EXP12SX and the PCIe3 SAS adapters or integrated SAS controllers is through SAS YO12 or X12 cables. All ends of the YO12 and X12 cables have mini-SAS HD narrow connectors.

The EXP24SX or EXP12SX includes redundant AC power supplies and two power cords.

Figure 1-10 shows the EXP24SX drawer.



Figure 1-10 The EXP24SX drawer

Figure 1-11 shows the EXP12SX drawer.



Figure 1-11 The EXP12SX drawer

1.7 IBM i Solution Editions for Power S914

The IBM i Solution Editions are designed to help you take advantage of the combined experience and expertise of IBM and independent software vendors (ISVs) in building business value with your IT investments. A qualifying purchase of software, maintenance, services, or training for a participating ISV solution is required when purchasing an IBM i Solution Edition.

The IBM Power S914 Solution Edition supports 4-core (#4928) and 6-core (#4927) configurations. For a list of participating ISVs, a registration form, and more details, see [IBM i Solution Editions](#).

Here are the requirements to be eligible to purchase an IBM Solution Edition for Power S914:

- ▶ The server must be ordered with a Solution Edition FC.
- ▶ The server must be a new purchase, not a miscellaneous execution system (MES) upgrade.
- ▶ The offering must include new or upgraded software licenses, or software maintenance from the ISV for the qualifying IBM server. Services and training for the qualifying server can also be provided. Contact the individual ISV for details.
- ▶ Proof of purchase (for example, a copy of the invoice) of the solution with a participating ISV must be provided to IBM on request. The proof must be dated within 90 days before or after the date of order of the qualifying server.
- ▶ The combined value (software, maintenance, services, and training) of the ISV purchase must be US \$6,500 or greater. Contact the individual ISV for details.
- ▶ Not eligible as an IBM Capacity BackUp (CBU).

1.8 IBM Capacity BackUp

The CBU designation enables you to temporarily transfer IBM i processor license entitlements and IBM i user license entitlements that are purchased for a primary machine to a secondary CBU-designated system for HADR operations. Temporarily transferring these resources instead of purchasing them for your secondary system can result in significant savings. Processor activations cannot be transferred.

If your primary or CBU machine is sold or discontinued from use, any temporary entitlement transfers must be returned to the machine on which they were originally acquired. For CBU registration, terms, and conditions, and more information, see [IBM Power Systems Capacity BackUp](#).

1.8.1 Power S922 (9009-22A) IBM Capacity BackUp offering

The CBU specify (#0444) is available only as part of a new server purchase. Certain system prerequisites must be met, and system registration and approval are required before the CBU specify feature can be applied on a new server. Standard IBM i terms and conditions do not allow either IBM i processor license entitlements or IBM i user license entitlements to be transferred permanently or temporarily. These entitlements remain with the machine for which they were ordered. When you register the association between your primary and on-order CBU system, you must agree to certain terms and conditions regarding the temporary transfer.

After a new CBU system is registered as a pair with the proposed primary system and the configuration is approved, you can temporarily move your optional IBM i processor license entitlement and IBM i user license entitlements from the primary system to the CBU system when the primary system is down or while the primary system processors are inactive. The CBU system can then support failover and role-swapping for a full range of test, disaster recovery, and high availability scenarios. Temporary entitlement transfer means that the entitlement is a property that is transferred from the primary system to the CBU system, which may remain in use on the CBU system while the registered primary and CBU systems are deployed for the high availability or disaster recovery operation. The intent of the CBU offering is to enable regular role-swap operations.

The primary systems for a Power S922 (9009-22A) CBU server with its IBM i P10 software tier can be a POWER9 or POWER8 server with a P10 or P20 software tie, as shown in the following list:

- ▶ Power S824 (8286-42A) server
- ▶ Power S814 6-core or 8-core (8286-41A) server
- ▶ Power S822 (8284-22A) server
- ▶ Power S924 (9009-42A) server
- ▶ Power S914 6-core or 8-core (9009-41A) server
- ▶ Power S922 (9009-22A) server

The primary machine must be in the same enterprise as the CBU system. The IBM i Solution Editions are not eligible for CBU status.

Before you can temporarily transfer IBM i processor license entitlements from the registered primary system, you must have more than one IBM i processor license on the primary machine and at least one IBM i processor license on the CBU system. An activated processor must be available on the CBU system to use the transferred entitlement. You can then transfer any IBM i processor entitlements above the minimum one, assuming the total IBM i workload on the primary system does not require the IBM i entitlement that you want to transfer during the time of the transfer. During this temporary transfer, the CBU system's internal records of its total number of IBM i processor license entitlements are not updated, and you might see IBM i license noncompliance warning messages from the CBU system. These warning messages in this situation do not mean that you are not in compliance.

Before you can temporarily transfer IBM i user entitlements, you must have more than the minimum number of IBM i user entitlements on a primary system. You can then transfer any IBM i user entitlements above the minimum, assuming the total IBM i users on the primary system do not require the IBM i entitlement that you want to transfer during the time of the transfer. The Power S824 and Power S924 servers do not have user entitlements that can be transferred, and only processor license entitlements can be transferred. The minimum number of IBM i users on the POWER9 and POWER8 with IBM i user entitlements are:

- ▶ Power S814 6-core or 8-core (8286-41A) server: Ten users
- ▶ Power S822 (8284-22A) server: Ten users
- ▶ Power S914 6-core or 8-core (9009-41A) server: Ten users
- ▶ Power S922 (9009-22A) server: Ten users

For example, if you have a Power S914 6-core server as your primary system with two IBM i processor license entitlements (one above the minimum) and 40 IBM i user entitlements (30 above the minimum), you can temporarily transfer up to one IBM i entitlement and up to 30 user entitlements. During this temporary transfer, the CBU system's internal records of its total number of IBM i processor and user license entitlements is not updated, and you might see IBM i license noncompliance warning messages from the CBU system.

1.8.2 Power S914 (9009-41A) IBM Capacity BackUp offering

The CBU specify (#0444) is available only as part of a new server purchase. Certain system prerequisites must be met, and system registration and approval are required before the CBU specify feature can be applied on a new server. Standard IBM i terms and conditions do not allow either IBM i processor license entitlements or IBM i user license entitlements to be transferred permanently or temporarily. These entitlements remain with the machine for which they were ordered. When you register the association between your primary and on-order CBU systems, you must agree to certain terms and conditions regarding the temporary transfer.

After a new CBU system is registered along with the proposed primary system and the configuration is approved, you can temporarily move your optional IBM i processor license entitlement and IBM i user license entitlements from the primary system to the CBU system when the primary system is down or while the primary system processors are inactive. The CBU system can then support failover and role-swapping for a full range of test, disaster recovery, and high availability scenarios. Temporary entitlement transfer means that the entitlement is a property that is transferred from the primary system to the CBU system, which may remain in use on the CBU system while the registered primary and CBU systems are in deployment for the high availability or disaster recovery operation. The intent of the CBU offering is to enable regular role-swap operations. The primary machine must be in the same enterprise as the CBU system. The IBM i Solution Editions are not eligible for CBU status.

The Power S914 server is available with six or eight cores in the P10 software tier and four cores in the P05 software tier.

Power S914 software tiers for IBM i on model 9009-41A

Here are the software tiers:

- ▶ The 4-core processor (#EP10, QPRCFEAT EP10) is IBM i software tier P05.
- ▶ The 6-core processor (#EP11, QPRCFEAT EP11) is IBM i software tier P10.
- ▶ The 8-core processor (#EP12, QPRCFEAT EP12) is IBM i software tier P10.

Power S914 CBU server in the P10 software tier

The primary systems for a Power S914 (9009-41A) CBU server with a IBM i P10 software tier can be a POWER8 or POWER9 server with a P10 or P20 software tier:

- ▶ Power S824 (8286-42A) server
- ▶ Power S814 6-core or 8-core (8286-41A) server
- ▶ Power S822 (8284-22A) server
- ▶ Power S924 (9009 42A) server
- ▶ Power S914 6-core or 8-core (9009-41A) server
- ▶ Power S922 (9009-22A) server

Before you can temporarily transfer IBM i user entitlements, you must have more than the minimum number of IBM i user entitlements on a primary server. You can then transfer any IBM i user entitlements above the minimum, assuming that the total IBM i users on the primary system do not require the IBM i entitlement that you want to transfer during the time of the transfer. The Power S924 and S824 servers do not have IBM i user entitlements to transfer, only processor entitlements. For a P10 primary, the minimum number of IBM i user entitlements on the eligible P10 POWER9 and POWER8 servers are:

- ▶ Power S814 6-core or 8-core (8286-41A) server: Ten users
- ▶ Power S822 (8284-22A) server: Ten users
- ▶ Power S914 6-core or 8-core (9009-41A) server: Ten users
- ▶ Power S922 (9009-22A) server: Ten users

Power S914 CBU server in the P05 software tier

The primary systems for a Power S914 (9009-41A) CBU server with a IBM i P05 software tier can be a POWER8 or POWER9 server with a P05 or P10 software tier:

- ▶ Power S814 (8286-41A) 4, 6, or 8 core server
- ▶ Power S822 (8284-22A) server
- ▶ Power S914 (9009-41A) 4, 6, or 8 core server
- ▶ Power S922 (9009-22A) server

Before you can temporarily transfer IBM i user entitlements, you must have more than the minimum number of IBM i user entitlements on a primary server. You can then transfer any IBM i user entitlements above the minimum, assuming that the total IBM i users on the primary system do not require the IBM i entitlement that you want to transfer during the time of the transfer. The minimum number of IBM i user entitlements on the P05 or P10 POWER9 and POWER8 with IBM i user entitlements are:

- ▶ Power S814 4-core (8286-41A) server: Five users
- ▶ Power S814 6-core or 8-core (8286-41A) server: Ten users
- ▶ Power S822 (8284-22A) server: Ten users
- ▶ Power S914 4-core (9009-41A) server: Five users
- ▶ Power S914 6-core or 8-core (9009-41A) server: Ten users
- ▶ Power S922 (9009-22A) server: Ten users

For example, if you have a 2-core server as your primary system with two IBM i processor license entitlements (one above the minimum) and 50 IBM i user entitlements (20 above the minimum), you can temporarily transfer up to one IBM i entitlement and up to 20 user entitlements. During this temporary transfer, the CBU system's internal records of its total number of IBM i processor and user license entitlements is not updated, and you might see IBM i license noncompliance warning messages from the CBU system.

1.8.3 Power S924 (9009-42A) IBM Capacity BackUp offering

The CBU specify (#0444) is available only as part of a new server purchase. Certain system prerequisites must be met, and system registration and approval are required before the CBU specify feature can be applied on a new server. Standard IBM i terms and conditions do not allow either IBM i processor license entitlements or IBM i user license entitlements to be transferred permanently or temporarily. These entitlements remain with the machine for which they were ordered. When you register the association between your primary and on-order CBU systems, you must agree to certain terms and conditions regarding the temporary transfer.

After a new CBU system is registered along with the proposed primary system and the configuration is approved, you can temporarily move your optional IBM i processor license entitlement and IBM i Enterprise Enablement (#5250) entitlements from the primary system to the CBU system when the primary system is down or while the primary system processors are inactive. The CBU system can then support failover and role-swapping for a full range of test, disaster recovery, and high availability scenarios. Temporary entitlement transfer means that the entitlement is a property that is transferred from the primary system to the CBU system, which may remain in use on the CBU system while the registered primary and CBU systems are in deployment for the high availability or disaster recovery operation. The intent of the CBU offering is to enable regular role-swap operations.

Before you can temporarily transfer Enterprise Enablement entitlements (#5250), you must have more than one Enterprise Enablement entitlement (#5250) on the primary server and at least one Enterprise Enablement entitlement (#5250) on the CBU system. You can then transfer the entitlements that are not required on the primary server during the time of transfer and that are above the minimum of one entitlement. The minimum number of permanent entitlements on the CBU is one; however, you are required to license all permanent workloads, such as replication workloads. If, for example, the replication workload uses four processor cores at peak workload, then you are required to permanently license four cores on the CBU.

For example, if you have a 12-core Power S824 server as your primary system with six IBM i processor license entitlements (five above the minimum) and two Enterprise Enablement entitlements (#5250) (one above the minimum), you can temporarily transfer up to five IBM i entitlements and one Enterprise Enablement entitlement (#5250). During the temporary transfer, the CBU system's internal records of its total number of IBM i processor entitlements are not updated, and you might see IBM i license noncompliance warning messages from the CBU system.

The CBU specify (#0444) is available only as part of a new server purchase. Certain system prerequisites must be met, and system registration and approval are required before the CBU specify feature can be applied on a new server. Standard IBM i terms and conditions do not allow either IBM i processor license entitlements or Enterprise Enablement entitlements (#5250) to be transferred permanently or temporarily. These entitlements remain with the machine for which they were ordered. When you register the association between your primary and on-order CBU systems, you must agree to certain terms and conditions regarding the temporary transfer.

The servers with P20 or higher software tiers do not have user entitlements that can be transferred, and only processor license entitlements can be transferred.

Here are the eligible primary servers for a Power S924 CBU:

- ▶ Power S824 (8286-42A) server
- ▶ Power S924 (9009-42A) server
- ▶ Power E870 (9119-MME) server
- ▶ Power E880 (9119-MHE) server
- ▶ Power E870C (9080-MME) server
- ▶ Power E880C (9080-MHE) server

Power S924 software tiers for IBM i on 9009-42A

Here are the Power S924 software tiers for IBM on model 9009-42A:

- ▶ The 8-core processor (#EP1E, QPRCRFEAT EP1E) is IBM i software tier P20.
- ▶ The 10-core processor (#EP1F, QPRCRFEAT EP1F) is IBM i software tier P20.
- ▶ The 11-core processor (#EP1H, QPRCRFEAT EP1H) is IBM i software tier P20.
- ▶ The 12-core processor (#EP1G, QPRCRFEAT EP1G) is IBM i software tier P20.

1.9 System racks

The Power S914, Power S922, and Power S924 servers are mounted in the 36U 7014-T00 (#0551), the 42U 7014-T42 (#0553), or the IBM 42U Slim Rack (7965-94Y) racks. These racks are built to the 19-inch EIA 310D standard.

Order information: The racking approach for the initial order must be either a 7014-T00, 7014-T42, or 7965-94Y. If an extra rack is required for I/O expansion drawers as an MES to an existing system, either an #0551, #0553, or #ER05 rack must be ordered.

You must leave 2U of space at either the bottom or top of the rack, depending on the client's cabling preferences, to allow for cabling to exit the rack.

If a system will be installed in a rack or cabinet that is not from IBM, ensure that the rack meets the requirements that are described in 1.9.10, "Original equipment manufacturer racks" on page 52.

Responsibility: The client is responsible for ensuring that the installation of the drawer in the preferred rack or cabinet results in a configuration that is stable, serviceable, safe, and compatible with the drawer requirements for power, cooling, cable management, weight, and rail security.

1.9.1 IBM 7014 Model T00 rack

The 1.8-meter (71-inch) Model T00 is compatible with past and present IBM Power Systems servers. The features of the T00 rack are as follows:

- ▶ Has 36U (EIA units) of usable space.
- ▶ Has optional removable side panels.
- ▶ Has optional side-to-side mounting hardware for joining multiple racks.
- ▶ Has increased power distribution and weight capacity.
- ▶ Supports both AC and DC configurations.
- ▶ Up to four PDUs can be mounted in the PDU bays (see Figure 1-13 on page 47), but others can fit inside the rack. For more information, see 1.9.7, “The AC power distribution unit and rack content” on page 45.
- ▶ For the T00 rack, three door options are available:
 - Front Door for 1.8 m Rack (#6068).

This feature provides an attractive black full height rack door. The door is steel, with a perforated flat front surface. The perforation pattern extends from the bottom to the top of the door to enhance ventilation and provide some visibility into the rack.
 - A 1.8 m Rack Acoustic Door (#6248).

This feature provides a front and rear rack door that is designed to reduce acoustic sound levels in a general business environment.
 - A 1.8 m Rack Trim Kit (#6263).

If no front door is used in the rack, this feature provides a decorative trim kit for the front.
- ▶ Ruggedized Rack Feature

For enhanced rigidity and stability of the rack, the optional Ruggedized Rack Feature (#6080) provides more hardware that reinforces the rack and anchors it to the floor. This hardware is designed primarily for use in locations where earthquakes are a concern. The feature includes a large steel brace or truss that bolts into the rear of the rack.

It is hinged on the left side so it can swing out of the way for easy access to the rack drawers when necessary. The Ruggedized Rack Feature also includes hardware for bolting the rack to a concrete floor or a similar surface, and bolt-in steel filler panels for any unoccupied spaces in the rack.
- ▶ Weights are as follows:
 - T00 base empty rack: 244 kg (535 lb).
 - T00 full rack: 816 kg (1795 lb).
 - Maximum weight of drawers is 572 kg (1260 lb).
 - Maximum weight of drawers in a zone 4 earthquake environment is 490 kg (1080 lb). This number equates to 13.6 kg (30 lb) per EIA.

Important: If more weight is added to the top of the rack, for example, adding #6117, the 490 kg (1080 lb) must be reduced by the weight of the addition. As an example, #6117 weighs approximately 45 kg (100 lb) so the new maximum weight of drawers that the rack can support in a zone 4 earthquake environment is 445 kg (980 lb). In the zone 4 earthquake environment, the rack must be configured starting with the heavier drawers at the bottom of the rack.

1.9.2 IBM 7014 Model T42 rack

The 2.0-meter (79.3-inch) Model T42 addresses the client requirement for a tall enclosure to house the maximum amount of equipment in the smallest possible floor space. The following features are for the model T42 rack (which differ from the model T00):

- ▶ The T42 rack has 42U (EIA units) of usable space (6U of extra space).
- ▶ The model T42 supports AC power only.
- ▶ Weights are as follows:
 - T42 base empty rack: 261 kg (575 lb)
 - T42 full rack: 930 kg (2045 lb)

The available door options for the Model T42 rack are shown in Figure 1-12.



Figure 1-12 Door options for the T42 rack

The door options are explained in the following list:

- ▶ The 2.0 m Rack Trim Kit (#6272) is used if no front door is used in the rack.
- ▶ The Front Door for a 2.0 m Rack (#6069) is made of steel, with a perforated flat front surface. The perforation pattern extends from the bottom to the top of the door to enhance ventilation and provide some visibility into the rack. This door is non-acoustic and has a depth of about 25 mm (1 in.).
- ▶ The 2.0 m Rack Acoustic Door (#6249) consists of a front and rear door to reduce noise by approximately 6 dB(A). It has a depth of approximately 191 mm (7.5 in.).

- ▶ The High-End Appearance Front Door (#6250) provides a front rack door with a field-installed Power 780 logo indicating that the rack contains a Power 780 system. The door is not acoustic and has a depth of about 90 mm (3.5 in.).

High end: For the High-End Appearance Front Door (#6250), use the High-End Appearance Side Covers (#6238) to make the rack appear as though it is a high-end server (but in a 19-inch rack format instead of a 24-inch rack).

- ▶ The #ERG7 provides an attractive black full height rack door. The door is steel, with a perforated flat front surface. The perforation pattern extends from the bottom to the top of the door to enhance ventilation and provide some visibility into the rack. The non-acoustic door has a depth of about 134 mm (5.3 in.).

Rear Door Heat Exchanger

To lead away more heat, a special door that is named the Rear Door Heat Exchanger (#6858) is available. This door replaces the standard rear door on the rack. Copper tubes that are attached to the rear door circulate chilled water, provided by the customer. The chilled water removes heat from the exhaust air being blown through the servers and attachments that are mounted in the rack. With industry standard quick couplings, the water lines in the door attach to the customer-supplied secondary water loop.

For more information about planning for the installation of the IBM Rear Door Heat Exchanger, see [IBM Knowledge Center](#).

1.9.3 IBM 42U Slim Rack 7965-94Y

The 2.0-meter (79-inch) Model 7965-94Y is compatible with past and present IBM Power Systems servers and provides an excellent 19-inch rack enclosure for your data center. Its 600 mm (23.6 in.) width that is combined with its 1100 mm (43.3 in.) depth plus its 42 EIA enclosure capacity provides great footprint efficiency for your systems and allows it to be easily located on standard 24-inch floor tiles.

The IBM 42U Slim Rack has a lockable perforated front steel door, providing ventilation, physical security, and visibility of indicator lights in the installed equipment within. In the rear, either a lockable perforated rear steel door (#EC02) or a lockable Rear Door Heat Exchanger (RDHX)(1164-95X) is used. Lockable optional side panels (#EC03) increase the rack's aesthetics, help control airflow through the rack, and provide physical security. Multiple 42U Slim Racks can be bolted together to create a rack suite (#EC04).

Up to six optional 1U PDUs can be placed vertically in the sides of the rack. More PDUs can be located horizontally, but they each use 1U of space in this position.

1.9.4 1.8 Meter Rack (#0551)

The 1.8 Meter Rack (#0551) is a 36 EIA unit rack. The rack that is delivered as #0551 is the same rack that is delivered when you order the 7014-T00 rack. The included features might vary. Certain features that are delivered as part of the 7014-T00 must be ordered separately with #0551.

1.9.5 2.0 Meter Rack (#0553)

The 2.0 Meter Rack (#0553) is a 42 EIA unit rack. The rack that is delivered as #0553 is the same rack that is delivered when you order the 7014-T42 rack. The included features might vary. Certain features that are delivered as part of the 7014-T42 must be ordered separately with #0553.

1.9.6 Rack (#ER05)

This feature provides a 19-inch, 2.0-meter high rack with 42 EIA units of total space for installing a rack-mounted central electronics complex or expansion units. The 600 mm wide rack fits within a data center's 24-inch floor tiles and provides better thermal and cable management capabilities. The following features are required on #ER05:

- ▶ #EC01 front door
- ▶ #EC02 rear door or #EC05 Rear Door Heat Exchanger (RDHX) indicator

PDUs on the rack are optional. Each #7196 and #7189 PDU uses one of six vertical mounting bays. Each PDU beyond four uses 1U of rack space.

If ordering Power Systems equipment in an MES order, use the equivalent rack #ER05 instead of 7965-94Y so IBM Manufacturing can ship the hardware in the rack.

1.9.7 The AC power distribution unit and rack content

For rack models T00 and T42, 12-outlet PDUs are available, which include the AC PDUs #9188 and #7188 and the AC Intelligent PDU+ #5889 and #7109. The Intelligent PDUs+ (#5889 and #7109) are identical to #9188 and #7188 PDUs, but are equipped with one Ethernet port, one console serial port, and one RS232 serial port for power monitoring.

The PDUs have 12 client-usable IEC 320-C13 outlets. There are six groups of two outlets that are fed by six circuit breakers. Each outlet is rated up to 10 amps, but each group of two outlets is fed from one 15 amp circuit breaker.

High-function PDUs provide more electrical power per PDU and offer better “PDU footprint” efficiency. In addition, they are intelligent PDUs that provide insight to actual power usage by receptacle and also provide remote power on/off capability for easier support by individual receptacle. The new PDUs can be ordered as #EPTJ, #EPTL, #EPTN, and #EPTQ.

High-function PDU FCs are shown in Table 1-16.

Table 1-16 Available high-function PDUs

PDUs	1-phase or 3-phase depending on country wiring standards	3-phase 208 V depending on country wiring standards
Nine C19 receptacles	EPTJ	EPTL
Twelve C13 receptacles	EPTN	EPTQ

In addition, the following high-function PDUs were announced in October 2019:

- ▶ High Function 9xC19 PDU plus (#ECJJ):

This is an intelligent, switched 200-240 volt AC Power Distribution Unit (PDU) plus with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack

making the nine C19 receptacles easily accessible. For comparison, this is most similar to the earlier generation #EPTJ PDU.

▶ High Function 9xC19 PDU plus 3-Phase (#ECJL):

This is an intelligent, switched 208 volt 3-phase AC Power Distribution Unit (PDU) plus with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the nine C19 receptacles easily accessible. For comparison, this is most similar to the earlier generation #EPTL PDU.

▶ High Function 12xC13 PDU plus (#ECJN):

This is an intelligent, switched 200-240 volt AC Power Distribution Unit (PDU) plus with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the twelve C13 receptacles easily accessible. For comparison, this is most similar to the earlier generation #EPTN PDU.

▶ High Function 12xC13 PDU plus 3-Phase (#ECJQ):

This is an intelligent, switched 208 volt 3-phase AC Power Distribution Unit (PDU) plus with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the twelve C13 receptacles easily accessible. For comparison, this is most similar to the earlier generation #EPTQ PDU.

Table 1-17 lists the feature codes for the high-function PDUs announced in October 2019.

Table 1-17 High-function PDUs available after October 2019

PDUs	1-phase or 3-phase depending on country wiring standards	3-phase 208 V depending on country wiring standards
Nine C19 receptacles	ECJJ	ECJL
Twelve C13 receptacles	ECJN	ECJQ

Four PDUs can be mounted vertically in the back of the T00 and T42 racks. Figure 1-13 shows the placement of the four vertically mounted PDUs. In the rear of the rack, two more PDUs can be installed horizontally in the T00 rack and three in the T42 rack. The four vertical mounting locations are filled first in the T00 and T42 racks. Mounting PDUs horizontally uses 1U per PDU and reduces the space that is available for other racked components. When mounting PDUs horizontally, the preferred approach is to use fillers in the EIA units that are occupied by these PDUs to facilitate proper air-flow and ventilation in the rack.




Figure 1-13 Power distribution unit placement and view

The PDU receives power through a UTG0247 power-line connector. Each PDU requires one PDU-to-wall power cord. Various power cord features are available for various countries and applications by varying the PDU-to-wall power cord, which must be ordered separately. Each power cord provides the unique design characteristics for the specific power requirements. To match new power requirements and save previous investments, these power cords can be requested with an initial order of the rack or with a later upgrade of the rack features.

Table 1-18 shows the available wall power cord options for the PDU and intelligent power distribution unit (iPDU) features, which must be ordered separately.

Table 1-18 Wall power cord options for the power distribution unit and iPDU features

Feature code	Wall plug	Rated voltage (V AC)	Phase	Rated amperage	Geography
#6653	IEC 309, 3P+N+G, 16A	230	3	16 amps/phase	Internationally available
#6489	IEC309 3P+N+G, 32A	230	3	32 amps/phase	EMEA
#6654	NEMA L6-30	200-208, 240	1	24 amps	US, Canada, LA, Japan
#6655	RS 3750DP (watertight)	200-208, 240	1	24 amps	US, Canada, LA, Japan
#6656	IEC 309, P+N+G, 32A	230	1	24 amps	EMEA
#6657	PDL	230-240	1	32 amps	Australia, New Zealand
#6658	Korean plug	220	1	30 amps	North and South Korea
#6492	IEC 309, 2P+G, 60A	200-208, 240	1	48 amps	US, Canada, LA, Japan
#6491	IEC 309, P+N+G, 63A	230	1	63 amps	EMEA

Notes: Ensure that the appropriate power cord feature is configured to support the power that is being supplied. Based on the power cord that is used, the PDU can supply 4.8 - 19.2 kVA. The power of all the drawers plugged into the PDU must not exceed the power cord limitation.

The Universal PDUs are compatible with previous models.

To better enable electrical redundancy, each server has two power supplies that must be connected to separate PDUs, which are not included in the base order.

For maximum availability, a preferred approach is to connect power cords from the same system to two separate PDUs in the rack, and to connect each PDU to independent power sources.

For detailed power requirements and power cord details about the 7014 and 7965-94Y racks, see the IBM Knowledge Center at:

<https://www.ibm.com/support/knowledgecenter/POWER9/p9hdx/POWER9welcome.htm>

1.9.8 Rack-mounting rules

Consider the following primary rules when you mount the system into a rack:

- ▶ The system is placed at any location in the rack. For rack stability, start filling a rack from the bottom.
- ▶ Any remaining space in the rack can be used to install other systems or peripheral devices, if the maximum permissible weight of the rack is not exceeded and the installation rules for these devices are followed.
- ▶ Before placing the system into the service position, be sure to follow the rack manufacturer's safety instructions regarding rack stability.

Order information: The racking approach for the initial order must be either a 7014-T00, 7014-T42, or 7965-94Y. If an extra rack is required for I/O expansion drawers as an Manufacturing Equipment Specification (MES) to an existing system, either an #0551, #0553, or #ER05 rack must be ordered.

You must leave 2U of space at either the bottom or top of the rack, depending on the client's cabling preferences, to allow for cabling to exit the rack.

1.9.9 Useful rack additions

This section highlights several rack addition solutions for IBM Power Systems rack-based systems.

IBM System Storage 7226 Model 1U3 Multi-Media Enclosure

The IBM System Storage™ 7226 Model 1U3 Multi-Media Enclosure can accommodate up to two tape drives, two RDX removable disk drive docking stations, or up to four DVD-RAM drives.

The IBM System Storage 7226 Multi-Media Enclosure supports LTO Ultrium and DAT160 Tape technology, DVD-RAM, and RDX removable storage requirements on the following IBM systems:

- ▶ IBM POWER6™ processor-based systems
- ▶ IBM POWER7® processor-based systems
- ▶ IBM POWER8 processor-based systems
- ▶ IBM POWER9 processor-based systems

The IBM System Storage 7226 Multi-Media Enclosure offers an expansive list of drive feature options, as shown in Table 1-19.

Table 1-19 Supported drive features for the 7226-1U3

Feature code	Description	Status
#5619	DAT160 SAS Tape Drive	Available
#EU16	DAT160 USB Tape Drive	Available
#1420	DVD-RAM SAS Optical Drive	Available
#1422	DVD-RAM Slim SAS Optical Drive	Available
#5762	DVD-RAM USB Optical Drive	Available
#5763	DVD Front USB Port Sled with DVD-RAM USB Drive	Available

Feature code	Description	Status
#5757	DVD RAM Slim USB Optical Drive	Available
#8248	LTO Ultrium 5 Half High Fibre Tape Drive	Available
#8241	LTO Ultrium 5 Half High SAS Tape Drive	Available
#8348	LTO Ultrium 6 Half High Fibre Tape Drive	Available
#8341	LTO Ultrium 6 Half High SAS Tape Drive	Available
#EU03	RDX 3.0 Removable Disk Docking Station	Available

Here are the option descriptions:

- ▶ **DAT160 160 GB Tape Drives:** With SAS or USB interface options and a data transfer rate up to 12 MBps (assumes 2:1 compression), the DAT160 drive is read/write compatible with DAT160, and DDS4 data cartridges.
- ▶ **LTO Ultrium 5 Half-High 1.5 TB SAS and Tape Drive:** With a data transfer rate up to 280 MBps (assuming a 2:1 compression), the LTO Ultrium 5 drive is read/write compatible with LTO Ultrium 5 and 4 data cartridges, and read-only compatible with Ultrium 3 data cartridges. By using data compression, an LTO-5 cartridge can store up to 3 TB of data.
- ▶ **LTO Ultrium 6 Half-High 2.5 TB SAS and Tape Drive:** With a data transfer rate up to 320 MBps (assuming a 2.5:1 compression), the LTO Ultrium 6 drive is read/write compatible with LTO Ultrium 6 and 5 media, and read-only compatibility with LTO Ultrium 4. By using data compression, an LTO-6 cartridge can store up to 6.25 TB of data.
- ▶ **DVD-RAM:** The 9.4 GB SAS Slim Optical Drive with an SAS and USB interface option is compatible with most standard DVD disks.
- ▶ **RDX removable disk drives:** The RDX USB docking station is compatible with most RDX removable disk drive cartridges when it is used in the same operating system. The 7226 offers the following RDX removable drive capacity options:
 - 500 GB (#1107)
 - 1.0 TB (#EU01)
 - 2.0 TB (#EU2T)

Removable RDX drives are in a rugged cartridge that inserts in to an RDX removable (USB) disk docking station (#1103 or #EU03). RDX drives are compatible with docking stations, which are installed internally in IBM POWER6, IBM POWER6+, POWER7, IBM POWER7+, POWER8, and POWER9 processor-based servers, where applicable.

Media that is used in the 7226 DAT160 SAS and USB tape drive features are compatible with DAT160 tape drives that are installed internally in IBM POWER6, POWER6+, POWER7, POWER7+, POWER8, and POWER9 processor-based servers.

Media that is used in LTO Ultrium 5 Half-High 1.5 TB tape drives are compatible with Half High LTO5 tape drives that are installed in the IBM TS2250 and TS2350 external tape drives, IBM LTO5 tape libraries, and half-high LTO5 tape drives that are installed internally in IBM POWER6, POWER6+, POWER7, POWER7+, POWER8, and POWER9 processor-based servers.

Figure 1-14 shows the IBM System Storage 7226 Multi-Media Enclosure.



Figure 1-14 IBM System Storage 7226 Multi-Media Enclosure

The IBM System Storage 7226 Multi-Media Enclosure offers customer-replaceable unit (CRU) maintenance service to help make the installation or replacement of new drives efficient. Other 7226 components are also designed for CRU maintenance.

The IBM System Storage 7226 Multi-Media Enclosure is compatible with most IBM POWER6, POWER6+, POWER7, POWER7+, POWER8, and POWER9 processor-based systems that offer current level AIX, IBM i, and Linux operating systems.

Unsupported: IBM i does not support 7226 USB tape devices but does support 7226 USB DVD devices.

For a complete list of host software versions and release levels that support the IBM System Storage 7226 Multi-Media Enclosure, see [System Storage Interoperation Center \(SSIC\)](#).

Note: Any of the existing 7216-1U2, 7216-1U3, and 7214-1U2 multimedia drawers are also supported.

Flat panel display options

The IBM 7316 Model TF4 is a rack-mountable flat panel console kit that can also be configured with the tray pulled forward and the monitor folded up, providing full viewing and keying capability for the HMC operator.

The Model TF4 is a follow-on product to the Model TF3 and offers the following features:

- ▶ A slim, sleek, and lightweight monitor design that occupies only 1U (1.75 in.) in a 19-inch standard rack.
- ▶ A 18.5-inch (409.8 mm x 230.4 mm) flat panel TFT monitor with truly accurate images and virtually no distortion.
- ▶ The ability to mount the IBM Travel Keyboard in the 7316-TF4 rack keyboard tray.

- Support for the IBM 1x8 Rack Console Switch (#4283) IBM Keyboard/Video/Mouse (KVM) switches.

#4283 is a 1x8 Console Switch that fits in the 1U space behind the TF4. It is a CAT5 based switch containing eight rack interface (ARI) ports for connecting either PS/2 or USB console switch cables. It supports chaining of servers that use IBM Conversion Options switch cable (#4269). This feature provides four cables that connect a KVM switch to a system, or can be used in a daisy-chain scenario to connect up to 128 systems to a single KVM switch. It also supports server-side USB attachments.

1.9.10 Original equipment manufacturer racks

The system can be installed in a suitable original equipment manufacturer (OEM) rack if that the rack conforms to the EIA-310-D standard for 19-inch racks. This standard is published by the Electrical Industries Alliance. For more information, see the IBM Knowledge Center at:

<https://www.ibm.com/support/knowledgecenter/POWER9/p9hdx/POWER9welcome.htm>

The website mentions the following key points:

- The front rack opening must be 451 mm wide \pm 0.75 mm (17.75 in. \pm 0.03 in.), and the rail-mounting holes must be 465 mm \pm 0.8 mm (18.3 in. \pm 0.03 in.) apart on-center (that is, the horizontal width between the vertical columns of holes on the two front-mounting flanges and on the two rear-mounting flanges). Figure 1-15 is a top view showing the specification dimensions.




Figure 1-15 Top view of the rack specification dimensions (not specific to IBM)

- ▶ The vertical distance between the mounting holes must consist of sets of three holes spaced (from bottom to top) 15.9 mm (0.625 in.), 15.9 mm (0.625 in.), and 12.67 mm (0.5 in.) on-center, making each three-hole set of vertical hole spacing 44.45 mm (1.75 in.) apart on center. Rail-mounting holes must be 7.1 mm ± 0.1 mm (0.28 in. ± 0.004 in.) in diameter. Figure 1-16 shows the top front specification dimensions.




Figure 1-16 Rack specification dimensions: Top front view

1.10 Hardware Management Console

This section describes the Hardware Management Consoles (HMCs) that are available for Power Systems servers.

1.10.1 New features

Here are some of the new features of the HMCs:

- ▶ New HMCs are now based on systems with POWER processors.
- ▶ Intel x86-based HMCs are supported but are no longer available.
- ▶ Virtual HMCs (vHMCs) are available for x86 and Power Systems virtual environments.

1.10.2 Hardware Management Console overview

Administrators can use the HMC, which is a dedicated appliance, to configure and manage system resources on IBM Power Systems servers. GUI, command-line interface (CLI), or REST API interfaces are available. The HMC provides basic virtualization management support for configuring logical partitions (LPARs) and dynamic resource allocation, including processor and memory settings for selected Power Systems servers.

The HMC also supports advanced service functions, including guided repair and verification, concurrent firmware updates for managed systems, and around-the-clock error reporting through IBM Electronic Service Agent (ESA) for faster support.

The HMC management features help improve server usage, simplify systems management, and accelerate provisioning of server resources by using PowerVM virtualization technology.

The HMC is available as a hardware appliance or as a vHMC. The Power S922, Power S914, and Power S924 servers support several service environments, including attachment to one or more HMCs or vHMCs. This is the default configuration for servers supporting multiple logical partitions with dedicated resource or virtual I/O.

Here are the HMCs for various hardware architectures:

- ▶ X86-based HMCs: 7042-CR7, CR8, or CR9
- ▶ POWER based HMC: 7063-CR1
- ▶ vHMC on x86 or Power Systems LPARs

Hardware support for customer replaceable units (CRUs) come standard with the HMC. In addition, users can upgrade this support level to IBM onsite support to be consistent with other Power Systems servers.

Note:

- ▶ An HMC or vHMC is required for the Power S922, Power S914, and Power S924 servers.
- ▶ Integrated Virtual Management (IVM) is no longer supported.

For more information about vHMC, see [Virtual HMC Appliance \(vHMC\) Overview](#).

Figure 1-17 shows the HMC model selections and tier updates.

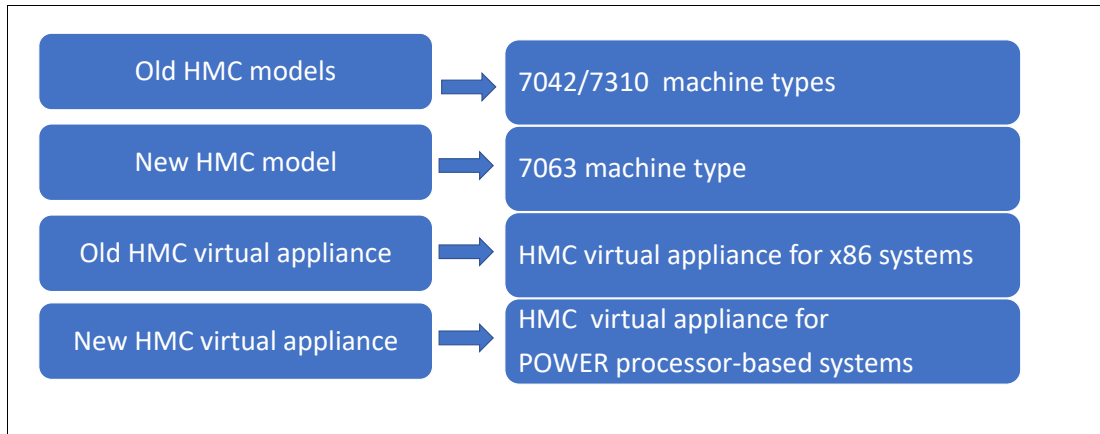


Figure 1-17 HMC model selections

1.10.3 Hardware Management Console code level

The HMC code must be running at Version 9 Release 1 modification 920 (V9R1M920) or later when you use the HMC with the Power S922, Power S914, and Power S924 servers.

If you are attaching an HMC to a new server or adding a function to an existing server that requires a firmware update, the HMC machine code might need to be updated to support the firmware level of the server. In a dual-HMC configuration, both HMCs must be at the same version and release of the HMC code.

To determine the HMC machine code level that is required for the firmware level on any server, go to [Fix Level Recommendation Tool \(FLRT\)](#) on or after the planned availability date for this product.

FLRT identifies the correct HMC machine code for the selected system firmware level.

Note:

- ▶ Access to firmware and machine code updates is conditional on entitlement and license validation in accordance with IBM policy and practice. IBM might verify entitlement through customer number, serial number electronic restrictions, or any other means or methods that are employed by IBM at its discretion.
- ▶ HMC V9 supports only the Enhanced+ version of the GUI. The Classic version is no longer available.
- ▶ HMC V9R1.911.0 added support for managing IBM OpenPOWER systems. The same HMC that is used to manage flexible service processor (FSP)-based enterprise systems can manage the baseboard management controller (BMC) based Power Systems AC and Power Systems LC servers. This support provides a consistent and consolidated hardware management solution.
- ▶ HMC V9 supports connections to servers that are based on IBM servers that are based on POWER9, POWER8, and POWER7 processors. There is no support in this release for servers that are based on POWER6 processors or earlier.

1.10.4 Two architectures of Hardware Management Console

There are now two options for the HMC hardware: The earlier Intel-based HMCs, and the newer HMCs that are based on an IBM POWER8 processor. The x86-based HMCs are no longer available for ordering, but are supported as an option for managing the Power S922, Power S914, and Power S924 servers.

You may use either architecture to manage the servers. You also may use one Intel-based HMC and one POWER8 based HMC if the software is at the same level.

It is a preferred practice to use the new POWER8 processor-based consoles for server management.

Intel-based HMCs

HMCs that are based on Intel processors that support V9 code are:

- ▶ 7042-CR9
- ▶ 7042-CR8
- ▶ 7042-CR7

7042-CR6 and earlier HMCs are not supported by the Power S922, Power S914, and Power S924 servers.

The 7042-CR9 has the following specifications:

- ▶ 2.4 GHz Intel Xeon processor E5-2620 V3
- ▶ 16 GB (1 x 16 GB) of 2.133 GHz DDR4 system memory
- ▶ 500 GB SATA SFF HDD
- ▶ SATA CD-RW and DVD-RAM
- ▶ Four Ethernet ports
- ▶ Six USB ports (two front and four rear)
- ▶ One PCIe slot

POWER8 processor-based HMC

The POWER8 processor-based HMC is machine type and model 7063-CR1. It has the following specifications:

- ▶ 1U base configuration
- ▶ IBM POWER8 120 W 6-core CPU
- ▶ 32 GB (4 x 8 GB) of DDR4 system memory
- ▶ Two 2-TB SATA LFF 3.5-inch HDD RAID 1 disks
- ▶ Rail bracket option for round hole rack mounts
- ▶ Two USB 3.0 hub ports in the front of the server
- ▶ Two USB 3.0 hub ports in the rear of the server
- ▶ Redundant 1 kW power supplies
- ▶ Four 10-Gb Ethernet Ports (RJ-45) (10 Gb/1 Gb/100 Mb)
- ▶ One 1-Gb Ethernet port for management (BMC)

All future HMC development will be done for the POWER8 processor-based 7063-CR1 model and its successors.

Note: System administrators can remotely start or stop a 7063-CR1 HMC by using `ipmitool` or WebUI.

1.10.5 Hardware Management Console connectivity to POWER9 processor-based systems

POWER9 processor-based servers and their predecessor systems that are managed by an HMC require Ethernet connectivity between the HMC and the server's service processor. Additionally, to perform an operation on an LPAR, initiate Live Partition Mobility (LPM), or perform PowerVM Active Memory Sharing operations, you must have an Ethernet link to the managed partitions. A minimum of two Ethernet ports are needed on the HMC to provide such connectivity.

For the HMC to communicate properly with the managed server, eth0 of the HMC must be connected to either the HMC1 or HMC2 ports of the managed server, although other network configurations are possible. You may attach a second HMC to the remaining HMC port of the server for redundancy. The two HMC ports must be addressed by two separate subnets.

Figure 1-18 shows a simple network configuration to enable the connection from the HMC to the server and to allow for dynamic LPAR operations. For more information about HMC and the possible network connections, see *IBM Power Systems HMC Implementation and Usage Guide*, SG24-7491.

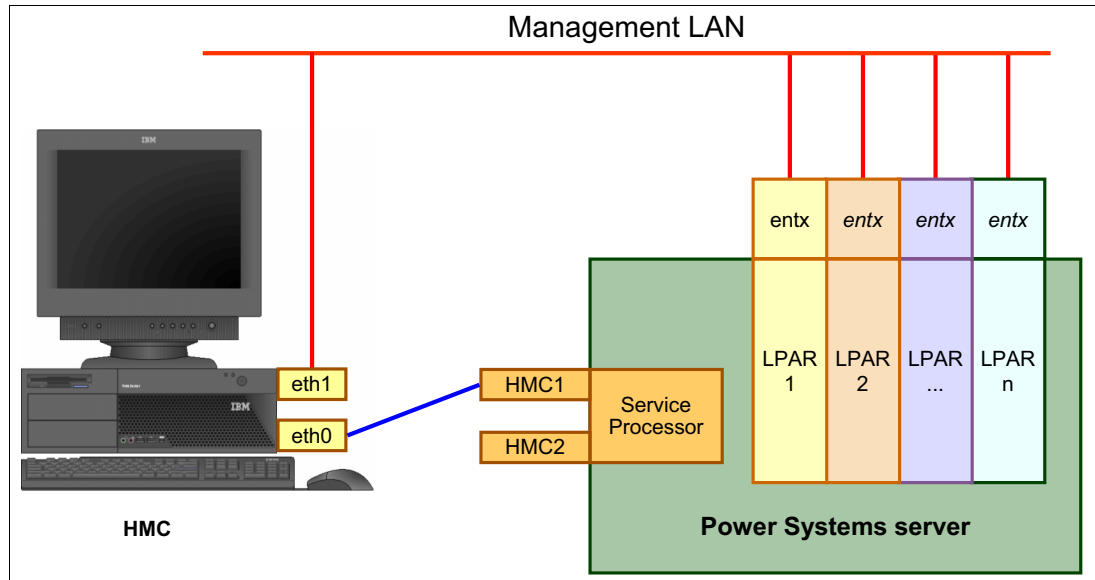


Figure 1-18 Network connections from the HMC to service processor and LPARs

By default, the service processor HMC ports are configured for dynamic IP address allocation. The HMC can be configured as a DHCP server, providing an IP address at the time that the managed server is powered on. In this case, the FSP is allocated an IP address from a set of address ranges that is predefined in the HMC software.

If the service processor of the managed server does not receive a DHCP reply before timeout, predefined IP addresses are set up on both ports. Static IP address allocation is also an option and can be configured by using the Advanced System Management Interface (ASMI) menus.

Notes: The two service processor HMC ports have the following features:

- ▶ 1 Gbps connection speed.
- ▶ Visible only to the service processor. They can be used to attach the server to an HMC or to access the ASMI options from a client directly from a client web browser.
- ▶ Use the following network configuration if no IP addresses are set:
 - Service processor eth0 (HMC1 port): 169.254.2.147 with netmask 255.255.255.0
 - Service processor eth1 (HMC2 port): 169.254.3.147 with netmask 255.255.255.0

1.10.6 High availability Hardware Management Console configuration

The HMC is an important hardware component. Although Power Systems servers and their hosted partitions can continue to operate when the managing HMC becomes unavailable, certain operations, such as dynamic LPAR, partition migration that uses PowerVM LPM, or the creation of a partition, cannot be performed without the HMC. Power Systems servers may have two HMCs attached to a system, which provides redundancy if one of the HMCs is unavailable.

To achieve HMC redundancy for a POWER9 processor-based server, the server must be connected to two HMCs:

- ▶ The HMCs must be running the same level of HMC code.
- ▶ The HMCs must use different subnets to connect to the service processor.
- ▶ The HMCs must be able to communicate with the server's partitions over a public network to allow for full synchronization and functionality.

Figure 1-19 shows one possible highly available HMC configuration that manages two servers. Each HMC is connected to one FSP port of each managed server.

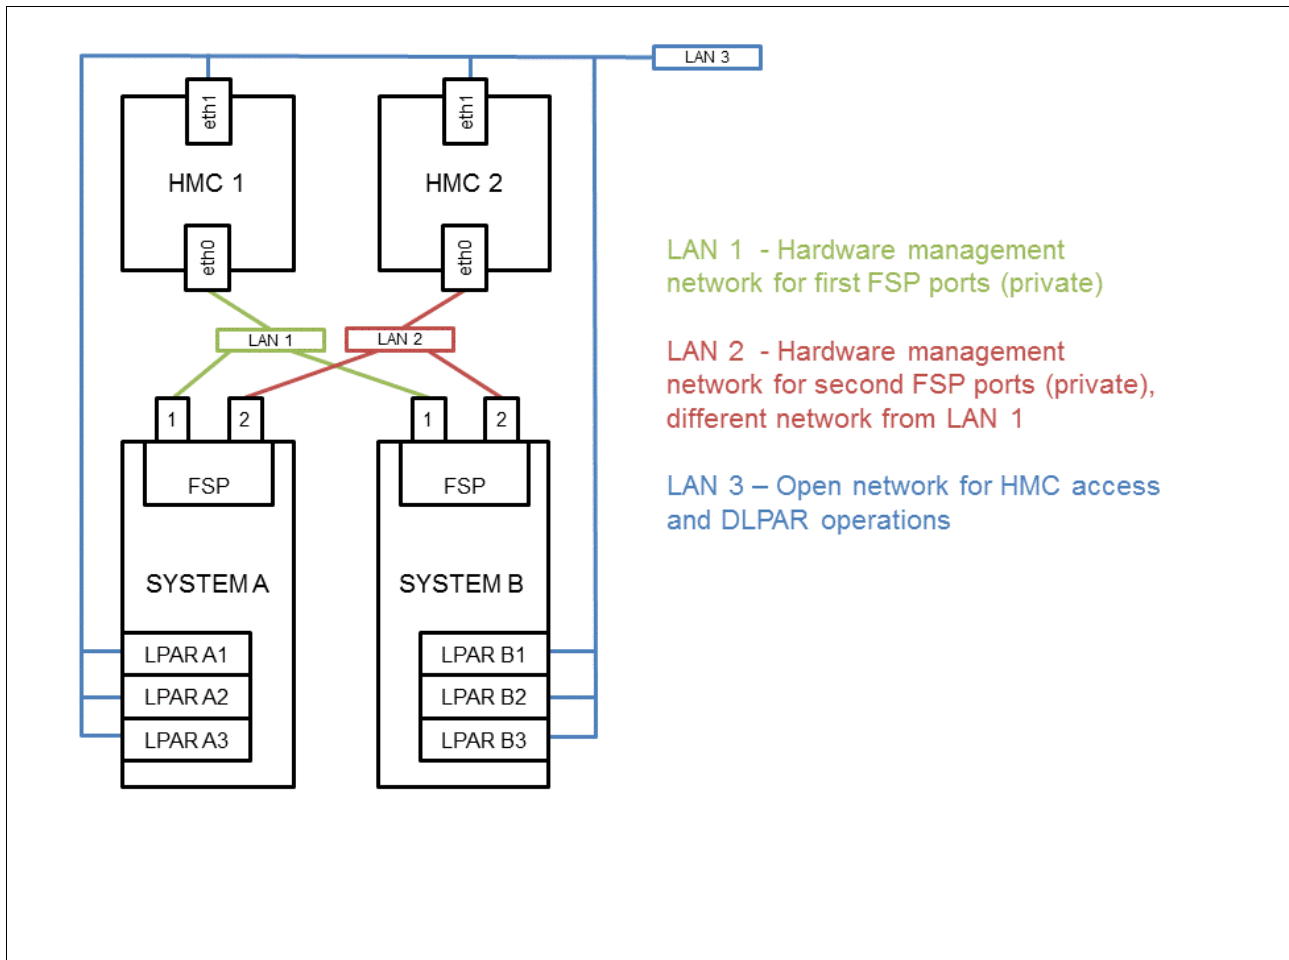


Figure 1-19 Highly available HMC networking example.

For simplicity, only the hardware management networks (LAN1 and LAN2) are highly available. However, the open network (LAN3) can be made highly available by using a similar concept and adding a second network between the partitions and HMCs.

For more information about redundant HMCs, see *IBM Power Systems HMC Implementation and Usage Guide*, SG24-7491.



Architecture and technical overview

This chapter describes the overall system architecture for the IBM Power System S922 (9009-22A), the IBM Power System S914 (9009-41A), and the IBM Power System S924 (9009-42A) servers. The bandwidths that are provided throughout the section are theoretical maximums that are used for reference.

The speeds that are shown are at an individual component level. Multiple components and application implementation are key to achieving the best performance.

Always do the performance sizing at the application workload environment level and evaluate performance by using real-world performance measurements and production workloads.

Figure 2-2 on page 61, Figure 2-3 on page 62, and Figure 2-1 on page 60 show the general architecture of the Power S922, Power S914, and Power S924 servers.

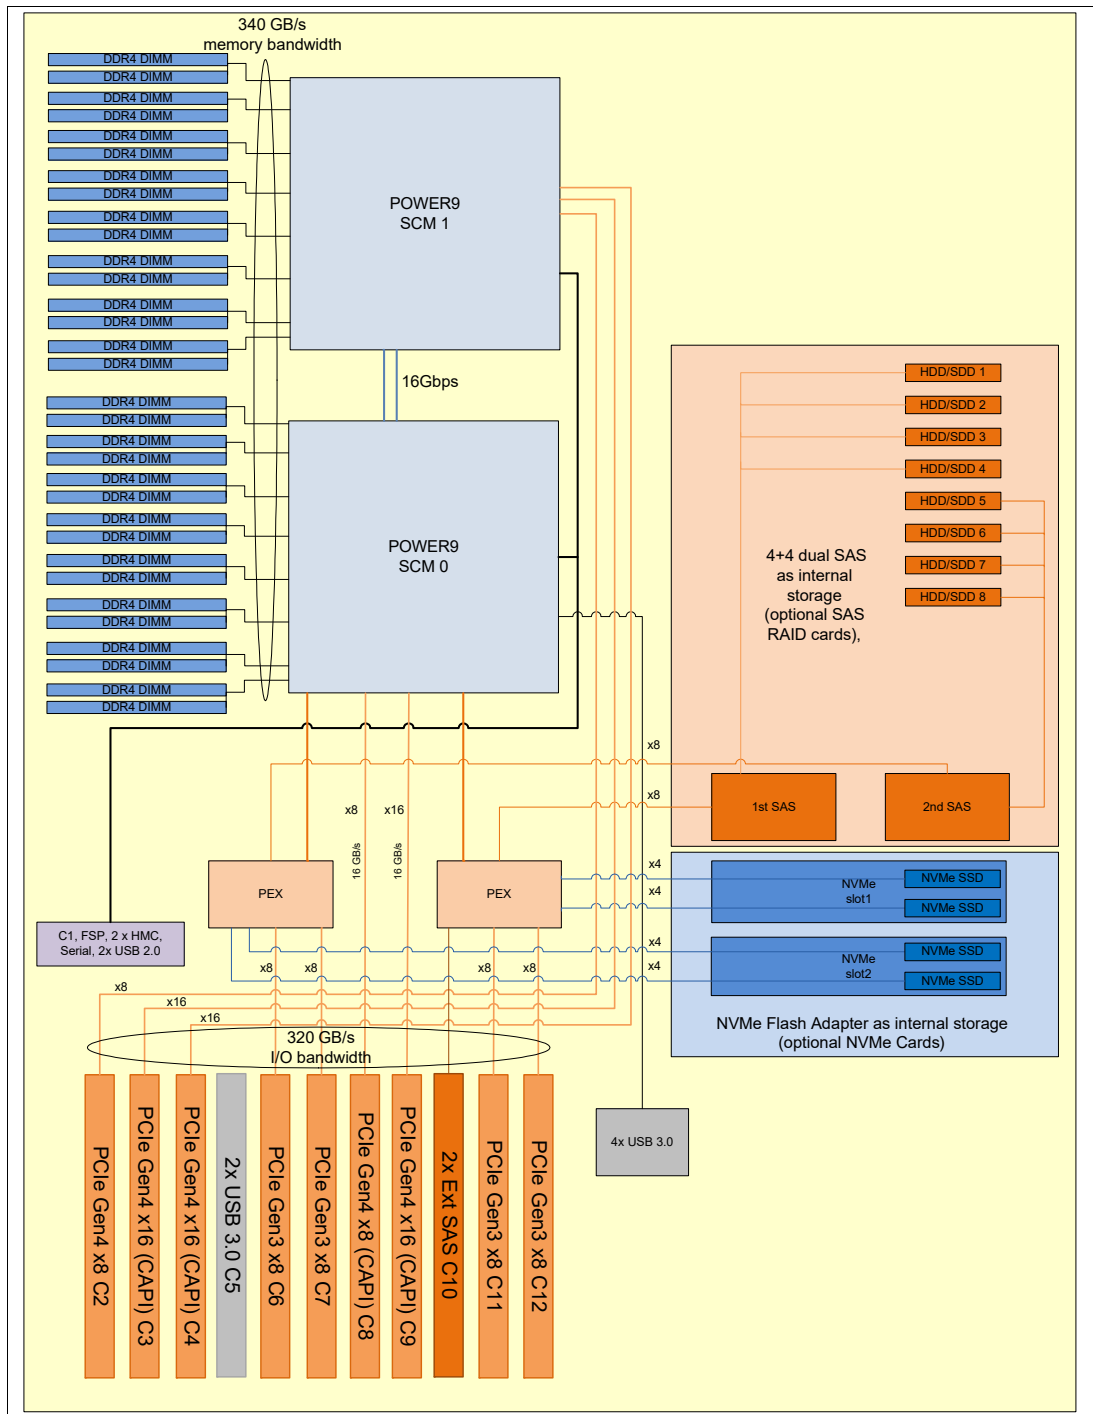


Figure 2-1 Power S922 logical system diagram

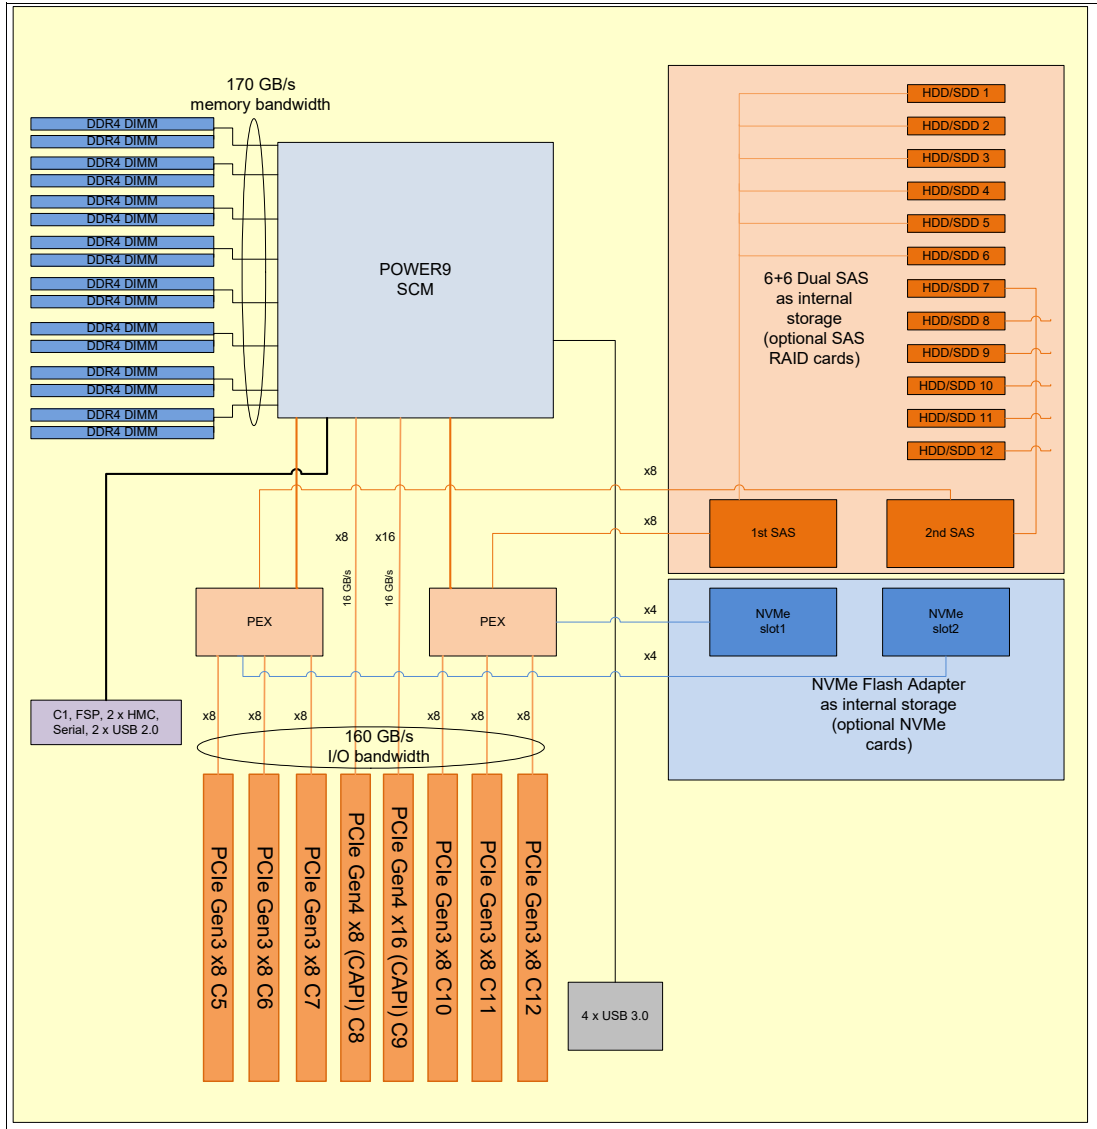


Figure 2-2 Power S914 logical system diagram

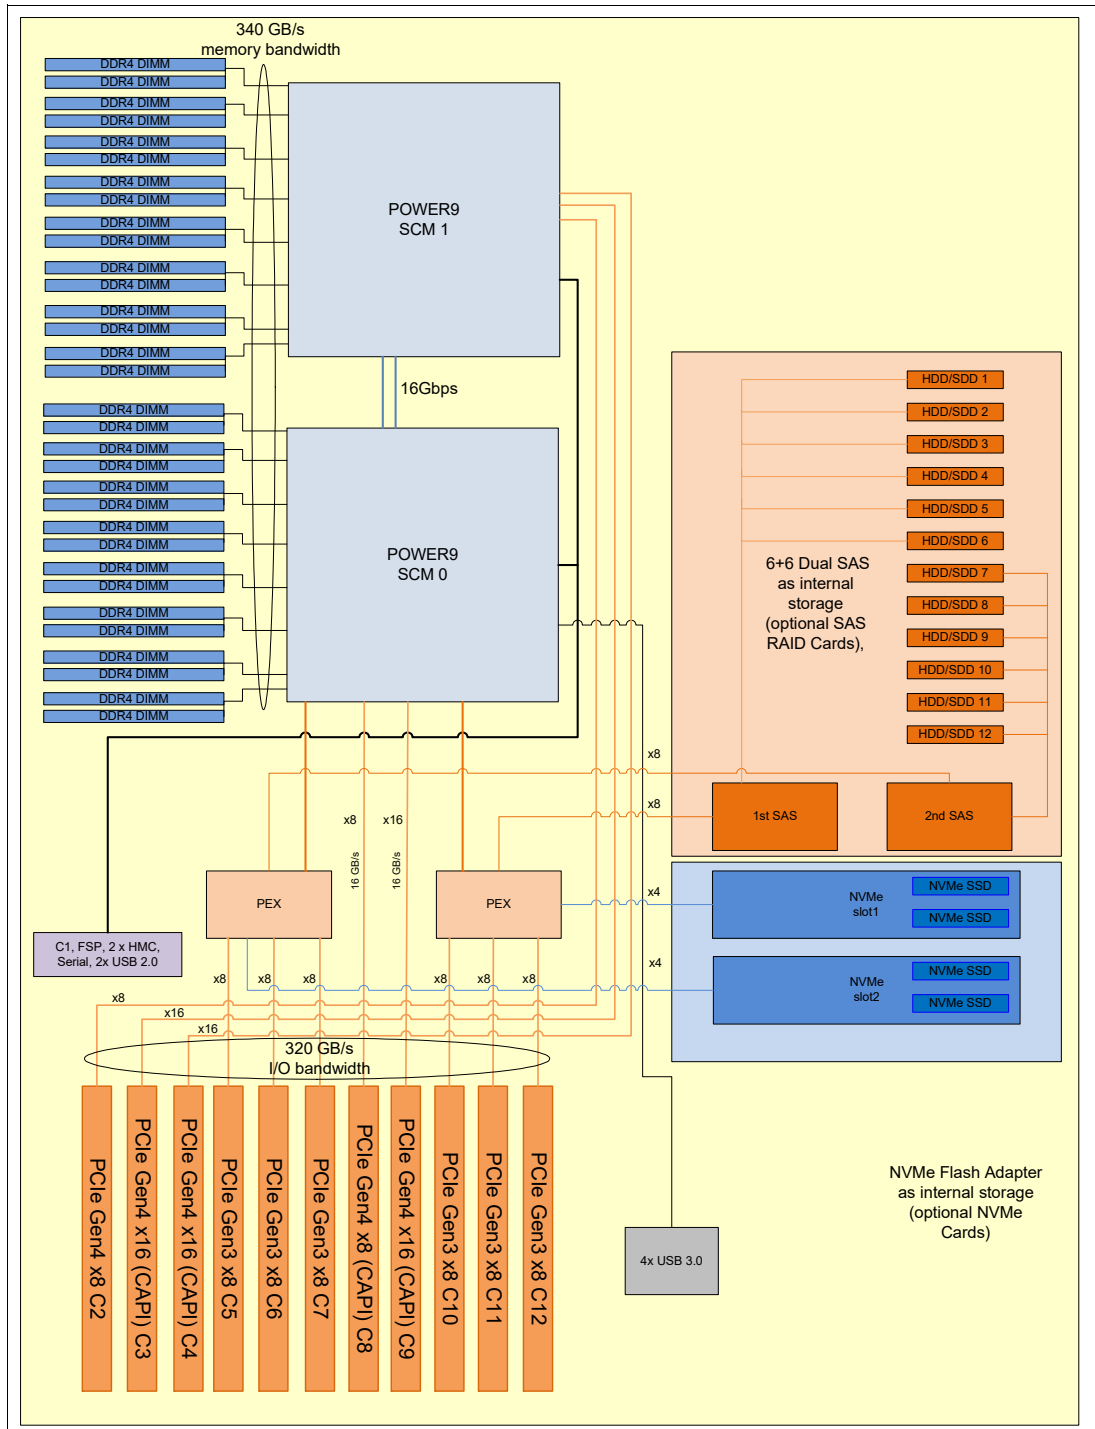


Figure 2-3 Power S924 logical system diagram

2.1 The IBM POWER9 processor

This section introduces the latest processor in the IBM Power Systems product family, and describes its main characteristics and features in general.

2.1.1 POWER9 processor overview

This is the architectural design of the POWER9 processor.

The servers are offered with various numbers of cores that are activated, and a selection of clock frequencies, so IBM can make offerings at several price points, and customers can select a particular server (or servers) to fit their budget and performance requirements.

The POWER9 processor is single-chip modules (SCMs) manufactured on the IBM 14-nm FinFET Silicon-On-Insulator (SOI) architecture. Each module is 68.5 mm x 68.5 mm and contains 8 billion transistors.

As shown in Figure 2-4, the chip contains 24 cores, two memory controllers, Peripheral Component Interconnect Express (PCIe) Gen4 I/O controllers, and an interconnection system that connects all components within the chip at 7 TBps. Each core has 512 KB of L2 cache, and 10 MB of L3 embedded DRAM (eDRAM). The interconnect also extends through module and system board technology to other POWER9 processors in addition to DDR4 memory and various I/O devices.

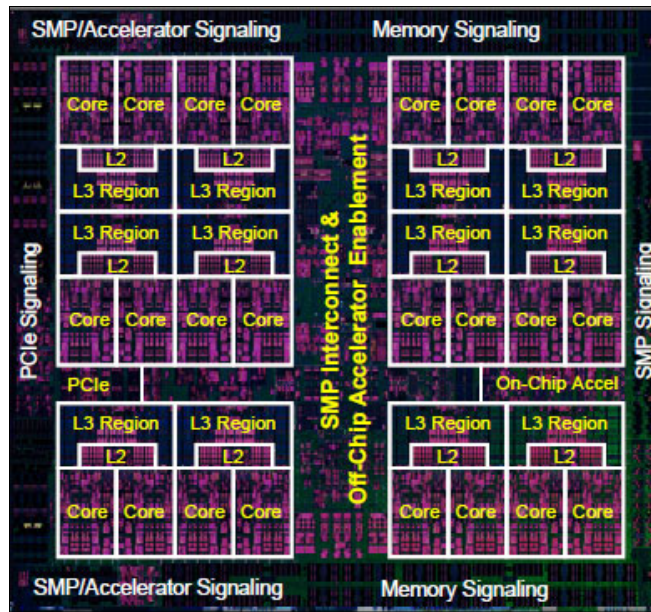


Figure 2-4 The 24-core POWER9 processor

POWER9 processor has eight memory channels, and each channel supports up to two DDR4 DIMM slots. The Power S914 server can support up to 1 TB of memory, and the Power S922 and Power S924 servers in a two-SCM configuration can support up to 4 TB of memory.

Limitation: The Power S914 server in a 4-core configuration is limited to 64 GB of memory.

2.1.2 POWER9 processor features

The POWER9 chip provides an embedded algorithm for the following features:

- ▶ External Interrupt Virtualization Engine. Reduces the code impact/path length and improves performance compared to the previous architecture.
- ▶ Gzip compression and decompression.
- ▶ PCIe Gen4 support.
- ▶ Two memory controllers that support direct-attached DDR4 memory.
- ▶ Cryptography: Advanced encryption standard (AES) engine.
- ▶ Random number generator (RNG).
- ▶ Secure Hash Algorithm (SHA) engine: SHA-1, SHA-256, and SHA-512, and Message Digest 5 (MD5).
- ▶ IBM Data Mover Tool.

Table 2-1 provides a summary of the POWER9 processor technology.

Note: The total values represent the maximum of 12 cores for the POWER9 architecture. Servers that are discussed in this paper have a maximum of 16 or 20 cores per module.

Table 2-1 Summary of the POWER9 processor technology

Technology	POWER9 processor
Die size	68.5 mm × 68.5 mm
Fabrication technology	<ul style="list-style-type: none">▶ 14-nm lithography▶ Copper interconnect▶ SOI▶ eDRAM
Maximum processor cores	12
Maximum execution threads core/module	8/96
Maximum L2 cache core/module	512 KB/6 MB
Maximum On-chip L3 cache core/module	10 MB/120 MB
Number of transistors	8 billion
Compatibility	With prior generation of POWER processor

2.1.3 POWER9 processor core

The POWER9 processor core is a 64-bit implementation of the IBM Power Instruction Set Architecture (ISA) Version 3.0, and has the following features:

- ▶ Multi-threaded design, which is capable of up to eight-way simultaneous multithreading (SMT)
- ▶ 64 KB, eight-way set-associative L1 instruction cache
- ▶ 64 KB, eight-way set-associative L1 data cache
- ▶ Enhanced prefetch, with instruction speculation awareness and data prefetch depth awareness

- ▶ Enhanced branch prediction that uses both local and global prediction tables with a selector table to choose the best predictor
- ▶ Improved out-of-order execution
- ▶ Two symmetric fixed-point execution units
- ▶ Two symmetric load/store units and two load units, all four of which can also run simple fixed-point instructions
- ▶ An integrated, multi-pipeline vector-scalar floating point unit for running both scalar and SIMD-type instructions, including the Vector Multimedia eXtension (VMX) instruction set and the improved Vector Scalar eXtension (VSX) instruction set, which is capable of up to 16 floating point operations per cycle (eight double precision or 16 single precision)
- ▶ In-core AES encryption capability
- ▶ Hardware data prefetching with 16 independent data streams and software control
- ▶ Hardware decimal floating point (DFP) capability

More information about Power ISA Version 3.0, see [OpenPOWER: IBM Power ISA Version 3.0B](#).

Figure 2-5 shows a picture of the POWER9 core, with some of the functional units highlighted.




Figure 2-5 POWER9 SMT4 processor core

2.1.4 Simultaneous multithreading

POWER9 processor advancements in multi-core and multi-thread scaling are remarkable. A significant performance opportunity comes from parallelizing workloads to enable the full potential of the microprocessor, and the large memory bandwidth. Application scaling is influenced by both multi-core and multi-thread technology.

SMT enables a single physical processor core to simultaneously dispatch instructions from more than one hardware thread context. With SMT, each POWER9 core can present eight hardware threads. Because there are multiple hardware threads per physical processor core, more instructions can run at the same time. SMT is primarily beneficial in commercial environments where the speed of an individual transaction is not as critical as the total number of transactions that are performed. SMT typically increases the throughput of workloads with large or frequently changing working sets, such as database servers and web servers.

Table 2-2 shows a comparison between the different POWER processors in terms of SMT capabilities that are supported by each processor architecture.

Table 2-2 SMT levels that are supported by POWER processors

Technology	Cores/system	Maximum SMT mode	Maximum hardware threads per partition
IBM POWER4	32	Single thread	32
IBM POWER5	64	SMT2	128
IBM POWER6	64	SMT2	128
IBM POWER7	256	SMT4	1024
IBM POWER8	192	SMT8	1536
IBM POWER9	192	SMT8	1535

2.1.5 Processor feature codes

The Power S922 (9009-22A) server is a 2-socket server with up to 20 cores. A system can be ordered with a single processor, and a second processor can be added as a miscellaneous execution system (MES) upgrade.

Table 2-3 shows the processor feature codes (FCs) for the Power S922 server.

Table 2-3 Processor feature codes specification for the Power S922 server

Number of cores	Frequency	Feature code
4 cores	2.8 - 3.8 GHz	#EP16
8 cores	3.4 - 3.9 GHz	#EP18
10 cores	2.9 - 3.8 GHz	#EP19

The Power S914 (9009-41A) server is the entry server that supports a one-processor socket with up to eight cores.

Table 2-4 shows the processor FCs for the Power S914 server.

Table 2-4 Processor feature codes specification for the Power S914 server

Number of cores	Frequency	Feature code
4 cores	2.3 - 3.8 GHz	#EP10
6 cores	2.3 - 3.8 GHz	#EP11
8 cores	2.8 - 3.8 GHz	#EP12

The Power S924 (9009-42A) server is a powerful 2-socket server with up to 24 cores. A system can be ordered with a single processor and a second processor can be added as an MES upgrade.

Table 2-5 shows the processor FCs for the Power S924 server.

Table 2-5 Processor feature codes specification for the Power S924 server

Number of cores	Frequency	Feature code
8 cores	3.8 - 4.00 GHz	#EP1E
10 cores	3.5 - 3.9 GHz	#EP1F
11 cores	3.45 - 3.9 GHz	#EP1H
12 cores	3.4 - 3.9 GHz	#EP1G

2.1.6 Memory access

The scale-out machines use industrial standard DDR4 DIMMs. Each POWER9 module has two memory controllers, which are connected to eight memory channels. Each memory channel can support up to two DIMMs. A single POWER9 module can support a maximum of 16 DDR4 DIMMs.

The speed of the memory depends on the DIMM size and placement.

Table 2-6 shows the DIMM speeds.

Table 2-6 DIMM speed

Registered DIMM (RDIMM) size	Mbps (1 DIMM per port)	Mbps (2 DIMMs per port)
8 GB	2666	2133
16 GB	2666	2133
32 GB	2400	2133
64 GB	2400	2133
128 GB	2400	2133

The Power S914 server can support up to 1 TB of memory. The Power S922 and Power S924 servers with a two-SCM configuration can operate up to 4 TB of memory.

Figure 2-6 shows an overview of the POWER9 direct attach memory.




Figure 2-6 Overview of POWER9 direct attach memory

2.1.7 On-chip L3 cache innovation and intelligent caching

Similar to POWER8, the POWER9 processor uses a breakthrough in material engineering and microprocessor fabrication to implement the L3 cache in eDRAM and place it on the processor die. The L3 cache is critical to a balanced design, as is the ability to provide good signaling between the L3 cache and other elements of the hierarchy, such as the L2 cache or SMP interconnect.

The on-chip L3 cache is organized into separate areas with differing latency characteristics. Each processor core is associated with a fast 10 MB local region of L3 cache (FLR-L3), but also has access to other L3 cache regions as a shared L3 cache. Additionally, each core can negotiate to use the FLR-L3 cache that is associated with another core, depending on the reference patterns. Data can also be cloned and stored in more than one core's FLR-L3 cache, again depending on the reference patterns. This intelligent cache management enables the POWER9 processor to optimize the access to L3 cache lines and minimize overall cache latencies.

Here are the L3 cache features on the POWER9 processor:

- ▶ Private 10-MB L3 cache/shared L3.1.
- ▶ 20-way set associative.
- ▶ 128-byte cache lines with 64-byte sector support.
- ▶ 10 EDRAM banks (interleaved for access overlapping).
- ▶ 64-byte wide data bus to L2 for reads.
- ▶ 64-byte wide data bus from L2 for L2 castouts.
- ▶ Eighty 1 Mb EDRAM macros that are configured in 10 banks, with each bank having a 64-byte wide data bus.
- ▶ All cache accesses have the same latency.
- ▶ 20-way directory that is organized as four banks, with up to four reads or two reads and two writes every two processor clock cycles to differing banks.

2.1.8 Hardware transactional memory

Transactional memory is an alternative to lock-based synchronization. It attempts to simplify parallel programming by grouping read and write operations and running them like a single operation. Transactional memory is like database transactions where all shared memory accesses and their effects are either committed together or discarded as a group. All threads can enter the critical region simultaneously. If there are conflicts in accessing the shared memory data, threads try accessing the shared memory data again or are stopped without updating the shared memory data. Therefore, transactional memory is also called a lock-free synchronization. Transactional memory can be a competitive alternative to lock-based synchronization.

Transactional memory provides a programming model that makes parallel programming easier. A programmer delimits regions of code that access shared data, and the hardware runs these regions atomically and in isolation, buffering the results of individual instructions and retrying execution if isolation is violated. Generally, transactional memory enables programs to use a programming style that is close to coarse-grained locking to achieve performance that is close to fine-grained locking.

Most implementations of transactional memory are based on software. The POWER9 processor-based systems provide a hardware-based implementation of transactional memory that is more efficient than the software implementations and requires no interaction with the processor core, therefore enabling the system to operate in maximum performance.

2.1.9 Coherent Accelerator Processor Interface 2.0

IBM Coherent Accelerator Processor Interface (CAPI) 2.0 is the evolution of CAPI and defines a coherent accelerator interface structure for attaching special processing devices to the POWER9 processor bus. As with the original CAPI, CAPI2 can attach accelerators that have coherent shared memory access with the processors in the server and share full virtual address translation with these processors by using standard PCIe Gen4 buses with twice the bandwidth compared to the previous generation.

Applications can have customized functions in Field Programmable Gate Arrays (FPGAs) and queue work requests directly into shared memory queues to the FPGA. Applications can also have customized functions by using the same effective addresses (pointers) they use for any threads running on a host processor. From a practical perspective, CAPI enables a specialized hardware accelerator to be seen as an extra processor in the system with access to the main system memory and coherent communication with other processors in the system.

Figure 2-7 shows a comparison of the traditional model, where the accelerator must go through the processor to access memory with CAPI.

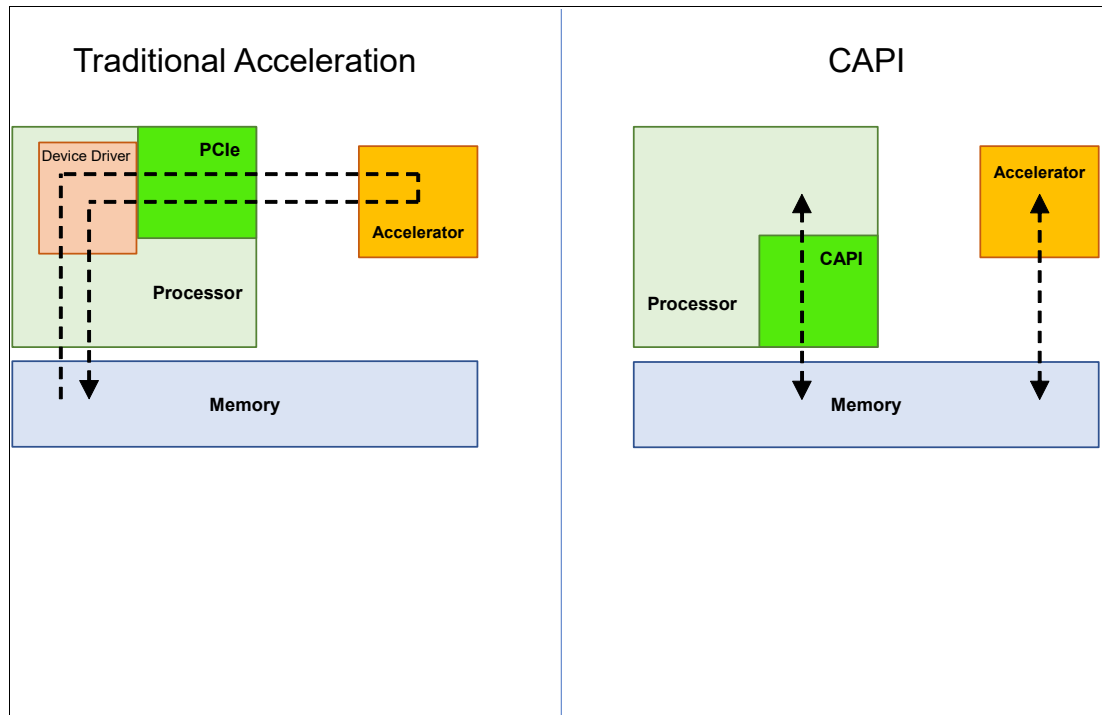


Figure 2-7 CAPI accelerator that is attached to the POWER9 processor

The benefits of using CAPI include the ability to access shared memory blocks directly from the accelerator, perform memory transfers directly between the accelerator and processor cache, and reduce the code path length between the adapter and the processors. This reduction in the code path length might occur because the adapter is not operating as a traditional I/O device, and there is no device driver layer to perform processing. CAPI also presents a simpler programming model.

The accelerator adapter implements the POWER Service Layer (PSL), which provides address translation and system memory cache for the accelerator functions. The custom processors on the system board, consisting of an FPGA or an ASIC, use this layer to access shared memory regions, and cache areas as though they were a processor in the system. This ability enhances the performance of the data access for the device and simplifies the programming effort to use the device. Instead of treating the hardware accelerator as an I/O device, it is treated as a processor, which eliminates the requirement of a device driver to perform communication. It also eliminates the need for direct memory access that requires system calls to the OS kernel. By removing these layers, the data transfer operation requires fewer clock cycles in the processor, improving the I/O performance.

The implementation of CAPI on the POWER9 processor enables hardware companies to develop solutions for specific application demands. Companies use the performance of the POWER9 processor for general applications and the custom acceleration of specific functions by using a hardware accelerator with a simplified programming model and efficient communication with the processor and memory resources.

2.1.10 Power management and system performance

The POWER9 scale-out models introduced new features for EnergyScale including new variable processor frequency modes that provide a significant performance boost beyond the static nominal frequency. The following modes can be modified or disabled.

The default performance mode depends on the server model. For the Power S914 server (9009-41A), Dynamic Performance mode is enabled by default, and the Power S922 (9009-22A) and Power S924 (9009-42A) servers have Maximum Performance mode enabled by default. The difference in the Power S914 setup is that some servers are used in the office workspace where the extra fan noise might be unacceptable. If acoustic concern is not an issue, you can change to Maximum Performance mode.

Disable all modes

The processor clock frequency is set to its nominal value, and the power that is used by the system remains at a nominal level. This option was the default for all systems before POWER9.

Static Power Save mode

Reduces the power consumption by lowering the processor clock frequency and the voltage to fixed values. This option also reduces the power consumption of the system while still delivering predictable performance.

Dynamic Power Performance mode

Causes the processor frequency to vary based on the processor use. During periods of high use, the processor frequency is set to the maximum value allowed, which might be above the nominal frequency. Additionally, the frequency is lowered below the nominal frequency during periods of moderate and low processor use.

Maximum Performance mode

The mode enables the system to reach the maximum frequency under certain conditions. The power consumption increases. The maximum frequency is approximately 20% better than nominal.

The controls for all of these modes are available on the Advanced System Management Interface (ASMI) and can be dynamically modified.

2.1.11 Comparison of the POWER9, POWER8, and POWER7+ processors

Table 2-7 shows comparable characteristics between the generations of POWER9, POWER8, and POWER7 processors.

Table 2-7 Comparison of technology for the POWER9 processor and prior generations

Characteristics	POWER9	POWER8	POWER7+
Technology	14 nm	22 nm	32 nm
Die size	68.5 mm x 68.5 mm	649 mm ²	567 mm ²
Number of transistors	8 billion	4.2 billion	2.1 billion
Maximum cores	12	12	8
Maximum SMT threads per core	8 threads	8 threads	4 threads

Characteristics	POWER9	POWER8	POWER7+
Maximum frequency	3.8 - 4.0 GHz	4.15 GHz	4.4 GHz
L2 Cache	512 KB shared between cores	512 KB per core	256 KB per core
L3 Cache	10 MB of FLR-L3 cache per two cores with each core having access to the full 120 MB of L3 cache, on-chip eDRAM	8 MB of FLR-L3 cache per core with each core having access to the full 96 MB of L3 cache, on-chip eDRAM	10 MB of FLR-L3 cache per core with each core having access to the full 80 MB of L3 cache, on-chip eDRAM
Memory support	DDR4	DDR3 and DDR4	DDR3
I/O bus	PCIe Gen4	PCIe Gen3	GX++

2.2 Memory subsystem

The Power S914 server is a one-socket system that supports a single POWER9 processor module. The server supports a maximum of 16 DDR4 DIMM slots. Memory features that are supported are 8 GB, 16 GB, 32 GB, and 64 GB in 6-core and 8-core configurations that offer a maximum system memory of 1 TB. Memory speeds vary depending on the DIMM size and modules placement, as shown in Table 2-8.

Note: If you use the 4-core memory #EP10, the maximum of 64 GB RAM is supported. The following memory configurations are supported:

- ▶ Quantity of 2 or 4 or 6 or 8 of #EM60 (8 GB DDR4 DRAM)
- ▶ Quantity of 2 or 4 of #EM62 (16 GB DDR4 DRAM)
- ▶ Quantity of 2 of #EM63 (32 GB DDR4 DRAM)

The Power S922 and Power S924 servers are two-socket servers that support up to two POWER9 processor modules. The servers support a maximum of 32 DDR4 DIMM slots, with 16 DIMM slots per installed processor. Memory features that are supported are 8 GB, 16 GB, 32 GB, 64 GB, and 128 GB, enabling a maximum system memory of 4 TB. Memory speeds vary depending on the DIMM size and modules placement, as shown in Table 2-8.

Table 2-8 POWER9 memory speed

RDIMM size	Mbps (1 DIMM per port)	Mbps (2 DIMMs per port)
8 GB	2400	2133
16 GB	2666	2133
32 GB	2400	2133
64 GB	2400	2133
128 GB	2400	2133

The maximum theoretical memory bandwidth for POWER9 processor module is 170 GBps. The total maximum theoretical memory bandwidth for a two-socket system is 340 GBps.

These servers support an optional feature that is called Active Memory Expansion that enables the effective maximum memory capacity to be much larger than the true physical memory on AIX. This feature runs innovative compression and decompression of memory content by using a dedicated coprocessor to provide memory expansion up to 125%, depending on the workload type and its memory usage.

For example, a server with 256 GB RAM physically installed can effectively be expanded over 512 GB RAM. This approach can enhance virtualization and server consolidation by allowing a partition to do more work with the same physical amount of memory or a server to run more partitions and do more work with the same physical amount of memory.

2.2.1 Memory placement rules

The following memory options are orderable:

- ▶ 8 GB DDR4 DRAM (#EM60)
- ▶ 16 GB DDR4 DRAM (#EM62)
- ▶ 32 GB DDR4 DRAM (#EM63)
- ▶ 64 GB DDR4 DRAM (#EM64)
- ▶ 128 GB DDR4 DRAM (#EM65)

Note: If you use the 4-core memory #EP10, the only memory sizes that are allowed is 8 GB (#EM60), 16 GB (#EM62) and 32 GB (#EM63).

All memory must be ordered in pairs, with a minimum of 32 GB for the Power S914, Power S924, and Power 922 servers that have a single processor module installed. 64 GB is the minimum for servers with two processors modules that are installed (the Power S924 and Power S922 servers) per server.

The supported maximum memory is as follows for the Power S914 server:

- ▶ One processor module installed (4-core): 64 GB (eight 8 GB DIMMS, four 16 GB DIMMs, or two 32 GB DIMMs)
- ▶ One processor module installed (6-core or 8-core): 1 TB (Sixteen 64 GB DIMMs)

The supported maximum memory is as follows for the Power S924 and the Power S922 servers:

- ▶ One processor module installed: 2 TB (Sixteen 128 GB DIMMs)
- ▶ Two processors modules installed: 4 TB (Thirty-two 128 GB DIMMs)

The basic rules for memory placement follow:

- ▶ Each FC equates to a single physical DIMM.
- ▶ All memory features must be ordered in pairs.
- ▶ All memory DIMMs must be installed in pairs.
- ▶ Each DIMM within a pair must be of the same capacity.

In general, the preferred approach is to install memory evenly across all processors in the system. Balancing memory across the installed processors enables memory access in a consistent manner and typically results in the best possible performance for your configuration. You should account for any plans for future memory upgrades when you decide which memory feature size to use at the time of the initial system order.

Figure 2-8 shows the physical memory DIMM topology for the Power S914 server.

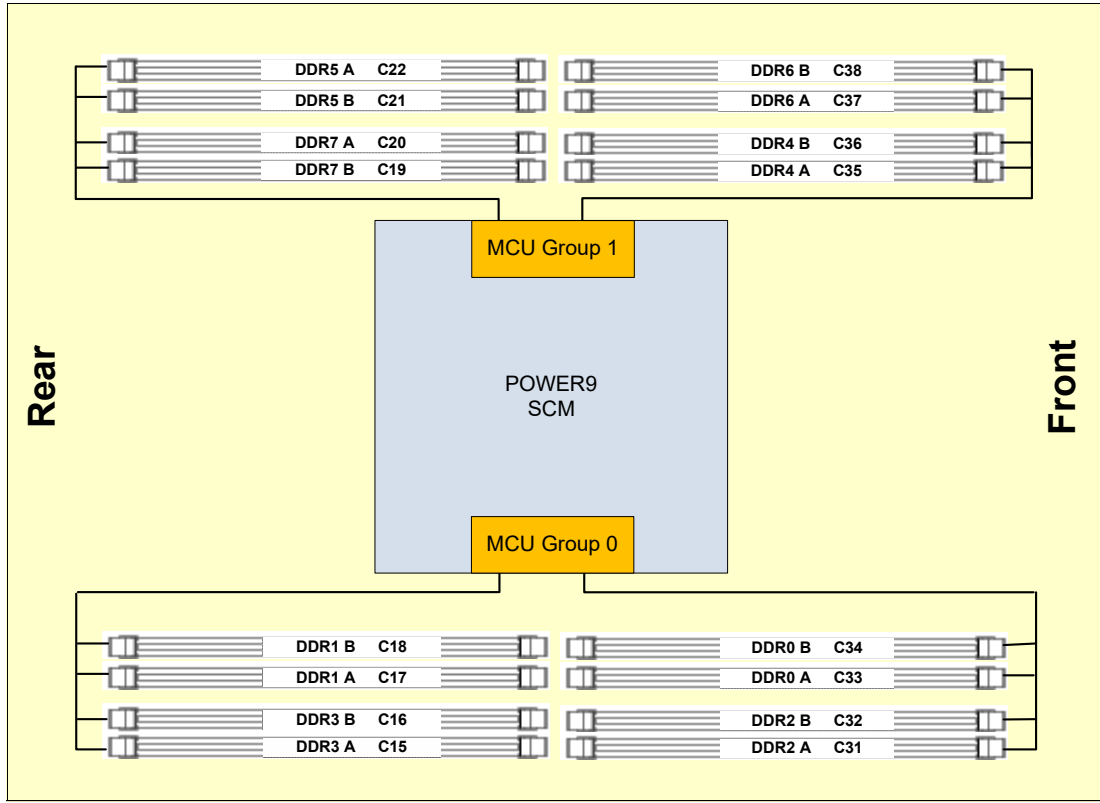


Figure 2-8 Memory DIMM topology for the Power S914 server

For systems with a single processor module that is installed, the plugging order for the memory DIMMs is as follows (see Figure 2-9):

- ▶ Pair installation:
 - The first DIMM pair is installed at Red 1 (C33 DDR0-A and C17 DDR1-A).
 - The second DIMM pair is installed at Gold 2 (C36 DDR4-A and C22 DDR5-A).
 - The third DIMM pair is installed at Cyan 3 (C31 DDR2-A and C15 DDR3-A).
 - The fourth DIMM pair is installed at Gray 4 (C38 DDR6-A and C20 DDR7-A).
- ▶ Quad installation:
 - Two fifth DIMM pairs (or quad) are installed at Red 5 (C34 DDR0-B and C18 DDR1-B) and at Cyan 5 (C32 DDR2-B and C16 DDR3-B).
 - Two sixth DIMM pairs (or quad) are installed at Gold 6 (C35 DDR4-B and C21 DDR5-B) and at Gray 6 (C37 DDR6-B and C19 DDR7-B).

MCU Group 0								MCU Group 1							
DDR0		DDR1		DDR2		DDR3		DDR4		DDR5		DDR6		DDR7	
A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B
C33	C34	C17	C18	C31	C32	C15	C16	C36	C35	C22	C21	C38	C37	C20	C19
1		1		3		3		2		2		4		4	
	5		5		5		5		6		6		6		6

Figure 2-9 DIMM plug sequence for the Power S914 server

More considerations:

- ▶ You may not mix 1R and 2R DIMMs on a single channel within an MCU group because they run at different DIMM data rates.

Table 2-9 lists the feature codes of the supported memory modules.

Table 2-9 Memory feature codes

Size	Feature code	Rank
8 GB	#EM60	1R
16 GB	#EM62	1R
32 GB	#EM63	2R
64 GB	#EM64	2R
128 GB	#EM65	2R

- ▶ DIMMs in the same color cells must be identical (same size and rank).

Note: Each color in Figure 2-9 represents a unique DIMM size and type.

Figure 2-10 shows the physical memory DIMM topology for the Power S922 and the Power S924 servers.

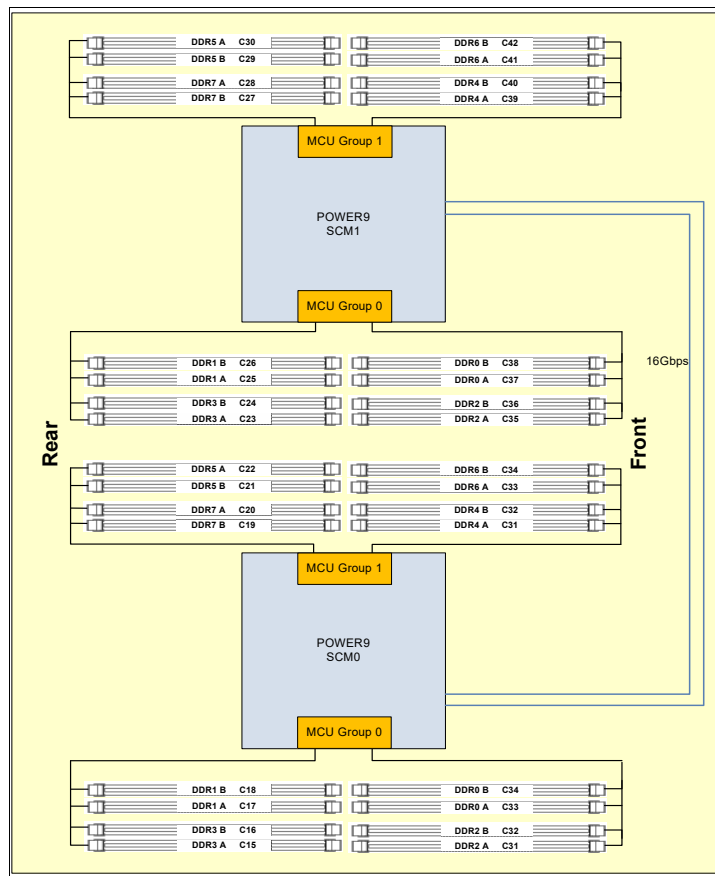


Figure 2-10 Memory DIMM topology for the Power S922 and the Power S924

For the Power S922 and Power S924 servers, the plugging order for the memory DIMMS is as follows (see Figure 2-11):

- ▶ Pair installation:
 - The first DIMM pair is installed at Red 1 (C33 DDR0-A and C17 DDR1-A) of SCM-0.
 - The second DIMM pair is installed at Green 2 (C41 DDR0-A and C25 DDR1-A) of SCM-1.
 - The third DIMM pair is installed at Gold 3 (C36 DDR4-A and C22 DDR5-A) of SCM-0.
 - The fourth DIMM pair is installed at Purple 4 (C44 DDR4-A and C30 DDR5-A) of SCM-1.
 - The fifth DIMM pair is installed at Cyan 5 (C31 DDR2-A and C15 DDR3-A) of SCM-0.
 - The sixth DIMM pair is installed at Pink 6 (C39 DDR2-A and C23 DDR3-A) of SCM-1.
 - The seventh DIMM pair is installed at Gray 7 (C38 DDR6-A and C20 DDR7-A) of SCM-0.
 - The eighth DIMM pair is installed at Yellow 8 (C46 DDR6-A and C28 DDR7-A) of SCM-1.
- ▶ Quad installation:
 - Two ninth DIMM pairs (or quad) are installed at Red 9 (C34 DDR0-B and C18 DDR1-B) and at Cyan 9 (C32DDR2-B and C16 DDR3-B) of SCM-0.
 - Two tenth DIMM pairs (or quad) are installed at Green 10 (C42 DDR0-B and C26 DDR1-B) and at Pink 10(C40 DDR2-B and C24 DDR3-B) of SCM-1.
 - Two eleventh DIMM pairs (or quad) are installed at Gold 11 (C35 DDR4-B and C21 DDR5-B) and at Gray 11(C37 DDR6-B and C19 DDR7-B) of SCM-0.
 - Two twelfth DIMM pairs (or quad) are installed at Purple 12 (C43 DDR4-B and C29 DDR5-B) and at Yellow 12(C45 DDR6-B and C27 DDR7-B) of SCM-1.

MCU Group 0								MCU Group 1								MCU Group 0								MCU Group 1							
DDR0		DDR1		DDR2		DDR3		DDR4		DDR5		DDR6		DDR7		DDR0		DDR1		DDR2		DDR3		DDR4		DDR5		DDR6		DDR7	
A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B
C33	C34	C17	C18	C31	C32	C15	C16	C36	C35	C22	C21	C38	C37	C20	C19	C41	C42	C25	C26	C39	C40	C23	C24	C44	C43	C30	C29	C46	C45	C28	C27
1		1								3		3				2		2						4		4				8	
				5		5						7		7				10		10						12		12			
	9		9							11		11										10		10				12			
				9		9						11		11								10		10				12			

Figure 2-11 DIMM plug sequence for Power S922 and Power S924 servers

More considerations:

- ▶ You may not mix 1R and 2R DIMMs on a single channel within an MCU group because they run at different DIMM data rates. For more information, see Table 2-9 on page 75.
- ▶ DIMMs at same color cells must be identical (same size and rank).

2.2.2 Memory bandwidth

The POWER9 processor has exceptional cache, memory, and interconnect bandwidths. The next sections show the bandwidth capabilities of the Power S914, Power S922, and Power S924 servers.

Power S922 server bandwidth

Table 2-10 shows the maximum bandwidth estimates for a single core on the Power S922 server.

Table 2-10 The Power S922 single core bandwidth maximum

Single core	Power S922 server	Power S922 server
	1 core @ 3.8 GHz (maximum)	1 core @ 3.9 GHz (maximum)
L1 (data) cache	364.8 GBps	374.4 GBps
L2 cache	364.8 GBps	374.4 GBps
L3 cache	243.2 GBps	249.6 GBps

For an entire Power S922 server that is populated with two processor modules, the overall bandwidths are shown in Table 2-11.

Table 2-11 The Power S922 total bandwidth maximum estimates

Total bandwidths	Power S922 server	Power S922 server
	16 cores @ 3.9 GHz (maximum)	20 cores @ 3.8 GHz (maximum)
L1 (data) cache	5990.4 GBps	7296 GBps
L2 cache	5990.4 GBps	7296 GBps
L3 cache	3993.6 GBps	4864 GBps
Total memory	340 GBps	340 GBps
PCIe Interconnect	320 GBps	320 GBps
X Bus SMP	16 GBps	16 GBps

Power S914 server bandwidth

The bandwidth figures for the caches are calculated as follows:

- ▶ L1 cache: In one clock cycle, four 16-byte load operations and two 16-byte store operations can be accomplished. The value varies depending on the clock of the core. The formula is as follows:
3.8 GHz core: $(4 * 16 \text{ B} + 2 * 16 \text{ B}) * 3.8 \text{ GHz} = 364.8 \text{ GBps}$
- ▶ L2 cache: In one clock cycle, one 64-byte load operation and two 16-byte store operations can be accomplished. The value varies depending on the clock of the core. The formula is as follows:
3.8 GHz core: $(1 * 64 \text{ B} + 2 * 16 \text{ B}) * 3.8 \text{ GHz} = 364.8 \text{ GBps}$
- ▶ L3 cache: One 32-byte load operation and one 32-byte store operation can be accomplished at one clock cycle. The formula is as follows:
3.8 GHz core: $(1 * 32 \text{ B} + 1 * 32 \text{ B}) * 3.8 \text{ GHz} = 243.2 \text{ GBps}$

Table 2-12 shows the maximum bandwidth estimates for a single core on the Power S914 server.

Table 2-12 The Power S914 single core bandwidth estimates

Single core	Power S 914 server 1 core @ 3.8 GHz (maximum)
L1 (data) cache	364.8 GBps
L2 cache	364.8 GBps
L3 cache	243.2 GBps

For an entire Power S914 system that is populated with one processor module, the overall bandwidths are what is shown in Table 2-13.

Table 2-13 The Power S914 total bandwidth maximum estimates

Total bandwidths	Power S914 server	Power S914 server	Power S914 server
	4 cores @ 3.8 GHz (maximum)	6 cores @ 3.8 GHz (maximum)	8 cores @ 3.8 GHz (maximum)
L1 (data) cache	1459.2 GBps	2188.8 GBps	2918.4 GBps
L2 cache	1459.2 GBps	2188.8 GBps	2918.4 GBps
L3 cache	972.8 GBps	1459.2 GBps	1945.6 GBps
Total memory	170 GBps	170 GBps	170 GBps
PCIe Interconnect	160 GBps	160 GBps	160 GBps

Power S924 server bandwidth

The bandwidth figures for the caches are calculated as follows:

- ▶ L1 cache: In one clock cycle, four 16-byte load operations and two 16-byte store operations can be accomplished. The value varies depending on the clock of the core. The formulas are as follows:
 - 3.9 GHz core: $(4 * 16 \text{ B} + 2 * 16 \text{ B}) * 3.9 \text{ GHz} = 374.4 \text{ GBps}$
 - 4.0 GHz core: $(4 * 16 \text{ B} + 2 * 16 \text{ B}) * 4.0 \text{ GHz} = 384 \text{ GBps}$
- ▶ L2 cache: In one clock cycle, one 64-byte load operation and two 16-byte store operations can be accomplished. The value varies depending on the clock of the core. The formulas are as follows:
 - 3.9 GHz core: $(1 * 64 \text{ B} + 2 * 16 \text{ B}) * 3.9 \text{ GHz} = 374.4 \text{ GBps}$
 - 4.0 GHz core: $(1 * 64 \text{ B} + 2 * 16 \text{ B}) * 4.0 \text{ GHz} = 384 \text{ GBps}$
- ▶ L3 cache: One 32-byte load operation and one 32-byte store operation can be accomplished at half-clock speed. The formula is as follows:
 - 3.9 GHz core: $(1 * 32 \text{ B} + 1 * 32 \text{ B}) * 3.9 \text{ GHz} = 249 \text{ GBps}$
 - 4.0 GHz core: $(1 * 32 \text{ B} + 1 * 32 \text{ B}) * 4.0 \text{ GHz} = 256 \text{ GBps}$

Processor modules for the Power S922 and Power S924 servers run with higher frequency than the Power S914 server.

Table 2-14 shows the maximum bandwidth estimates for a single core on the Power S924 server.

Table 2-14 The Power S924 single core bandwidth maximum estimates

Single core	Power S924 server	Power S924 server
	1 core @ 3.9 GHz (maximum)	1 core @ 4.0 GHz (maximum)
L1 (data) cache	374.4 GBps	384 GBps
L2 cache	374.4 GBps	384 GBps
L3 cache	249 GBps	256 GBps

For an entire Power S924 server that is populated with two processor modules, the overall bandwidths are shown in Table 2-15.

Table 2-15 The Power S924 total bandwidth maximum estimates

Total bandwidths	Power S924 server	Power S924 server	Power S924 server
	16 cores @ 4.0 GHz (maximum)	20 cores @ 3.9 GHz (maximum)	24 cores @ 3.9 GHz (maximum)
L1 (data) cache	6144 GBps	7488 GBps	8985.6 GBps
L2 cache	6144 GBps	7488 GBps	8985.6 GBps
L3 cache	4096 GBps	4992 GBps	5990.4 GBps
Total memory	340 GBps	340 GBps	340 GBps
PCIe Interconnect	320 GBps	320 GBps	320 GBps
X Bus SMP	16 GBps	16 GBps	X Bus SMP

Note: There are several POWER9 design points to consider when comparing hardware designs that use SMP communication bandwidths as a unique measurement. POWER9 provides:

- ▶ More cores per socket leading to lower inter-CPU communication.
- ▶ More RAM density (up to 2 TB per socket) that leads to less inter-CPU communication.
- ▶ Greater RAM bandwidth for less dependence on an L3 cache.
- ▶ Intelligent hypervisor scheduling that places RAM usage close to the CPU.
- ▶ New SMP routing so that multiple channels are available when congestion occurs.

2.3 System bus

This section provides information about the internal system buses.

The Power S914, Power S924, and Power S922 servers have internal I/O connectivity through PCIe Gen4 and Gen3 (PCI Express Gen4/Gen3 or PCIe Gen4/Gen3) slots, and also external connectivity through SAS adapters.

The internal I/O subsystem on the Power S914, Power S924, and Power S922 servers is connected to the PCIe controllers on a POWER9 processor in the system. An IBM Power System server in a two-socket configuration has a bus that has 80 PCIe G4 lanes running at a maximum of 16 Gbps full-duplex, which provides 320 GBps of I/O connectivity to the PCIe slots, SAS internal adapters, and USB ports. The Power S914 server with one processor module provides 160 GBps I/O bandwidth (maximum).

Some PCIe slots are connected directly to the PCIe Gen4 buses on the processors, and PCIe Gen3 devices are connected to these buses through PCIe Gen3 Switches. For more information about which slots are connected directly to the processor and which ones are attached to a PCIe Gen3 Switch (referred as PEX), see Figure 2-3 on page 62.

Figure 2-12 compares the POWER8 and POWER9 I/O buses architecture.

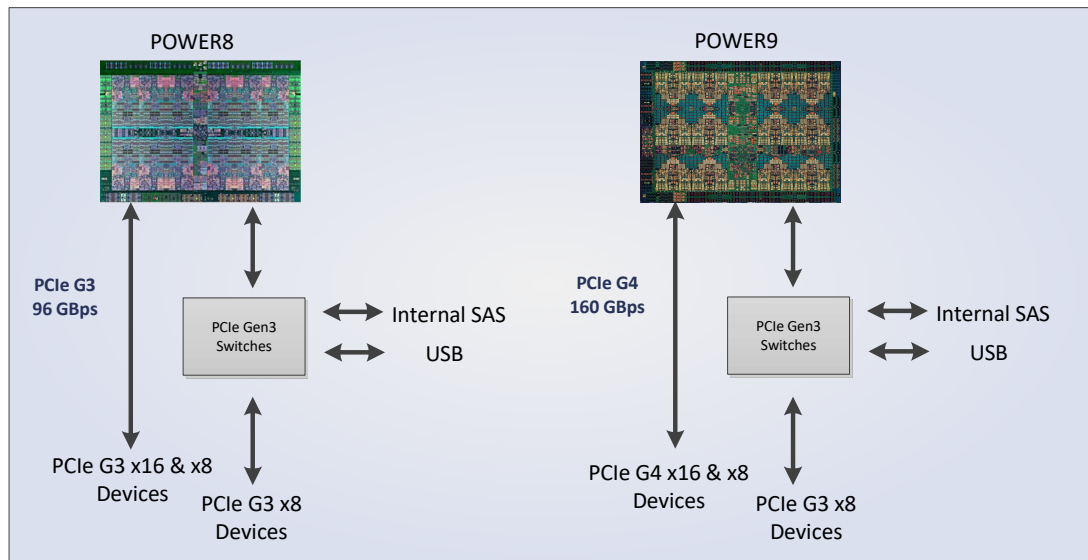


Figure 2-12 Comparison of POWER and POWER9 I/O buses architectures

Table 2-16 lists the I/O bandwidth of Power S914, Power S924, and Power S922 processor configurations.

Table 2-16 I/O bandwidth

I/O	I/O bandwidth (maximum theoretical)
Total I/O bandwidth	Power S914 server with one processor: <ul style="list-style-type: none"> ▶ 80 GBps simplex ▶ 160 GBps duplex Power S924 and S922 servers with two processors: <ul style="list-style-type: none"> ▶ 160 GBps simplex ▶ 320 GBps duplex

For PCIe Interconnect, each POWER9 processor module has 40 PCIe lanes running at 16 Gbps full-duplex. The bandwidth formula is calculated as follows:

$$\text{Forty lanes} * 2 \text{ processors} * 16 \text{ Gbps} * 2 = 320 \text{ GBps}$$

2.4 Internal I/O subsystem

The internal I/O subsystem is on the system board, which supports PCIe slots. PCIe adapters on the Power S922, Power S914, and Power S924 servers are hot-pluggable.

All PCIe slots support enhanced error handling (EEH). PCI EEH-enabled adapters respond to a special data packet that is 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 restart.

2.4.1 Slot configuration

The various slot configurations are described in this section. We combined the Power S914 and Power S924 servers into a single section.

Slot configuration for the Power S922 server

The Power S922 server provides PCIe Gen3 and PCIe Gen4 slots. The number of PCIe slots that are available on the Power S922 server depends on the number of installed processors. Table 2-17 provides information about the PCIe slots in the Power S922 server.

Table 2-17 PCIe slot locations and descriptions for the Power S922 server

Slot availability	Description	Adapter size
Two slots (P1-C6 and P1-C12)	PCIe Gen3 x8	Half-height and half-length
Two slots (P1-C7 and P1-C11)	PCIe Gen3 x8	Half-height and half-length
Three slots (P1-C3 ^a , P1-C4 ^a , and P1-C9)	PCIe Gen4 x16	Half-height and half-length
Two slots (P1-C2 ^a and P1-C8)	PCIe Gen4 x8 with x16 connector	Half-height and half-length

a. The slot is available when the second processor slot is populated.

Table 2-18 lists the PCIe adapter slot locations and details for the Power S922 server.

Table 2-18 PCIe slot locations and details for the Power S922 server

Location code	Description	Slot capabilities		
		CAPI	Single Root I/O Virtualization (SR-IOV)	I/O adapter enlarged capacity enablement order ^a
P1-C2 ^b	PCIe Gen4 x8 with x16 connector	No	Yes	5
P1-C3 ^b	PCIe Gen4 x16	Yes	Yes	2

Location code	Description	Slot capabilities		
		CAPI	Single Root I/O Virtualization (SR-IOV)	I/O adapter enlarged capacity enablement order ^a
P1-C4 ^b	PCIe Gen4 x16	Yes	Yes	3
P1-C6	PCIe Gen3 x8 with x16 connector	No	Yes	6
P1-C7	PCIe Gen3 x8	No	Yes	10
P1-C8 ^b	PCIe Gen4 x8 with x16 connector	Yes	Yes	4
P1-C9 ^b	PCIe Gen4 x16	Yes	Yes	1
P1-C11	PCIe Gen3 x8 (default LAN slot)	No	Yes	11
P1-C12	PCIe Gen3 x8 with x16 connector	No	Yes	7

- a. Enabling the I/O adapter enlarged capacity option affects only Linux partitions.
b. A high-performance slot that is directly connected to the processor module. The connectors in these slots are differently colored than the slots in the PCIe3 switches.

Figure 2-13 shows the rear view of the Power S922 server with the location codes for the PCIe adapter slots.




Figure 2-13 Rear view of a rack-mounted Power S922 server with PCIe slots location codes

Slot configurations for Power S914 and Power S924 servers

The Power S914 and Power S924 servers provide PCIe Gen3 and PCIe Gen4 slots. The number of PCIe slots that are available on the Power S924 server depends on the number of installed processors.

Table 2-19 provides information about the PCIe slots in the Power S914 and Power S924 servers.

Table 2-19 PCIe slot locations and descriptions for the Power S914 and Power S924 servers

Slot availability	Description	Adapter size
Two slots (P1-C6 and P1-C12)	PCIe Gen3 x8 with x16 connector	Full-height and half-length
Four slots (P1-C5, P1-C7, P1-C10, and P1-C11)	PCIe Gen3 x8	Full-height and half-length
Three slots (P1-C3 ^a , P1-C4, and P1-C9)	PCIe Gen4 x16	Full-height and half-length
Two slots (P1-C2 and P1-C8)	PCIe Gen4 x8 with x16 connector	Full-height and half-length

a. The slot is available when the second processor slot of the Power S924 server is populated.

Table 2-20 lists the PCIe adapter slot locations and details for the Power S914 and Power S924 servers.

Table 2-20 PCIe slot locations and details for the Power S914 and Power S924 servers

Location code	Description	Slot capabilities		
		CAPI	SR-IOV	I/O adapter enlarged capacity enablement order ^a
P1-C2 ^b	PCIe Gen4 x8 or NVLink slot	No	Yes	N/A (S914) 5 (S924)
P1-C3 ^b	PCIe Gen4 x16	Yes	Yes	N/A (S914) 2 (S924)
P1-C4 ^b	PCIe Gen4 x16	Yes	Yes	N/A (S914) 3 (S924)
P1-C5	PCIe Gen3 x8	No	Yes	5 (S914) 8 (S924)
P1-C6	PCIe Gen3 x8 with x16 connector	No	Yes	3 (S914) 6 (S924)
P1-C7	PCIe Gen3 x8	No	Yes	7 (S914) 10(S924)
P1-C8 ^b	PCIe Gen4 x8	Yes	Yes	2 (S914) 4 (S924)
P1-C9 ^b	PCIe Gen4 x16	Yes	Yes	1 (S914) 1 (S924)
P1-C10	PCIe Gen3 x8	No	Yes	6 (S914) 9 (S924)

Location code	Description	Slot capabilities		
		CAPI	SR-IOV	I/O adapter enlarged capacity enablement order ^a
P1-C11	PCIe Gen3 x8 (default LAN slot)	No	Yes	8 (S914) 11 (S924)
P1-C12	PCIe Gen3 x8 with x16 connector	No	Yes	4 (S914) 7(S924)

a. Enabling the I/O adapter enlarged capacity option affects only Linux partitions.

b. A high-performance slot that is directly connected to the processor module. The connectors in these slots are differently colored than the slots in the PCIe3 switches.

Figure 2-14 shows the rear view of the Power S924 server with the location codes for the PCIe adapter slots.




Figure 2-14 Rear view of a rack-mounted Power S924 server with the PCIe slots location codes

2.4.2 System ports

The system board has one serial port that is called a *system port*. The one system port is RJ45 and is supported by AIX and Linux for attaching serial devices, such as an asynchronous device, for example, a console. If the device does not have an RJ45 connection, a converter cable such as #3930 can provide a 9-pin D-shell connection.

2.5 Peripheral Component Interconnect adapters

This section covers the various types and functions of the PCI adapters that are supported by the Power S914, Power S922, and Power S924 servers.

Important: There is no FCoE support on POWER9 systems.

2.5.1 Peripheral Component Interconnect Express

PCIe uses a serial interface and enables point-to-point interconnections between devices (by 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 Gen4, which are capable of 16 GBps simplex (32 GBps duplex) on a single x16 interface. PCIe Gen4 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 size connector 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 the slot.

IBM POWER9 processor-based servers can support two different form factors of PCIe adapters:

- ▶ PCIe low-profile (LP) cards, which are used with the Power S922 PCIe slots. These cards are not compatible with Power S914 and Power S924 servers because of their low height, but there are similar cards in other form factors.
- ▶ PCIe full height and full high cards are not compatible with the Power S922 server and are designed for the following servers:
 - Power S914 server
 - Power S924 server

Before adding or rearranging adapters, use the [IBM System Planning Tool \(SPT\)](#) to validate the new adapter configuration.

If you are installing a new feature, ensure that you have the software that is required to support the new feature and determine whether there are any existing update prerequisites to install. To do this, go to the [IBM Power Systems Prerequisite website](#).

The following sections describe the supported adapters and provide tables of orderable feature numbers. The tables indicate operating system support (AIX, IBM i, and Linux) for each of the adapters.

Note: The maximum number of adapters in each case may require the server to have an external I/O drawer.

2.5.2 LAN adapters

To connect the Power S914, Power S922, and Power S924 servers to a local area network (LAN), you can use the LAN adapters that are supported in the PCIe slots of the system unit.

Table 2-21 lists the LAN adapters that are available for the Power S922 server.

Table 2-21 Available LAN adapters for Power S922 servers.

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0W	2CC4	PCIe2 2-port 10/1 GbE BaseT RJ45 Adapter	0	12	AIX, IBM i, and Linux
EN0U	2CC3	PCIe2 4-port (10 Gb+1 GbE) Copper SFP+RJ45 Adapter	0	12	AIX, IBM i, and Linux
EN0S	2CC3	PCIe2 4-Port (10 Gb+1 GbE) SR+RJ45 Adapter	0	12	AIX, IBM i, and Linux
5899	576F	PCIe2 4-port 1 GbE Adapter	0	12	AIX, IBM i, and Linux
EN0X	2CC4	PCIe2 LP 2-port 10/1 GbE BaseT RJ45 Adapter	0	9	AIX, IBM i, and Linux
EN0V	2CC3	PCIe2 LP 4-port (10 Gb+1 GbE) Copper SFP+RJ45 Adapter	0	9	AIX, IBM i, and Linux
EN0T	2CC3	PCIe2 LP 4-Port (10 Gb+1 GbE) SR+RJ45 Adapter	0	9	AIX, IBM i, and Linux
5260	576F	PCIe2 LP 4-port 1 GbE Adapter	0	9	AIX, IBM i, and Linux
EC2S	58FA	PCIe3 2-Port 10 Gb NIC & ROCE SR/Cu Adapter	0	4	AIX, IBM i, and Linux
EC37	57BC	PCIe3 LP 2-port 10 GbE NIC&RoCE SFP+ Copper Adapter	0	9	AIX, IBM i, and Linux
EC38	57BC	PCIe3 2-port 10 GbE NIC&RoCE SFP+ Copper Adapter	0	12	AIX, IBM i, and Linux
EC2U	58FB	PCIe3 2-Port 25/10 Gb NIC & ROCE SR/Cu Adapter	0	4	AIX, IBM i, and Linux
EC3B	57BD	PCIe3 2-Port 40 GbE NIC RoCE QSFP+ Adapter	0	12	AIX, IBM i, and Linux
EC3B	57BD	PCIe3 2-Port 40 GbE NIC RoCE QSFP+ Adapter	0	12	AIX, IBM i, and Linux
EN0K	2CC1	PCIe3 4-port (10 Gb FCoE & 1 GbE) SFP+Copper & RJ45	0	12	AIX, IBM i, and Linux
EN0H	2B93	PCIe3 4-port (10 Gb Fibre Channel over Ethernet (FCoE) & 1 GbE) SR & RJ45	0	12	AIX, IBM i, and Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN15	2CE3	PCIe3 4-port 10 GbE SR Adapter	0	12	AIX, IBM i, and Linux
EC3T	2CEB	PCIe3 LP 1-port 100 Gb EDR IB Adapter x16	0	3	Linux
EC2R	58FA	PCIe3 LP 2-Port 10 Gb Network Interface Card (NIC) & ROCE SR/Cu Adapter	0	8	AIX, IBM i, and Linux
EC37	57BC	PCIe3 LP 2-port 10 GbE NIC&RoCE SFP+ Copper Adapter	0	9	AIX, IBM i, and Linux
EC3E	2CEA	PCIe3 LP 2-port 100 Gb EDR IB Adapter x16	0	3	Linux
EC3L	2CEC	PCIe3 LP 2-port 100 GbE (NIC & RoCE) QSFP28 Adapter x16	0	3	AIX, IBM i, and Linux
EC2T	58FB	PCIe3 LP 2-Port 25/10 Gb NIC & ROCE SR/Cu Adapter	0	8	AIX, IBM i, and Linux
EC3A	57BD	PCIe3 LP 2-Port 40 GbE NIC RoCE QSFP+ Adapter	0	8	AIX, IBM i, and Linux
EC3A	57BD	PCIe3 LP 2-Port 40 GbE NIC RoCE QSFP+ Adapter	0	8	AIX, IBM i, and Linux
EN0J	2B93	PCIe3 LP 4-port (10 Gb FCoE & 1 GbE) SR & RJ45	0	9	AIX, IBM i, and Linux
EN0L	2CC1	PCIe3 LP 4-port(10 Gb FCoE & 1 GbE) SFP+Copper & RJ45	0	9	AIX, IBM i, and Linux
EC62	2CF1	PCIe4 LP 1-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux
EC64	2CF2	PCIe4 LP 2-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux
EC67	2CF3	PCIe4 LP 2-port 100 Gb ROCE EN LP adapter	0	3	AIX, IBM i, and Linux

Table 2-22 lists the available LAN adapters for a Power S914 server.

Table 2-22 Available LAN adapters in Power S914 servers.

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0W	2CC4	PCIe2 2-port 10/1G bE BaseT RJ45 Adapter	0	12	AIX, IBM i, and Linux
EN0U	2CC3	PCIe2 4-port (10 Gb+1 GbE) Copper SFP+RJ45 Adapter	0	12	AIX, IBM i, and Linux
EN0S	2CC3	PCIe2 4-Port (10 Gb+1 GbE) SR+RJ45 Adapter	0	12	AIX, IBM i, and Linux
5899	576F	PCIe2 4-port 1 GbE Adapter	0	13	AIX, IBM i, and Linux
EC3U	2CEB	PCIe3 1-port 100 Gb EDR IB Adapter x16	0	1	Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC2S	58FA	PCIe3 2-Port 10 Gb NIC & ROCE SR/Cu Adapter	0	8	AIX, IBM i, and Linux
EC38	57BC	PCIe3 2-port 10 GbE NIC&RoCE SFP+ Copper Adapter	0	12	AIX, IBM i, and Linux
EC3F	2CEA	PCIe3 2-port 100 Gb EDR IB Adapter x16	0	1	Linux
EC3M	2CEC	PCIe3 2-port 100 GbE (NIC & RoCE) QSFP28 Adapter x16	0	1	AIX, IBM i, and Linux
EC2U	58FB	PCIe3 2-Port 25/10 Gb NIC & ROCE SR/Cu Adapter	0	8	AIX, IBM i, and Linux
EC3B	57BD	PCIe3 2-Port 40 GbE NIC RoCE QSFP+ Adapter	0	12	AIX, IBM i, and Linux
EN0K	2CC1	PCIe3 4-port (10 Gb FCoE & 1 GbE) SFP+Copper & RJ45	0	12	AIX, IBM i, and Linux
EN0H	2B93	PCIe3 4-port (10 Gb FCoE & 1 GbE) SR & RJ45	0	12	AIX, IBM i, and Linux
EN15	2CE3	PCIe3 4-port 10 GbE SR Adapter	0	12	AIX, IBM i, and Linux
EC63	2CF1	PCIe4 1-port 100 Gb EDR InfiniBand CAPI adapter	0	1	Linux
EC65	2CF2	PCIe4 2-port 100 Gb EDR InfiniBand CAPI adapter	0	1	Linux
EC66	2CF3	PCIe4 2-port 100 Gb ROCE EN adapter	0	1	AIX, IBM i, and Linux

Table 2-23 lists the available LAN adapters for a Power S924 server.

Table 2-23 Available LAN adapters in Power S924 servers.

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0W	2CC4	PCIe2 2-port 10/1 GbE BaseT RJ45 Adapter	0	25	AIX, IBM i, and Linux
EN0U	2CC3	PCIe2 4-port (10 Gb+1 GbE) Copper SFP+RJ45 Adapter	0	25	AIX, IBM i, and Linux
EN0S	2CC3	PCIe2 4-Port (10 Gb+1 GbE) SR+RJ45 Adapter	0	25	AIX, IBM i, and Linux
5899	576F	PCIe2 4-port 1 GbE Adapter	0	26	AIX, IBM i, and Linux
EC3U	2CEB	PCIe3 1-port 100 Gb EDR IB Adapter x16	0	3	Linux
EC2S	58FA	PCIe3 2-Port 10 Gb NIC & ROCE SR/Cu Adapter	0	13	AIX, IBM i, and Linux
EC38	57BC	PCIe3 2-port 10 GbE NIC&RoCE SFP+ Copper Adapter	0	25	AIX, IBM i, and Linux
EC3F	2CEA	PCIe3 2-port 100 Gb EDR IB Adapter x16	0	3	Linux
EC3M	2CEC	PCIe3 2-port 100 GbE (NIC & RoCE) QSFP28 Adapter x16	0	3	AIX, IBM i, and Linux
EC2U	58FB	PCIe3 2-Port 25/10 Gb NIC & ROCE SR/Cu Adapter	0	13	AIX, IBM i, and Linux
EC3B	57BD	PCIe3 2-Port 40 GbE NIC RoCE QSFP+ Adapter	0	25	AIX, IBM i, and Linux
EN0K	2CC1	PCIe3 4-port (10 Gb FCoE & 1 GbE) SFP+Copper & RJ45	0	25	AIX, IBM i, and Linux
EN0H	2B93	PCIe3 4-port (10 Gb FCoE & 1 GbE) SR & RJ45	0	25	AIX, IBM i, and Linux
EN15	2CE3	PCIe3 4-port 10 GbE SR Adapter	0	25	AIX, IBM i, and Linux
EC63	2CF1	PCIe4 1-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux
EC65	2CF2	PCIe4 2-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux

2.5.3 Graphics accelerator adapters

An adapter can be configured to operate in either 8-bit or 24-bit color modes. The adapter supports both analog and digital monitors.

Table 2-24 on page 90 lists the available graphics accelerator adapter for the Power S922 server.

Table 2-24 The graphics accelerator adapter that is supported in the Power S922 server

Feature code	CCIN	Description	Minimum	Maximum	OS support
5269	5269	PCIe LP POWER GXT145 Graphics Accelerator	0	6	AIX and Linux

Table 2-25 lists the available graphics accelerator adapter for the Power S914 server.

Table 2-25 The graphics accelerator adapter that is supported in the Power S914 server

Feature code	CCIN	Description	Minimum	Maximum	OS support
5748	5269	POWER GXT145 PCI Express Graphics Accelerator	0	4	AIX and Linux

Table 2-26 lists the available graphics accelerator adapter for the Power S924 server.

Table 2-26 The graphics accelerator card that is supported in the Power S924 server

Feature code	CCIN	Description	Minimum	Maximum	OS support
5748	5269	POWER GXT145 PCI Express Graphics Accelerator	0	7	AIX and Linux

2.5.4 SAS adapters

Table 2-27 lists the SAS adapters that are available for the Power S922 server.

Table 2-27 The PCIe SAS adapters that are available for the Power S922 server

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ0J	57B4	PCIe3 RAID SAS Adapter Quad-port 6 Gb x8	0	8	AIX, IBM i, and Linux
EJ0L	57CE	PCIe3 12 GB Cache RAID SAS Adapter Quad-port 6 Gb x8	0	19	AIX, IBM i, and Linux
EJ0M	57B4	PCIe3 LP RAID SAS Adapter Quad-Port 6 Gb x8	0	7	AIX, IBM i, and Linux
EJ10	57B4	PCIe3 SAS Tape/DVD Adapter Quad-port 6 Gb x8	0	12	AIX, IBM i, and Linux
EJ11	57B4	PCIe3 LP SAS Tape/DVD Adapter Quad-port 6 Gb x8	0	7	AIX, IBM i, and Linux
EJ14	57B1	PCIe3 12 GB Cache RAID PLUS SAS Adapter Quad-port 6 Gb x8	0	8	AIX, IBM i, and Linux

Table 2-28 lists the SAS adapters that are available for the Power S914 server.

Table 2-28 The PCIe SAS adapters that are available for the Power S914 server

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ0J	57B4	PCIe3 RAID SAS Adapter Quad-port 6 Gb x8	0	10	AIX, IBM i, and Linux
EJ0L	57CE	PCIe3 12 GB Cache RAID SAS Adapter Quad-port 6 Gb x8	0	19	AIX, IBM i, and Linux
EJ10	57B4	PCIe3 SAS Tape/DVD Adapter Quad-port 6 Gb x8	0	12	AIX, IBM i, and Linux
EJ14	57B1	PCIe3 12 GB Cache RAID PLUS SAS Adapter Quad-port 6 Gb x8	0	8	AIX, IBM i, and Linux

Table 2-29 lists the SAS adapters that are available for Power S924 servers.

Table 2-29 The PCIe SAS adapters that are available for Power S924 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ0J	57B4	PCIe3 RAID SAS Adapter Quad-port 6 Gb x8	0	19	AIX, IBM i, and Linux
EJ0L	57CE	PCIe3 12 GB Cache RAID SAS Adapter Quad-port 6 Gb x8	0	19	AIX, IBM i, and Linux
EJ10	57B4	PCIe3 SAS Tape/DVD Adapter Quad-port 6 Gb x8	0	24	AIX, IBM i, and Linux
EJ14	57B1	PCIe3 12 GB Cache RAID PLUS SAS Adapter Quad-port 6 Gb x8	0	16	AIX, IBM i, and Linux

2.5.5 Fibre Channel adapters

The servers support direct or SAN connection to devices that use Fibre Channel adapters.

Note: If you are attaching a device or switch with an SC type fiber connector, then an LC-SC 50-Micron Fiber Converter Cable (#2456) or an LC-SC 62.5-Micron Fiber Converter Cable (#2459) is required.

Table 2-30 summarizes the available Fibre Channel adapters for Power S922 servers. They all have LC connectors.

Table 2-30 The PCIe Fibre Channel adapters that are available for Power S922 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0A	577F	PCIe3 16 Gb 2-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN0B	577F	PCIe3 LP 16 Gb 2-port Fibre Channel Adapter	0	8	AIX, IBM i, and Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0F	578D	PCIe2 LP 8 Gb 2-Port Fibre Channel Adapter	0	8	AIX, IBM i and Linux ^a
EN0G	578D	PCIe2 8 Gb 2-Port Fibre Channel Adapter	0	12	AIX, IBM i and Linux ^a
EN0Y	N/A	PCIe2 LP 8 Gb 4-port Fibre Channel Adapter	0	8	AIX, IBM i, and Linux
EN12		PCIe2 8 Gb 4-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN1A	578F	PCIe3 32 Gb 2-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN1B	578F	PCIe3 LP 32 Gb 2-port Fibre Channel Adapter	0	8	AIX, IBM i, and Linux
EN1C	578E	PCIe3 16 Gb 4-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN1D	578E	PCIe3 LP 16 Gb 4-port Fibre Channel Adapter	0	8	AIX, IBM i, and Linux

a. IBM i requires VIOS

Table 2-31 summarizes the available Fibre Channel adapters for Power S914 servers. They all have LC connectors.

Table 2-31 The PCIe Fibre Channel adapters that are available for Power S914 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0A	577F	PCIe3 16 Gb 2-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN0G	5789	PCIe2 8 Gb 2-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN12		PCIe2 8 Gb 4-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN1A	578F	PCIe3 32 Gb 2-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN1C	578E	PCIe3 16 Gb 4-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux

Table 2-32 summarizes the available Fibre Channel adapters for Power S924 servers. They all have LC connectors.

Table 2-32 The PCIe Fibre Channel adapters that are available for Power S924 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0A	577F	PCIe3 16 Gb 2-port Fibre Channel Adapter	0	25	AIX, IBM i, and Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EN0G	578D	PCIe2 8 Gb 2-Port Fibre Channel Adapter	0	25	AIX, IBM i, and Linux
EN12		PCIe2 8 Gb 4-port Fibre Channel Adapter	0	12	AIX, IBM i, and Linux
EN1A	578F	PCIe3 32 Gb 2-port Fibre Channel Adapter	0	25	AIX, IBM i, and Linux
EN1C	578E	PCIe3 16 Gb 4-port Fibre Channel Adapter	0	25	AIX, IBM i, and Linux

Note: The usage of N_Port ID Virtualization (NPIV) through the Virtual I/O Server (VIOS) requires an NPIV-capable Fibre Channel adapter, such as the #EN0A.

2.5.6 InfiniBand host channel adapter

The InfiniBand Architecture (IBA) is an industry-standard architecture for server I/O and inter-server communication. It was developed by the InfiniBand Trade Association (IBTA) to provide the levels of reliability, availability, performance, and scalability that are necessary for present and future server systems with levels better than can be achieved by using bus-oriented I/O structures.

InfiniBand is an open set of interconnect standards and specifications. The main InfiniBand specification is published by the IBTA and is available at the [IBTA website](#).

InfiniBand is based on a switched fabric architecture of serial point-to-point links, where these InfiniBand links can be connected to either host channel adapters (HCAs), which are used primarily in servers, or target channel adapters (TCAs), which are used primarily in storage subsystems.

The InfiniBand physical connection consists of multiple byte lanes. Each individual byte lane is a four-wire, 2.5, 5.0, or 10.0 Gbps bidirectional connection. Combinations of link width and byte lane speed allow for overall link speeds of 2.5 - 120 Gbps. The architecture defines a layered hardware protocol and also a software layer to manage initialization and the communication between devices. Each link can support multiple transport services for reliability and multiple prioritized virtual communication channels.

For more information about InfiniBand, see *HPC Clusters Using InfiniBand on IBM Power Systems Servers*, SG24-7767.

A connection to supported InfiniBand switches is accomplished by using the QDR optical cables (#3290 and #3293).

Table 2-33 lists the InfiniBand adapters that are available for Power S922 servers.

Table 2-33 InfiniBand adapters available for Power S922 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3E	2CEA	PCIe3 LP 2-port 100 Gb EDR InfiniBand Adapter x16	0	3	Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3T	2CEB	PCIe3 LP 1-port 100 Gb EDR InfiniBand Adapter x16	0	3	Linux
EC62	2CF1	PCIe4 LP 1-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux
EC64	2CF2	PCIe4 LP 2-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux

Table 2-34 lists the InfiniBand adapters that are available for Power S914 servers.

Table 2-34 InfiniBand adapters available for Power S914 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3u		PCIe3 1-port 100 Gb EDR IB Adapter x16	0	1	Linux
EC63	2CF1	PCIe4 1-port 100 Gb EDR InfiniBand CAPI adapter	0	1	Linux
EC65	2CF2	PCIe4 2-port 100 Gb EDR InfiniBand CAPI adapter	0	1	Linux

Table 2-35 lists the InfiniBand adapters available for Power S924 servers.

Table 2-35 InfiniBand adapters available for Power S924 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3u		PCIe3 1-port 100 Gb EDR IB Adapter x16	0	1	Linux
EC63	2CF1	PCIe4 1-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux
EC65	2CF2	PCIe4 2-port 100 Gb EDR InfiniBand CAPI adapter	0	3	Linux

2.5.7 Cryptographic coprocessor

The cryptographic coprocessor card that is supported for the Power S922 server is shown in Table 2-36.

Table 2-36 The cryptographic coprocessor available for Power S9222 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ33	4767	PCIe3 Crypto Coprocessor BSC-Gen3 4767	0	12	AIX, IBM i, and Linux

The cryptographic coprocessor cards that are supported for the Power S914 server are shown in Table 2-37.

Table 2-37 Cryptographic coprocessors available for Power S914 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ32	4767	PCIe3 Crypto Coprocessor no BSC 4767	0	7	AIX, IBM i, and Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ33	4767	PCIe3 Crypto Coprocessor BSC-Gen3 4767	0	6	AIX, IBM i, and Linux

The cryptographic coprocessor cards that are supported for the Power S924 server are shown in Table 2-38.

Table 2-38 Cryptographic coprocessors available for Power S924 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EJ32	4767	PCIe3 Crypto Coprocessor no BSC 4767	0	10	AIX, IBM i, and Linux
EJ33	4767	PCIe3 Crypto Coprocessor BSC-Gen3 4767	0	18	AIX, IBM i, and Linux

2.5.8 Coherent Accelerator Processor Interface adapters

The CAPI-capable adapters that are available for Power S922 servers are shown in Table 2-39.

Table 2-39 CAPI-capable adapters available for Power S922 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3L	2CEC	PCIe3 LP 2-port 100 GbE (NIC & RoCE) QSFP28 Adapter x16	0	3	AIX, IBM i, and Linux
EC62	2CF1	PCIe4 LP 1-port 100 Gb EDR InfiniBand CAPI adapter	0	3	AIX, IBM i, and Linux
EC64	2CF2	PCIe4 LP 2-port 100 Gb EDR InfiniBand CAPI adapter	0	3	AIX, IBM i, and Linux

The CAPI-capable adapters that are available for Power S914 servers are shown in Table 2-40.

Table 2-40 CAPI-capable adapters available for Power S914 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3M	2CEC	PCIe3 2-port 100 GbE (NIC & RoCE) QSFP28 Adapter x16	0	1	AIX, IBM i, and Linux
EC63	2CF1	PCIe4 1-port 100 Gb EDR InfiniBand CAPI adapter	0	1	AIX, IBM i, and Linux
EC65	2CF2	PCIe4 2-port 100 Gb EDR InfiniBand CAPI adapter	0	1	AIX, IBM i, and Linux

The CAPI-capable adapters that are available for Power S924 servers are shown in Table 2-41 on page 96.

Table 2-41 CAPI-capable adapters available for Power S924 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC3M	2CEC	PCIe3 2-port 100 GbE (NIC & RoCE) QSFP28 Adapter x16	0	3	AIX, IBM i, and Linux
EC63	2CF1	PCIe4 1-port 100 Gb EDR InfiniBand CAPI adapter	0	3	AIX, IBM i, and Linux
EC65	2CF2	PCIe4 2-port 100 Gb EDR InfiniBand CAPI adapter	0	3	AIX, IBM i, and Linux

2.5.9 USB adapters

The USB adapters that are available for Power S922 servers are shown in Table 2-39 on page 95.

Table 2-42 USB adapters available for Power S922 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC45	58F9	PCIe2 LP 4-Port USB 3.0 Adapter	0	8	AIX, IBM i, and Linux
EC46	58F9	PCIe2 4-Port USB 3.0 Adapter	0	12	AIX, IBM i, and Linux

The USB adapters that are available for Power S914 servers are shown in Table 2-40 on page 95.

Table 2-43 USB adapters available for Power S914 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC46	58F9	PCIe2 4-Port USB 3.0 Adapter	0	12	AIX, IBM i, and Linux
EJ32	4767	PCIe3 Crypto Coprocessor no BSC 4767	0	7	AIX, IBM i, and Linux

The USB adapters that are available for Power S924 servers are shown in Table 2-41.

Table 2-44 USB adapters available for Power S924 servers

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC46	58F9	PCIe2 4-Port USB 3.0 Adapter	0	25	AIX, IBM i, and Linux
EJ32	4767	PCIe3 Crypto Coprocessor no BSC 4767	0	10	AIX, IBM i, and Linux

2.6 Internal storage

The internal storage on the Power S922, S914, and Power S924 servers depends on the DASD/Media backplane that is used. The servers support various DASD/Media backplanes:

Power S922 servers support:

#EJ1F	Eight SFF-3 with optional split card (#EJ1H)
#EJ1G	Expanded function Storage Backplane 8 SFF-3 Bays/Single IOA with write cache
#EC59	Optional PCIe3 NVMe carrier card with two M.2 module slots

Note: If #EC59 is ordered, a minimum of one #ES14 Mainstream 400 GB solid-state drive (SSD) NVMe M.2 module must be ordered.

Power S914 and Power S924 servers support:

#EJ1C	Twelve SFF-3 Bays with an optional split card (#EJ1E)
#EJ1D	Eighteen SFF-3 Bays/Dual IOA with Write Cache
#EJ1M	Twelve SFF-3 Bays/RDX Bays with write cache and optional external SAS port
#EC59	Optional PCIe3 Non-Volatile Memory express (NVMe) carrier card with two M.2 module slots

IBM i operating system performance: Clients with write-sensitive disk/hard disk drive (HDD) workloads should upgrade from the base storage backplane (#EJ1C/#EJ1E) to the expanded function storage backplanes (#EJ1M/#EJ1D) to gain the performance advantage of write cache

2.6.1 Backplane (#EJ1C)

This FC is the base storage backplane with an integrated SAS controller for SAS bays in the system unit. SAS bays are 2.5-inch or SFF and use drives that are mounted on a carrier/tray that is specific to the system unit (SFF-3).

The high-performance SAS controller provides RAID 0, RAID 5, RAID 6, or RAID 10 support for either HDDs or SSDs. JBOD support for HDD is also supported. The controller has no write cache.

For servers that support split backplane capability, add #EJ1E. For write cache performance, use #EJ1D or #EJ1M instead of this backplane.

Both 5xx and 4-KB sector HDDs/SSDs are supported. 5xx and 4-KB drives cannot be mixed in the same array.

This FC provides a storage backplane with one integrated SAS adapter with no cache, running 12 SFF-3 SAS bays in the system unit and one RDX bay in the system unit.

Supported operating systems:

- ▶ Red Hat Enterprise Linux
- ▶ SUSE Linux Enterprise Server
- ▶ Ubuntu Server

- ▶ AIX
- ▶ IBM i

The internal connections to the physical disks are shown in Figure 2-15 on page 98.

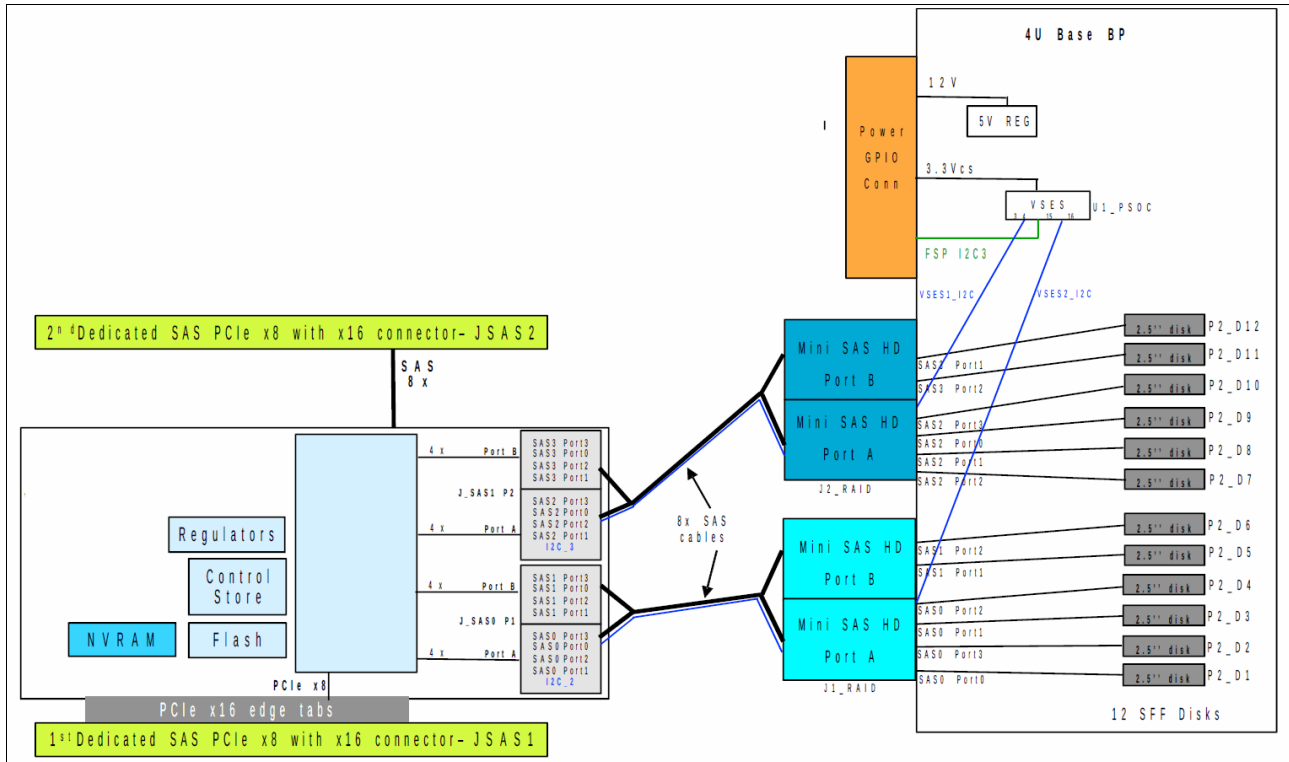


Figure 2-15 #EJ1C connections

2.6.2 Split backplane option (#EJ1E)

This FC modifies the base storage backplane cabling and adds a second, high-performance SAS controller. The existing 12 SFF-3 SAS bays are cabled to be split into two sets of six bays, each with one SAS controller. Both SAS controllers are in integrated slots and do not use a PCIe slot.

The high-performance SAS controllers each provides RAID 0, RAID 5, RAID 6, or RAID 10 support. JBOD support for HDDs is also supported. There is no write cache on either controller.

Both 5x and 4-KB sector HDDs/SSDs are supported. 5x and 4-KB drives cannot be mixed in the same array.

This FC provides a second integrated SAS adapter with no cache and internal cables to provide two sets of six SFF-3 bays in the system unit.

You must have an #EJ1C backplane to use this FC.

Supported operating systems:

- ▶ Red Hat Enterprise Linux
- ▶ SUSE Linux Enterprise Server
- ▶ Ubuntu Server
- ▶ AIX

► IBM i

The internal connections to the physical disks are shown in Figure 2-16 on page 99.

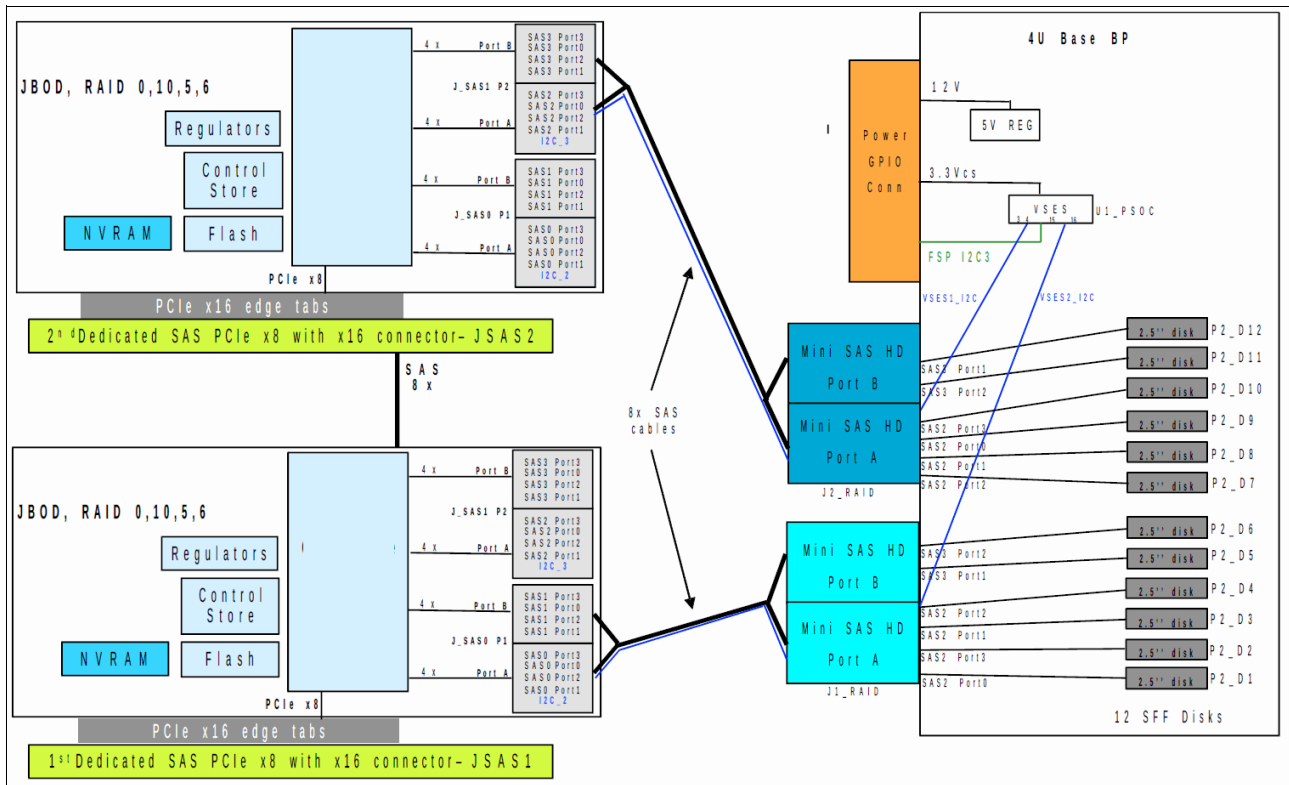


Figure 2-16 #EJ1E connections

2.6.3 Backplane (#EJ1D)

This FC is an expanded function storage backplane with dual integrated SAS controllers with write cache and an optional external SAS port. High-performance controllers run SFF-3 SAS bays in the system unit. Dual controllers (also called dual I/O adapters or paired controllers) and their write cache are placed in integrated slots and do not use PCIe slots. The write cache augments the controller's high performance for workloads with writes, especially for HDDs. A 1.8-GB physical write cache is used with compression to provide up to 7.2-GB cache capacity. The write cache contents are protected against power loss by using flash memory and super capacitors, which removes the need for battery maintenance.

The high-performance SAS controllers provide RAID 0, RAID 5, RAID 6, RAID 10, RAID 5T2, RAID 6T2, or RAID 10T2 support. Patented Active/Active configurations with at least two arrays are supported.

The Easy Tier function is supported so that the dual controllers can automatically move hot data to attached SSDs and cold data to attached HDDs for AIX and Linux, and VIOS environments.

SFF or 2.5-inch drives are mounted on a carrier/ tray that is specific to the system unit (SFF-3). The backplane has 18 SFF-3 bays.

This backplane enables two SAS ports (#EJ0W) at the rear of the system unit to support the attachment of one EXP24S/EXP24SX I/O drawer in mode1 to hold HDDs or SSDs.

#EJ0W is an optional feature with #EJ1M, and one 8 PCIe slot is used by #EJ0W.

This backplane does not support a split backplane. For a split backplane, use #EJ1C plus #EJ1E.

Both 5xx and 4-KB sector HDDs/SSDs are supported. 5xx and 4-KB drives cannot be mixed in the same array.

This FC provides a storage backplane with a pair of integrated SAS adapters with write cache, with an optional external SAS port running up to:

- ▶ A set of 18 SFF-3 SAS bays in the system unit
- ▶ Two SAS ports at the rear of the system unit to connect to a single EXP24S/EXP24SX I/O drawer

Supported operating systems:

- ▶ Red Hat Enterprise Linux
- ▶ SUSE Linux Enterprise Server
- ▶ Ubuntu Server
- ▶ AIX
- ▶ IBM i

2.6.4 Expanded Function Backplane (#EJ1M)

This FC is an expanded function storage backplane with dual integrated SAS controllers with write cache and an optional external SAS port. High-performance controllers run SFF-3 SAS bays and an RDX bay in the system unit. Dual controllers (also called dual I/O adapters or paired controllers) and their write cache are placed in integrated slots and do not use PCIe slots. A write cache augments the controller's high performance for workloads with writes, especially for HDDs. A 1.8-GB physical write cache is used with compression to provide up to 7.2-GB cache capacity. The write cache contents are protected against power loss by using flash memory and super capacitors, which removes the need for battery maintenance.

The high-performance SAS controllers provide RAID 0, RAID 5, RAID 6, RAID 10, RAID 5T2, RAID 6T2, or RAID 10T2 support. Patented Active/Active configurations with at least two arrays are supported.

The Easy Tier function is supported so that the dual controllers can automatically move hot data to attached SSDs and cold data to attached HDDs for AIX and Linux, and VIOS environments.

SFF or 2.5-inch drives are mounted on a carrier/ tray that is specific to the system unit (SFF-3). The backplane has 12 SFF-3 bays.

This backplane also enables two SAS ports (#EJ0W) at the rear of the system unit, which support the attachment of one EXP24S/EXP24SX I/O drawer in mode1, which holds HDDs or SSDs.

#EJ0W is an optional feature with #EJ1M, and one 8 PCIe slot is used by #EJ0W.

This backplane does not support a split backplane. For a split backplane, use the #EJ1C with #EJ1E backplane features.

Both 5xx and 4-KB sector HDDs/SSDs are supported. 5xx and 4-KB drives cannot be mixed in the same array.

This FC provides an expanded function storage backplane with a pair of integrated SAS adapters with a write cache, and an optional external SAS port running up to:

- ▶ A set of 12 SFF-3 SAS bays in the system unit
- ▶ One RDX bay in the system unit
- ▶ Two SAS ports at the rear of the system unit to connect to a single EXP24S/EXP24SX I/O drawer

Supported operating systems:

- ▶ Red Hat Enterprise Linux
- ▶ SUSE Linux Enterprise Server
- ▶ Ubuntu Server
- ▶ AIX
- ▶ IBM i

2.6.5 NVMe support

This section provides information about the available NVMe carriers and adapters.

NVMe adapters

Table 2-45 provides a list of available NVMe cards for the S922 server.

Table 2-45 Available NVMe adapters for the S922

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC5G	58FC	PCIe3 LP 1.6 TB SSD NVMe Adapter	0	5	AIX and Linux
EC5C	58FD	PCIe3 LP 3.2 TB SSD NVMe adapter	0	5	AIX and Linux
EC5E	58FE	PCIe3 LP 6.4 TB SSD NVMe adapter	0	5	AIX and Linux

Table 2-46 provides a list of available NVMe cards for the S914 server.

Table 2-46 Available NVMe adapters for the S914

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC5B	58FC	PCIe3 LP 1.6 TB SSD NVMe Adapter	0	3	AIX and Linux
EC5D	58FD	PCIe3 LP 3.2 TB SSD NVMe adapter	0	3	AIX and Linux
EC5F	58FE	PCIe3 LP 6.4 TB SSD NVMe adapter	0	3	AIX and Linux
EC6V	58FC	PCIe3 x8 1.6 TB NVMe Flash Adapter for IBM i	0	3	IBM i
EC6X	58FD	PCIe3 x8 3.2 TB NVMe Flash Adapter for IBM i	0	3	IBM i
EC6Z	58FE	PCIe3 x8 6.4 TB NVMe Flash Adapter for IBM i	0	3	IBM i

Table 2-47 provides a list of available NVMe cards for the S924 server.

Table 2-47 Available NVMe adapters for the S924

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC5B	58FC	PCIe3 LP 1.6 TB SSD NVMe Adapter	0	5	AIX and Linux

Feature code	CCIN	Description	Minimum	Maximum	OS support
EC5D	58FD	PCIe3 LP 3.2 TB SSD NVMe adapter	0	5	AIX and Linux
EC5F	58FE	PCIe3 LP 6.4 TB SSD NVMe adapter	0	5	AIX and Linux
EC6V	58FC	PCIe3 x8 1.6 TB NVMe Flash Adapter for IBM i	0	3	IBM i
EC6X	58FD	PCIe3 x8 3.2 TB NVMe Flash Adapter for IBM i	0	3	IBM i
EC6Z	58FE	PCIe3 x8 6.4 TB NVMe Flash Adapter for IBM i	0	3	IBM i

PCIe3 NVMe carrier card with two M.2 module slots (#EC59)

The NVMe option offers fast start times, and is ideally suited to housing the rootvg of VIOS partitions.

#EC59 is a carrier card for 400 GB Mainstream SSD NVMe module (#ES14). The maximum quantity is two of #ES14 per #EC59.

This FC provides an PCIe3 NVMe card with two M.2 module slots.

Supported operating systems:

- ▶ SUSE Linux Enterprise Server 12 Service Pack 3 or later
- ▶ SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 12 Service Pack 3 or later
- ▶ Red Hat Enterprise Linux
- ▶ Ubuntu Server
- ▶ AIX

If the NVMe carrier card #EC59 is selected, you do not have to order disk units. If you do not order SAN Boot (#EC59 or #0837), then you must order at least one disk unit. If you do not order HDD/SSD/SAN boot (#0837), then #EC59 (with at least one of #ES14) is the load source.

Figure 2-17 shows the location of the #EC59 feature code in a Power S922 server.



Figure 2-17 Two #EC59 feature codes in a Power S922 server

Figure 2-18 on page 103 shows an #EC59 feature code.




Figure 2-18 #EC59 feature code with the cover open showing two #ES14 modules fitted

Each NVMe device (#ES14) is a separate PCIe endpoint, which means that each NVMe device can be assigned to a different logical partition (LPAR) or Virtual I/O Server (VIOS). At the operating system level, each #ES14 device appears to the operating system as an individual disk. For example, in AIX, the device might appear as hdisk0.

Tip: If two #EC59 feature codes are configured, each with two #ES14 feature codes, it is possible to have the rootvg of the first VIOS mirrored to an #ES14 in each #EC59, and the second VIOS could be mirrored to the other two modules, which provides excellent performance and resilience.

Or, each #ES14 could be assigned to a separate partition or VIOS server as a boot device.

400 GB SSD NVMe M.2 module (#ES14)

This feature code is a 400 GB Mainstream SSD formatted in 4096 byte sectors (4 KB). The drive is mounted on the PCIe NVMe carrier card with 2 M.2 SOCKETS (#EC59). The Drive Write Per Day (DWPD) rating is 1 calculated over a 5-year period. Approximately 1,095 TB of data can be written over the life of the drive, but depending on the nature of the workload, this number might be larger. As a preferred practice, use this device for boot support and non-intensive workloads.

Note: You must order, at a minimum, one #ES14 module with each #EC59 that you order. The maximum quantity is two #ES14 per #EC59.

Using this FC for other functions beyond boot support and non-intensive workloads might result in throttled performance and high temperatures that lead to timeouts and critical thermal warnings.

Supported operating systems:

- ▶ SUSE Linux Enterprise Server 12 Service Pack 3 or later
- ▶ SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 12 Service Pack 3 or later
- ▶ Red Hat Enterprise Linux
- ▶ Ubuntu Server
- ▶ AIX

Note:

- ▶ Assignment of a #ES14 feature code to the VIOS is supported.
- ▶ IBM i is *not* supported (IBM i does not support this feature as a Virtual Target Device to IBM i through VIOS).

Figure 2-19 shows two #ES14 NVMe modules.




Figure 2-19 Two NVMe modules (#ES14)

2.6.6 Backplane (#EJ1F)

The backplane option provides SFF-3 SAS bays in the system unit. These 2.5-inch or SFF SAS bays can contain SAS drives (HDDs or SSDs) that are mounted on a Gen3 tray or carrier. Thus, the drives are designated SFF-3. SFF-1 or SFF-2 drives do not fit in an SFF-3 bay. All SFF-3 bays support concurrent maintenance or hot-plug capability.

This backplane option uses leading-edge, integrated SAS RAID controller technology that is designed and patented by IBM. A custom-designed PowerPC based ASIC chip is the basis of these SAS RAID controllers, and provides RAID 0, RAID 5, RAID 6, and RAID 10 functions with HDDs and SSDs. Internally, SAS ports are implemented and provide plentiful bandwidth. The integrated SAS controllers are placed in dedicated slots and do not reduce the number of available PCIe slots.

The Storage Backplane option (#EJ1F) provides eight SFF-3 bays and one SAS controller with zero write cache.

Optionally, by adding the Split Backplane (#EJ1H), a second integrated SAS controller with no write cache is provided, and the eight SFF-3 bays are logically divided into two sets of four bays. Each SAS controller independently runs one of the four-bay sets of drives.

This backplane option supports HDDs or SSDs or a mixture of HDDs and SSDs in the SFF-3 bays. Mixing HDDs and SSDs applies even within a single set of four bays of the split backplane option. If you are mixing HDDs and SSDs, they *must* be in separate arrays (unless you use the Easy Tier function).

This backplane option can offer different drive protection options: RAID 0, RAID 5, RAID 6, or RAID 10. RAID 5 requires a minimum of three drives of the same capacity. RAID 6 requires a minimum of four drives of the same capacity. RAID 10 requires a minimum of two drives. Hot-spare capability is supported by RAID 5, RAID 6, or RAID 10.

RAID 5 and RAID 6 result in more drive write activity than mirroring or unprotected drives.

This backplane option is supported by AIX and Linux, and VIOS. As a preferred practice, the drives should be protected.

2.6.7 Expanded Function Storage Backplane (#EJ1G)

In addition to supporting HDDs and SSDs in the SFF-3 SAS bays, the Expanded Function Storage Backplane (#EJ1G) supports the optional attachment of an EXP12SX/EXP24SX drawer. All bays are accessed by both of the integrated SAS controllers. The bays support concurrent maintenance (hot-plug).

2.6.8 RAID support

There are multiple protection options for HDD/SSD drives in the Power S91, Power S922, and Power S924 servers, whether they are contained in the SAS SFF bays in the system unit or drives in disk-only I/O drawers. Although protecting drives is always preferred, AIX and Linux users can choose to leave a few or all drives unprotected at their own risk, and IBM supports these configurations.

Drive protection

HDD/SSD drive protection can be provided by AIX, IBM i, and Linux, or by the HDD/SSD hardware controllers.

Apart from the #EC59 option, all of the storage backplanes offer RAID. The default storage backplanes (#EJ1C for the Power S914 and Power S924 servers, and #EJ1F and #EJ1G for the Power S922 server) contain one SAS HDD/SSD controller and provide support for JBOD and RAID 0, 5, 6, and 10 for AIX or Linux. A secondary non-redundant controller is added when you use #EJ1E for the Power S914 and Power S924 servers or #EJ1H for the Power S922 server, so each of the six disk bays has a separated disk controller.

When you choose the optional #EJ1D, #EJ1M, or #EJ1H storage backplane, the controller is replaced by a pair of high-performance RAID controllers with dual integrated SAS controllers with 1.8 GB of physical write cache. High-performance controllers run 18 SFF-3 SAS bays with 1.8-inch SSD bays. Dual controllers (also called dual I/O adapters or paired controllers) and their write cache are placed in integrated slots and do not use PCIe slots. Patented active/active configurations with at least two arrays are supported.

The write cache, which is responsible for increasing write performance by caching data before it is written to the physical disks, can have its data compression capabilities activated, which provides up to 7.2 GB effective cache capacity. The write cache contents are protected against power loss by flash memory and super capacitors, which removes the need for battery maintenance.

The high-performance SAS controllers provide RAID 0, RAID 5, RAID 6, and RAID 10 support, and the Easy Tier variants (RAID 5T2, RAID 6T2, and RAID 10T2) if the server has both HDDs and SSDs installed.

The Easy Tier function is supported, so the dual controllers can automatically move hot data to an attached SSD and cold data to an attached HDD for AIX and Linux, and VIOS environments. To learn more about Easy Tier, see 2.6.9, “Easy Tier” on page 106.

AIX and Linux can use disk drives that are formatted with 512-byte blocks when they are mirrored by the operating system. These disk drives must be reformatted to 528-byte sectors when used in RAID arrays. Although a small percentage of the drive's capacity is lost, extra

data protection, such as error-correcting code (ECC) and bad block detection, is gained in this reformatting. For example, a 300 GB disk drive, when reformatted, provides approximately 283 GB. IBM i always uses drives that are formatted to 528 bytes. SSDs are always formatted with 528-byte sectors.

Supported RAID functions

The base hardware supports RAID 0, 5, 6, and 10. When more features are configured, the server supports hardware RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2:

- ▶ RAID 0 provides striping for performance, but does not offer any fault tolerance.

The failure of a single drive results in the loss of all data on the array. This version of RAID increases I/O bandwidth by simultaneously accessing multiple data paths.

- ▶ RAID 5 uses block-level data striping with distributed parity.

RAID 5 stripes both data and parity information across three or more drives. Fault tolerance is maintained by ensuring that the parity information for any given block of data is placed on a drive that is separate from the ones that are used to store the data itself. This version of RAID provides data resiliency if a single drive fails in a RAID 5 array.

- ▶ RAID 6 uses block-level data striping with dual distributed parity.

RAID 6 is the same as RAID 5 except that it uses a second level of independently calculated and distributed parity information for more fault tolerance. A RAID 6 configuration requires N+2 drives to accommodate the additional parity data, making it less cost-effective than RAID 5 for equivalent storage capacity. This version of RAID provides data resiliency if one or two drives fail in a RAID 6 array. When you work with large capacity disks, RAID 6 enables you to sustain data parity during the rebuild process.

- ▶ RAID 10 is a striped set of mirrored arrays.

It is a combination of RAID 0 and RAID 1. A RAID 0 stripe set of the data is created across a two-disk array for performance benefits. A duplicate of the first stripe set is then mirrored on another two-disk array for fault tolerance. This version of RAID provides data resiliency if a single drive fails, and it can provide resiliency for multiple drive failures.

RAID 5T2, RAID 6T2, and RAID 10T2 are RAID levels with Easy Tier enabled. They require that both types of disks exist on the system under the same controller (HDDs and SSDs), and that both types are configured under the same RAID type.

2.6.9 Easy Tier

With a standard backplane (#EJ1C or #EJ1F), the server can handle both HDDs and SSDs that are attached to its storage backplane if they are on separate arrays.

The high-function backplane (#EJ1D) can handle both types of storage in two different ways:

- ▶ Separate arrays: SSDs and HDDs coexist on separate arrays, just like the Standard SAS adapter can.
- ▶ Easy Tier: SSDs and HDDs coexist under the same array.

When the SSDs and HDDs are under the same array, the adapter can automatically move the most accessed data to faster storage (SSDs) and less accessed data to slower storage (HDDs). This is called *Easy Tier*.

There is no need for coding or software intervention after the RAID is configured correctly. Statistics on block accesses are gathered every minute, and after the adapter realizes that some portion of the data is being frequently requested, it moves this data to faster devices. The data is moved in chunks of 1 MB or 2 MB called *bands*.

From the operating system point-of-view, there is just a regular array disk. From the SAS controller point-of-view, there are two arrays with parts of the data being serviced by one tier of disks and parts by another tier of disks.

Figure 2-20 shows an Easy Tier array.




Figure 2-20 Easy Tier array

The Easy Tier configuration is accomplished through a standard operating system SAS adapter configuration utility. Figure 2-21 and Figure 2-22 show two examples of tiered array creation for AIX.

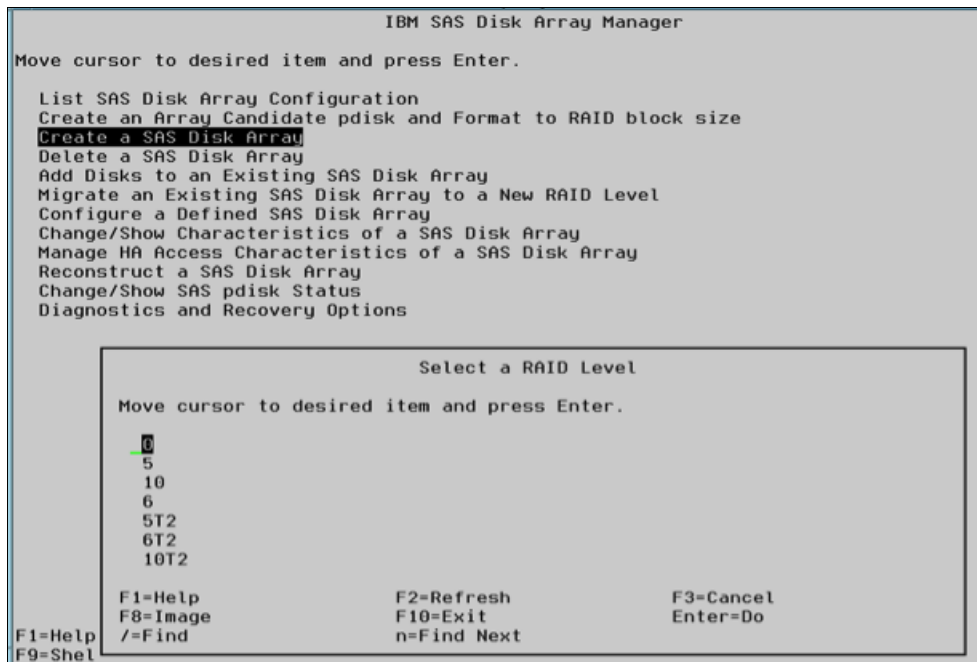


Figure 2-21 Array type selection panel on AIX RAID Manager

Name	Resource	State	Description	Size	
sissas1	FEFFFFFF	Primary	PCIe3 12GB Cache RAID SAS Adapter Quad-port 6Gb x8		
sissas0	FEFFFFFF	HA Linked	Remote adapter SN 00325001		
hdisk1	FC0000FF	Optimal	RAID 5T2 Array (N/N)	773.5GB	← RAID 5T2
pdisk0	000400FF	Active	Array Member	139.6GB	RAID 5 SSD - 2 + 1 x 177.8 GB RAID 5 HDD - 3 + 1 x 139.6 GB
pdisk1	000401FF	Active	Array Member	139.6GB	
pdisk2	000402FF	Active	Array Member	139.6GB	
pdisk3	000403FF	Active	Array Member	139.6GB	
pdisk7	000407FF	Active	SSD Array Member	177.8GB	
pdisk6	000406FF	Active	SSD Array Member	177.8GB	
pdisk8	000408FF	Active	SSD Array Member	177.8GB	
hdisk2	FC0100FF	Optimal	RAID 6T2 Array (N/N)	1090GB	
pdisk10	00040AFF	Active	SSD Array Member	387.9GB	RAID 6 SSD - 3 + 2 x 387.9 GB RAID 6 HDD - 4 + 2 x 139.6 GB
pdisk11	00040BFF	Active	SSD Array Member	387.9GB	
pdisk4	000404FF	Active	Array Member	139.6GB	
pdisk20	000414FF	Active	SSD Array Member	387.9GB	
pdisk21	000415FF	Active	SSD Array Member	387.9GB	
pdisk9	000409FF	Active	SSD Array Member	177.8GB	
pdisk5	000405FF	Active	Array Member	139.6GB	
pdisk12	00040CFF	Active	Array Member	139.6GB	
pdisk13	00040DFF	Active	Array Member	139.6GB	
pdisk14	00040EFF	Active	Array Member	139.6GB	
pdisk15	00040FFF	Active	Array Member	139.6GB	
hdisk3	FC0200FF	Optimal	RAID 10T2 Array (0/0)	666.6GB	← RAID 10T2
pdisk22	000416FF	Active	SSD Array Member	387.9GB	RAID 10 SSD - 1 + 1 x 387.9 GB RAID 10 HDD - 2 + 2 x 139.6 GB
pdisk23	000417FF	Active	SSD Array Member	387.9GB	
pdisk16	000410FF	Active	Array Member	139.6GB	
pdisk17	000411FF	Active	Array Member	139.6GB	
pdisk18	000412FF	Active	Array Member	139.6GB	
pdisk19	000413FF	Active	Array Member	139.6GB	

Figure 2-22 Tiered arrays (RAID 5T2, RAID 6T2, and RAID 10T2) example on AIX RAID Manager

To support Easy Tier, make sure that the server is running at least the following minimum levels:

- ▶ VIOS 2.2.3.3 with interim fix IV56366 or later
- ▶ AIX 7.1 TL3 SP3 or later

- ▶ AIX 6.1 TL9 SP3 or later
- ▶ RHEL 6.5 or later
- ▶ SUSE Linux Enterprise Server 11 SP3 or later

IBM i supports Easy Tier only through VIOS.

2.6.10 External SAS ports

The Power S914 and Power S924 DASD backplanes (#EJ1D and #EJ1M) offer a connection to an external SAS port.

Power S914 and Power S924 servers that use the high-performance RAID feature support two external SAS ports. The external SAS ports are used for expansion to an external SAS drawer.

More drawers and the IBM System Storage 7226 Tape and DVD Enclosure Express (Model 1U3) can be attached by installing more SAS adapters.

Note: Only one SAS drawer is supported from the external SAS port. More SAS drawers can be supported through SAS adapters.

2.6.11 Media drawers

IBM multimedia drawers, such as the 7226-1U3 or 7214-1U2, or tape units, such as the TS2240, TS2340, TS3100, TS3200, and TS3310, can be connected by using external SAS ports.

2.6.12 External DVD drives

There is a trend to use good quality USB flash drives rather than DVD drives. Being mechanical, DVD drives are less reliable than solid-state technology.

If you feel that you do need a DVD drive, IBM offers a stand-alone external USB unit (#EUA5), which is shown in Figure 2-23.



Figure 2-23 Stand-alone USB DVD drive with cable (#EUA5)

Note: If you use an external/stand-alone USB drive, which does not have its own power supply, you should use a USB socket at the front of the system to ensure enough current is available.

2.6.13 RDX removable disk drives

These Power Systems servers support RDX removable disk drives, which are commonly used for quick backups.

If #EJ1C or #EJ1M is configured, an internal bay is available and can be populated by #EU00.

There is also an external/stand-alone RDX unit (#EUA4).

Various disk drives are available, as shown in Table 2-48.

Table 2-48 RDX disk drives

Feature code	Part number	Description
#1107	46C5379	500 GB Removable Disk Drive
#EU01	46C2335	1 TB Removable Disk Drive
#EU2T	46C2975	2 TB Removable Disk Drive

2.7 External IO subsystems

This section describes the PCIe Gen3 I/O expansion drawer that can be attached to the Power S922, Power S914, and Power S924 servers.

2.7.1 Peripheral Component Interconnect Express Gen3 I/O expansion drawer

The PCIe Gen3 I/O expansion (EMX0) drawer is a 4U high, PCI Gen3-based, and rack-mountable I/O drawer. It offers two PCIe fan-out modules (#EMXF or #EMXG). The PCIe fan-out module provides six PCIe Gen3 full-high, full-length slots (two x16 and four x8). The PCIe slots are hot-pluggable.

The PCIe fan-out module has two CXP ports, which are connected to two CXP ports on a PCIe Optical Cable Adapter (#EJ05, #EJ07, or #EJ08, depending on the server that is selected). A pair of active optical CXP cables (AOCs) or a pair of CXP copper cables are used for this connection.

Concurrent repair and add/removal of PCIe adapters is done by Hardware Management Console (HMC) guided menus or by operating system support utilities.

A BSC is used to house the full-high adapters that go into these slots. The BSC is the same BSC that is used with the previous generation server's #5802/5803/5877/5873 12X attached I/O drawers.

Figure 2-24 shows the back view of the PCIe Gen3 I/O expansion drawer.




Figure 2-24 Rear view of the PCIe Gen3 I/O expansion drawer

2.7.2 PCIe Gen3 I/O expansion drawer optical cabling

I/O drawers are connected to the adapters in the system node with data transfer cables:

- ▶ 3M Optical Cable Pair for PCIe3 Expansion Drawer (#ECC7)
- ▶ 10M Optical Cable Pair for PCIe3 Expansion Drawer (#ECC8)
- ▶ 3M Copper CXP Cable Pair for PCIe3 Expansion Drawer (#ECCS)

Cable lengths: Use the 3.0 m cables for intra-rack installations. Use the 10.0 m cables for inter-rack installations.

Limitation: You cannot mix copper and optical cables on the same PCIe Gen3 I/O drawer. Both fan-out modules use copper cables or both use optical cables.

A minimum of one PCIe3 Optical Cable Adapter for PCIe3 Expansion Drawer is required to connect to the PCIe3 6-slot fan-out module in the I/O expansion drawer. The fan-out module has two CXP ports. The top CXP port of the fan-out module is cabled to the top CXP port of the PCIe3 Optical Cable Adapter. The bottom CXP port of the fan-out module is cabled to the bottom CXP port of the same PCIe3 Optical Cable Adapter.

To set up the cabling correctly, follow these steps:

1. Connect an optical cable or copper CXP cable to connector T1 on the PCIe3 optical cable adapter in your server.
2. Connect the other end of the optical cable or copper CXP cable to connector T1 on one of the PCIe3 6-slot fan-out modules in your expansion drawer.
3. Connect another cable to connector T2 on the PCIe3 optical cable adapter in your server.
4. Connect the other end of the cable to connector T2 on the PCIe3 6-slot fan-out module in your expansion drawer.
5. Repeat steps 1 - 4 for the other PCIe3 6-slot fan-out module in the expansion drawer, if required.

Drawer connections: Each fan-out module in a PCIe3 Expansion Drawer can be connected only to a single PCIe3 Optical Cable Adapter for PCIe3 Expansion Drawer.

Figure 2-25 shows the connector locations for the PCIe Gen3 I/O expansion drawer.




Figure 2-25 Connector locations for the PCIe Gen3 I/O expansion drawer

Figure 2-26 shows typical optical cable connections.

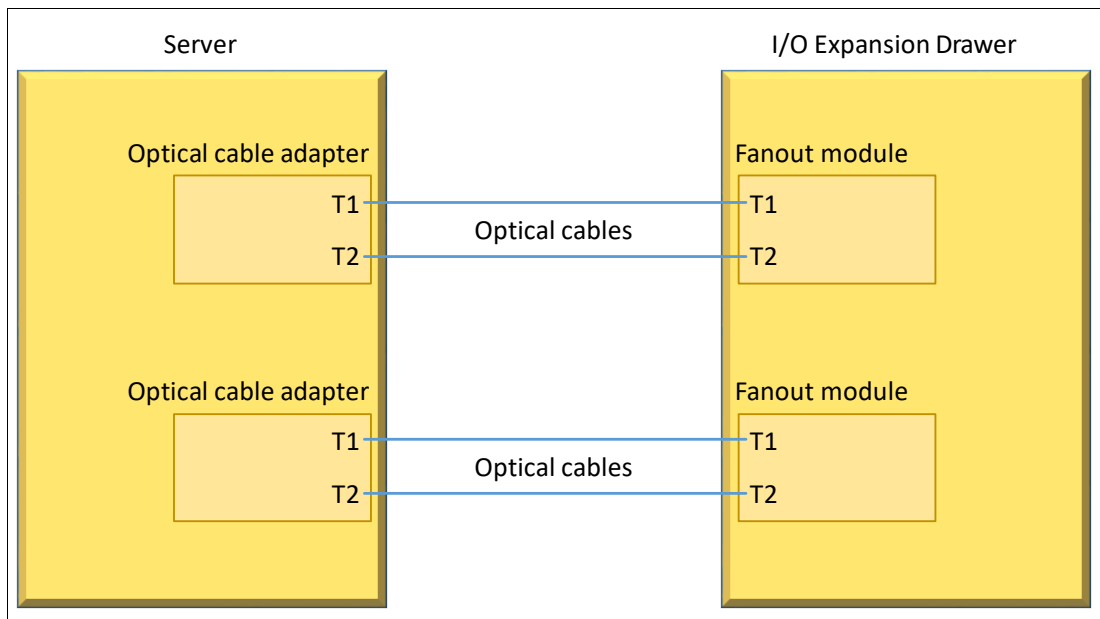


Figure 2-26 Typical optical cable connections

General rules for the PCI Gen3 I/O expansion drawer configuration

The PCIe3 Optical Cable Adapter for PCIe3 Expansion Drawer (#EJ05) is supported in slots P1-C4 and P1-C9 for the Power S922 system. This is a double-wide adapter that requires two adjacent slots. If #EJ05 is installed in this slot, the external SAS port is not allowed in the system.

The PCIe3 cable adapter for the PCIe3 EMX0 expansion drawer (#EJ08) is supported in P1-C9 for the Power S914 and Power S924 single processor systems. It is supported in P1-C9, P1-C3, and P1-C4 in the Power S924 double processor systems.

Table 2-49 shows PCIe adapter slot priorities and maximum adapters that are supported in the Power S922, Power S914, and Power S924 systems.

Table 2-49 PCIe adapter slot priorities and maximum adapters that are supported

System	Feature code	Slot priorities	Maximum number of adapters supported
Power S922 (1 processor)	#EJ05	9	1
Power S922 (2 processor)	#EJ05	9/10, 3/4	2
Power S914	#EJ08	9	1
Power S924 (1 processor)	#EJ08	9	1
Power S924 (2 processor)	#EJ08	9, 3, 4	3

2.7.3 PCIe Gen3 I/O expansion drawer system power control network cabling

There is no system power control network (SPCN) that is used to control and monitor the status of power and cooling within the I/O drawer. SPCN capabilities are integrated into the optical cables.

2.8 External disk subsystems

This section describes the following external disk subsystems that can be attached to the Power S922, Power S914, and Power S924 servers:

- ▶ EXP24SX SAS Storage Enclosure (#ESLS) and EXP12SX SAS Storage Enclosure (#ESLL)
- ▶ IBM Storage

2.8.1 EXP24SX SAS Storage Enclosure and EXP12SX SAS Storage Enclosure

The EXP24SX is a storage expansion enclosure with 24 2.5-inch SFF SAS bays. It supports up to 24 hot-swap HDDs or SSDs in only 2 EIA of space in a 19-inch rack. The EXP24SX SFF bays use SFF gen2 (SFF-2) carriers/ trays that are identical to the carrier/trays in the previous EXP24S drawer. With AIX and Linux, or VIOS, the EXP24SX can be ordered with four sets of six bays (mode 4), two sets of 12 bays (mode 2), or one set of 24 bays (mode 1). With IBM i, one set of 24 bays (mode 1) is supported.

There can be no mixing of HDDs and SSDs in the same mode 1 drawer. HDDs and SSDs can be mixed in a mode 2 or mode 4 drawer, but they cannot be mixed within a logical split of the drawer. For example, in a mode 2 drawer with two sets of 12 bays, one set can hold SSDs and one set can hold HDDs, but you cannot mix SSDs and HDDs in the same set of 12 bays.

The EXP12SX is a storage expansion enclosure with twelve 3.5-inch large form factor (LFF) SAS bays. It supports up to 12 hot-swap HDDs in only 2 EIA of space in a 19-inch rack. The EXP12SX SFF bays use LFF gen1 (LFF-1) carriers/trays. 4-KB sector drives (4096 or 4224) are supported. With AIX and Linux, and VIOS, the EXP12SX can be ordered with four sets of three bays (mode 4), two sets of six bays (mode 2), or one set of 12 bays (mode 1). Only 4-KB sector drives are supported in the EXP12SX drawer.

Four mini-SAS HD ports on the EXP24SX or EXP12SX are attached to PCIe Gen3 SAS adapters or attached to an integrated SAS controller in the Power S922, S914, or S924 systems. The following PCIe3 SAS adapters support the EXP24SX and EXP 12SX:

- ▶ PCIe3 RAID SAS Adapter Quad-port 6 Gb x8 (#EJ0J, #EJ0M, #EL3B, or #EL59)
- ▶ PCIe3 12 GB Cache RAID Plus SAS Adapter Quad-port 6 Gb x8 (#EJ14)

Earlier generation PCIe2 or PCIe1 SAS adapters are not supported by the EXP24SX.

The attachment between the EXP24SX or EXP12SX and the PCIe3 SAS adapters or integrated SAS controllers is through SAS YO12 or X12 cables. All ends of the YO12 and X12 cables have mini-SAS HD narrow connectors. The cable options are:

- ▶ X12 cable: 3-meter copper (#ECDJ)
- ▶ YO12 cables: 1.5-meter copper (#ECDT) or 3-meter copper (#ECDU)
- ▶ 3M 100 GbE Optical Cable QSFP28 (AOC) (#EB5R)
- ▶ 5M 100 GbE Optical Cable QSFP28 (AOC) (#EB5S)
- ▶ 10M 100 GbE Optical Cable QSFP28 (AOC) (#EB5T)
- ▶ 15M 100 GbE Optical Cable QSFP28 (AOC) (#EB5U)
- ▶ 20M 100 GbE Optical Cable QSFP28 (AOC) (#EB5V)
- ▶ 30M 100 GbE Optical Cable QSFP28 (AOC) (#EB5W)
- ▶ 50M 100 GbE Optical Cable QSFP28 (AOC) (#EB5X)
- ▶ 100M 100 GbE Optical Cable QSFP28 (AOC) (#EB5Y)

There are six SAS connectors at the rear of the EXP24SX and EXP12SX to which SAS adapters or controllers are attached. They are labeled T1, T2, and T3; there are two T1, two T2, and two T3 connectors.

- ▶ In mode 1, two or four of the six ports are used. Two T2 ports are used for a single SAS adapter, and two T2 and two T3 ports are used with a paired set of two adapters or a dual adapters configuration.
- ▶ In mode 2 or mode 4, four ports are used, two T2s and two T3s, to access all SAS bays.

Figure 2-27 shows connector locations for the EXP24SX and EXP12SX storage enclosures.




Figure 2-27 Connector locations for the EXP24SX and EXP12SX storage enclosures

For more information about SAS cabling, see the “Connecting an ESSL or ESLS storage enclosure to your system” topic in [IBM Knowledge Center](#).

The EXP24SX and EXP12SX drawers have many high-reliability design points:

- ▶ SAS bays that support hot-swap.
- ▶ Redundant and hot-plug power and fan assemblies.
- ▶ Dual power cords.
- ▶ Redundant and hot-plug ESMs.
- ▶ Redundant data paths to all drives.
- ▶ LED indicators on drives, bays, ESMs, and power supplies that support problem identification.
- ▶ Through the SAS adapters/controllers, drives that can be protected with RAID and mirroring and hot-spare capability.

2.8.2 IBM Storage

The IBM Storage Systems products and offerings provide compelling storage solutions with superior value for all levels of business, from entry-level to high-end storage systems. For more information about the various offerings, see [Data Storage Solutions](#).

The following sections highlight a few of the offerings.

IBM Flash Storage

The next generation of IBM Flash Storage delivers the extreme performance and efficiency you need to succeed, with a new pay-as-you-go option to reduce your costs and scale-on-demand. For more information, see [Flash Storage and All Flash Arrays](#).

IBM DS8880 hybrid storage

IBM DS8880 Hybrid Storage is a family of storage systems that includes the IBM DS8886 storage system for high-performance functionality in a dense, expandable package, and the IBM DS8884 storage system to provide advanced functionality for consolidated systems or multiple platforms in a space-saving design. IBM DS8880 systems combine resiliency and intelligent flash performance to deliver microsecond application response times and more than six-nines availability. For more information, see [IBM DS8880 hybrid storage - Overview](#).

IBM XIV Storage System

IBM XIV® Gen3 is a high-end, grid-scale storage system that excels in tuning-free consistent performance, extreme ease of use, and exceptional data economics, including inline, field-proven IBM Real-time Compression. IBM XIV is ideal for hybrid cloud, offering predictable service levels for dynamic workloads, simplified scale management, including in multi-tenant environments, flexible consumption models, and robust cloud automation and orchestration through OpenStack, the RESTful API, and VMware. It offers security and data protection through hot encryption, advanced mirroring and self-healing, and investment protection with perpetual licensing. For more information, see [IBM XIV Storage System - Overview](#).

IBM Storwize V7000

IBM Storwize® V7000 is an enterprise-class storage solution that offers the advantages of IBM Spectrum® Virtualize software. It can help you lower capital and operational storage costs with heterogeneous data services while optimizing performance with flash storage. IBM Storwize V7000 enables you to take advantage of hybrid cloud technology without replacing your current storage. For more information, see [IBM Storwize V7000 - Overview](#).

IBM Storwize V5000

IBM Storwize V5000 is a flexible storage solution that offers extraordinary scalability from the smallest to the largest system without disruption. Built with IBM Spectrum Virtualize software, it can help you lower capital and operational storage costs with heterogeneous data services. IBM Storwize V5000 is an easily customizable and upgradeable solution for better investment protection, improved performance, and enhanced efficiency. For more information, see [IBM Storwize V5000 - Overview](#).

2.9 Operating system support

The Power S922, Power S914, and Power S924 servers support the following operating systems:

- ▶ AIX
- ▶ IBM i (by using the VIOS)
- ▶ Linux

In addition, the VIOS can be installed in special partitions that provide support to other partitions running AIX, IBM i, or Linux operating systems for using features such as virtualized I/O devices, PowerVM Live Partition Mobility (LPM), or PowerVM Active Memory Sharing.

For more information about the software that is available on IBM Power Systems, see [IBM Power Systems Software](#).

2.9.1 AIX operating system

The following sections describe the various levels of AIX operating system support.

IBM periodically releases maintenance packages (service packs or technology levels) for the AIX operating system. For more information about these packages, downloading, and obtaining the CD-ROM, see [Fix Central](#).

The Fix Central website also provides information about how to obtain the fixes that are included on CD-ROM.

The Service Update Management Assistant (SUMA), which can help you automate the task of checking and downloading operating system downloads, is part of the base operating system. For more information about the `suma` command, see [IBM Knowledge Center](#).

The following minimum levels of AIX support the Power S922, Power S914, and Power S924 servers:

- ▶ If you are installing the AIX operating system LPAR with any I/O configuration:
 - AIX Version 7.2 with the 7200-02 Technology Level and Service Pack 7200-02-02-1810, or later
 - AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-02-1810, or later
 - AIX Version 6.1 with the 6100-09 Technology Level and Service Pack 6100-09-11-1810, or later
 - AIX Version 7.2 with the 7200-01 Technology Level and Service Pack 7200-01-04-1806, or later
 - AIX Version 7.2 with the 7200-00 Technology Level and Service Pack 7200-00-06-1806, or later
 - AIX Version 7.1 with the 7100-04 Technology Level and Service pack 7100-04-06-1806, or later
- ▶ If you are installing the AIX operating system Virtual I/O only LPAR:
 - AIX Version 7.2 with the 7200-02 Technology Level and Service Pack 7200-02-01-1732, or later
 - AIX Version 7.2 with the 7200-01 Technology Level and Service Pack 7200-01-01-1642, or later
 - AIX Version 7.2 with the 7200-00 Technology Level and Service Pack 7200-00-01-1543, or later
 - AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-01-1731, or later
 - AIX Version 7.1 with the 7100-04 Technology Level and Service Pack 7100-04-01-1543, or later
 - AIX Version 6.1 with the 6100-09 Technology Level and Service Pack 6100-09-06-1543, or later (AIX 6.1 service extension required.)

2.9.2 IBM i

IBM i is supported on the Power S922, Power S914, and Power S924 servers with the following minimum required levels:

- ▶ IBM i 7.3 TR4
- ▶ IBM i 7.2 TR8

IBM periodically releases maintenance packages (service packs or technology levels) for the IBM i operating system. For more information about these packages, downloading, and obtaining the CD-ROM, see [IBM Fix Central](#).

For compatibility information for hardware features and the corresponding AIX and IBM i Technology Levels, see [IBM Prerequisites](#).

2.9.3 Linux operating system

Linux is an open source, cross-platform operating system that runs on numerous platforms from embedded systems to mainframe computers. It provides an UNIX -like implementation across many computer architectures.

The supported versions of Linux on the Power S922, Power S914, and Power S924 servers are as follows:

- ▶ If you are installing the Linux operating system LPAR:
 - Red Hat Enterprise Linux 7 for Power LE Version 7.4, or later (POWER8 mode).
 - SUSE Linux Enterprise Server 12 Service Pack 3, or later.
 - Ubuntu Server 16.04.4, or later (POWER8 mode).
- ▶ If you are installing the Linux operating systems LPAR in non-production SAP implementations:
 - SUSE Linux Enterprise Server 12 Service Pack 3, or later.
 - SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 12 Service Pack 3, or later.
 - Red Hat Enterprise Linux 7 for Power LE, version 7.4, or later (POWER8 mode).
 - Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 7 for Power LE version 7.4, or later (POWER8 mode). Linux supports almost all of the Power Systems I/O, and the configurator verifies support on order.

Service and productivity tools

Service and productivity tools are available in a YUM repository that you can use to download, and then install, all the recommended packages for your Red Hat, SUSE Linux, or Fedora distribution. The packages are available from [Service and productivity tools for Linux on Power servers](#).

To learn about developing on the IBM Power Architecture®, find packages, get access to cloud resources, and discover tools and technologies, see [Linux on IBM Power Systems Developer Portal](#).

The IBM Advance Toolchain for Linux on Power is a set of open source compilers, runtime libraries, and development tools that users can use to take leading-edge advantage of IBM POWER hardware features on Linux. For more information, see [Advanced Toolchain for Linux on Power](#).

For more information about SUSE Linux Enterprise Server, see [SUSE Linux Enterprise Server](#).

For more information about Red Hat Enterprise Linux, see [Red Hat Enterprise Linux](#).

2.9.4 Virtual I/O Server

The minimum required level of VIOS for the Power S922, Power S914, and Power S924 servers is VIOS 2.2.6.21.

IBM regularly updates the VIOS code. To find information about the latest updates, see [Fix Central](#).

2.10 POWER9 reliability, availability, and serviceability capabilities

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

The elements of RAS can be described as follows:

- 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 nondisruptively the faults are repaired.

Table 2-50 provides a list of the Power Systems RAS capabilities by operating system. The HMC is an optional feature on scale-out Power Systems servers.

Table 2-50 Selected reliability, availability, and serviceability features by operating system

RAS feature	AIX	IBM i	Linux
Processor			
First failure data capture (FFDC) for fault detection/error isolation	X	X	X
Dynamic Processor Deallocation	X	X	X
I/O subsystem			
PCI Express bus enhanced error detection	X	X	X
PCI Express bus enhanced error recovery	X	X	X
PCI Express card hot-swap	X	X	X
Memory availability			
Memory Page Deallocation	X	X	X
Special Uncorrectable Error Handling	X	X	X
Fault detection and isolation			

RAS feature	AIX	IBM i	Linux
Storage Protection Keys	X	Not used by OS	Not used by OS
Error log analysis	X	X	X
Serviceability			
Boot-time progress indicators	X	X	X
Firmware error codes	X	X	X
Operating system error codes	X	X	X
Inventory collection	X	X	X
Environmental and power warnings	X	X	X
Hot-swap DASD / media	X	X	X
Dual disk controllers / Split backplane	X	X	X
EED collection	X	X	X
SP/OS "Call Home" on non-HMC configurations	X	X	X
IO adapter/device stand-alone diagnostic tests with PowerVM	X	X	X
SP mutual surveillance with IBM POWER Hypervisor	X	X	X
Dynamic firmware update with HMC	X	X	X
Service Agent Call Home Application	X	X	X
Service Indicator LED support	X	X	X
System dump for memory, POWER Hypervisor, and SP	X	X	X
IBM Knowledge Center / IBM Systems Support Site service publications	X	X	X
System Service/Support Education	X	X	X
Operating system error reporting to HMC Service Focal Point (SFP) application	X	X	X
RMC secure error transmission subsystem	X	X	X
Healthcheck scheduled operations with HMC	X	X	X
Operator panel (real or virtual [HMC])	X	X	X
Concurrent Op Panel Display Maintenance	X	X	X
Redundant HMCs	X	X	X
High availability clustering support	X	X	X
Repair and Verify Guided Maintenance with HMC	X	X	X
PowerVM Live Partition / Live Application Mobility With PowerVM Enterprise Edition	X	X	X
EPOW			
EPOW errors handling	X	X	X

2.11 Manageability

Several functions and tools help with manageability so that you can efficiently and effectively manage your system.

2.11.1 Service user interfaces

The service user interface enables support personnel or the client to communicate with the service support applications in a server by using a console, interface, or terminal. Delivering a clear, concise view of available service applications, the service interface enables the support team to manage system resources and service information in an efficient and effective way. Applications that are available through the service interface are carefully configured and placed to give service providers access to important service functions.

Various service interfaces are used, depending on the state of the system and its operating environment. Here are the primary service interfaces:

- ▶ Light Path, which provides indicator lights to help a service technical find a component in need of service.
- ▶ Service processor.
- ▶ ASMI.
- ▶ Operator panel.
- ▶ An operating system service menu, which obtains error codes directly from the hardware.
- ▶ Service Focal Point (SFP) on the HMC.

Service processor

The service processor is a controller that is running its own operating system. It is a component of the service interface card.

The service processor operating system has specific programs and device drivers for the service processor hardware. The host interface is a processor support interface that is connected to the POWER processor. The service processor is always working, regardless of the main system unit's state. The system unit can be in the following states:

- ▶ Standby (power off)
- ▶ Operating, ready to start partitions
- ▶ Operating with running logical partitions (LPARs)

The service processor is used to monitor and manage the system hardware resources and devices. The service processor checks the system for errors, ensuring that the connection to the management console for manageability purposes is functioning, and accepting ASMI Secure Sockets Layer (SSL) network connections. The service processor can view and manage the machine-wide settings by using the ASMI, which enables complete system and partition management from the HMC.

Analyzing a system that does not start: The flexible service processor (FSP) can analyze a system that does not start. Reference codes and detailed data are available in the ASMI and are transferred to the HMC.

The service processor uses two Ethernet ports that run at 1-Gbps speed. Consider the following information:

- ▶ Both Ethernet ports are visible only to the service processor and can be used to attach the server to an HMC or to access the ASMI. The ASMI options can be accessed through an HTTP server that is integrated into the service processor operating environment.
- ▶ Both Ethernet ports support only auto-negotiation. Customer-selectable media speed and duplex settings are not available.
- ▶ Both Ethernet ports have a default IP address, as follows:
 - Service processor eth0 (HMC1 port) is configured as 169.254.2.147.
 - Service processor eth1 (HMC2 port) is configured as 169.254.3.147.
 - DHCP using the HMC for the HMC management networks is also possible.

The following functions are available through the service processor:

- ▶ Call Home
- ▶ ASMI
- ▶ Error information (error code, part number, and location codes) menu
- ▶ View of guarded components
- ▶ Limited repair procedures
- ▶ Generate dump
- ▶ LED Management menu
- ▶ Remote view of ASMI menus
- ▶ Firmware update through a USB key

Advanced System Management Interface

ASMI is the interface to the service processor that enables you to manage the operation of the server, such as auto-power restart, and to view information about the server, such as the error log and Vital Product Data (VPD). Various repair procedures require connection to the ASMI.

The ASMI is accessible through the management console. It is also accessible by using a web browser on a system that is connected directly to the service processor (in this case, either a standard Ethernet cable or a crossed cable) or through an Ethernet network. ASMI can also be accessed from an ASCII terminal, but this is available only while the system is in the platform powered-off mode.

Figure 2-28 shows a method of opening ASMI on a particular server by using the HMC GUI.

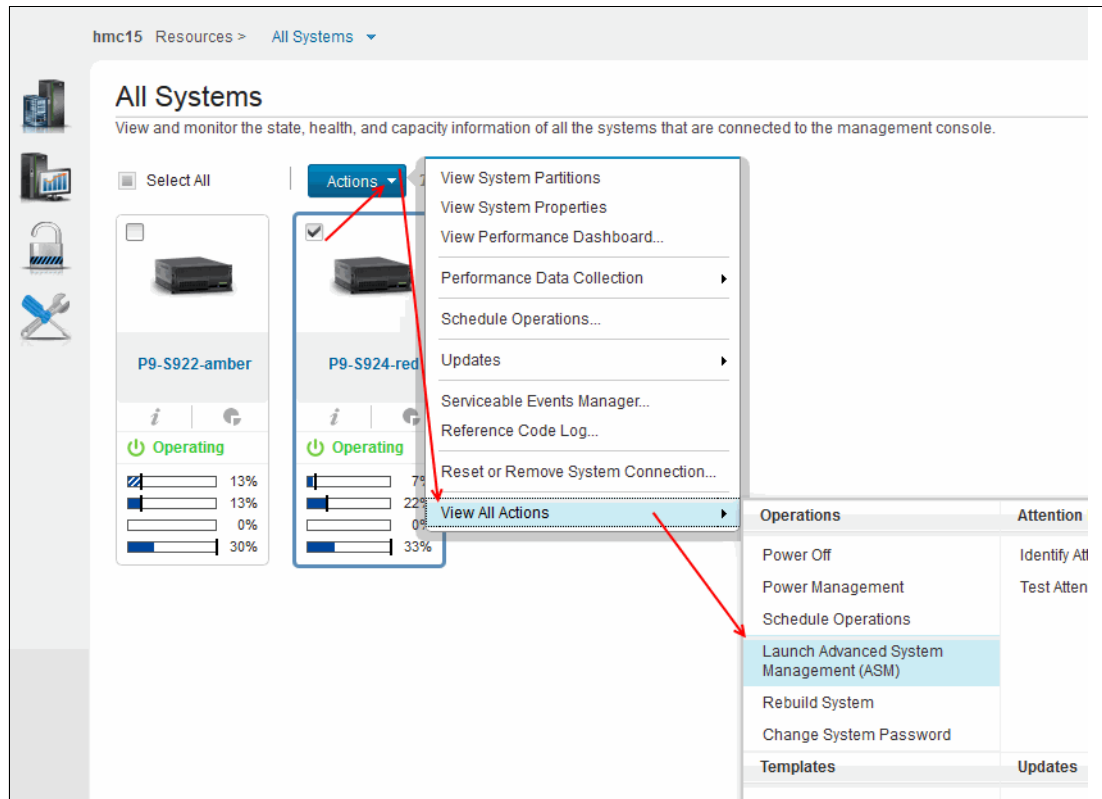


Figure 2-28 Starting the Advanced System Management Interface through the HMC GUI

You are prompted for confirmation about which FSP to use, and then a login window opens, as shown in Figure 2-29.

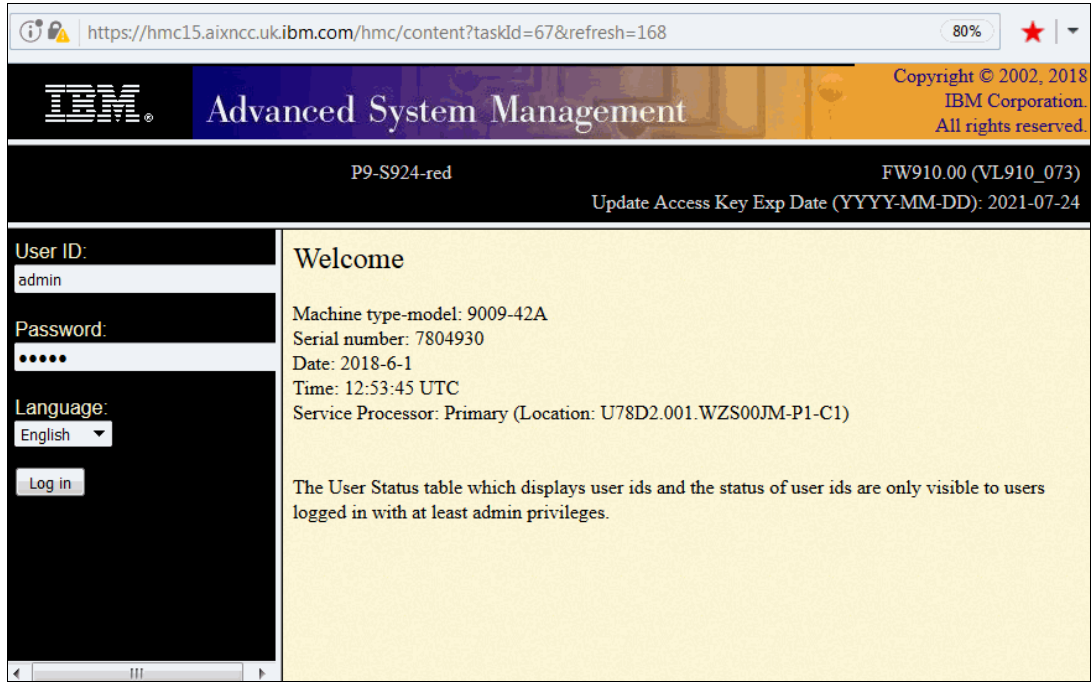


Figure 2-29 The ASMI login window

After you are logged in (the default credentials are admin/admin), you see the menu that is shown in Figure 2-30.




Figure 2-30 An ASMI window

Tip: If you click **Expand all menus**, as shown in the red ring, you can then use the search function (Ctrl+f) in your browser to find quickly menu items.

Use the ASMI to change the service processor IP addresses or to apply certain security policies and prevent access from unwanted IP addresses or ranges.

You might be able to use the service processor's default settings. In that case, accessing the ASMI is not necessary. To access ASMI, use one of the following methods:

- ▶ Use a management console.
 - If configured to do so, the management console connects directly to the ASMI for a selected system from this task.
 - To connect to the ASMI from a management console, complete the following steps:
 - a. Open **Systems Management** from the navigation pane.
 - b. From the work window, select one of the managed systems.
 - c. From the System Management tasks list, click **Operations** → **Launch Advanced System Management (ASMI)**.

- ▶ Use a web browser.

At the time of writing, the supported web browsers are Microsoft Internet Explorer (Version 10.0.9200.16439), Mozilla Firefox ESR (Version 24), and Chrome (Version 30). Later versions of these browsers might work, but are not officially supported. The JavaScript language and cookies must be enabled, and TLS 1.2 might need to be enabled.

The web interface is available during all phases of system operation, including the initial program load (IPL) and run time. However, several of the menu options in the web interface are unavailable during IPL or run time to prevent usage or ownership conflicts if the system resources are in use during that phase. The ASMI provides an SSL web connection to the service processor. To establish an SSL connection, open your browser go to the following address:

`https://<ip_address_of_service_processor>`

Note: To make the connection through Internet Explorer, click **Tools Internet Options**. Clear the **Use TLS 1.0** check box, and click **OK**.

- ▶ Use an ASCII terminal.

The ASMI on an ASCII terminal supports a subset of the functions that are provided by the web interface and is available only when the system is in the platform powered-off mode. The ASMI on an ASCII console is not available during several phases of system operation, such as the IPL and run time.

- ▶ Command-line start of the ASMI

Either on the HMC itself or when properly configured on a remote system, it is possible to start ASMI web interface from the HMC command line. Open a terminal window on the HMC or access the HMC with a terminal emulation and run the following command:

```
asmmenu --ip <ip address>
```

On the HMC itself, a browser window opens automatically with the ASMI window and, when configured properly, a browser window opens on a remote system when issued from there.

The operator panel

The service processor provides an interface to the operator panel, which is used to display system status and diagnostic information.

The operator panel is formed of two parts: One is always installed, and the second might be optional.

The part that is always installed provides LEDs and sensors:

- ▶ Power LED:
 - Color: Green.
 - Off: Enclosure is off (AC cord is not connected).
 - On Solid: Enclosure is powered on.
 - On Blink: Enclosure is in the standby-power state.
- ▶ Enclosure Identify LED. Color: Blue.
- ▶ Enclosure Identify LED:
 - Color: Blue.
 - Off: Normal.
 - On Solid: Identify State.

- ▶ System Fault LED:
 - Color: Amber.
 - Off: Normal.
 - On Solid: Check Error Log.
- ▶ System Roll-up LED:
 - Color: Amber.
 - Off: Normal.
 - On Solid: Fault.
- ▶ Power Button.
- ▶ System Reset Switch.
- ▶ Two Thermal Sensors.
- ▶ One Pressure/Altitude Sensor.

The LCD operator panel is optional in some systems, but there must be at least one in a server in a rack containing any IBM Power S914, Power S922, or Power S924 server. The panel can be accessed by using the switches on the front panel.

Here are several of the operator panel features:

- ▶ A 2 x 16 character LCD display
- ▶ Increment, decrement, and “enter” buttons

The following functions are available through the operator panel:

- ▶ Error information
- ▶ Generate dump
- ▶ View machine type, model, and serial number
- ▶ Limited set of repair functions

The System Management Services (SMS) error log is accessible through the SMS menus. This error log contains errors that are found by partition firmware when the system or partition is starting.

The service processor’s error log can be accessed on the ASMI menus.

You can also access the system diagnostics from a Network Installation Management (NIM) server.

IBM i and its associated machine code provide dedicated service tools (DSTs) as part of the IBM i licensed machine code (Licensed Internal Code) and System Service Tools (SSTs) as part of IBM i. DSTs can be run in dedicated mode (no operating system is loaded). DSTs and diagnostic tests are a superset of those available under SSTs.

The IBM i End Subsystem (**ENDSBS *ALL**) command can shut down all IBM and customer applications subsystems except for the controlling subsystem QTCL. The Power Down System (**PWRDWSYS**) command can be set to power down the IBM i partition and restart the partition in DST mode.

You can start SSTs during normal operations, which keep all applications running, by using the IBM i Start Service Tools (**STRSST**) command (when signed on to IBM i with the appropriately secured user ID).

With DSTs and SSTs, you can look at various logs, run various diagnostic tests, or take several kinds of system memory dumps or other options.

Depending on the operating system, the following service-level functions are what you typically see when you use the operating system service menus:

- ▶ Product activity log
- ▶ Trace Licensed Internal Code
- ▶ Work with communications trace
- ▶ Display/Alter/Dump
- ▶ Licensed Internal Code log
- ▶ Main storage memory dump manager
- ▶ Hardware service manager
- ▶ Call Home/Customer Notification
- ▶ Error information menu
- ▶ LED management menu
- ▶ Concurrent/Non-concurrent maintenance (within scope of the OS)
- ▶ Managing firmware levels:
 - Server
 - Adapter
- ▶ Remote support (access varies by OS)

Service Focal Point on the Hardware Management Console

Service strategies become more complicated in a partitioned environment. The Manage Serviceable Events task in the management console can help streamline this process.

Each LPAR reports errors that it detects and forwards the event to the SFP application that is running on the management console, without determining whether other LPARs also detect and report the errors. For example, if one LPAR reports an error for a shared resource, such as a managed system power supply, other active LPARs might report the same error.

By using the Manage Serviceable Events task in the management console, you can avoid long lists of repetitive Call Home information by recognizing that they are repeated errors and consolidating them into one error.

In addition, you can use the Manage Serviceable Events task to initiate service functions on systems and LPARs, including the exchanging of parts, configuring connectivity, and managing memory dumps.

2.11.2 IBM Power Systems Firmware maintenance

The IBM Power Systems Client-Managed Microcode is a methodology that enables you to manage and install microcode updates on Power Systems servers and their associated I/O adapters.

Firmware entitlement

With the new HMC Version V8R8.1.0.0 and Power Systems servers, the firmware installations are restricted to entitled servers. The customer must be registered with IBM and entitled by 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 some extra time. The Power Systems Firmware is relocated from the public repository to the access control repository. The I/O firmware remains on the public repository, but the server must be entitled for installation. When the `ls1ic` command is run to display the firmware levels, a new value, `update_access_key_exp_date`, is added. The HMC GUI and the ASMI menu show the Update access key expiration date.

When the system is no longer entitled, the firmware updates fail. Some new System Reference Code (SRC) packages are available:

- ▶ E302FA06: Acquisition entitlement check failed
- ▶ E302FA08: Installation entitlement check failed

Any firmware release that was made available during the entitled time frame can still be installed. For example, if the entitlement period ends on 31 December 2014, and a new firmware release is release before the end of that entitlement period, then it can still be installed. If that firmware is downloaded after 31 December 2014, 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.

Note: The update access key expiration date requires a valid entitlement of the system to perform firmware updates.

You can find an update access key at [IBM CoD Home](#).

For more information, go to [IBM Entitled Software Support](#).

Firmware updates

System firmware is delivered as a release level or a service pack. Release levels support the general availability (GA) of new functions or features, and new machine types or models. Upgrading to a higher release level is disruptive to customer operations. IBM intends to introduce no more than two new release levels per year. These release levels will be supported by service packs. Service packs are intended to contain only firmware fixes and not introduce new functions. A *service pack* is an update to an existing release level.

If the system is managed by a management console, you use the management console for firmware updates. By using the management console, you can take advantage of the CFM option when concurrent service packs are available. CFM is the IBM Power Systems Firmware updates that can be partially or wholly concurrent or nondisruptive. With the introduction of CFM, IBM is increasing its clients' opportunity to stay on a given release level for longer periods. Clients that want maximum stability can defer until there is a compelling reason to upgrade, such as the following reasons:

- ▶ A release level is approaching its end of service date (that is, it has been available for about a year, and soon service will not be supported).
- ▶ Move a system to a more standardized release level when there are multiple systems in an environment with similar hardware.
- ▶ A new release has a new function that is needed in the environment.
- ▶ A scheduled maintenance action causes a platform restart, which provides an opportunity to also upgrade to a new firmware release.

The updating and upgrading of system firmware depends on several factors, such as whether the system is stand-alone or managed by a management console, the firmware that is installed, and what operating systems are running on the system. These scenarios and the associated installation instructions are comprehensively outlined in the firmware section of [Fix Central](#).

You might also want to review the preferred practice white papers that are found at [Service and support best practices for Power Systems](#).

Firmware update steps

The system firmware consists of service processor microcode, Open Firmware microcode, and SPCN microcode.

The firmware and microcode can be downloaded and installed either from an HMC, from a running partition, or from USB port number 1 on the rear, if that system is not managed by an HMC.

Power Systems Firmware has a permanent firmware boot side (A side) and a temporary firmware boot side (B side). New levels of firmware must be installed first on the temporary side to test the update's compatibility with existing applications. When the new level of firmware is approved, it can be copied to the permanent side.

For access to the initial websites that address this capability, see [Support for IBM Systems](#). For Power Systems, select the **Power** link.

Although the content under the Popular links section can change, click the **Firmware and HMC updates** link to go to the resources for keeping your system's firmware current.

If there is an HMC to manage the server, the HMC interface can be used to view the levels of server firmware and power subsystem firmware that are installed and that are available to download and install.

Each IBM Power Systems server has the following levels of server firmware and power subsystem firmware:

- ▶ Installed level

This level of server firmware or power subsystem firmware is installed and will be installed into memory after the managed system is powered off and then powered on. It is installed on the temporary side of system firmware.

- ▶ Activated level

This level of server firmware or power subsystem firmware is active and running in memory.

- ▶ Accepted level

This level is the backup level of the server or power subsystem firmware. You can return to this level of server or power subsystem firmware if you decide to remove the installed level. It is installed on the permanent side of system firmware.

Figure 2-31 shows the different levels in the HMC.

EC Number	LIC Type	Machine Type/Model/Serial Number	Update Access Key Expiration	Installed Level	Activated Level	Accepted Level	Unactivated Deferred Level	Platform IPL Level	Update Control
01VL910	Managed System Primary	9009-42A*7804930	07/24/2021	FW910.00 (073)	FW910.00 (073)	FW910.00 (073)		FW910.00 (073)	Management Console

Figure 2-31 Firmware levels

IBM provides the CFM function on selected Power Systems servers. This function supports applying nondisruptive system firmware service packs to the system concurrently (without requiring a restart operation to activate changes). For systems that are not managed by an HMC, the installation of system firmware is always disruptive.

The concurrent levels of system firmware can, on occasion, contain fixes that are known as *deferred*. These deferred fixes can be installed concurrently but are not activated until the next IPL. Deferred fixes, if any, are identified in the Firmware Update Descriptions table of the firmware document. For deferred fixes within a service pack, only the fixes in the service pack that cannot be concurrently activated are deferred.

Table 2-51 shows the file-naming convention for system firmware.

Table 2-51 Firmware naming convention

PPNNSSS_ FFF_ DDD			
PP	Package identifier	01	-
NN	Platform and class	VL	Low end
SSS	Release indicator		
FFF	Current fix pack		
DDD	Last disruptive fix pack		

The following example uses the convention:

01VL910,73,73 = POWER9 Entry Systems Firmware for 9009-41A, 9009-22A, and 9009-42A

An installation is disruptive if the following statements are true:

- ▶ The release levels (SSS) of the currently installed and the new firmware differ.
- ▶ The service pack level (FFF) and the last disruptive service pack level (DDD) are equal in the new firmware.

Otherwise, an installation is concurrent if the service pack level (FFF) of the new firmware is higher than the service pack level that is installed on the system and the conditions for disruptive installation are not met.

2.11.3 Concurrent firmware maintenance improvements

Since POWER6, firmware service packs (updates) are concurrently applied and take effect immediately. Occasionally, a service pack is shipped where most of the features can be concurrently applied, but because changes to some server functions (for example, changing initialization values for chip controls) cannot occur during operation, a patch in this area required a system restart for activation.

With the Power-On Reset Engine (PORE), the firmware can now dynamically power off processor components, change the registers, and reinitialize while the system is running, without discernible impact to any applications running on a processor, which potentially allows concurrent firmware changes in POWER9, which in earlier designs required a restart to take effect.

Activating new firmware functions requires the installation of a firmware release level (upgrades). This process is disruptive to server operations, and requires a scheduled outage and full server restart.

2.11.4 Electronic Services and Electronic Service Agent

IBM transformed its delivery of hardware and software support services to help you achieve higher system availability. Electronic Services is a web-enabled solution that offers an exclusive, no additional charge enhancement to the service and support that is available for IBM servers. These services provide the opportunity for greater system availability with faster problem resolution and preemptive monitoring. The Electronic Services solution consists of two separate, but complementary, elements:

- ▶ Electronic Services news page
- ▶ Electronic Service Agent

Electronic Services news page

The Electronic Services news page is a single internet entry point that replaces the multiple entry points that traditionally are used to access IBM internet services and support. With the news page, you can gain easier access to IBM resources for assistance in resolving technical problems.

Electronic Service Agent

The Electronic Service Agent (ESA) is software that is on your server. It monitors events and transmits system inventory information to IBM on a periodic, client-defined timetable. The ESA automatically reports hardware problems to IBM.

Early knowledge about potential problems enables IBM to deliver proactive service that can result in higher system availability and performance. In addition, information that is collected through the ESA is made available to IBM Support Services Representatives (IBM SSRs) when they help answer your questions or diagnose problems. Installation and use of ESA for problem reporting enables IBM to provide better support and service for your IBM server.

To learn how Electronic Services can work for you, see [IBM Electronic Services](#) (an IBM ID is required).

Here are some of the benefits of Electronic Services:

- ▶ Increased uptime

The ESA tool enhances the warranty or maintenance agreement by providing faster hardware error reporting and uploading system information to IBM Support, which can translate to less time that is wasted monitoring the symptoms, diagnosing the error, and manually calling IBM Support to open a problem record.

Its 24x7 monitoring and reporting mean no more dependence on human intervention or off-hours customer personnel when errors are encountered in the middle of the night.

- ▶ Security

The ESA tool is designed to be secure in monitoring, reporting, and storing the data at IBM. The ESA tool securely transmits either through the internet (HTTPS or VPN) or modem, and can be configured to communicate securely through gateways to provide customers a single point of exit from their site.

Communication is one way. Activating ESA does not enable IBM to call into a customer's system. System inventory information is stored in a secure database, which is protected behind IBM firewalls. It is viewable only by the customer and IBM. The customer's business applications or business data is never transmitted to IBM.

- ▶ More accurate reporting

Because system information and error logs are automatically uploaded to the IBM Support center with the service request, customers are not required to find and send system information, decreasing the risk of misreported or misdiagnosed errors.

When inside IBM, problem error data is run through a data knowledge management system, and knowledge articles are appended to the problem record.

- ▶ Customized support

By using the IBM ID that you enter during activation, you can view system and support information by selecting **My Systems** at [Electronic Support](#).

My Systems provides valuable reports of installed hardware and software by using information that is collected from the systems by ESA. Reports are available for any system that is associated with the customers IBM ID. Premium Search combines the function of search and the value of ESA information, providing advanced search of the technical support knowledge base. Using Premium Search and the ESA information that was collected from your system, your clients can see search results that apply specifically to their systems.

For more information about how to use the power of IBM Electronic Services, contact your IBM SSR, or see [Electronic Support](#).

Service Event Manager

The Service Event Manager (SEM) enables the user to decide which of the Serviceable Events are called home by the ESA. It is possible to lock certain events. Some customers might not allow data to be transferred outside their company. After the SEM is enabled, the analysis of the possible problems might take longer.

- ▶ The SEM can be enabled by running the following command:

```
chhmc -c sem -s enable
```

- ▶ You can disable SEM mode and specify what state in which to leave the Call Home feature by running the following commands:

```
chhmc -c sem -s disable --callhome disable
```

```
chhmc -c sem -s disable --callhome enable
```

The easiest way to set up the ESA is by using the wizard in the HMC GUI, which is shown in Figure 2-32.




Figure 2-32 Accessing the Electronic Service Agent setup wizard

The wizard guides the user through the necessary steps, including entering details about the location of the system, the contact details, and other details. The user can select which HMC or HMCs should be used, as shown in Figure 2-33.

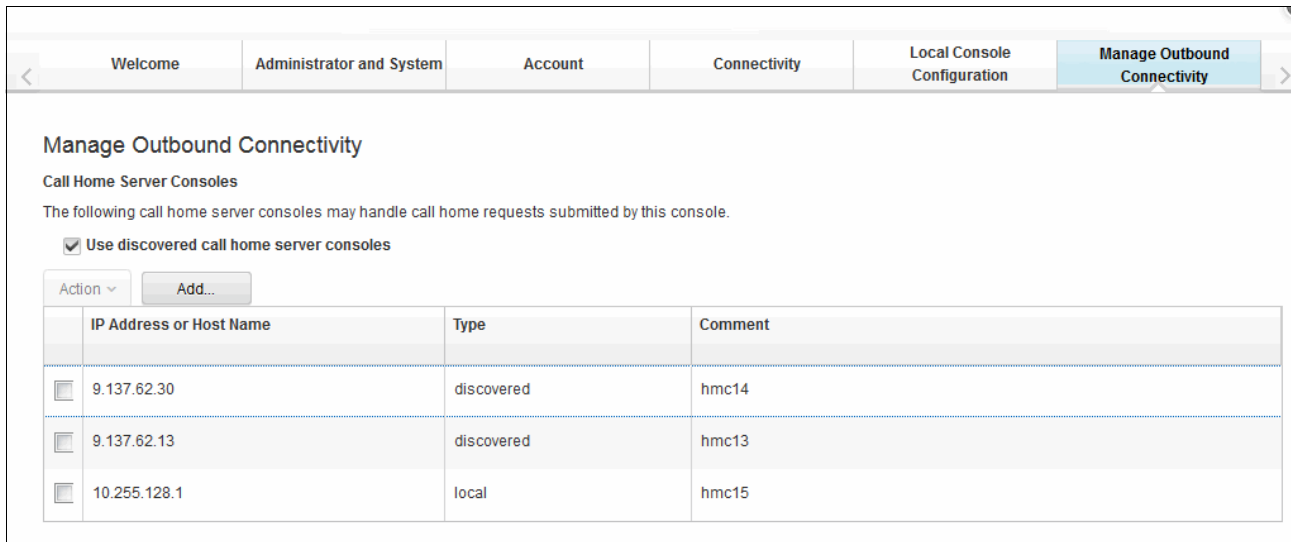


Figure 2-33 Managing which HMCs are used for Call Home

If you select a server, you can easily open the SEM, as shown in Figure 2-34.

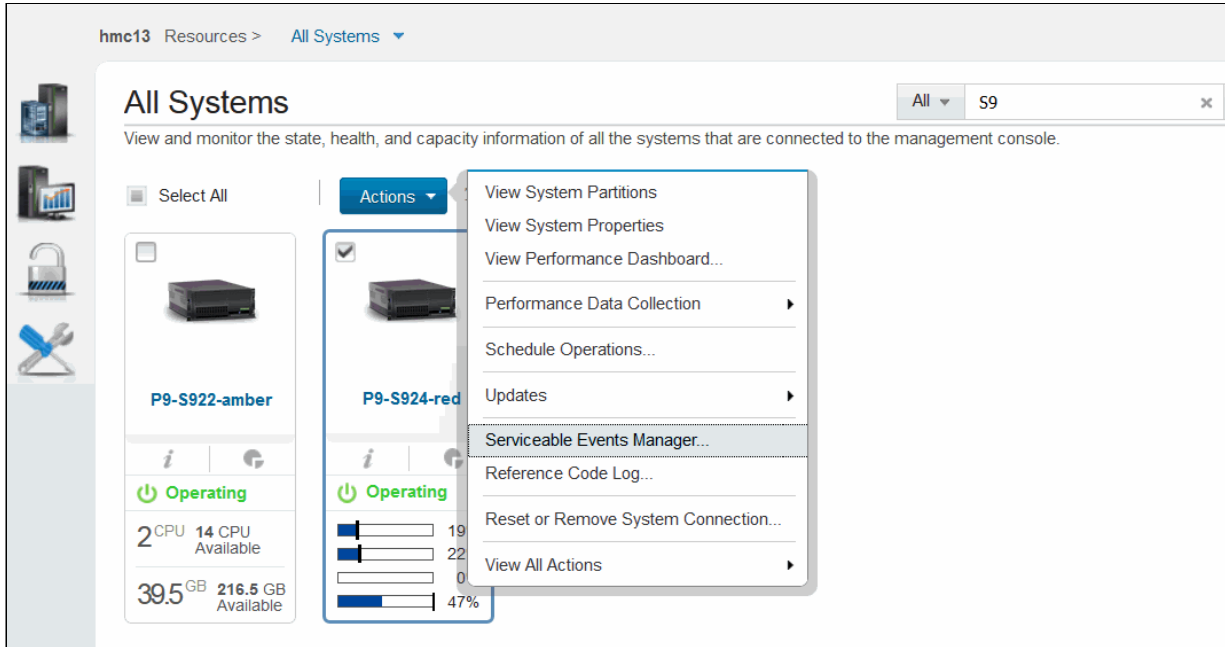


Figure 2-34 Opening the Serviceable Events Manager for a server

The user can then navigate the SEM menus to see the events for this server, as shown in Figure 2-35.




Figure 2-35 Serviceable events for this server



Virtualization

Virtualization is a key factor for productive and efficient use of IBM Power Systems servers. In this chapter, you find a brief description of virtualization technologies that are available for POWER9. The following IBM Redbooks publications provide more information about the virtualization features:

- ▶ *IBM PowerVM Best Practices*, SG24-8062
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *IBM PowerVM Virtualization Active Memory Sharing*, REDP-4470
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- ▶ *IBM Power Systems SR-IOV: Technical Overview and Introduction*, REDP-5065

3.1 POWER9 Hypervisor

Power Systems servers that are combined with PowerVM technology offer key capabilities that can help you consolidate and simplify your IT environment:

- ▶ Improve server usage and share I/O resources to reduce the total cost of ownership (TCO) and better use IT assets.
- ▶ Improve business responsiveness and operational speed by dynamically reallocating resources to applications as needed to better match changing business needs or handle unexpected changes in demand.
- ▶ Simplify IT infrastructure management by making workloads independent of hardware resources so that you can make business-driven policies to deliver resources that are based on time, cost, and service-level requirements.

Combined with features in the POWER9 processors, the IBM POWER Hypervisor delivers functions that enable other system technologies, including logical partitioning (LPAR) technology, virtualized processors, IEEE VLAN-compatible virtual switch, virtual SCSI adapters, virtual Fibre Channel adapters, and virtual consoles. The POWER Hypervisor is a basic component of the system's firmware and offers the following functions:

- ▶ Provides an abstraction between the physical hardware resources and the LPARs that use them.
- ▶ Enforces partition integrity by providing a security layer between LPARs.
- ▶ Controls the dispatch of virtual processors to physical processors.
- ▶ Saves and restores all processor state information during a logical processor context switch.
- ▶ Controls hardware I/O interrupt management facilities for LPARs.
- ▶ Provides virtual LAN channels between LPARs that help reduce the need for physical Ethernet adapters for inter-partition communication.
- ▶ Monitors the service processor and performs a reset or reload if it detects the loss of the service processor, notifying the operating system if the problem is not corrected.

The POWER Hypervisor is always active, regardless of the system configuration and also when not connected to the managed console. It requires memory to support the resource assignment to the LPARs on the server. The amount of memory that is required by the POWER Hypervisor firmware varies according to several factors:

- ▶ Memory that is required for hardware page tables (HPT).
- ▶ Memory that is required to support I/O devices.
- ▶ Memory that is required for virtualization.

Memory that is required for hardware page tables

Each partition on the system has its own HPT that contributes to hypervisor memory usage. The HPT is used by the operating system to translate from effective addresses to physical real addresses in the hardware. This translation from effective to real addresses allows multiple operating systems to run simultaneously in their own logical address space. Whenever a virtual processor for a partition is dispatched on a physical processor, the hypervisor indicates to the hardware the location of the partition HPT that should be used when translating addresses.

The amount of memory for the HPT is based on the maximum memory size of the partition and the HPT ratio. The default HPT ratio is either 1/64th of the maximum (for IBM i partitions) or 1/128th (for AIX, Virtual I/O Server (VIOS), and Linux partitions) of the maximum memory size of the partition. AIX, VIOS, and Linux use larger page sizes (16 and 64 KB) instead of using 4 KB pages. Using larger page sizes reduces the overall number of pages that must be tracked, so the overall size of the HPT can be reduced. As an example, for an AIX partition with a maximum memory size of 256 GB, the HPT would be 2 GB.

When defining a partition, the maximum memory size that is specified should be based on the amount of memory that can be dynamically added to the dynamic partition (DLPAR) without having to change the configuration and restart the partition.

In addition to setting the maximum memory size, the HPT ratio can also be configured. The `hpt_ratio` parameter for the `chsyscfg` Hardware Management Console (HMC) command can be issued to define the HPT ratio that is used for a partition profile. The valid values are 1:32, 1:64, 1:128, 1:256, or 1:512. Specifying a smaller absolute ratio (1/512 is the smallest value) decreases the overall memory that is assigned to the HPT. Testing is required when changing the HPT ratio because a smaller HPT might incur more CPU consumption because the operating system might need to reload the entries in the HPT more frequently. Most customers have chosen to use the IBM provided default values for the HPT ratios.

Memory usage for I/O devices

In support of I/O operations, the hypervisor maintains structures that are called the Translation Control Entries (TCEs), which provides an information path between I/O devices and partitions. The TCEs provide the address of the I/O buffer, indication of read versus write requests, and other I/O-related attributes. There are many TCEs in use per I/O device, so multiple requests can be active simultaneous to the same physical device. To provide better affinity, the TCEs are spread across multiple processor chips or drawers to improve performance while accessing the TCEs. For physical I/O devices, the base amount of space for the TCEs is defined by the hypervisor based on the number of I/O devices that are supported. A system that supports high-speed adapters can also be configured to allocate more memory to improve I/O performance. Linux is the only operating system that uses these additional TCEs so that the memory can be freed for use by partitions if the system is using only AIX or IBM i.

Memory usage for virtualization features

Virtualization requires more memory to be allocated by the POWER hypervisor for hardware statesave areas and various virtualization technologies. For example, on POWER9 systems, each processor core supports up to eight simultaneous multithreading (SMT) threads of execution, and each thread contains over 80 different registers. The POWER hypervisor must set aside save areas for the register contents for the maximum number of virtual processors that is configured. The greater the number of physical hardware devices, the greater the number of virtual devices, the greater the amount of virtualization, and the more hypervisor memory is required. For efficient memory consumption, wanted and maximum values for various attributes (processors, memory, and virtual adapters) should be based on business needs and not set to values that are significantly higher than actual requirements.

Predicting memory that is used by the POWER hypervisor

The IBM System Planning Tool (SPT) is a resource that can be used to estimate the amount of hypervisor memory that is required for a specific server configuration. After the SPT executable file is downloaded and installed, you can define a configuration by selecting the correct hardware platform, selecting the installed processors and memory, and defining partitions and partition attributes. SPT can estimate the amount of memory that will be assigned to the hypervisor, which assists you when you change an existing configuration or deploy new servers.

The POWER Hypervisor provides the following types of virtual I/O adapters:

- ▶ Virtual SCSI
- ▶ Virtual Ethernet
- ▶ Virtual Fibre Channel
- ▶ Virtual (TTY) console

3.1.1 Virtual SCSI

The POWER Hypervisor provides a virtual SCSI mechanism for the virtualization of storage devices. The storage virtualization is accomplished by using two paired adapters: a virtual SCSI server adapter and a virtual SCSI client adapter.

3.1.2 Virtual Ethernet

The POWER Hypervisor provides a virtual Ethernet switch function that allows partitions either fast and secure communication on the same server without any need for physical interconnection or connectivity outside of the server if a Layer 2 bridge to a physical Ethernet adapter is set in one VIOS partition, also known as Shared Ethernet Adapter (SEA).

3.1.3 Virtual Fibre Channel

A virtual Fibre Channel adapter is a virtual adapter that provides client LPARs with a Fibre Channel connection to a storage area network through the VIOS partition. The VIOS partition provides the connection between the virtual Fibre Channel adapters on the VIOS partition and the physical Fibre Channel adapters on the managed system.

3.1.4 Virtual (TTY) console

Each partition must have access to a system console. Tasks such as operating system installation, network setup, and various problem analysis activities require a dedicated system console. The POWER Hypervisor provides the virtual console by using a virtual TTY or serial adapter and a set of Hypervisor calls to operate on them. Virtual TTY does not require the purchase of any additional features or software, such as the PowerVM Edition features.

3.2 POWER processor modes

Although they are not virtualization features, the POWER processor modes are described here because they affect various virtualization features.

On Power System servers, partitions can be configured to run in several modes, including the following modes:

- ▶ POWER7 compatibility mode

This is the mode for POWER7+ and POWER7 processors, implementing Version 2.06 of the IBM Power Instruction Set Architecture (ISA). For more information, see [IBM Knowledge Center](#).

- ▶ POWER8 compatibility mode

This is the native mode for POWER8 processors implementing Version 2.07 of the IBM Power ISA. For more information, see [IBM Knowledge Center](#).

- ▶ POWER9 compatibility mode

This is the native mode for POWER9 processors implementing Version 3.0 of the IBM Power ISA. For more information, see [IBM Knowledge Center](#).

Figure 3-1 shows the available processor modes on a POWER9 system.

The screenshot shows a configuration window for a POWER9 system. The window has several tabs: General, Processors (selected), Memory, I/O, Virtual Adapters, Power Controlling, and Settings. The main content area is titled "Detailed below are the current processing settings for this partition profile." and is divided into several sections:

- Processing mode:** Radio buttons for "Dedicated" and "Shared" (selected).
- Processing units:** Fields for "Total managed system processing units : 16.00", "Minimum shared processing units : 0.1", "Desired shared processing units : 4.0", "Maximum shared processing units : 8.0", and "Shared processor pool: DefaultPool (0)".
- Virtual processors:** Fields for "Minimum processing units required for each virtual processor : 0.10", "Newer operating system levels support : 0.05", "Minimum virtual processors : 1.0", "Desired virtual processors : 8.0", and "Maximum virtual processors : 16.0".
- Sharing mode:** A checked checkbox for "Uncapped" and a "Weight : 128" field with up/down arrows.
- Processor compatibility mode:** A dropdown menu currently set to "default".

At the bottom left are "OK", "Cancel", and "Help" buttons. A dropdown menu is open below the "Processor compatibility mode" field, showing the following options: "default", "POWER7", "POWER8", and "POWER9_base".

Figure 3-1 POWER9 processor modes

Processor compatibility mode is important when Live Partition Mobility (LPM) migration is planned between different generation of servers. An LPAR that potentially might be migrated to a machine that is managed by a processor from another generation must be activated in a specific compatibility mode.

Table 3-1 shows an example where the processor mode must be selected when a migration from POWER9 to POWER8 is planned.

Table 3-1 Processor compatibility modes for a POWER9 to POWER8 migration

Source environment POWER9 server		Destination environment POWER8 server			
		Active migration		Inactive migration	
Wanted processor compatibility mode	Current processor compatibility mode	Wanted processor compatibility mode	Current processor compatibility mode	Wanted processor compatibility mode	Current processor compatibility mode
POWER9	POWER9	Fails because the wanted processor mode is not supported on destination.		Fails because the wanted processor mode is not supported on destination.	
POWER9	POWER8	Fails because the wanted processor mode is not supported on destination.		Fails because the wanted processor mode is not supported on destination.	
Default	POWER9	Fails because the wanted processor mode is not supported on destination.		Default	POWER8
POWER8	POWER8	POWER8	POWER8	POWER8	POWER8
Default	POWER8	Default	POWER8	Default	POWER8
POWER7	POWER7	POWER7	POWER7	POWER7	POWER7

3.3 Active Memory Expansion

Active Memory Expansion is an optional feature for the Power S914 and the Power S924 servers (#4794) or for the Power S922 server (#4793) in the e-Config tool.

This feature enables memory expansion on the system. By using compression and decompression of memory, content can effectively expand the maximum memory capacity, providing more server workload capacity and performance.

Active Memory Expansion is a technology that allows the effective maximum memory capacity to be much larger than the true physical memory maximum. Compression and decompression of memory content can allow memory expansion up to 125% for AIX partitions, which in turn enables a partition to perform more work or support more users with the same physical amount of memory. Similarly, it can allow a server to run more partitions and do more work for the same physical amount of memory.

Note: The Active Memory Expansion feature is not supported by IBM i and the Linux operating system.

3.4 Single Root I/O Virtualization

Single Root I/O Virtualization (SR-IOV) is an extension to the Peripheral Component Interconnect Express (PCIe) specification that allows multiple operating systems to simultaneously share a PCIe adapter with little or no runtime involvement from a hypervisor or other virtualization intermediary.

SR-IOV is PCI standard architecture that enables PCIe adapters to become self-virtualizing. It enables adapter consolidation, through sharing, much like logical partitioning enables server consolidation. With an adapter capable of SR-IOV, you can assign virtual *slices* of a single physical adapter to multiple partitions through logical ports; all of this is done without a VIOS.

POWER9 provides the following SR-IOV enhancements:

- ▶ Faster ports: 10 Gb, 25 Gb, 40 Gb, and 100 Gb
- ▶ More Virtual Functions (VFs) per port: Sixty VFs per port (120 VFs per adapter) for 100-Gb adapters
- ▶ vNIC and vNIC failover support for Linux

For more information, see *IBM Power Systems SR-IOV: Technical Overview and Introduction*, REDP-5065.

3.5 PowerVM

The PowerVM platform is the family of technologies, capabilities, and offerings that delivers industry-leading virtualization on IBM Power Systems servers. It is the umbrella branding term for Power Systems virtualization (logical partitioning, IBM Micro-Partitioning®, POWER Hypervisor, VIOS, LPM, and more). As with Advanced Power Virtualization in the past, PowerVM is a combination of hardware enablement and software.

The Power S914, Power S922, and Power S924 servers come with PowerVM Enterprise Edition (#5228) by default. Furthermore, a temporary PowerVM Enterprise license (#ELPM) is delivered for an old server to support a seamless move to POWER9 at no additional cost.

Logical partitions

LPARs and virtualization increase the usage of system resources and add a level of configuration possibilities.

Logical partitioning is the ability to make a server that is run as though it were two or more independent servers. When you logically partition a server, you divide the resources on the server into subsets called LPARs. You can install software on an LPAR, and the LPAR runs as an independent logical server with the resources that you have allocated to the LPAR. LPAR is the equivalent of a virtual machine (VM).

You can assign processors, memory, and input/output devices to LPARs. You can run AIX, IBM i, Linux, and the VIOS in LPARs. The VIOS provides virtual I/O resources to other LPARs with general-purpose operating systems.

LPARs share a few system attributes, such as the system serial number, system model, and processor feature code (FC). All other system attributes can vary from one LPAR to another.

Micro-Partitioning

When you use the Micro-Partitioning technology, you can allocate fractions of processors to an LPAR. An LPAR that uses fractions of processors is also known as a *shared processor partition* or *micropartition*. Micropartitions run over a set of processors that is called a *shared processor pool*, and virtual processors are used to let the operating system manage the fractions of processing power that are assigned to the LPAR. From an operating system perspective, a virtual processor cannot be distinguished from a physical processor, unless the operating system is enhanced to determine the difference. Physical processors are abstracted into virtual processors that are available to partitions.

On the POWER9 processors, a partition can be defined with a processor capacity as small as 0.05processing units. The shared processor partitions are dispatched and time-sliced on the physical processors under the control of the POWER Hypervisor. The shared processor partitions are created and managed by the HMC.

The Power S914 server supports up to eight cores in a single system. Here are the maximum numbers:

- ▶ Eight dedicated partitions
- ▶ 160 micropartitions (maximum of 20 micropartitions per physical active core)

The Power S922 server supports up to 20 cores in a single system. Here are the maximum numbers:

- ▶ 20 dedicated partitions
- ▶ 400 micropartitions (maximum of 20 micropartitions per physical active core)

The Power S924 server supports up to 24 cores in a single system. Here are the maximum numbers:

- ▶ 24 dedicated partitions
- ▶ 480 micropartitions (maximum of 20 micropartitions per physical active core)

The maximum amounts are supported by the hardware, but the practical limits depend on application workload demands.

Processing mode

When you create an LPAR, you can assign entire processors for dedicated use, or you can assign partial processing units from a shared processor pool. This setting defines the processing mode of the LPAR.

Dedicated mode

In dedicated mode, physical processors are assigned as a whole to partitions. The SMT feature in the POWER9 processor core allows the core to run instructions from two, four, or eight independent software threads simultaneously.

Shared dedicated mode

On POWER9 processor technology-based servers, you can configure dedicated partitions to become processor donors for idle processors that they own, allowing for the donation of spare CPU cycles from dedicated processor partitions to a shared processor pool. The dedicated partition maintains absolute priority for dedicated CPU cycles. Enabling this feature can help increase system usage without compromising the computing power for critical workloads in a dedicated processor.

Shared mode

In shared mode, LPARs use virtual processors to access fractions of physical processors. Shared partitions can define any number of virtual processors (the maximum number is 20 times the number of processing units that are assigned to the partition). The POWER Hypervisor dispatches virtual processors to physical processors according to the partition's processing units entitlement. One processing unit represents one physical processor's processing capacity. All partitions receive a total CPU time equal to their processing unit's entitlement. The logical processors are defined on top of virtual processors. So, even with a virtual processor, the concept of a logical processor exists, and the number of logical processors depends on whether SMT is turned on or off.

3.5.1 Multiple shared processor pools

Multiple shared processor pools (MSPPs) are supported on POWER9 processor-based servers. This capability allows a system administrator to create a set of micropartitions with the purpose of controlling the processor capacity that can be used from the physical shared processor pool.

Micropartitions are created and then identified as members of either the default processor pool or a user-defined shared processor pool. The virtual processors that exist within the set of micropartitions are monitored by the POWER Hypervisor, and processor capacity is managed according to user-defined attributes.

If the Power Systems server is under heavy load, each micropartition within a shared processor pool is assured of its processor entitlement, plus any capacity that it might be allocated from the reserved pool capacity if the micropartition is uncapped.

If certain micropartitions in a shared processor pool do not use their capacity entitlement, the unused capacity is ceded and other uncapped micropartitions within the same shared processor pool are allocated the additional capacity according to their uncapped weighting. In this way, the entitled pool capacity of a shared processor pool is distributed to the set of micropartitions within that shared processor pool.

All Power Systems servers that support the MSPPs capability have a minimum of one (the default) shared processor pool and up to a maximum of 64 shared processor pools.

3.5.2 Virtual I/O Server

The VIOS is part of PowerVM. It is specific appliance that allows the sharing of physical resources between LPARs to allow more efficient usage (for example, consolidation). In this case, the VIOS owns the physical resources (SCSI, Fibre Channel, network adapters, or optical devices) and allows client partitions to share access to them, thus minimizing the number of physical adapters in the system. The VIOS eliminates the requirement that every partition owns a dedicated network adapter, disk adapter, and disk drive. The VIOS supports OpenSSH for secure remote logins. It also provides a firewall for limiting access by ports, network services, and IP addresses.

Figure 3-2 shows an overview of a VIOS configuration.

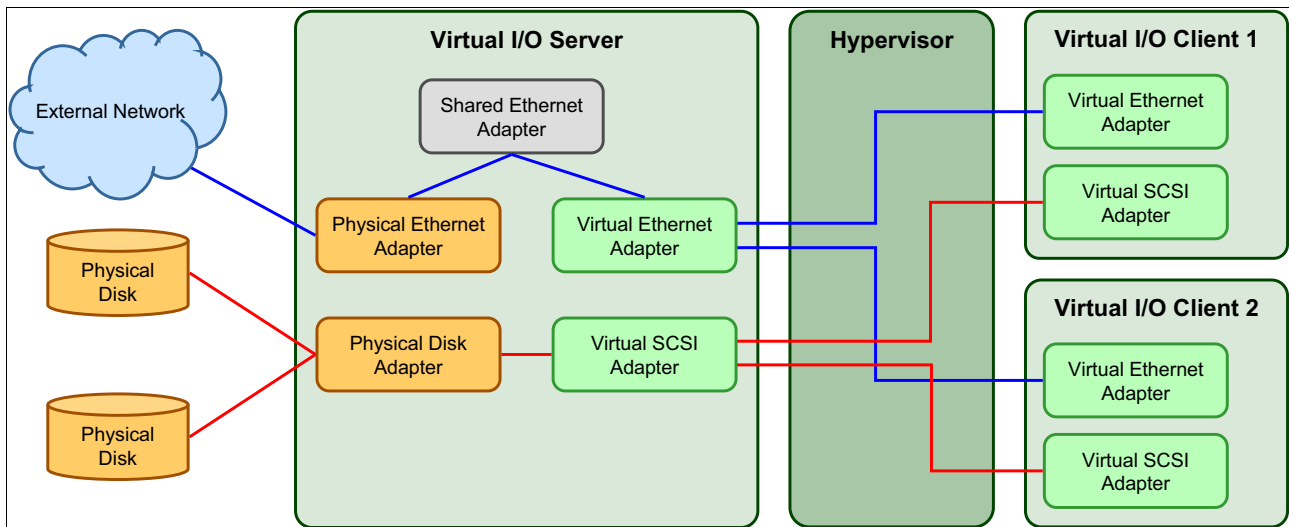


Figure 3-2 Architectural view of the VIOS

It is a preferred practice to run two VIOSes per physical server.

Shared Ethernet Adapter

A SEA can be used to connect a physical Ethernet network to a virtual Ethernet network. The SEA provides this access by connecting the POWER Hypervisor VLANs with the VLANs on the external switches. Because the SEA processes packets at Layer 2, the original MAC address and VLAN tags of the packet are visible to other systems on the physical network. IEEE 802.1 VLAN tagging is supported.

By using the SEA, several client partitions can share one physical adapter, and you can connect internal and external VLANs by using a physical adapter. The SEA service can be hosted only in the VIOS, not in a general-purpose AIX or Linux partition, and acts as a Layer 2 network bridge to securely transport network traffic between virtual Ethernet networks (internal) and one or more (Etherchannel) physical network adapters (external). These virtual Ethernet network adapters are defined by the POWER Hypervisor on the VIOS.

Virtual SCSI

Virtual SCSI is used to see a virtualized implementation of the SCSI protocol. Virtual SCSI is based on a client/server relationship. The VIOS LPAR owns the physical resources and acts as a server or, in SCSI terms, a target device. The client LPARs access the virtual SCSI backing storage devices that are provided by the VIOS as clients.

The virtual I/O adapters (virtual SCSI server adapter and a virtual SCSI client adapter) are configured by using a managed console or through the Integrated Virtualization Manager on smaller systems. The virtual SCSI server (target) adapter is responsible for running any SCSI commands that it receives. It is owned by the VIOS partition. The virtual SCSI client adapter allows a client partition to access physical SCSI and SAN-attached devices and LUNs that are assigned to the client partition. The provisioning of virtual disk resources is provided by the VIOS.

N_Port ID Virtualization

N_Port ID Virtualization (NPIV) is a technology that allows multiple LPARs to access independent physical storage through the same physical Fibre Channel adapter. This adapter is attached to a VIOS partition that acts only as a pass-through, managing the data transfer through the POWER Hypervisor.

Each partition has one or more virtual Fibre Channel adapters, each with their own pair of unique worldwide port names, enabling you to connect each partition to independent physical storage on a SAN. Unlike virtual SCSI, only the client partitions see the disk.

For more information and requirements for NPIV, see *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

3.5.3 Live Partition Mobility

LPM allows you to move a running LPAR from one system to another without disruption. Inactive partition mobility allows you to move a powered-off LPAR from one system to another one.

LPM provides systems management flexibility and improves system availability by:

- ▶ Avoiding planned outages for hardware upgrade or firmware maintenance.
- ▶ Avoiding unplanned downtime. With preventive failure management, if a server indicates a potential failure, you can move its LPARs to another server before the failure occurs.

For more information and requirements for NPIV, see *IBM PowerVM Live Partition Mobility*, SG24-7460.

3.5.4 Active Memory Sharing

Active Memory Sharing provides system memory virtualization capabilities, allowing multiple partitions to share a common pool of physical memory.

The physical memory of an IBM Power System can be assigned to multiple partitions in either dedicated or shared mode. A system administrator can assign some physical memory to a partition and some physical memory to a pool that is shared by other partitions. A single partition can have either dedicated or shared memory:

- ▶ With a pure dedicated memory model, the system administrator's task is to optimize available memory distribution among partitions. When a partition suffers degradation because of memory constraints and other partitions have unused memory, the administrator can manually issue a dynamic memory reconfiguration.
- ▶ With a shared memory model, the system automatically decides the optimal distribution of the physical memory to partitions and adjusts the memory assignment based on partition load. The administrator reserves physical memory for the shared memory pool, assigns partitions to the pool, and provides access limits to the pool.

3.5.5 Active Memory Deduplication

In a virtualized environment, the systems might have a considerable amount of duplicated information that is stored on RAM after each partition has its own operating system, and some of them might even share kinds of applications. On heavily loaded systems, this behavior might lead to a shortage of the available memory resources, forcing paging by the Active Memory Sharing partition operating systems, the Active Memory Deduplication pool, or both, which might decrease overall system performance.

Active Memory Deduplication allows the POWER Hypervisor to map dynamically identical partition memory pages to a single physical memory page within a shared memory pool. This way enables a better usage of the Active Memory Sharing shared memory pool, increasing the system's overall performance by avoiding paging. Deduplication can cause the hardware to incur fewer cache misses, which also leads to improved performance.

Active Memory Deduplication depends on the Active Memory Sharing feature being available, and it uses CPU cycles that are donated by the Active Memory Sharing pool's VIOS partitions to identify deduplicated pages. The operating systems that are running on the Active Memory Sharing partitions can suggest to the POWER Hypervisor that some pages (such as frequently referenced read-only code pages) are good for deduplication.

3.5.6 Remote Restart

Remote Restart is a high availability option for partitions. If there is an error that causes a server outage, a partition that is configured for Remote Restart can be restarted on a different physical server. At times, it might take longer to start the server, in which case the Remote Restart function can be used for faster reprovisioning of the partition. Typically, this can be done faster than restarting the server that failed and then restarting the partitions. The Remote Restart function relies on technology similar to LPM where a partition is configured with storage on a SAN that is shared (accessible) by the server that will host the partition.

The HMC version V9R1 brings following enhancements to the Remote Restart feature.

- ▶ Remote restart a partition with reduced or minimum CPU/memory on the target system.
- ▶ Remote restart by choosing a different virtual switch on the target system.
- ▶ Remote restart the partition without turning on the partition on the target system.
- ▶ Remote restart the partition for test purposes when the source-managed system is in the Operating or Standby state.
- ▶ Remote restart through the REST API.

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM PowerAI: Deep Learning Unleashed on IBM Power Systems Servers*, SG24-8409
- ▶ *IBM Power System AC922 Introduction and Technical Overview*, REDP-5472
- ▶ *IBM Power System AC922 Introduction and Technical Overview*, REDP-5494
- ▶ *IBM Power System L922 Introduction and Technical Overview*, REDP-5496
- ▶ *IBM Power Systems LC921 and LC922 Introduction and Technical Overview*, REDP-5495
- ▶ *IBM Power Systems H922 and H924 Introduction and Technical Overview*, REDP-5498
- ▶ *IBM Power System S822LC for High Performance Computing Introduction and Technical Overview*, REDP-5405
- ▶ *IBM PowerVM Best Practices*, SG24-8062
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590

You can search for, view, download, or order these documents and other Redbooks publications, Redpapers, web docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ IBM Fix Central website
<http://www.ibm.com/support/fixcentral/>
- ▶ IBM Knowledge Center
<http://www.ibm.com/support/knowledgecenter/>
- ▶ IBM Knowledge Center: IBM Power Systems Hardware
<http://www-01.ibm.com/support/knowledgecenter/api/redirect/powersys/v3r1m5/index.jsp>
- ▶ IBM Knowledge Center: Migration combinations of processor compatibility modes for active Partition Mobility
<http://www-01.ibm.com/support/knowledgecenter/api/redirect/powersys/v3r1m5/topic/p7hc3/iphc3pcmcombosact.htm>

- ▶ IBM Portal for OpenPOWER - POWER9 Monza Module
https://www.ibm.com/systems/power/openpower/tgcmDocumentRepository.xhtml?aliasId=POWER9_Monza
- ▶ IBM Power Systems
<http://www.ibm.com/systems/power/>
- ▶ IBM Storage website
<http://www.ibm.com/systems/storage/>
- ▶ IBM Systems Energy Estimator
<http://www-912.ibm.com/see/EnergyEstimator/>
- ▶ IBM System Planning Tool website
<http://www.ibm.com/systems/support/tools/systemplanningtool/>
- ▶ NVIDIA Tesla V100
<https://www.nvidia.com/en-us/data-center/tesla-v100/>
- ▶ NVIDIA Tesla V100 Performance Guide
<http://images.nvidia.com/content/pdf/volta-marketing-v100-performance-guide-us-r6-web.pdf>
- ▶ OpenCAPI
<http://opencapi.org/technical/use-cases/>
- ▶ OpenPOWER Foundation
<https://openpowerfoundation.org/>
- ▶ Power Systems Capacity on Demand website
<http://www.ibm.com/systems/power/hardware/cod/>
- ▶ Support for IBM Systems website
<http://www.ibm.com/support/entry/portal/Overview?brandind=Hardware~Systems~Power>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5497-00

ISBN 0738456934

Printed in U.S.A.

Get connected

