

IBM Power11 Scale-Out Servers

Introduction and Overview

Tim Simon

Henry Vo

Larry Bolhuis

David Harlow

Jean-Manuel Lenez

Jaqui Lynch

Michael Miller

Dean Mussari

Steve Palazzolo

Gerd Reinhardt



IBM Power

 **Hybrid Cloud**



IBM Redbooks

IBM Power11 Scale-Out Servers: Introduction and Overview

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

First Edition (October 2025)

This edition applies to IBM Power11 models 9824-42A, 9856-42H, 9824-22A, and 9856-22H.

© **Copyright International Business Machines Corporation 2025. All rights reserved.**

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

Contents

Notices	ix
Trademarks	x
Preface	xi
Authors	xii
Now you can become a published author, too!	xiv
Comments welcome	xv
Stay connected to IBM Redbooks	xv
Chapter 1. Overview of IBM Power11 scale-out servers	1
1.1 Power11 processor-based servers overview	2
1.2 IBM Power roadmaps	3
1.2.1 IBM POWER processor	4
1.2.2 IBM AIX roadmap	5
1.2.3 IBM i roadmap	5
1.2.4 PowerVM VIOS roadmap	6
1.2.5 Red Hat Enterprise Linux and IBM Power	7
1.2.6 SUSE Enterprise Linux and IBM Power	7
1.3 Power11 platform improvements from the Power10 platform	8
1.3.1 New Power11 processor-based server portfolio	9
1.4 The Power11 platform: Trusted, autonomous, and modern	13
1.4.1 Building a trusted infrastructure with Quantum-Safe Encryption	14
1.4.2 Protecting, detecting, and recovering with IBM Power Cyber Vault	19
1.4.3 Quantum-safe compliance with one-click inventory discovery	22
1.4.4 Zero planned downtime and end-to-end orchestration for maintenance events ..	23
1.4.5 Automated data collection for faster error resolution	27
1.4.6 Increasing efficiency and reducing energy consumption through smart scheduling options	28
1.4.7 Consolidating Oracle 19c workloads	32
1.4.8 Day 1 availability of the Power11 platform in IBM Power Virtual Server with hybrid billing	33
1.4.9 Off-chip AI acceleration with the IBM Spyre Accelerator	34
1.4.10 Red Hat OpenShift with support for an expanded software ecosystem	36
1.5 Introducing the Power11 processor	36
1.5.1 New Power11 processor details	37
1.5.2 Comparing Power11, Power10, and Power9 processors	38
1.5.3 Chip packaging options	39
1.5.4 Spare core support	42
1.6 Operating system support	43
1.6.1 AIX operating system	43
1.6.2 IBM i operating system	44
1.6.3 Linux operating system	44
1.6.4 VIOS	45
1.7 Firmware and Hardware Management Console	45
1.7.1 HMC requirements	46
1.8 Rack support	46
1.8.1 New rack considerations	47
Chapter 2. IBM Power11 scale-out servers	49

2.1	Power11 scale-out servers	50
2.1.1	Family overview	50
2.2	Power S1122 (9824-22A) server overview	51
2.2.1	Processors	51
2.2.2	Memory	52
2.2.3	PCIe slots	52
2.2.4	NVMe drives	53
2.2.5	I/O expansion drawers	53
2.2.6	Power S1122 system configuration	55
2.2.7	Comparing Power S1122, IBM Power S1022, and IBM Power S1022s servers	58
2.3	Power S1124 (9824-42A) server overview	60
2.3.1	Processors	61
2.3.2	Memory	62
2.3.3	PCIe slots	62
2.3.4	NVMe drive support	63
2.3.5	I/O expansion drawers	63
2.3.6	Power S1124 system minimum configuration	65
2.3.7	Comparing Power S1124 and Power S1024	68
2.4	Processor and memory activation on Power S1122 and Power S1124 systems	69
2.4.1	Advanced memory functions on Power11 scale-out servers	70
2.5	Environmental and physical specifications	71
Chapter 3. I/O subsystem		75
3.1	Internal I/O	76
3.1.1	Internal PCIe Gen 5 subsystem	77
3.1.2	Internal NVMe storage subsystem	80
3.1.3	NVMe backplane	81
3.1.4	NVMe JBOF card	81
3.1.5	NVMe JBOF to backplane cabling	83
3.2	Enhancing I/O scalability with expansion drawers	85
3.2.1	PCIe expansion card	85
3.2.2	Supported I/O drawers	90
3.2.3	Non-supported drawers	91
3.2.4	PCIe Gen 4 I/O Expansion Drawer	91
3.3	Available adapter list	98
3.3.1	Storage area network adapters	98
3.3.2	Network adapters	99
3.3.3	Transceiver (small form-factor pluggable) replacement support for FC and Ethernet Adapter Features	101
3.3.4	SAS and tape adapters	101
3.4	NVMe U.2 devices	102
3.4.1	USB adapters	104
3.4.2	Cryptographic adapters	105
3.4.3	Graphical adapters	107
3.5	Other device support	108
3.5.1	No longer supported	108
3.5.2	PCIe adapters that require increased cooling	109
Chapter 4. Artificial intelligence support		113
4.1	On-chip support	114
4.2	AI acceleration with IBM Spyre Accelerator	114
4.2.1	Deploying AI in the enterprise	115
4.2.2	Fit-for-purpose AI acceleration	116

4.2.3 IBM Spyre Accelerator	117
4.3 Solutions	119
4.3.1 Barriers to scaling AI and the IBM Power technology advantage	120
4.3.2 IBM Power AI-ready portfolio	120
4.3.3 Strategic outcomes that are enabled by IBM Power	121
4.3.4 Uniqueness of the platform	121
Chapter 5. Automation and management	123
5.1 Hardware Management Console overview	124
5.1.1 HMC options	124
5.1.2 BMC network connectivity rules for the 7063-CR2 HMC	126
5.1.3 High availability HMC configuration	127
5.1.4 HMC code level requirements	127
5.1.5 Keeping your HMC up to date	128
5.1.6 New features	128
5.1.7 Using the automated maintenance tool	129
5.2 Ansible	131
5.3 Terraform	135
5.4 PowerVC	137
Chapter 6. Operating systems	141
6.1 Subscription licensing	142
6.1.1 AIX subscription licensing	142
6.1.2 IBM i subscription licensing	143
6.2 AIX	143
6.2.1 AIX 7.3 key features	144
6.2.2 Supported levels	144
6.2.3 AIX maintenance levels	145
6.2.4 Licensing	145
6.2.5 New AIX editions	146
6.3 IBM i	148
6.3.1 Introducing IBM i	148
6.3.2 Supported levels	149
6.3.3 New IBM i features that are supported by the Power11 platform	150
6.3.4 IBM i software tiers	150
6.3.5 IBM i licensing	151
6.3.6 IBM i subscription bundles for P20 and P30	152
6.3.7 IBM i Software Maintenance Agreement	155
6.3.8 Capacity BackUp and dedicated backup machine licensing models for IBM i	155
6.3.9 Upgrading an existing IBM i license to Version 7.6	157
6.3.10 IBM i simplification	157
6.4 Linux on Power	157
6.4.1 Supported distributions	158
6.4.2 Licensing	162
6.5 Red Hat OpenShift	164
6.6 Introducing Kernel-based Virtual Machine support	166
6.7 PowerVM Virtual I/O Server	167
6.8 Setting your LPAR compatibility mode	167
Chapter 7. Enterprise solutions	169
7.1 High availability and disaster recovery solutions	170
7.1.1 PowerHA SystemMirror for AIX	170
7.1.2 IBM PowerHA SystemMirror for i	170
7.1.3 VM Recovery Manager	171

7.2 IBM Db2	171
7.3 Oracle	172
7.3.1 Running Oracle Standard Edition 2 on IBM Power	173
7.4 SAP HANA	175
7.5 Banking	176
7.6 Healthcare	178
7.6.1 Epic	179
Chapter 8. Servicing IBM Power11 scale-out servers	181
8.1 IBM maintenance	182
8.2 IBM Expert Care	182
8.3 IBM tools and interfaces	183
8.4 Baseboard Management Controller card	184
8.5 ASMI	184
8.5.1 More ASMI functions	186
8.6 Entitled System Support	188
8.7 System firmware	188
8.7.1 Update Access Keys	198
Chapter 9. Virtualization	205
9.1 PowerVM	206
9.1.1 IBM PowerVM Hypervisor	206
9.1.2 Multiple shared processor pools	210
9.1.3 Virtual I/O Server	210
9.1.4 Live Partition Mobility	212
9.1.5 Active Memory Expansion	212
9.1.6 Remote Restart	213
9.1.7 POWER processor modes	213
9.1.8 Single Root I/O Virtualization	214
9.1.9 Virtualization features resources	214
9.1.10 Resource Groups	215
9.2 KVM support	222
Chapter 10. Hybrid cloud solutions	225
10.1 IBM Power Private Cloud with Shared Utility Capacity	226
10.1.1 IBM Cloud Management Console	227
10.2 Red Hat OpenShift	228
10.2.1 Red Hat OpenShift on IBM Power	228
10.2.2 Red Hat OpenShift AI on Power servers	231
10.3 IBM Power Private Cloud Rack Solution	232
10.3.1 Key features and components	232
10.3.2 Available configurations	233
10.3.3 IBM Cloud Pak	234
10.3.4 Other cloud enablement solutions	236
10.4 IBM Power Virtual Server	236
10.4.1 IBM Power Virtual Server options	237
10.4.2 IBM Power Virtual Server in the cloud	239
10.4.3 IBM Power Virtual Server Private Cloud	240
10.4.4 Introducing Power11 processor-based servers in the cloud	241
Abbreviations and acronyms	245
Related publications	247
IBM Redbooks	247

Online resources	247
Help from IBM	248

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](https://www.ibm.com/legal/copytrade.shtml) 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 <https://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®	IBM Security®	PowerVM®
Concert®	IBM Spectrum®	Rational®
FICON®	IBM Spyre™	Redbooks®
HashiCorp®	IBM Turbonomic®	Redbooks (logo)  ®
IBM®	IBM Watson®	Satellite®
IBM API Connect®	IBM watsonx®	Spyre™
IBM Cloud®	IBM Z®	Storwize®
IBM Cloud Pak®	Instana®	SystemMirror®
IBM Cloud Satellite®	Micro-Partitioning®	Terraform®
IBM Concert®	Passport Advantage®	Tivoli®
IBM FlashSystem®	Power Architecture®	Turbonomic®
IBM Garage®	Power8®	watsonx®
IBM Instana®	Power9®	watsonx.ai®
IBM Research®	PowerHA®	WebSphere®

The following terms are trademarks of other companies:

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

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

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

Red Hat, Ansible, Fedora, OpenShift, 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, VMware vSphere, 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

IBM Power11 scale-out servers are designed to meet the evolving demands of modern enterprises. They offer an integrated, end-to-end infrastructure stack that delivers high performance, robust security, and operational agility. These systems support autonomous IT operations (ITOps), helping organizations reduce risk, streamline management, and accelerate innovation across hybrid and multicloud environments.

With a strong focus on business continuity, Power11 processor-based servers provide zero planned downtime (ZPD) for maintenance, automated failover by using spare processor cores, and rapid ransomware detection and recovery through IBM Power Cyber Vault. These servers also improve productivity and efficiency by simplifying infrastructure operations and supporting flexible consumption models such as Capacity on Demand (CoD) and Power Enterprise Pools (PEP) 2.0.

Power11 scale-out servers are designed for the future with quantum-safe security capabilities. They include advanced cryptographic protections such as Secure Boot and Live Partition Mobility (LPM), which are hardened against quantum-era threats. The servers also support the latest FIPS 140-3 certified cryptographic modules, helping safeguard sensitive data and ensure long-term compliance in a complex threat landscape.

For organizations adopting artificial intelligence (AI), Power11 processor-based servers are AI-ready by design. They feature built-in hardware acceleration for machine learning and inferencing workloads. Their seamless integration with IBM Cloud® enables rapid scaling and consistent performance, enabling enterprises to deploy new applications and AI use cases with confidence.

This IBM Redbooks® publication provides a hardware architecture analysis and highlights the changes, new technologies, and major features that are introduced in these systems.

This publication is intended for the following professionals seeking a deeper understanding of IBM Power server products:

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

This book expands the IBM Power documentation set by offering a desktop reference with a detailed technical description of the Power11 scale-out server models.

Authors

This book was produced by a team of specialists from around the world working with IBM Redbooks.

Tim Simon is an IBM Redbooks® Project Leader who is based in Tulsa, Oklahoma, US. He has more than 40 years of experience at IBM®, primarily in technical sales, and has built a distinguished career helping clients design and implement IBM solutions that align with their business goals. He holds a Bachelor of Science degree in Mathematics from Towson University in Maryland. Tim's expertise spans a broad range of IBM technologies, with deep experience in developing customer solutions that use IBM Power, IBM Storage, and IBM Z® platforms.

Henry Vo is an IBM Redbooks Project Leader with 11 years of experience at IBM. He brings strong technical expertise in business problem-solving, risk and root-cause analysis, and the development of technical strategies that align with business objectives. Throughout his career at IBM, Henry has held various roles, including Project Manager, ST/FT/ETE Tester, Back-End Developer, and Department of Labor (DOL) Agent for New York. He is certified in both IBM z and IBM Power technologies. Henry earned his master's degree in management information systems from the University of Texas at Dallas in 2012 and continues to apply his diverse skill set to drive innovation and deliver value across IBM projects.

Larry Bolhuis is an IBM i and cloud infrastructure advisor for Service Express, an IBM Platinum Business Partner who is based in Grand Rapids, Michigan. He brings 37 years of experience with IBM i systems, following earlier work on the IBM S/34 and IBM S/36 platforms. Since 1988, Larry has volunteered with COMMON, serving in nearly every role from session manager to President. He has contributed to more than 30 IBM certifications and holds the Advanced Technical Expert badge. Larry is internationally recognized as a speaker at COMMON and other user groups across North America and Europe. He has been named an IBM Champion for 6 consecutive years.

David Harlow is a Senior Solutions Architect with Converge Technology Solutions, Inc., who is based in Spring Hill, Tennessee. He has extensive pre-sales experience in architecture, system design, planning, and project management. David specializes in implementing large-scale, mission-critical projects, including production system upgrades, new installations, Power servers with IBM i upgrades and consolidation, IBM i with Virtual I/O Server (VIOS), and IBM Storwize® storage solutions. He also focuses on business continuity and disaster recovery (DR). David consistently delivers complex projects on time and ahead of schedule. He previously contributed to IBM Redbooks, including *IBM BladeCenter PS703 and PS704 Technical Overview and Introduction*, REDP-4744.

Jean-Manuel Lenez is a Presales Engineer with IBM Switzerland, serving in this role since 1999. He specializes in UNIX, IBM Power, IBM AIX®, and IBM i server technologies, along with related products such as IBM PowerVM®, IBM PowerHA®, PowerSC, Linux on Power, and IBM Cloud. Jean-Manuel leads projects with major customers on topics including AI, deep learning, SAP HANA, server consolidation, and high availability and disaster recovery (HADR).

Jaqui Lynch is an independent consultant who is based in the US, with more than 48 years of experience across diverse platforms and operating systems (OSs). Her background includes mainframes, UNIX systems, midrange systems, and personal workstations. Since moving to UNIX in 1993, she has focused on IBM Power technologies beginning in 1999. Jaqui spent over 22 years as an IBM customer before joining three IBM Business Partners. Her consulting expertise includes enterprise architecture, virtualization, performance optimization, security, and solution delivery on Power running AIX and Linux. She is a frequent speaker at IBM and industry conferences and contributes regularly to IBM TechChannel magazine. Jaqui holds a master's degree in philosophy from the University of Illinois Springfield.

Michael Miller is an IBM Champion and Enterprise Architect with Peller Technologies (formerly Mainline Information Systems), an IBM Platinum Business Partner. He has 25 years of experience with IBM i and works extensively with IBM Power platforms, covering presales architecture through post-sales implementation. Michael holds multiple IBM certifications in both sales and technical domains, including a Professional Cloud Architect credential. He is a recognized expert in IBM Cloud and IBM Power Virtual Server (PowerVS), collaborating closely with IBM developers to shape the future of PowerVS. He has spoken at conferences such as IBM TechXchange and COMMON POWERUp. He holds a Bachelor of Science degree from Western Kentucky University.

Dean Mussari is an IBM Power Brand Technical Specialist serving the US National Market. He brings over 35 years of experience with IBM servers and storage solutions, primarily in large-scale retail environments. Dean specializes in IBM Power, with a strong focus on IBM i. Before joining IBM, he built a distinguished career delivering enterprise infrastructure solutions that enhance performance and reliability. He holds a master's degree in computer science from Loyola University Chicago.

Steve Palazzolo is an IBM Power Brand Technical Specialist serving the US National Market. He has more than 36 years of experience designing and implementing IBM server and storage solutions, with a background in large-scale retail environments. Before joining IBM, Steve led an IBM Power implementation team at a managed service provider. His expertise centers on Power servers, particularly in IBM i environments, and IBM FlashSystem® storage solutions. Steve holds a Bachelor of Arts degree from Barat College of DePaul University.

Gerd Reinhardt is a Hardware Product Engineer for IBM Power who is based in Ehningen, Baden-Württemberg, Germany. He has over 32 years of experience at IBM, including more than 27 years specializing in hardware support for IBM Power and its predecessors IBM System i5, IBM iSeries, and IBM AS/400. Gerd has held various roles in both front-end and back-end support. He holds a certificate in electronics, an advanced technical college qualification, and a diploma in business administration from the Academy of Bad Harzburg. He has also been recognized by the Trustforte Corporation (New York, NY) as having attained the equivalent of a Bachelor of Science degree in Information Technology. Gerd is IBM-accredited at the Senior Level within the Product Services Profession. His areas of expertise include IBM i, IBM Power hardware, and PowerVM. He has contributed to IBM technical documentation and co-authored several IBM Redbooks, including the IBM Systems Handbook and the IBM System Builder series.

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

Nicole Nett
IBM Chief Architect Power Servers, IBM Austin

Lyudmila Simeonova
IBM AIX I/O Device Drivers/PowerVM VIOS, IBM Bulgaria

Milen Paiakov
IBM Power Courseware Developer, IBM Bulgaria

Ruby Zgabay
IBM Power Systems Chief Engineer, IBM Austin

Kaveh Naderi
IBM Power System Chief Architect, IBM Austin

Paul Bastide
IBM Senior Software Engineer, IBM Cambridge

Tanya Koleva
IBM Technology Lifecycle Services, IBM Bulgaria

Manjunath Kumatagi
IBM Senior Software Engineer, Red Hat OpenShift on Power, IBM Bangalore

Manoj Kumar
IBM Chief Engineer, Red Hat OpenShift, Red Hat OpenShift AI, IBM Turbonomic®, IBM Instana® on Power, Austin, TX

Pablo Pereira Ruiz
Senior UNIX/AIX Administrator/Support en Interamericana de Cómputos S.A., Uruguay

To the authors of *IBM Power11 E1150 Introduction and Technical Overview*, SG24-8589 and *IBM Power E1180: Introduction and Overview*, SG24-8587 for sharing content freely for this publication.

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 IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

<https://www.linkedin.com/groups/2130806>

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

<https://www.redbooks.ibm.com/subscribe>

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

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



Overview of IBM Power11 scale-out servers

IBM Power11 processor-based servers represent an advancement in enterprise computing. Building on the strengths of Power10 processor-based servers, Power11 processor-based servers introduce up to 25% more cores per chip, higher clock speeds, and improved energy efficiency. It continues IBM's focus on reliability, availability, and serviceability (RAS) and integrates quantum-safe security features to make critical workloads future-ready. The processor supports demanding enterprise applications, particularly those involving artificial intelligence (AI), analytics, and hybrid cloud environments.

Power11 processor-based servers use Integrated Stack Capacitor (ISC) technology and advanced 2.5D packaging to boost performance and efficiency. These innovations, when combined with enhanced thermal management, such as more efficient fans and heat sinks, improve system density, cooling effectiveness, and overall compute capability.

A key feature of Power11 processor-based servers is continued support for AI workloads through the Matrix Math Accelerator (MMA) architecture and the IBM Spyre™ Accelerator, which is optimized for generative AI and complex model inference. Power11 processor-based servers are a strong alternative to traditional GPU-heavy AI infrastructures. In addition, Power11 processor-based servers strengthen their virtualization capabilities with deeper Kernel-based Virtual Machine (KVM) integration, improving compatibility with Linux-native tools and hybrid cloud platforms. It is a versatile choice for enterprises seeking scalable, AI-ready infrastructure.

The following topics are covered in this chapter:

- ▶ Power11 processor-based servers overview
- ▶ IBM Power roadmaps
- ▶ Power11 platform improvements from the Power10 platform
- ▶ The Power11 platform: Trusted, autonomous, and modern
- ▶ Introducing the Power11 processor
- ▶ Operating system support
- ▶ Firmware and Hardware Management Console
- ▶ Rack support

1.1 Power11 processor-based servers overview

For more than 35 years, IBM Power servers have been a cornerstone of enterprise computing, delivering unmatched performance, reliability, and availability for mission-critical workloads worldwide. From the original IBM RS/6000 systems that were powered by the Power1 processor to the advanced Power11 platforms, the IBM Power platform evolved to meet the demands of modern IT infrastructure.

The IBM Power Architecture® follows a consistent innovation cadence, introducing a new generation of processors approximately every 3 years. Each iteration delivers enhancements, such as more cores and threads, improved energy efficiency, increased memory bandwidth, and expanded I/O capabilities, helping ensure that the platform remains at the forefront of enterprise computing.

The Power11 platform is built on three foundational principles:

- ▶ **Trusted**

Delivers continuous business operations with up to 99.9999% availability¹ (on the IBM Power E1180), zero planned downtime (ZPD) for maintenance², quantum-safe cryptography, and accelerated system recovery.

- ▶ **Autonomous**

Boosts operational efficiency through intelligent automation, dynamic performance tuning, and reduced manual intervention.

- ▶ **Modern**

Enables AI-infused application deployment across hybrid environments, whether on-premises and in the cloud, by using off-chip AI acceleration and seamless integration with platforms such as Red Hat OpenShift and IBM watsonx®.

The Power11 platform is segmented into three primary tiers to address a wide range of workload requirements:

- ▶ **Entry-level**

IBM Power S1122 and IBM Power S1124, which are optimized for departmental and edge workloads.

- ▶ **Midrange**

IBM Power E1150, which is ideal for enterprise consolidation and scalable virtualization.

- ▶ **High-end**

IBM Power E1180, which is designed for large-scale, mission-critical, and AI-intensive workloads.

The Power11 platform introduces several key advancements across the portfolio:

- ▶ Up to six nines (99.9999%) availability on high-end systems
- ▶ Enhanced security with Quantum-Safe Encryption (QSE) and Secure Boot improvements

¹ IBM Power11 E1180 is designed to be the most resilient server in the history of the IBM Power platform. On average, Power E1180 servers are estimated to meet a 99.9999% stand-alone availability standard, as defined in section "Enterprise Single Server Availability Standard" in [Introduction to IBM Power Reliability, Availability, and Serviceability for Power10 and Power11 processor-based servers Using IBM PowerVM](#).

² Based on IBM internal testing of system upgrade scenarios, many tasks, such as Virtual I/O Server (VIOS) updates, hot-plug adapter changes, I/O adapter firmware updates, and concurrent system firmware updates can be completed in place. Some tasks, such as non-concurrent system firmware and hardware maintenance, might require Live Partition Mobility (LPM) support.

- ▶ Performance gains of up to 15% per core and up to 25% per thread through combined hardware and software optimizations
- ▶ Significant throughput improvements in midrange and entry-level systems
- ▶ Up to 50% increase in memory bandwidth by using DDR5 and advanced memory controllers

Also, the Power11 platform introduces serial number-preserving upgrades from Power10 processor-based servers (available starting Q4 2025), enabling seamless migrations without disrupting asset tracking or licensing.

This book focuses on Power11 scale-out servers, which are the entry-level offerings in the Power11 platform. It delivers a balanced combination of performance, scalability, and cost efficiency, making it an ideal platform for enterprise workloads, virtualization, and AI integration.

1.2 IBM Power roadmaps

IBM continues to advance its Power Architecture to address the evolving demands of AI, hybrid cloud, and mission-critical enterprise workloads. Central to IBM's hybrid cloud strategy, the Power platform is optimized for Red Hat OpenShift, enabling containerized workloads and cloud-native application development.

IBM Power Virtual Server (PowerVS) extends these capabilities into the cloud so that clients can run Power server workloads on-premises or in IBM Cloud while benefiting from flexible consumption models and meeting data residency and compliance requirements. This hybrid approach helps ensure seamless workload portability and supports modernization initiatives.

The Power platform also maintains a robust ecosystem of operating systems (OSs), including IBM AIX (a UNIX like OS with a roadmap extending beyond 2035); IBM i, which is an integrated OS that is known for its security and stability; and Linux on Power, which supports distributions such as Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server, making it ideal for open-source and cloud-native environments.

Complementing its technology stack, IBM built a strategic network of alliances that expanded the reach and capabilities of the Power platform. Notably, PowerVS is a certified platform for RISE with SAP, enabling high-performance, reliable SAP workload deployment in hybrid cloud environments. Red Hat technologies such as Red Hat OpenShift and Ansible are deeply integrated into Power servers, supporting automation and DevOps practices. Also, IBM collaborates with major cloud providers, such as Microsoft, Amazon Web Services (AWS), and Google Cloud to deliver hybrid and multicloud solutions, enabling Power workloads to extend seamlessly into public cloud infrastructures.

This section describes the IBM roadmap for the evolution of the IBM Power infrastructure.

1.2.1 IBM POWER processor

The IBM Power roadmap continues to evolve with the upcoming Power11 processor-based generation, which is scheduled for launch in Q3 2025. Building on the innovations of the Power10 platform, the Power11 platform introduces enhancements across the processor architecture, packaging, and energy efficiency. A key advancement is the integration of an extra silicon layer for improved energy management, enabling higher performance with reduced power consumption. This generation also increases core strength and count so that it can handle more demanding workloads. IBM is refining its Matrix Math Accelerators (MMAs), first introduced in the Power10 platform, to further accelerate AI inferencing directly on the chip, eliminating the need for external GPUs in many scenarios.

Looking beyond the Power11 platform, the IBM roadmap emphasizes full-stack innovation and adaptability to emerging technologies. The Power11 platform is designed to be memory-neutral, supporting both DDR4 and DDR5 through the Open Memory Interface (OMI), and introduces improved thermal infrastructure through advanced packaging and cooling technologies. These developments enhance performance and reduce operational costs in data centers. Future iterations of the Power Architecture are expected to continue this trend, focusing on hybrid cloud readiness, AI integration, and support for open-source ecosystems. IBM's long-term vision includes deeper collaboration with the OpenPOWER community and a renewed emphasis on openness, flexibility, and sustainability in enterprise computing.

Figure 1-1 illustrates the generations of the POWER processor family and the continued plans for improvement. With the release of the Power11 platform, the next generation of POWER processors is already in development, and more processor designs are in the pipeline.

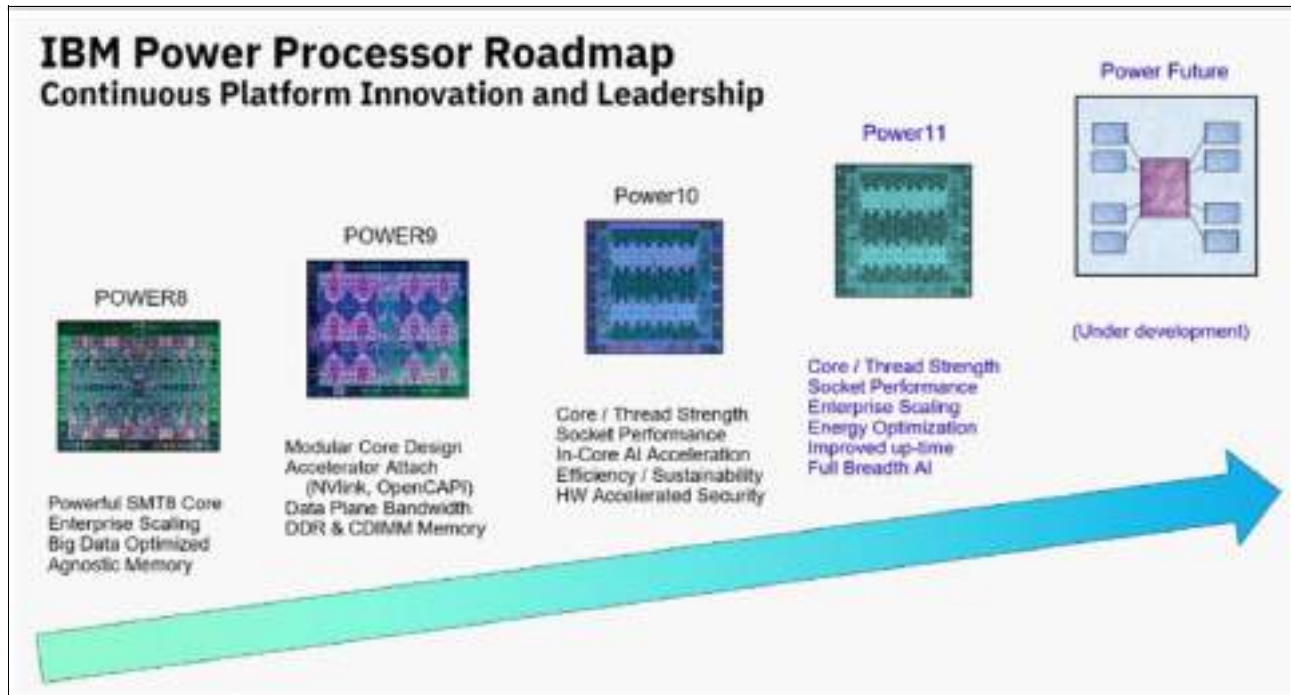


Figure 1-1 POWER processor roadmap

1.2.2 IBM AIX roadmap

The IBM AIX OS is a proprietary UNIX like OS that is built on open standards and designed for Power servers.

For more than 38 years, AIX has powered mission-critical applications and databases in industries with high-performance computing (HPC) needs, such as finance, retail, healthcare, government, manufacturing, and insurance. Paired with Power11 processor-based servers, AIX 7.3 (the latest version) expands its reach into new markets and workloads to deliver unmatched performance, security, scalability, and reliability. It supports digital transformation through flexible subscription models that are tailored to business needs and enables the adoption of technologies such as hybrid cloud, AI, and cloud-native applications.

AIX binary compatibility enables applications to run unchanged and without recompilation on the newest release, safeguarding IT investments in the platform.

IBM is committed to delivering an AIX release roadmap of further innovations and purposeful support that extends beyond 2035, helping ensure a stable platform for enterprise workloads.

Figure 1-2 shows the current roadmap.




Figure 1-2 Support timeline for future generations of AIX

For more information on IBM's strategy and roadmap for IBM AIX see [Strategy and Roadmap for the IBM AIX Operating System](#).

Additional information about AIX can be found in 6.2, "AIX" on page 143.

1.2.3 IBM i roadmap

Since 2008, IBM has consistently released new versions of IBM i approximately every 3 years, with IBM i 7.6 being the latest. Each release introduces new features, workload support, and language enhancements while maintaining compatibility with earlier versions so that customers can preserve and extend their existing investments. IBM i is recognized as one of the most secure, self-optimizing, and scalable OSs in the world. IBM continues to deliver updates and enhancements throughout each release lifecycle through Program Temporary Fixes (PTFs) and Technology Refreshes (TRs).

A key differentiator in the enterprise OS market is IBM's long-term commitment to IBM i, which is demonstrated by its publicly available roadmap, which currently extends to 2035 and is updated annually, underscoring IBM's dedication to the platform's future.

Figure 1-3 shows the IBM i roadmap.

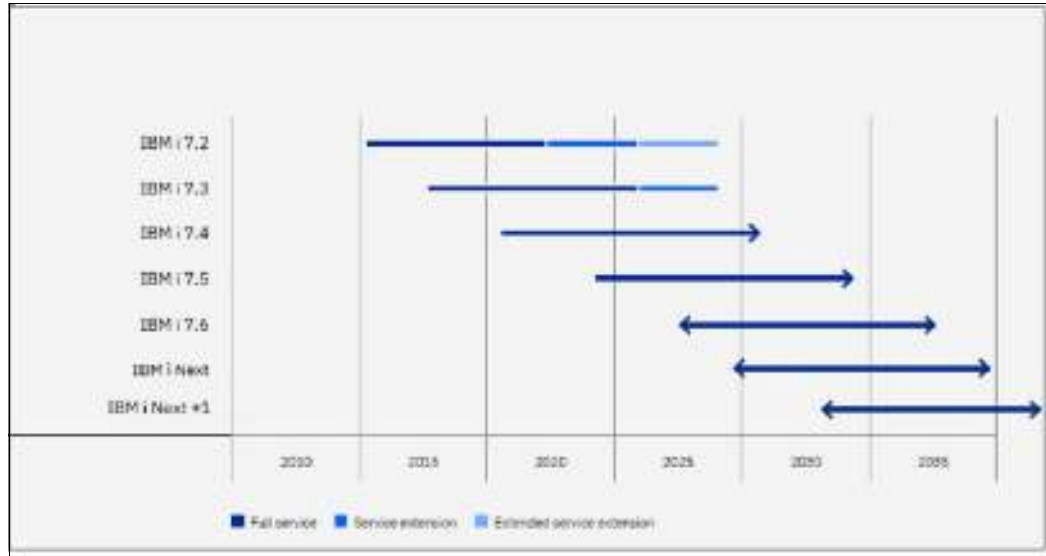


Figure 1-3 IBM i roadmap

1.2.4 PowerVM VIOS roadmap

PowerVM VIOS is a critical component of the IBM Power virtualization stack, enabling shared access to physical I/O resources for client logical partitions (LPARs). Like all IBM software, VIOS follows a defined product lifecycle that includes general availability, support phases, and end of service (EOS) milestones.

Table 1-1 shows the VIOS release schedule as of July 2025.

Table 1-1 VIOS release schedule

VIOS release	Release date	End of fix support ^a	Latest Fix Pack ^b	Next Fix Pack ^c
4.1.1	Dec-24	31 December 2027 (estimated)	VIOS_FP_4.1.1.0	23 July 2025
4.1.0	Nov-23	30 November 2026 (estimated)	VIOS_FP_4.1.0.30	23 July 2025
3.1.4	Dec-22	30 April 2026 (estimated)	VIOS_FP_3.1.4.50	23 July 2025
3.1.3	Sept-21	30-Sept-24	VIOS_FP_3.1.3.40	None
3.1.2	Nov-20	30-Nov-23	VIOS_FP_3.1.2.60	None

- The end of the maintenance period for a VIOS Release. Fix Packs and interim fixes will not be created for a VIOS release after the end of fix support.
- The most recent Fix Pack that is available for the VIOS release and links to that Fix Pack on Fix Central.
- The target availability date for the next Fix Pack for the VIOS Release. After the final Fix Pack for a VIOS release is available, further fixes for the release are provided as interim fixes.

Note: The EOS date for a VIOS release is not included in this table. EOS dates are available on the [IBM Software Lifecycle](#) page.

For more information about the PowerVM VIOS roadmap, see the [PowerVM VIOS Lifecycle Information](#) and [System to PowerVM Virtual I/O Server maps](#).

1.2.5 Red Hat Enterprise Linux and IBM Power

Red Hat affirmed its strong commitment to the IBM Power platform, particularly as both companies align around hybrid cloud, AI, and open-source innovation. This partnership is rooted in their shared vision of delivering enterprise-grade solutions that are scalable, secure, and optimized for modern workloads.

At the Red Hat Summit 2025, Red Hat emphasized its dedication to supporting IBM Power through continued enhancements to RHEL on Power servers. These enhancements include performance tuning for Power10 and Power11 processors and deeper integration with Red Hat OpenShift for containerized and cloud-native applications. These efforts help ensure that Power users can fully use Red Hat's open hybrid cloud technologies across on-premises and cloud environments.

Red Hat also plays a pivotal role in enabling AI workloads on Power, with support for AI inference and acceleration technologies that align with IBM's hardware innovations, such as MMAs and the upcoming IBM Spyre AI Accelerator. This synergy enables enterprises to deploy AI models efficiently on a Power infrastructure by using Red Hat's open-source toolchains and platforms.

In addition, Red Hat and IBM continue to collaborate on automation and DevOps through tools such as Ansible Automation Platform, which is optimized for managing Power environments. This integration simplifies operations, enhances consistency, and accelerates application delivery across hybrid infrastructures.

For more information about RHEL support on IBM Power, see "Red Hat Enterprise Linux" on page 158.

1.2.6 SUSE Enterprise Linux and IBM Power

The partnership between SUSE Linux and Power servers is a long-standing, strategic collaboration that is focused on delivering enterprise-grade Linux solutions for mission-critical workloads, hybrid cloud, and SAP environments.

SUSE and IBM have worked together for more than two decades to provide robust, scalable, and secure Linux solutions on IBM Power infrastructure. Their joint efforts focus on the following items:

- ▶ SUSE Linux Enterprise Server for Power
- ▶ SUSE Linux Enterprise Server for SAP Applications, optimized for SAP HANA on Power
- ▶ Hybrid and multicloud enablement through container and virtualization technologies

SUSE's open-source model aligns with IBM's strategy to provide flexible, vendor-neutral solutions. This approach enables customers to modernize their IT environments while avoiding vendor lock-in.

For more information about SUSE Enterprise Linux, see "SUSE Linux Enterprise Server" on page 161.

1.3 Power11 platform improvements from the Power10 platform

Power11 scale-out servers introduce a comprehensive suite of advanced capabilities, all rooted in IBM's core strength: a full-stack, end-to-end design. By tightly integrating every layer of the infrastructure from the Power11 processor and system architecture to firmware, OSs, and cloud services, IBM Power delivers a synergistic platform that is built for autonomous IT. This holistic approach drives measurable business outcomes across three foundational pillars: unprecedented business continuity, enhanced productivity and efficiency, and accelerated growth and scalability.

Unprecedented business continuity

Power11 scale-out servers redefine operational resilience with features that are designed to eliminate downtime and help ensure continuous service delivery:

- ▶ **Planned downtime elimination**
Using technologies such as live updates, rolling upgrades, and autonomous fixes, Power11 processor-based servers enable maintenance without taking applications offline, targeting zero hours of planned downtime.
- ▶ **Spare core technology**
At the silicon level, the system includes spare processor cores that can be dynamically activated in response to hardware failures. This proactive fault-tolerance mechanism enables seamless failover without requiring system or application restarts, preserving compute capacity and system integrity. For more information about spare core technology, see 1.5.4, "Spare core support" on page 42.
- ▶ **IBM Power Cyber Vault**
Advanced threat detection can identify ransomware attacks in under a minute. Combined with automated recovery mechanisms, this feature enhances cyber resiliency. It is offered in collaboration with IBM Storage and IBM Expert Labs.
- ▶ **Quantum-safe protection**
The Power11 platform incorporates quantum-safe cryptography to secure system restarts and LPM, helping ensure future-ready data protection.
- ▶ **Crypto compliance**
Support for the IBM 4770 Crypto Card enables FIPS 140-3 Level 4 certification, enhancing compliance with stringent security standards.
- ▶ **Automated diagnostics**
Intelligent data collection accelerates error resolution, saving up to eight hours per support ticket and reducing time to recovery.

Enhanced productivity and efficiency

Power11 scale-out servers are designed to maximize IT efficiency and output:

- ▶ **Autonomous IT operations (ITOps)**
With built-in automation and AI-driven management, the system reduces manual intervention and operational impact.
- ▶ **Integrated monitoring and management**
Tools such as Hardware Management Console (HMC) and IBM Power Virtualization Center (PowerVC) provide centralized control, enabling streamlined operations across hybrid environments.

Accelerated growth and scalability

Power11 scale-out servers support dynamic scaling and hybrid cloud integration to meet evolving business demands:

- ▶ PowerVS

Immediate availability of the Power11 platform in IBM Cloud enables seamless extension of workloads to the cloud while maintaining consistent performance, security, and OS compatibility.

- ▶ Hybrid cloud agility

Enterprises can scale resources on demand, accelerate development and testing, and optimize costs through a consumption-based model.

- ▶ AI acceleration with MMA

Power11 processor cores feature MMA, a hardware acceleration unit that is optimized for AI and machine learning workloads. By running matrix operations directly on-chip, MMA reduces latency and boosts throughput, making Power11 scale-out servers ideal for deep learning inference and training tasks.

Technical enhancements

Power11 processor-based servers deliver enhanced performance through higher core frequencies, more cores per server, and improved memory bandwidth with reduced memory latency.

Table 1-2 shows some of the benefits that these improvements can provide to your enterprise.

Table 1-2 Power11 platform performance benefits

Improvement	Benefit
Up to 15 - 25% core performance increase with both hardware and software planned improvements	Exceed service-level agreements (SLAs) for response times and batch windows.
30 - 45% system performance increase with the new larger 30-Core Power11 Dual-Chip Modules (DCMs)	Helps ensure that increasing business workload demands can continue to be met with improved system performance.
Memory technology improvements with up to 50% higher data rates	Enhance performance and system scalability for memory-intensive workloads by optimizing both memory bandwidth and latency.

1.3.1 New Power11 processor-based server portfolio

The new Power11 processor-based server portfolio consists of the following offerings, providing clients with options that are designed to meet specific business requirements.

High-end: Power E1180

At the top of the portfolio is the Power E1180, a high-end enterprise server that is built for the most demanding workloads. It delivers exceptional performance, scalability, and resiliency, making it ideal for large-scale enterprise resource planning (ERP) systems, core banking, and high-throughput analytics. With support for massive memory footprints, advanced RAS features, and robust virtualization capabilities, the Power E1180 is designed for continuous availability and workload consolidation.

Midrange: Power E1150

The Power E1150 serves as the midrange workhorse of the Power11 platform. It brings many of the high-end capabilities of the Power E1180, such as spare core technology, pervasive memory encryption, and quantum-safe security into a more compact and cost-effective form factor. The Power E1150 is ideal for organizations seeking a balance between performance, scalability, and operational efficiency, especially in hybrid cloud environments.

Entry-level scale-out: Power S1124 and Power S1122

For smaller enterprises or distributed environments, IBM offers the Power S1124, IBM Power L1124, Power S1122, and IBM Power L1122 servers. These entry-level, scale-out systems are optimized for cost-effective deployment of AIX, IBM i, and Linux workloads. The Power L1122 and Power L1124 are targeted at the Linux marketplace but support up to 25% of their activated cores for AIX or IBM i workloads. These systems are suitable for departmental applications, edge computing, and cloud-native workloads. Despite their compact size, they benefit from the same Power11 processor technology, helping ensure consistent performance and security across the portfolio.

Figure 1-4 shows the new Power11 processor-based server portfolio.



Figure 1-4 The Power11 portfolio

Architectural technology Improvements

The Power11 portfolio introduces innovations across both the product lineup and the underlying core technologies, all aimed at delivering enhanced performance and greater value. One of the most notable advancements is in the processor and memory architecture, where the Power11 platform delivers substantial improvements over the previous the Power10 platform. These enhancements form the foundation for the next level of scalability, efficiency, and intelligent workload optimization.

Increased core counts and spare cores

Improved thermal dissipation solutions provide optimal cooling, supporting higher core densities and sustained performance.

Increased processor frequencies

ISC) technology improves power delivery and stability at the silicon level, enabling higher performance under sustained workloads. Refined algorithms dynamically adjust processor frequency based on workload characteristics, maximizing performance while maintaining energy efficiency.

Increased memory frequencies

The Power11 architecture introduces higher memory frequencies and increased memory bandwidth, resulting in substantial improvements in overall system performance. By supporting faster memory speeds and optimizing the memory subsystem, the Power11 platform boosts memory throughput while reducing latency. These enhancements provide quicker data access and processing, which is especially critical for memory-intensive workloads such as real-time analytics, large-scale databases, and AI applications. The combination of higher bandwidth and reduced latency enables smoother performance, better responsiveness, and more efficient usage of compute resources across the entire system.

Table 1-3 provides a listing of the portfolio with a comparison to the Power10 offerings.

Table 1-3 Portfolio overview

Segment	Models	Footprint capacity	Power10 comparison
High-end (4S+)	Power E1180	<ul style="list-style-type: none"> ▶ 4 nodes (CPC) each 5U with 4 processor modules each ▶ Up to 64 TB DDR5-based memory 	<ul style="list-style-type: none"> ▶ 15-core moving to 16-core ▶ Same memory capacity ▶ Introducing Resource Groups to maximize system usage ▶ In-place upgrades with serial number preservation
Mid-Range (4S4U)	Power E1150	<ul style="list-style-type: none"> ▶ 2 - 4 processors in a 4U form factor ▶ Up to 30 cores per module ▶ Up to 16 TB DDR5-based memory 	<ul style="list-style-type: none"> ▶ Up to 25% more core capacity ▶ Same memory capacity ▶ Introducing Resource Groups to maximize system usage ▶ In-place upgrades with serial number preservation
Scale-out 2S4U	Power S1124 Power L1124	<ul style="list-style-type: none"> ▶ Up to 2 processors in a 4U form factor ▶ Up to 30 cores per module ▶ Up to 8 TB DDR5-based memory 	<ul style="list-style-type: none"> ▶ Up to 25% more core capacity and same memory capacity ▶ Introducing Resource Groups to maximize system usage

Segment	Models	Footprint capacity	Power10 comparison
Scale-out 2S2U	Power S1122 Power L1122	<ul style="list-style-type: none"> ▶ Integrating the entry Single-Chip Module (eSCM) modules into the Power S1122 family with 4 or 10 processors per eSCM ▶ Up to 30 cores per dual core module ▶ 2 processors in a 2U form factor ▶ Up to 4 TB DDR5-based memory 	<ul style="list-style-type: none"> ▶ Up to 50% more core capacity ▶ Same memory ▶ Introducing Resource Groups to maximize system usage

Investment protection

Power11 processor-based servers are designed with strong investment protection in mind, helping ensure that customers can evolve their infrastructure without sacrificing prior investments. A key example is the ability to perform same-serial-number upgrades from the IBM Power E1080 to the new Power E1180 and from the IBM Power E1050 to the Power E1150. This upgrade path enables organizations to retain their existing system identity, simplifying software licensing, asset tracking, and operational continuity.

As part of this upgrade process, customers can migrate existing memory modules and I/O adapters from their current systems, preserving valuable hardware investments. This compatibility reduces the cost and complexity of migrating to the latest technology while delivering the performance, scalability, and security benefits of the Power11 architecture.

Also, IBM supports Power Enterprise Pools (PEP), which include both Power10 scale-out servers and Power11 scale-out servers. This mixed-system support enables organizations to gradually migrate workloads from older to newer systems within the same resource pool. It enables dynamic sharing of processor and memory activations across both generations, providing flexibility in capacity planning and workload placement. This approach protects existing investments and enables a smooth, nondisruptive path to modernization.

Portfolio simplification

One of the most notable changes in the Power11 platform is the streamlining of the scale-out server lineup. IBM reduced the number of scale-out models, simplifying customer choices and reducing complexity in deployment and support. With this focused approach organizations can select the right system for their needs while benefiting from the full-stack integration and innovation that IBM Power is known for.

Figure 1-5 on page 13 illustrates how customers can migrate from their existing IBM Power9® or Power10 processor-based servers to the new Power11 platform.

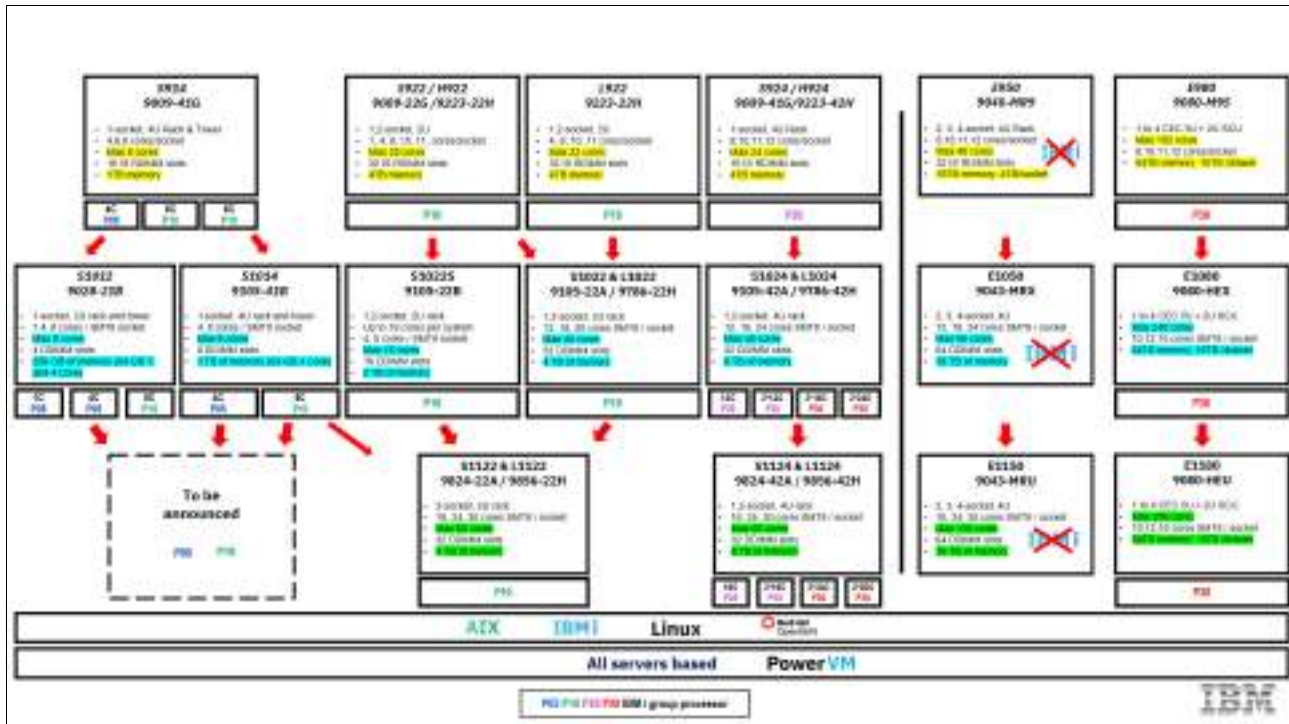


Figure 1-5 Migration path from the Power9 platform to the Power11 platform

1.4 The Power11 platform: Trusted, autonomous, and modern

The goal of a Power11 processor-based server is autonomous IT, where the server thinks, heals, and scales on its own. This approach supports business agility so that enterprises can respond quickly to market changes and customer demands, deploy services rapidly, and reduce time to market for products.

You can accomplish this goal by using the automation that is built into the platform and combined with AI-infused workflows. This automation spans the day-0 to day-2 IT lifecycle for Power11 processor-based servers and includes capabilities such as Infrastructure as Code (IaC) with Terraform, application configuration management with Red Hat Ansible, and enhancements to existing platform tools such as the HMC and IBM Cloud Management Console (IBM CMC). On the AI side, you can use AI-infused workflows with IBM watsonx and new tools such as IBM Concert®.

This section describes some of the technologies and solutions that make the Power11 platform a trusted, autonomous, and modern platform to support your business requirements.

1.4.1 Building a trusted infrastructure with Quantum-Safe Encryption

Infrastructure that is built by using Power servers benefits from robust security technologies that are integrated into both the hardware and software stacks. Power servers provide advanced security features at every level of the system to ensure comprehensive protection for sensitive data and applications. These features include advanced encryption technologies, Secure Boot capabilities, and integrated firmware updates. In addition, Power servers use IBM expertise in securing mission-critical workloads, making them a preferred choice for organizations that require a secure environment for their digital assets.

Workloads on Power11 processor-based servers benefit from improved cryptographic accelerator performance compared to previous generations. The Power11 processor supports accelerated cryptographic algorithms such as Advanced Encryption Standard (AES), SHA-2, and SHA-3, which results in higher per-core performance for these algorithms. This enhancement enables features such as AIX Logical Volume Encryption to operate with minimal impact on system performance.

The processor-core technology of the Power11 platform includes integrated security protections that provide the following benefits:

- ▶ Improved cryptographic performance
Integrated cryptographic support reduces the performance impact of encrypting and decrypting your data so that you can make encryption pervasive to protect all critical data.
- ▶ Increased application security
Hardened defenses against return-oriented programming attacks.
- ▶ Simplified hybrid cloud security
Hybrid cloud security administration through a single interface.
- ▶ Enhanced virtual machine (VM) isolation
Industry-leading VM isolation technology that defends against attacks that use OS or application vulnerabilities in the VM to access other VMs or the host system.

Encryption technologies and their applications

The Power11 platform emphasizes comprehensive security throughout its design and offers multiple encryption options. Key options include Transparent Memory Encryption (TME), Fully Homomorphic Encryption (FHE), and QSE.

Transparent Memory Encryption

TME encrypts data in memory to protect it from unauthorized access and tampering during run time. Operating at the hardware level, TME uses the Power11 processor cryptographic engines to perform encryption and decryption tasks efficiently. TME provides pervasive protection of data in memory with minimal impact on system performance because encryption and decryption are managed at the chip level. Its integration into normal operations is seamless and automatic.

Fully Homomorphic Encryption

FHE enables computations to be performed directly on encrypted data without decrypting it first. This capability helps ensure that sensitive data remains confidential even during processing. FHE operates at the software level and uses sophisticated mathematical algorithms to enable computations on ciphertexts.

Implementing FHE requires specialized libraries and frameworks. However, FHE is computationally intensive and can introduce a performance impact compared to conventional hardware-only encryption methods because of the complexity of the algorithms.

Quantum-Safe Encryption

QSE is designed to resist quantum attacks and secure data against the computational capabilities of future quantum computers that might break current cryptographic algorithms. QSE uses cryptographic algorithms that are believed to be resistant to quantum attacks, such as lattice-based, hash-based, and multivariate quadratic equation-based cryptography.

Many quantum-safe algorithms are still undergoing testing and standardization to help ensure that they provide robust security against future quantum advancements. QSE is typically used to secure long-term data, sensitive communications, and critical infrastructure.

Table 1-4 shows the relevant features and differences in these technologies.

Table 1-4 Key differences

Feature	TME	FHE	QSE
Encryption scope	Secures data in memory.	Enables computations on encrypted data.	Designed to resist future quantum computing threats.
Implementation level	Implemented in hardware.	Implemented by using a combination of hardware and software.	Implemented by using a combination of hardware and software.
Performance impact	Hardware that is accelerated through the Power11 cryptographic engines and designed to have minimal performance impact.	Involves substantial computational processing.	The impact of QSE varies. Some quantum-safe algorithms might introduce a performance impact; This impact is a subject of ongoing research.
Use cases	Used to protect data in memory.	Used for performing secure computations on sensitive data without decrypting it.	Provide long-term data protection and secure communications in the future.

Quantum safety compliance

QSE, also known as Post-Quantum Cryptography (PQC), refers to encryption methods that are secure against both classical and quantum computers. As quantum computers advance, they might pose a threat to existing cryptographic systems, potentially compromising their security. QSE is essential for protecting sensitive data, communication channels, and user identities in the age of quantum computing.

The urgency of adopting QSE stems from two primary concerns:

- ▶ Harvest-now, decrypt-later attacks.
Advanced quantum computers might enable adversaries to intercept and store encrypted communications today and decrypt them later, even before reaching Q-Day. (Q-Day is the anticipated point when quantum supremacy becomes widespread and many current encryption algorithms are no longer effective.)
- ▶ Complex migration
Migrating to QSE might require more than a decade because of the complexities of organizational structures and IT infrastructure.

Organizations should start evaluating and implementing QSE solutions immediately to help ensure continued protection and maintain stakeholder trust. Delaying QSE adoption might have severe consequences. Cryptographic systems that remain unchanged might be compromised if there is a successful quantum attack, exposing sensitive data and risking confidential business transactions and individual privacy. Financial institutions, critical infrastructure providers, and government agencies would face challenges in maintaining operational integrity and confidentiality. Therefore, prioritizing QSE implementation is crucial for long-term cybersecurity resilience.

The Power11 platform is designed to support these quantum-safe algorithms, helping ensure robust security as quantum computing advances. Power11 support for quantum-safe features provides the following benefits:

- ▶ Data encryption breakage protection
 - Risk
Quantum computers might break widely used cryptographic algorithms such as RSA, Elliptic Curve Cryptography (ECC), and traditional Diffie-Hellman key exchange protocols. Shor's algorithm, for example, could efficiently factor large integers and solve discrete logarithms, compromising the security of these algorithms.
 - Protection
The Power11 platform supports quantum-safe algorithms such as lattice-based, hash-based, code-based, and multivariate quadratic equation-based cryptography, which are believed to be resistant to quantum attacks. The crypto engines in the Power11 platform enhance the performance of these algorithms, helping ensure secure encryption and key exchange processes with minimal performance degradation.
- ▶ Secure communications
 - Risk
Quantum computers might intercept and decrypt secure communications, undermining protocols that currently rely on classical encryption methods.
 - Protection
The Power11 platform secures communication channels with quantum-resistant protocols, helping ensure data confidentiality even in the presence of quantum adversaries. End-to-end encryption is maintained throughout the data lifecycle, from storage to transmission, by using algorithms that are resistant to quantum attacks.

- ▶ Data integrity and authenticity
 - Risk

Quantum computers might forge digital signatures or tamper with data.
 - Protection

The Power11 platform supports quantum-safe digital signature algorithms such as eXtended Merkle Signature Scheme (XMSS) and Unbalanced Oil and Vinegar (UOV), providing strong security against quantum attacks.
- ▶ Long-term data protection
 - Risk

Sensitive data that is stored today might be harvested and decrypted in the future as quantum computers become more powerful, threatening long-term confidentiality.
 - Protection

Implementing QSE methods helps ensure that data remains secure over time, even as quantum computing capabilities evolve. The Power11 architecture supports updates to cryptographic libraries and protocols, enabling the adoption of new quantum-safe algorithms as they are developed and standardized.
- ▶ Physical and memory attack protection
 - Risk

Physical attacks on memory, such as cold start attacks, might expose sensitive data if not adequately protected.
 - Protection

Power11 TME helps ensure that data in memory is encrypted, protecting it from physical attacks during run time.

Quantum-safe algorithms that are supported by the Power11 platform

The following quantum-safe algorithms are supported by the Power11 platform:

- ▶ Lattice-based cryptography
 - Algorithms

ML-DSA (Dilithium), ML-KEM (Kyber), and NTRUEncrypt
 - Characteristics

Secure against quantum attacks. Based on lattice problems such as Learning With Errors (LWE) and Ring-Learning With Errors (Ring-LWE). Relatively efficient for hardware and software implementations.
- ▶ Hash-based cryptography
 - Algorithms

Merkle Signature Scheme (MSS), XMSS, and SPHINCS+
 - Characteristics

Secured through hash functions, though produces larger signatures and keys

- ▶ Code-based cryptography
 - Algorithms
 - McEliece Cryptosystem, Bit Flipping Key Encapsulation, and Hamming Quasi-Cyclic
 - Characteristics
 - Quantum-resistant based on decoding random linear codes, although the public keys can be large.
- ▶ Multivariate quadratic equations
 - Algorithms
 - UOV and Rainbow
 - Characteristics
 - Secure against solving systems of multivariate quadratic equations. Efficient in signature generation but might involve larger key sizes.

Power11 implementation

Power11 processors support quantum-safe algorithms in the following ways:

- ▶ Crypto engines
 - Multiple engines per core enable efficient running of cryptographic operations.
- ▶ Software updates
 - The architecture enables updates to cryptographic libraries, helping ensure the integration of new quantum-safe algorithms as they become standardized.

The Power11 platform design and capabilities help ensure robust security against future quantum threats by using hardware acceleration and flexible software updates, maintaining high security standards as the cryptographic landscape evolves.

Encryption enablement in hardware

There are two options for accelerating encryption in an Power11 processor-based server:

- ▶ Use the built-in encryption acceleration that is integrated into the Power11 processor.
- ▶ Use a Peripheral Component Interconnect Express (PCIe)-based encryption accelerator that IBM Power supports.

On-chip encryption support in the Power11 platform

The Power11 processor is designed to support future encryption technologies, including FHE and Quantum-Safe Cryptography to prepare for the quantum era. The Power11 processor Instruction Set Architecture (ISA) is tailored for these solutions' software libraries, which are available or will soon be available in the corresponding open source communities.

Workloads on the Power11 platform benefit from cryptographic algorithm acceleration, which enables higher per-core performance than Power10 processor-based servers for algorithms such as AES, SHA-2, and SHA-3. Features such as AIX Logical Volume Encryption can be activated with minimal performance impact because of this performance enhancement.

With four times as many AES encryption engines, Power11 processor technology is designed to deliver noticeably faster encryption performance. Power11 processor-based servers are more advanced than Power10 processor-based servers, with updates for the most stringent standards of today and future cryptographic standards, including post-quantum and FHE. The Power11 platform also introduces more improvements to container security. By using hardware features for a seamless user experience, TME simplifies encryption and supports end-to-end security without compromising performance.

1.4.2 Protecting, detecting, and recovering with IBM Power Cyber Vault

Regulatory frameworks such as the US Securities and Exchange Commission cybersecurity disclosure rules and the Digital Operational Resilience Act are no longer distant challenges but active requirements. More regulations are coming, each demanding higher levels of risk management and governance, mechanisms to promptly detect anomalous activities, and resilience practices to minimize the risk of data corruption or loss. The bar is rising, and the spotlight is now on every organization's security posture.

These frameworks go beyond checklists. They signal a shift: cybersecurity is no longer a compliance add-on but a foundational element of operational continuity, customer trust, and business strategy.

Traditional siloed security solutions

Most organizations still operate with fragmented security architectures that are collections of tools, processes, and teams that function in silos. Although some integration has occurred over time, loosely connecting security domains still leaves critical gaps in visibility, context, and response coordination.

A siloed approach results in the following issues:

- ▶ Limited understanding of threats across domains (for example, linking identity misuse with network anomalies).
- ▶ Reactive security postures that are measured by containment speed rather than business outcomes.
- ▶ Barriers to collaboration and innovation as teams struggle to securely share and access data.
- ▶ Missed opportunities to embed privacy and security into customer-facing products and services.

Even more concerning, a lack of contextual awareness across tools and data streams can open the door to full-scale compromises, such as ransomware attacks that paralyze operations.

Security was never meant to be siloed. It became that way out of necessity and the involved architectures. However, that model is no longer sufficient. Forward-looking enterprises are rethinking their cybersecurity strategy as a defense mechanism and a business enabler. They are integrating tools, processes, and insights across domains, from access and identity to threat detection, data governance, and compliance. They are moving from isolated defenses to connected intelligence.

Zero trust: a strategic shift

Many organizations are adopting zero trust frameworks as the foundation for their next-generation security programs. Built on the principles of least privilege access, continuous verification, and assumed breaches, zero trust offers a compelling blueprint.

However, the blueprint alone is not enough. Applying zero trust effectively requires the following points:

- ▶ A deep understanding of your business risk surface
- ▶ Alignment with regulatory and operational requirements
- ▶ Integration with your existing cloud and hybrid infrastructure
- ▶ Cultural and organizational readiness to shift to continuous governance and access control

In essence, zero trust is not a product. It is an operating model for security that aligns with the way modern businesses function, that is, dynamically, digitally, and in a distributed manner.

Today's threats: evolving, disruptive, and costly

Cyberattacks today are fast-moving, complex, and deeply disruptive. A typical incident flow includes the following steps:

- ▶ Service disruption is detected.
- ▶ Immediate shutdowns are initiated to contain the spread.
- ▶ A search for uncompromised backups begins.
- ▶ Slow and costly recovery efforts start.

According to recent reports, 78% of organizations take more than 100 days to fully recover from a major cyber incident. which includes downtime, lost trust, revenue, and opportunity³.

IBM provides a unified cyber resilient solution that includes Power11 processor-based servers, IBM FlashSystem storage, advanced software, and services from IBM Technology Expert Labs. This solution helps organizations respond to evolving cyberthreats and regulatory standards with IBM Power Cyber Vault.

IBM Power Cyber Vault

IBM Power Cyber Vault keeps your business running by combining the following features:

- ▶ Power Cyber Vault protected data
Automated immutable images are stored in the IBM Power Cyber Vault to protect client snapshots and backups.
- ▶ Continuous threat detection
Near real-time file system monitoring, IBM FlashSystem inline detection, and recovery data scanning. Threat detection time can be as low as 60 seconds.
- ▶ Rapid clean room response
Automated response to attacks creates IBM Power Cyber Vault clean rooms for testing and validation, enabling accelerated recovery.
- ▶ Cyber Vault Assisted Recovery
Automated multi-level testing and validation of IBM Power Cyber Vault images to accelerate system recovery. Organizations can return to full operations in hours instead of months.

³ Source:

<https://www.cohesity.com/content/dam/cohesity/resource-assets/research-reports/cyber-resilience-global-survey-report-en.pdf>

IBM Power Cyber Vault solution components

IBM Power Cyber Vault is a complete cyber resilient solution that can be deployed in as little as 30 days. It includes the following components:

- ▶ Power11 processor-based servers
 - Industry-leading threat resistance
 - PowerSC for security and compliance tracking
 - Built-in Secure Boot and memory encryption
 - Scalable, flexible architecture from entry-level to enterprise systems
- ▶ IBM FlashSystem Storage
 - Immutable copies for data protection
 - Ransomware detection for real-time threat identification
 - Flash-speed recovery for rapid restoration of snapshot data
 - Storage Insights Pro for comprehensive reporting
- ▶ IBM Technology Expert Labs – IBM Power Cyber Vault Deployment
 - A design workshop that is tailored to client-specific requirements
 - IBM Power Cyber Vault deployment to customize and enable the unified solution
- ▶ IBM Expertise Connect
 - Primary technical point of contact for IBM Power Cyber Vault
 - Dedicated to client success
 - Available postinstallation to guide client strategy and direction

IBM Power Cyber Vault implementation

The implementation of the solution consists of several phases, each designed to support risk mitigation and enable the fastest possible recovery:

- ▶ Identify
 - Assess risks and design mitigation solutions
 - IBM Security® and Resilience Assessment.
 - Provides an overview of the organization, processes, and strategies for operational and cyber resilience.
 - Delivers a prioritized list of recommendations to improve overall resilience.
 - IBM Cyber Vault Solution and Design Workshop enables the creation of a detailed solution based on client cyber resiliency goals, infrastructure, and required services.
- ▶ Protect
 - Safeguarded Cyber Vault copies and backups
 - Provides client-specific scheduling for creating and protecting immutable snapshots and backups to meet Recovery Point Objective (RPO) and Recovery Time Objective (RTO) requirements.
 - Snapshots. IBM FlashSystem and Storage Defender CSM save protected system snapshots into the Cyber Vault on a regular schedule.
 - Orchestration for crash consistency.
 - Application-specific customization is available.
- ▶ Detect
 - Real-time threat detection
 - Monitors real-time threats and cyberattack vectors to enable rapid response by using PowerSC Real-Time Compliance orchestration.

- Zero Trust Execution detects, prevents, and reports unapproved executable files and applications on Power11 processor-based servers.
 - IBM FlashSystem integrated ransomware detection alerts on active attacks and reports to PowerSC through IBM Storage Insights Pro.
 - Security Information and Event Management (SIEM) integration accepts alerts and threat reports from SIEM systems.
- ▶ Respond and recover
- Real-time threat response
- When threats are detected, the response phase initiates auditing and recovery.
- On receiving an integrity alert, IBM Power Cyber Vault automatically creates an immutable backup of the VM within seconds and spins up a copy in the clean room within minutes.
 - Runs multiple integrity checks in the clean room, including platform-specific tests.
 - Identifies corrupted copies and finds the newest clean copies.
 - After administrator approval, recovery to production meets RPO and RTO requirements.

1.4.3 Quantum-safe compliance with one-click inventory discovery

IBM PowerSC is a comprehensive security and compliance solution that is designed for virtualized environments running on Power servers with AIX, Linux, or IBM i OSs. It integrates deeply with the IBM Power platform to provide end-to-end protection through features such as security and compliance automation, trusted boot, trusted firewall, and trusted logging.

PowerSC helps organizations meet stringent regulatory standards such as PCI DSS, HIPAA, and SOX by automating the monitoring, auditing, and enforcement of security policies. It also includes advanced capabilities such as multi-factor authentication, intrusion detection, and fix management. With tools like PowerSC Trusted Surveyor, it helps ensure consistent network configuration and compliance across dynamic virtual environments, reducing administrative impact while maintaining strong security postures.

IBM is actively involved in developing and adopting quantum-safe cryptographic standards, including collaboration with NIST on its PQC Standardization project. PowerSC supports a seamless transition to NIST's forthcoming standards and offers capabilities to help verify and manage QSE.

The Quantum Safety Analysis feature evaluates the quantum-safe status of your VMs to identify areas of concern and assist with compliance with emerging quantum-safe cryptographic standards:

- ▶ Endpoint scanning
 - PowerSC scans endpoints on your system to identify cryptographic artifacts and vulnerabilities that are related to quantum safety. Scans can also be scheduled.
- ▶ Encryption inventory
 - Creates an inventory of where different encryption algorithms are implemented on your system.

- ▶ Cipher strength assessment
Estimates the strength of ciphers, certificates, and keys on your AIX endpoints and categorizes them as weak, strong, quantum-safe, or unclassified. PowerSC is a security and compliance management tool, not a cryptographic module validation tool.
- ▶ Quantum Safety Analysis report
Displays discovered cryptographic elements and their assessed strengths in the PowerSC GUI.

1.4.4 Zero planned downtime and end-to-end orchestration for maintenance events

Planned downtime remains one of the biggest challenges in ITOps, especially for a critical infrastructure such as the Power11 platform. IBM addresses this challenge by delivering *Autonomous IT* through extreme automation and AI-infused workflows that are embedded directly into the platform. This approach reduces operational effort, minimizes disruptions, and improves resilience.

Today's enterprise IT environments face increasing pressure to help ensure infrastructure reliability and security without disrupting operations. Key challenges include the following items:

- ▶ Operational complexity
Maintaining firmware, I/O adapters, and VIOSs is highly specialized and intricate.
- ▶ Service interruptions
Routine maintenance often requires downtime, which impacts application availability and business continuity.
- ▶ Security risks
Delayed updates expose systems to vulnerabilities, such as Common Vulnerabilities and Exposures (CVEs), which weaken compliance and increase risk.
- ▶ Resource drain
A significant amount of time is spent on maintenance tasks, such as fix evaluation, compatibility testing, and scheduling.
- ▶ High downtime costs
With downtime averaging \$336,000 per hour, annual maintenance-related outages can exceed \$2 million per system⁴.

These issues force organizations to make a difficult tradeoff: stay current and risk downtime, or delay updates and increase exposure. This approach is unsustainable in today's always-on digital economy.

⁴ Source:

<https://itic-corp.com/itic-2024-hourly-cost-of-downtime-part-2/#:~:text=ITIC's%2011th%20annual%20Hourly%20Cost,costs%20their%20company%20over%20%24100%2C000>

Defining infrastructure

In this context, *infrastructure* refers to the Power11 platform, which includes firmware, I/O components, VIOS, and the OS. To maintain performance and security, clients must complete the following tasks:

- ▶ Apply fixes and updates regularly.
- ▶ Remain on supported versions.
- ▶ Respond promptly to security advisories.
- ▶ Conduct preventive maintenance.

However, these actions typically require planned downtime. The core dilemma is whether businesses should accept regular service interruptions to stay current or defer updates and risk security and performance.

The Power11 platform: A platform for autonomous IT

The Power11 platform is designed to resolve this dilemma by enabling continuous business operations while keeping infrastructure secure, up to date, and high performing. The Power11 platform introduces a transformative approach to infrastructure lifecycle management that is purpose-built for Autonomous IT. The platform operates with minimal human intervention by using embedded intelligence to monitor, diagnose, and remediate issues in real time. This capability enables an infrastructure that is inherently resilient, adaptive, and capable of maintaining continuous service availability.

At the core of the Power11 platform is the *Automated Maintenance Framework*, which is orchestrated by IBM Concert, an AI-driven engine that automates the full maintenance lifecycle. This framework includes the following capabilities:

- ▶ Unified update orchestration for system firmware, VIOSs, and I/O adapters
- ▶ LPM to provide uninterrupted application availability during updates
- ▶ Prevalidation checks to confirm system readiness and compatibility
- ▶ Flexible update sourcing from IBM repositories, SFTP, NFS, USB, or HMC local file systems

Maintenance operations can run autonomously or manually through the HMC, providing both automation and administrative control. This architecture transforms maintenance from a disruptive, manual process into a seamless background operation.

IBM Concert: AI-powered application management for the modern enterprise

IBM Concert addresses the persistent challenge of siloed application data by using AI to deliver intelligent, prioritized recommendations. By streamlining issue identification and resolution, Concert reduces the mean time to resolution (MTTR) for critical risk factors such as the following ones:

- ▶ CVEs
- ▶ Certificate management issues
- ▶ Application compliance challenges

What sets IBM Concert apart

Built on IBM watsonx, IBM Concert uses advanced generative AI to analyze complex environments and deliver actionable insights:

- ▶ Data-neutral integration
IBM Concert supports a wide range of data sources across networks, infrastructure, and application architectures.
- ▶ Hybrid cloud by design
IBM Concert seamlessly supports hybrid cloud environments to meet the requirements of modern enterprise IT.
- ▶ AI-infused orchestration for proactive maintenance
IBM Concert orchestrates maintenance operations through intelligent workflows and end-to-end automation, including the following capabilities:
 - Inventory and risk discovery
Automatically identifies assets and uncovers hidden risks across the environment.
 - Risk analysis and remediation planning
Delivers prioritized, actionable plans to address vulnerabilities and compliance gaps.
 - Automated risk resolution
Enables automated actions to remediate issues, reducing manual effort and downtime.

Automation can run directly from the IBM Concert solution or from the HMC, providing flexibility and administrative control.

Key benefits of Power11 autonomous maintenance

The Power11 integrated solution for ZPD delivers measurable benefits across operational, security, and financial dimensions:

- ▶ Business continuity
Updates apply without interrupting workloads, preserving uptime and service-level commitments.
- ▶ Operational efficiency
Automation streamlines the maintenance lifecycle, reducing manual effort and administrative impact.
- ▶ Risk reduction
AI-driven orchestration provides timely updates, minimizing exposure to vulnerabilities and compliance risks.
- ▶ Productivity gains
IT staff benefit from a simplified, intelligent interface that reduces time on task from weeks to minutes.
- ▶ Cost optimization
Eliminating downtime and manual labor reduces the total cost of infrastructure operations.

Together, these capabilities position the Power11 platform as a foundational platform for secure, scalable, and self-managing enterprise IT infrastructure that aligns with the demands of modern digital transformation.

The new maintenance model

From the IT administrator's perspective, maintenance is now part of a unified, intelligent flow that includes the following steps:

- ▶ Inventory discovery and risk detection
- ▶ Smart planning
- ▶ Automated download and application of updates
- ▶ Validation
- ▶ VM (LPAR) evacuation and reintegration

This approach minimizes disruptions, helps ensure service-level agreement (SLA) compliance, and enables faster, more timely maintenance. Systems remain secure, stable, and aligned with evolving requirements without downtime or performance impact.

The Power11 built-in Automated Maintenance Tool enables platform maintenance without affecting applications. This automation can run through IBM Concert orchestration or be initiated manually from the HMC, independent of IBM Concert.

IBM Concert includes the following capabilities:

- ▶ Ability to update system firmware, VIOSs, and I/O adapter microcode from a single update flow, supporting both concurrent and disruptive updates
- ▶ Validation for LPM and VIOS redundancy (VIOS maintenance readiness check)
- ▶ Ability to automatically migrate partitions and return them as part of the update process:
 - Option to return to the source system or remain on the target system
 - Option to evacuate all LPARs or select a subset of LPARs and define the order
 - Option to choose the order of updates
 - Option to download only or download and update
- ▶ Support for multiple update sources:
 - IBM website (preferred for end-to-end automation)
 - SFTP server
 - HMC file system
 - NFS server
 - USB

The new tool appears as a Licensed Capability of the server, as shown in Figure 1-6 on page 27.

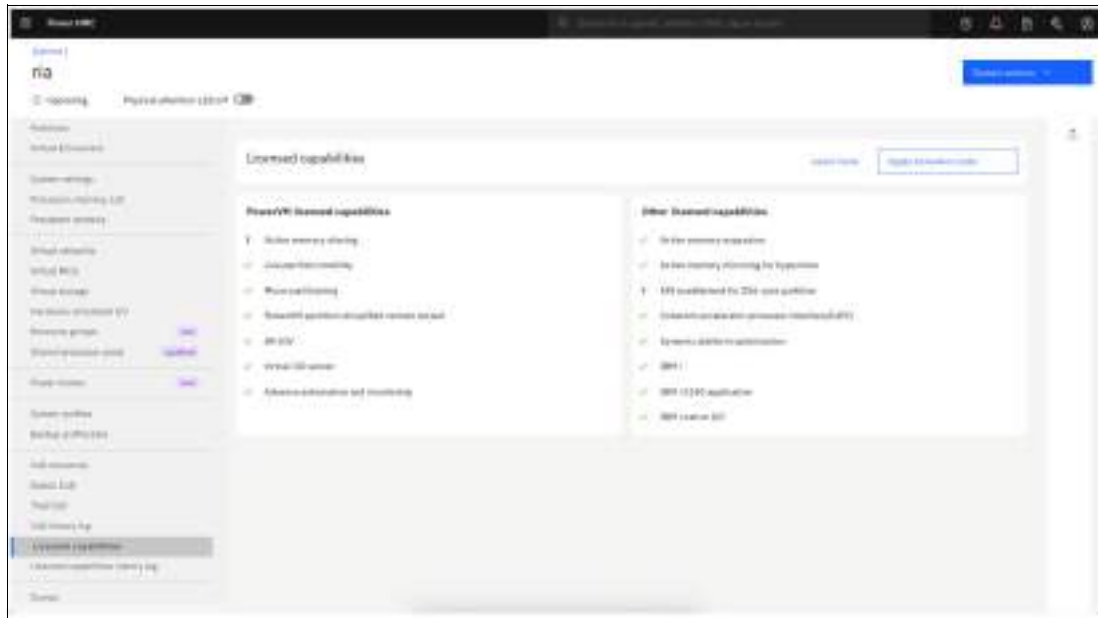


Figure 1-6 Power server Licensed Capabilities

For more information about using the Automated Maintenance Tool, see 5.1.7, “Using the automated maintenance tool” on page 129.

1.4.5 Automated data collection for faster error resolution

Enterprises are increasingly adopting hybrid cloud environments, with 82% reporting usage as of 2024 (IBM Cloud Survey). This shift introduces operational complexity, particularly in managing infrastructure and resolving issues. Manual error resolution continues to overwhelm IT teams, with 60% of organizations citing insufficient staffing for incident management. Moreover, the financial impact of downtime is escalating with over 90% of midsize and large enterprises reporting that a single hour of downtime now exceeds \$300,000 in cost⁵.

To address these challenges, IBM introduced enhanced support capabilities in the Power11 HMC. These capabilities include a streamlined case creation process and automated log collection through Call Home. You can initiate support cases directly from the HMC, and the system automatically gathers and transmits relevant First Failure Data Capture (FFDC) logs to IBM Support.

Here are some of the supported log types:

- ▶ System firmware (dumps)
- ▶ HMC (pedbg)
- ▶ VIOS (snap)
- ▶ Cloud connector logs
- ▶ Partition migration logs

Call Home configuration is required to create a case in CSP and upload logs. If Call Home is not configured, the HMC creates a problem record, but it does not create a case. Logs remain stored in the HMC and are not uploaded automatically.

⁵ Source: <https://uptimeinstitute.com/2019-data-center-industry-survey-results>

Note: As a best practice, enable Call Home because it is also required for Autonomous Error Resolution (AER) data collection.

If the HMC is a Virtual HMC and the problem is an HMC issue, the system collects logs when the issue occurs. However, it does not initiate a Call Home, and the logs are not sent automatically.

Use case: automated support workflow

Consider Louis, a system administrator at a large insurance company that runs critical back-end applications on IBM Power infrastructure. During a routine maintenance task involving LPM, Louis encounters a migration validation error. Previously, resolving such issues required manually opening a support case, coordinating with IBM Support, and collecting FFDC data. This process might take 5 - 8 hours before investigation begins.

With the Power11 platform, you can now open a support case directly from the HMC. The system automatically collects and submits the necessary diagnostic data, eliminating manual steps and reducing turnaround time. IBM Support uses AI-enhanced diagnostic tools that are trained on IBM Power servers to accelerate root cause analysis and resolution.

Benefits of the enhanced support model

Implementing this new automated data collection capability with applied AI workflows in IBM Support provides the following benefits:

- ▶ **Reduced time to resolution**
Automated log collection and AI-assisted diagnostics accelerate issue resolution.
- ▶ **Improved productivity**
System administrators save hours by avoiding manual data gathering and coordination.
- ▶ **Reduced operational risk**
Faster resolution minimizes downtime and its associated business impact.
- ▶ **Enhanced support experience**
A simplified, integrated workflow improves responsiveness and reduces stress for IT teams.

1.4.6 Increasing efficiency and reducing energy consumption through smart scheduling options

In recent years, sustainability became a central strategic priority in the IT industry due to regulatory pressure and business value. Over the past 5 years, climate commitments among companies increased by 800%, reflecting broader awareness and the recognition that sustainable practices can enhance profitability. At the time of writing, 83% of organizations are investing in research and development for low-carbon products and services, and products with sustainability attributes are achieving revenue growth rates more than 25% higher than traditional products⁶.

⁶ Source:

<https://www.pwc.com/us/en/services/esg/library/assets/pwc-sustainability-decarbonization-2025.pdf>

Although 95% of companies established operational sustainability goals, only 41% achieved measurable progress⁷. Many organizations face challenges in starting the implementation process. This situation highlights the need for explicit guidance, stronger governance frameworks, and better integration of sustainability metrics into core business operations.

IBM Power servers consistently improve energy efficiency per watt with each generation, and the Power11 platform continues this trend. The Power11 platform enhances efficiency, reducing energy use, carbon emissions, and the data center energy footprint.

Clients upgrading from the Power9 platform to the Power11 platform can reduce energy consumption by up to 60% for the same performance. The Power11 platform outperforms x86 systems, delivering twice the performance per watt. Compared to the Power10 platform, the Power11 platform offers up to 33% better performance per watt, with the Power S1122 model showing the largest gain⁸.

Figure 1-7 illustrates natural IT efficiency improvements from the Power9 platform to the Power11 platform, offering up to a 3:1 consolidation ratio, reducing energy use and carbon emissions by approximately 60%, and cutting the data center footprint by two-thirds.

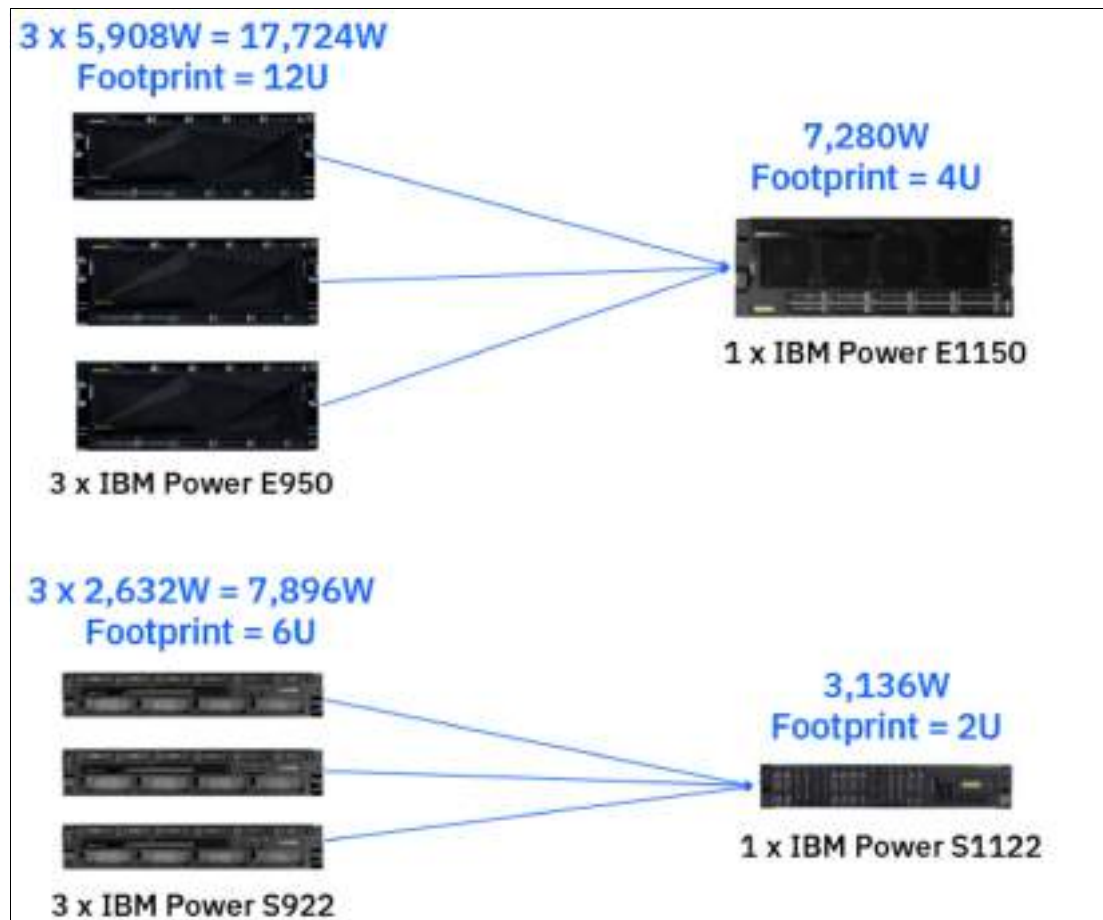


Figure 1-7 Power9 platform to Power11 platform energy consumption reduction

In addition to the base energy savings that are delivered by the Power11 platform, new features were introduced to address client feedback and further enhance operational efficiency.

⁷ Ibid

⁸ Source: <https://techchannel.com/power11/server-models/>

Partition level monitoring

At the time of writing, the HMC provides monitoring and reporting capabilities to track energy usage, carbon emissions, and other environmental factors for existing Power servers. The Power11 platform extends these functions by delivering insights at the partition level, as shown in Figure 1-8.



Figure 1-8 Partition level monitoring

Energy efficient mode

The Power11 platform introduces a new Energy efficient mode that, when enabled, can reduce energy consumption by up to 30%, with only an approximate 10% reduction in performance⁹. The impact varies depending on system configuration. You can dynamically configure the power mode through the HMC GUI, as illustrated in Figure 1-9.

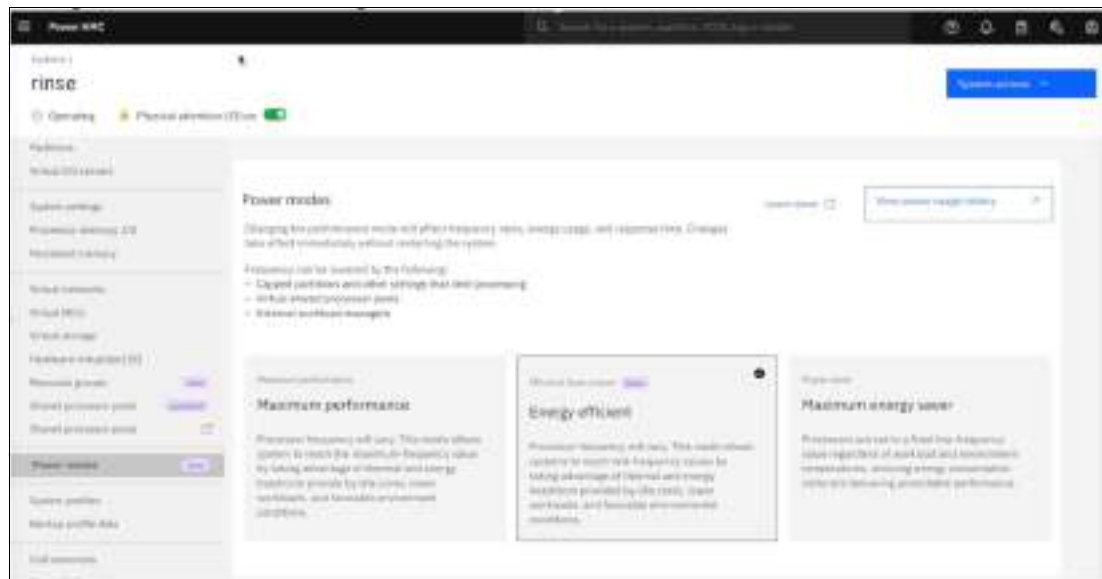


Figure 1-9 Power modes

The environmental dashboard shows that in maximum performance mode, energy consumption remains almost constant even when core usage decreases.

⁹ Based on IBM measurements on servers comparing Maximum Performance Mode to Energy efficient mode while running compute-based, disk-based, and memory-based workloads.

Figure 1-10 illustrates energy usage when the processor operates in maximum performance mode.

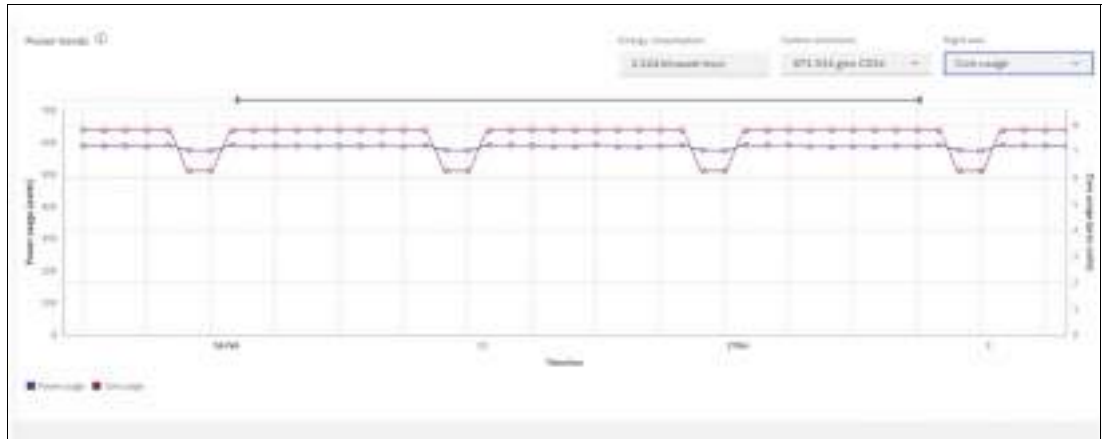


Figure 1-10 Energy consumption in max performance mode

Switching to Energy efficient mode, which is applied instantly, results in only a 10% reduction in performance while decreasing energy consumption by up to 30%, as shown in Figure 1-11.



Figure 1-11 Energy efficient mode

Power mode scheduling

If system usage follows a predictable pattern, you can use the new scheduling power mode feature to optimize energy efficiency. This feature enables the system to activate Energy efficient mode during designated time periods.

As illustrated in Figure 1-12, Energy efficient mode is engaged over the weekend, starting at 8:00 PM on Friday and ending at 5:00 AM on Monday.

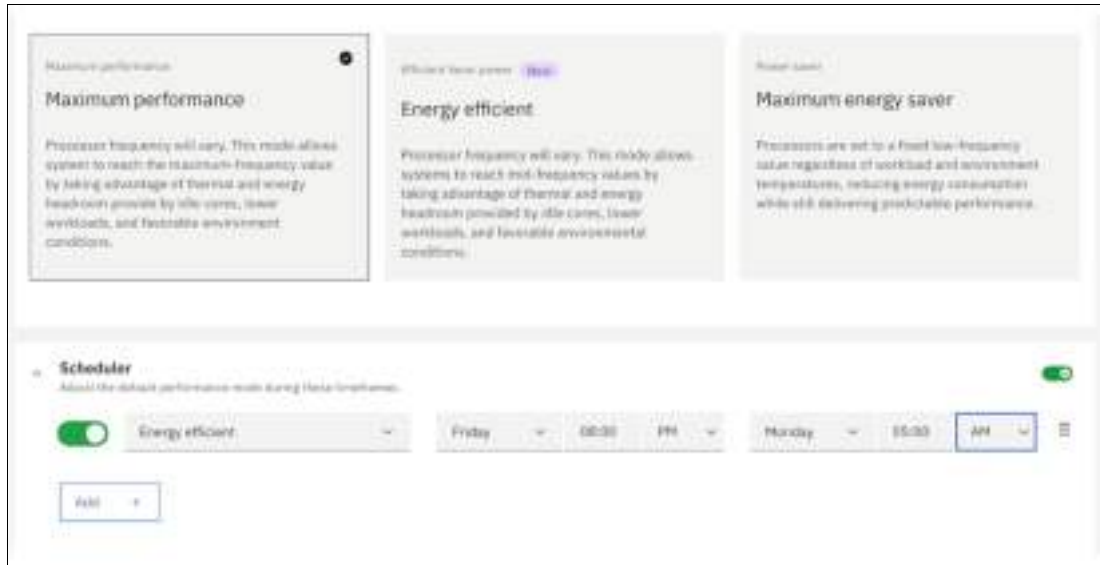


Figure 1-12 Power mode scheduling

The new environmental dashboard enables energy savings whenever opportunities arise. Automated scheduling capabilities allow the system to migrate between power modes efficiently.

1.4.7 Consolidating Oracle 19c workloads

At the time of writing, Oracle Database 19c is the latest long-term release, and Version 19.12 is the most recent Release Update for the AIX OS. Oracle Database 19c introduces numerous enhancements over previous versions. Oracle Database versions for IBM AIX are compatible with later AIX Technology Levels (TLs) within the same major version. For certification details, see the “Certification” section at the [My Oracle Support portal](#).

Starting with Oracle Database 11g Release 2 (11.2.0.4), Oracle Database Instant Client is supported on Linux on Power (32-bit and 64-bit). This support continues with Oracle Database 19c. Beginning with Oracle Database 12c Release 1, Little-Endian support was introduced for the Instant Client on Linux on Power. As of Oracle Database 12c Release 2, only Little-Endian is supported.

Oracle Database Instant Client is a lightweight version of the full Oracle Database Client binary files. It is available on Linux on Power, enabling applications that are deployed on Linux on Power to connect to Oracle Database instances regardless of the platform on which the database is deployed (AIX, IBM Power, or other platforms).

Although systems based on IBM Power8® processor technology reached end of support, many clients have not completed a TR. IBM demonstrated that migrating Oracle Database workloads from the Power8 platform to the Power11 platform can reduce processor core requirements by up to 50%¹⁰. This reduction directly reduces the number of required Oracle licenses, decreasing total cost of ownership (TCO).

¹⁰ Source: Based on performance results published <https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed>

The Power11 processor delivers substantial performance improvements, enabling up to a 25% increase in Oracle Database 19c workload consolidation compared to the Power10 platform¹¹. These gains enable organizations to optimize infrastructure efficiency, reduce software licensing costs through improved core usage, and enhance overall system throughput while maintaining enterprise-grade RAS features.

Oracle is committed to release their latest products on IBM Power. For more information, see the [Oracle statement](#) about support for 23ai on IBM Power.

1.4.8 Day 1 availability of the Power11 platform in IBM Power Virtual Server with hybrid billing

As hybrid cloud adoption reaches 82% among enterprises and the number of applications is projected to exceed 1 billion by 2028¹², organizations face increasing pressure to manage complexity, address skill shortages, and control infrastructure costs. The Power11 platform, combined with PowerVS, offers a secure, scalable, and high-performance hybrid cloud platform for mission-critical workloads.

A key differentiator is the day-one availability of the Power11 platform in both public and private PowerVS environments. This immediate access enables you to deploy virtual servers for development, testing, production, and disaster recovery (DR) without delay. With provisioning times under 10 minutes, you can accelerate time-to-value and reduce migration risk through a consistent architecture and user experience across on-premises and cloud environments.

Using the Power11 platform in PowerVS also unlocks advanced enterprise-grade security features, which include quantum-safe firmware signing, encrypted LPM, and a reduced number of hypervisor vulnerabilities (CVEs) compared to x86 alternatives. These capabilities are essential for organizations operating in regulated industries or managing sensitive data.

In addition, PowerVS integrates natively with IBM watsonx services through IBM Satellite® Connector, enabling you to seamlessly infuse AI into your workloads and accelerate digital transformation initiatives.

PowerVS now offers three generations of Power servers, providing flexibility for clients with diverse requirements:

- ▶ The Power9 platform
 - Soft landing for clients moving to cloud, with substantial support life remaining (through October 2028).
 - Lowest cost option with special pricing on existing capacity.
 - Clients can continue to provision and use the Power9 platform.
- ▶ The Power10 platform
 - Ideal for business-critical applications such as SAP RISE.
 - Broad capacity availability across data centers.
 - Excellent price and performance for AIX and IBM i workloads.
 - Clients can continue to provision and use the Power10 platform.

¹¹ The actual performance improvement depends on the specific configuration of the Power11 processor-based server. Not all configurations achieve a 25% improvement. For more information, see <https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed>.

¹² Source: <https://gitnux.org/hybrid-cloud-statistics/>

- ▶ The Power11 platform
 - Immediate access to the Power11 platform without waiting for on-premises delivery.
 - Provision in less than 10 minutes.
 - Suitable for development, testing, and application qualification.
 - The first cloud environment to provide the Power11 platform in both IBM Cloud and on-premises (Private).

From a business value perspective, PowerVS delivers up to 25% faster SAP workload migration compared to x86 platforms¹³, which is driven by architectural consistency and superior compute performance. Flexible consumption models, shared processor pools (SPPs), and tiered storage options help reduce TCO. Automated DR, secure backup, and nondisruptive maintenance further enhance operational efficiency. The platform also supports regulatory compliance with standards such as SOC 2, HIPAA, and PCI-DSS.

Common use cases include:

- ▶ Application modernization through access to more than 250 IBM Cloud services
- ▶ Business resiliency with warm DR environments
- ▶ Data center optimization by consolidating or exiting physical infrastructure
- ▶ SAP transformation initiatives with unmatched availability and scalability for large-scale HANA deployments

PowerVS is available in 22 IBM Cloud data centers worldwide for public cloud use and can also be deployed in dedicated private cloud environments at client sites to meet stringent regulatory or operational requirements. The Power11 platform will initially be available in four global data centers and will be fully enabled for automated DR operations. This broad availability enables organizations to adopt the Power11 platform in alignment with their business goals and compliance mandates.

1.4.9 Off-chip AI acceleration with the IBM Spyre Accelerator

The IBM Spyre Accelerator is a low-power, high-efficiency hardware module that is designed to complement Power11 processor-based servers for AI inference, analytics, and memory-intensive workloads. This feature is planned for support on the Power E1180 in the fourth quarter of 2025.

The card uses a system-on-chip architecture that is optimized for enterprise AI workloads. It includes 32 low-power AI cores and supports multi-precision operations for inference and training, including FP16/8 and INT8/4. The card is enabled for foundation models and is planned to support the Red Hat software stack (Open Compute Project Open Systems for AI). It will also support AI frameworks and libraries such as PyTorch and vLLM.

Using 5 nm technology, the IBM Spyre Accelerator delivers high performance with low power consumption, requiring only 75 W. These characteristics make it ideal for organizations seeking to accelerate AI workloads while maintaining energy efficiency.

Figure 1-13 on page 35 shows the IBM Spyre Accelerator.

¹³ Source: <https://www.ibm.com/new/product-blog/rise-with-sap-migration-ibm-power-virtual-server>



Figure 1-13 IBM Spyre Accelerator

Each IBM Spyre Accelerator operates at only 75 W, reducing the power footprint compared to traditional GPU-based accelerators, which often exceed 300 W per device. The IBM Spyre Accelerator is planned for support only inside the accelerator expansion drawer. Each expansion drawer supports up to eight IBM Spyre Accelerator cards.

Expansion drawer support on the Power E1180:

- ▶ At initial availability, a maximum of one accelerator expansion drawer is supported on a Power11 processor-based server.
- ▶ Support for more expansion drawers is planned for 2026.

Note: The IBM Spyre Accelerator requires an updated fan-out card in the expansion drawer.

Unlike GPUs, which typically rely on high-bandwidth memory that is integrated on the card, the IBM Spyre Accelerator is designed to access and use the system's main DDR5 memory modules. This design enables the use of terabytes of high-performance memory that is available within the Power11 processor-based server. It provides two primary advantages:

- ▶ Reduced latency data access through a shared memory model
- ▶ Elimination of memory duplication, avoiding the need to copy data between host and device memory, as is common with discrete GPUs

Using system RAM simplifies software development and memory management while contributing to a more energy-efficient and thermally balanced system architecture, which is an increasingly important factor in modern data centers and edge environments. This design makes the IBM Spyre Accelerator suitable for workloads that require frequent access to large memory spaces with minimal energy cost, such as real-time analytics, AI inferencing, and high-throughput data preprocessing.

Here are the minimum supported software stack levels for general availability (4Q 2025):

- ▶ FW1110.10
- ▶ HMC1111
- ▶ RHEL 9.6
- ▶ Spyre software stack container
- ▶ Red Hat OpenShift AI (technology preview in 4Q 2025, GA in 1Q 2026)

For more information about use cases and solutions with IBM Spyre Accelerator, see Chapter 4, “Artificial intelligence support” on page 113.

1.4.10 Red Hat OpenShift with support for an expanded software ecosystem

Red Hat OpenShift on IBM Power provides a robust, enterprise-grade container platform that enables you to modernize applications, deploy AI workloads, and integrate with a broad software ecosystem. Running Red Hat OpenShift on Power11 processor-based servers enables organizations to use the platform’s performance, scalability, and security while maintaining architectural consistency across hybrid cloud environments.

A key advantage of Red Hat OpenShift on Power is its support for an expanded software ecosystem, which includes IBM Cloud Pak® offerings, such as IBM Cloud Pak for Data and IBM Cloud Pak for Integration; Red Hat Runtimes like Quarkus and Spring Boot; and a wide range of open-source AI and ML frameworks, including PyTorch, TensorFlow, and ONNX. Many independent software vendor (ISV) applications are now certified for the Power Architecture, enabling financial institutions and other enterprises to run containerized workloads alongside traditional AIX or IBM i systems. This colocation reduces latency, improves data gravity, and simplifies integration.

Red Hat OpenShift on IBM Power also supports seamless hybrid cloud integration, particularly when deployed on IBM PowerVS. This integration provides access to IBM Cloud services such as IBM Cloud Object Storage, IBM Key Protect, IBM Event Streams (Kafka), and IBM API Connect®. These capabilities are essential for hybrid architectures where AI models may be trained on-premises and deployed in the cloud or vice versa, depending on data locality and compliance requirements.

For developers and IT architects, Red Hat OpenShift on Power offers a rich set of tools and resources. The Red Hat OpenShift Container Catalog provides pre-built images and Helm charts that are optimized for Power, while IaC templates that use IBM Terraform® simplify cluster provisioning on PowerVS. IBM also supports the ecosystem with community channels, technical workshops, and IBM Garage® engagements to accelerate adoption.

By combining the strengths of IBM Power hardware with the flexibility of Red Hat OpenShift and the breadth of the container ecosystem, organizations can build and scale AI-driven applications with confidence, agility, and enterprise-grade resilience.

1.5 Introducing the Power11 processor

The Power11 processor is designed to deliver higher clock speeds and up to 25% more cores per processor chip compared to IBM Power10 processor-based servers. The Power11 platform builds on the key capabilities that were introduced with the Power10 platform, including stronger RAS features, improved energy efficiency, enhanced energy management, and advanced quantum-safe security.

Beyond the processor itself, here are more improvements to packaging, the memory architecture, and AI acceleration:

► Packaging innovation

The Power11 processor uses new ISC technology and advanced 2.5D packaging, which are combined with cooling innovations such as improved heat sinks and more efficient fans. These enhancements optimize energy delivery, strengthen thread and core performance, and increase overall system capacity.

► Enhanced system architecture

Power11 processor-based servers feature a robust memory architecture that is based on DDR5 DDIMMs and enhanced OMI technology. These improvements deliver higher memory reliability, capacity, and bandwidth. Because OMI is technology-neutral, Power11 processor-based servers also support OMI DDR4 memory that is migrated from Power10 high-end systems, protecting client investments in memory technology.

► AI acceleration

The Power11 platform continues to support a wide range of enterprise AI use cases through the MMA architecture. Improvements in processor core strength and system capacity enhance MMA performance for inference workloads. In addition, IBM plans to incorporate the IBM Spyre Accelerator into Power11 offerings to provide more AI inferencing capabilities. Together, IBM POWER processors and the IBM Spyre Accelerator enable next-generation infrastructure to scale demanding AI workloads for business.

For more information about AI and IBM Spyre Accelerator, see Chapter 4, “Artificial intelligence support” on page 113.

1.5.1 New Power11 processor details

Here are some details on the new Power11 processor design:

► Technology and packaging

- 654 mm² 7 nm refined Samsung process (30 billion devices).
- Advanced 2.5D packaging technology optimizes space and improves signal integrity.
- ISC technology enhances power delivery and performance by integrating capacitors directly into the chip package.
- Supports single-chip or dual-chip sockets.

► Computational capabilities

- Up to 16 SMT8 cores (2 MB L2 cache per core).
- Up to 128 MB L3 cache with low-latency non-uniform cache architecture management.
- Enterprise performance focus.
 - 3x core performance compared to the Power9 platform.
 - 2x thread strength compared to the Power9 platform.
 - 4x L2 cache and 4x MMU per core compared to the Power9 platform.
 - 4x cryptographic engine per core compared to the Power9 platform.

- AI computational focus.
 - 2x general single instruction, multiple data (SIMD) per core compared to the Power9 platform.
 - 4x matrix SIMD per core compared to the Power9 platform.
 - New AI instructions and data types.
- ▶ Robust data plane
 - 2 TBps raw bandwidth (32 GTps) with PowerAXON and OMI signaling.
 - SMP interconnect supports up to 16 sockets.
 - 2x OMI memory bandwidth compared to the Power10 platform.
 - 64 TB OMI DDR large system memory capacity.

Figure 1-14 shows the Power11 processor.



Figure 1-14 Power11 processor diagram

Each of the cores has an L1 cache with 64 KB for the data cache and 96 KB for the instruction cache.

1.5.2 Comparing Power11, Power10, and Power9 processors

Table 1-5 shows a comparison of the Power11 processor to previous generations.

Table 1-5 Comparing the Power11 processor technology to previous processor generations

Characteristics	Power11 processor	Power10 processor	Power9 processor
Technology	7 nm	7 nm	14 nm
Die size	654 mm ²	602 mm ²	693 mm ²
Processor module size	68.5 mm x 77.5 mm	68.5 mm x 77.5 mm	68.5 mm x 68.5 mm

Characteristics	Power11 processor	Power10 processor	Power9 processor
Number of transistors	30 billion ^a	18 billion	8 billion
Maximum cores	16	15	12
Maximum hardware threads per core	8	8	8
Maximum static frequency and high-performance frequency range	3.8 - 4.4 GHz	3.75 - 4.15 GHz	3.9 - 4.0 GHz
L2 cache per core	2048 KB	2048 KB	512 KB
L3 cache	8 MB of L3 cache per core with each core having access to the full 128 MB of L3 cache, on-chip high-efficiency SRAMs	8 MB of L3 cache per core with each core having access to the full 120 MB of L3 cache, on-chip high-efficiency SRAM	10 MB of L3 cache per core with each core having access to the full 120 MB of L3 cache, on-chip eDRAM
Memory technology	DDR5 and DDR4 ^b	DDR4 and DDR5 ^c	DDR4 and DDR3
I/O bus	PCIe Gen 5	PCIe Gen 5	PCIe Gen 4

a. IBM Power recently changed its methodology for counting transistors to be more in line with the industry.

b. DDR4 only as a migration from the Power10 platform during the same serial number upgrade.

c. DDR5 support was added after GA.

1.5.3 Chip packaging options

To provide flexibility and scalability, the Power11 processor is packaged in several ways. For the high-end Power E1180, IBM uses an SCM, which provides one chip per socket and supports up to 16 sockets per system. For the midrange Power E1150, the Power11 processor is packaged with two chips per socket in a Dual-Chip Module (DCM) and supports up to four sockets per system. The DCM version is also used in the scale-out system line, including the Power S1122 and Power L1122, and Power S1124 and Power L1124, with up to two sockets per system. A third packaging option, the eSCM, is used in the entry-level processor features of the Power S1122 (the 4-core and 10-core options). The eSCM is a modified DCM in which the second chip in the module does not provide processor capability and only adds I/O functions.

Dual-Chip Module

The Power E1150 and Power S1122 systems use the Power11 processor architecture with a DCM design. This configuration enables each socket to contain two processor chips, which increases core density and performance.

Figure 1-15 shows the DCM layout with two fully operational Power11 processors that are combined into a single socket.

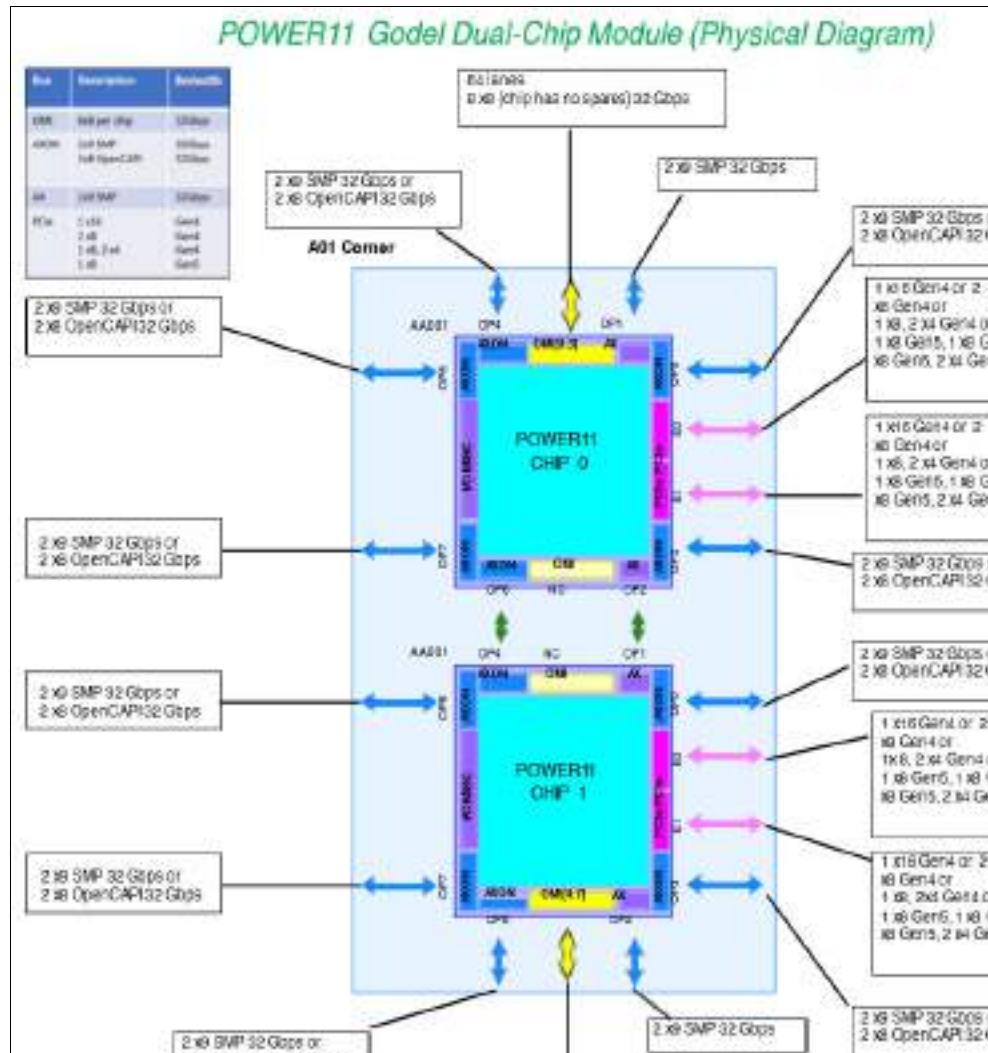


Figure 1-15 Golden dual chip (physical diagram)

In the Power S1122 and Power S1124 systems, the DCM plays a central role in delivering scalable performance. These systems are designed for midrange enterprise applications and benefit from Power11 innovations, including the DCM's support for high memory bandwidth and PCIe Gen5 connectivity. The DCM enables the systems to handle demanding workloads while maintaining a compact footprint, making them ideal for data centers with space and power constraints. The integration of OMI and DDIMMs further enhances memory performance and reliability, which is critical for mission-critical applications.

Both systems use the DCM to support IBM's goals of flexibility and efficiency in hybrid cloud environments. The DCM boosts raw performance and supports advanced features such as TME and dynamic resource allocation. These capabilities are essential for modern IT infrastructures that require both security and adaptability. Whether deployed in traditional data centers or as part of a hybrid cloud strategy, the DCM-equipped Power S1122 and Power S1124 systems provide a solid foundation for scalable, secure, and high-performance computing.

Entry Single-Chip Module

The eSCM is a derivative of the DCM in which all active compute cores run on the first chip, and the second chip provides more PCIe connectivity, acting as a switch.

The entry-level versions of the Power S1122 (the 4-core and 10-core modules) use the eSCM, and all other processor features use the DCM format. All Power L1122, Power S1124, and Power L1124 processor modules use the DCM design.

Table 1-6 shows this configuration.

Table 1-6 Using eSCM and DCM in the Power S1122 and Power L1122

Feature Code 22A / 22H	Processor type	Cores	Frequency range	Supported configurations
EP3Y / EP4J	DCM	30 cores + 2 spare	2.4 - 3.95 GHz	60 cores
EP3H / EP4B	DCM	24 cores + 2 spare	2.65 - 4.15 GHz	48 cores
EP3X / EP4C	DCM	16 cores + 2 spare	3.00 - 4.2 GHz	32 cores
ERGQ / NA	eSCM	10 cores	3.05 - 4.00 GHz	20 cores
ERGR / NA	eSCM	4 cores	3.60 - 4.00GHz	8 cores

Figure 1-16 shows the main differences between the eSCM and DCM structures:

- ▶ All active cores are on chip 0, and no active cores are on chip 1.
- ▶ Chip 1 works with chip 0 as a switch to provide more I/O connections.
- ▶ All active OMI interfaces are on chip 0, and no active OMI interfaces are on chip 1.

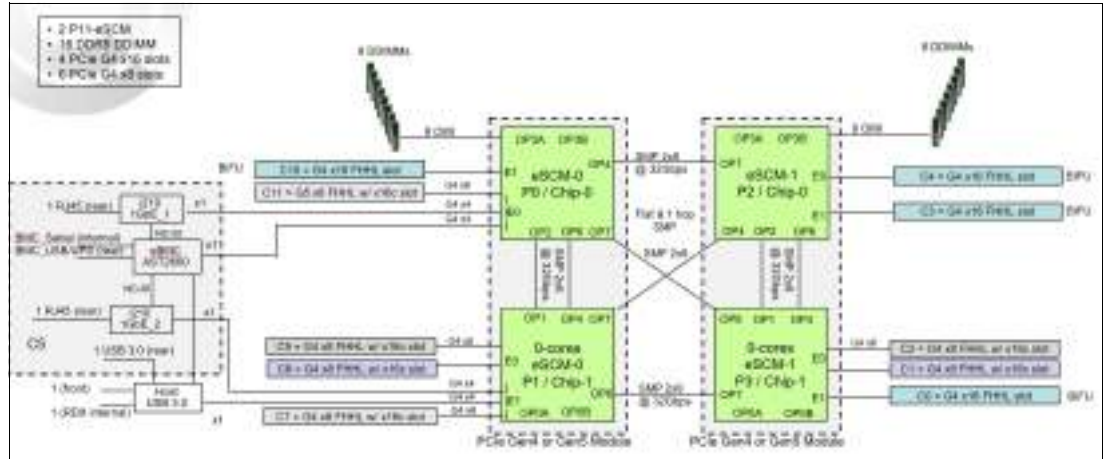


Figure 1-16 Entry Single-Chip Module logical diagram

The eSCM internal chip-to-chip connectivity, the SMP links across the eSCM in two-socket configurations, and the PCIe Gen5 bus structure are identical to the Power11 DCM implementation.

As with the Power11 DCM, 36 X-bus lanes are used for two chip-to-chip connections. These eSCM internal connections use the interface ports OP2 and OP6 on chip 0 and OP1 and OP4 on chip 1:

- ▶ Two 9 OP2 lanes of chip 0 connect to two 9 OP1 lanes of chip 1.
- ▶ Two 9 OP6 lanes of chip 0 connect to two 9 OP4 lanes of chip 1.

The eSCM module internal chip-to-chip links provide a theoretical maximum full-duplex bandwidth of 256 GBps.

To connect the two eSCMs in two-socket configurations, the interface ports OP4 and OP7 on chip 0 and OP6 and OP7 on chip 1 of the processor module are active and implement direct chip-to-chip SMP connections between the two eSCM modules.

As with the DCM, the eSCM offers differential PCIe Gen5 with a total of 64 lanes. Each chip of the eSCM provides 32 PCIe Gen5 lanes, which are grouped into two PCIe host bridges (E0 and E1) with 16 PCIe Gen5 lanes each:

- ▶ E0 and E1 on chip 0
- ▶ E0 and E1 on chip 1

Because only I/O is supported on chip 0 in the eSCM, the maximum number of DDIMMs that IS supported is eight per socket, for a maximum of 16 DDIMMs per system. In comparison, the DCM processor design supports up to 32 DDIMMs.

1.5.4 Spare core support

In previous generations of IBM POWER processors, IBM provided a proactive fault-tolerance mechanism to handle certain predictive core errors. When a suspect core was identified, PowerVM could automatically unconfigure it and if spare capacity was available (either unused or unlicensed) substitute these resources to maintain system processing capacity. This failover occurred without requiring a system or application restart, helping ensure continued operation. This feature remains available in the Power11 processor.

Because the substituted resources were not guaranteed to be on the same processor chip as the failing core, the spare might not share affinity as the original core. In addition, because the spare core was not a dedicated resource, it eventually might be needed or licensed. Therefore, the processor containing the suspect core was flagged for repair. As a result, a planned system outage was scheduled to replace the affected processor module and restore full system integrity.

Spare cores can also be used for processor core failures other than predictive errors. In these cases, the processor module is called out for repair. However, if the spare is successfully deployed, system owners might decide not to take down the system because full processor capacity and performance remain available.

With the Power11 platform, spare cores (one spare in each SCM and two spares in each DCM) are allocated on each processor module. The exceptions are the 16-core SCM, which allocates all cores for client usage, and the eSCM modules, which allocate no spares. These cores are used only to substitute for a failing core on the same processor module, which resolves affinity issues and helps ensure that the spare core is not otherwise used. As a result, predictive processor core failures in the Power11 platform do not require system downtime to repair the processor module.

1.6 Operating system support

The Power11 scale-out servers support the following OSs:

- ▶ AIX
- ▶ IBM i
- ▶ Linux
- ▶ Red Hat OpenShift

In addition, you can install the VIOS in special partitions that provide virtualization of I/O capabilities, such as network and storage connectivity. You can install multiple VIOS partitions to provide support and services to other partitions running AIX, IBM i, or Linux, such as virtualized devices and LPM capabilities.

For more information about the OSs and other software that are available on Power, see [IBM Power](#). The minimum supported levels of IBM AIX, IBM i, and Linux are described in the following sections.

IBM will announce support for newer versions of the OSs over time. Keep your OSs up to date to enable new security capabilities and support new features.

1.6.1 AIX operating system

The Power S1122 and Power S1124 servers support the minimum levels of the AIX OS that are described in this section:

- ▶ When using virtual I/O

The Power S1122 and Power S1124 servers support the following minimum levels of the AIX OS when installed with virtual I/O:

- AIX Version 7.3 with TL 7300-02 and Service Pack (SP) 7300-02-02-2420
- AIX Version 7.3 with TL 7300-03 and SP 7300-03-00-2446
- AIX Version 7.2 with TL 7200-05 and SP 7200-05-08-2420

- ▶ When using direct I/O connectivity

The Power S1122 and Power S1124 servers support the following minimum levels of the AIX OS when installed by using direct I/O connectivity:

- AIX Version 7.3 with TL 7300-02 and SP 7300-02-04
- AIX Version 7.3 with TL 7300-03 and SP 7300-03-01
- AIX Version 7.3 with TL 7300-04 and SP 7300-04-00
- AIX Version 7.2 with TL 7200-05 and SP 7200-05-10

Important:

- ▶ AIX 7.2 instances can use physical and virtual I/O adapters, but must run in an LPAR in IBM Power9 compatibility mode.
- ▶ AIX 7.3 instances can use physical and virtual I/O adapters, and can run in an LPAR in native Power11 mode.

1.6.2 IBM i operating system

The Power S1122 and Power S1124 support the following minimum levels of the IBM i OS both natively and as a client of VIOS:

- ▶ IBM i 7.4 TR12
- ▶ IBM i 7.5 TR6
- ▶ IBM i 7.6

Note: IBM i support is provided at a cost-effective Power10 software tier, even though the Power S1122 has two sockets. There are limitations on the maximum partition size, and all I/O must be virtualized through VIOS. VIOS is required, and IBM i partitions must be set to restricted I/O. Each IBM i partition supports up to four cores (real or virtual). Multiple IBM i partitions can run concurrently, and each partition can have up to four cores.

Native IBM i support is available on the Power S1122 (MTM 9824-22A) only when configured with two 4-core processors (Feature Code ERGR), with a maximum of eight active cores. This configuration is available at a IBM i Power10 software tier and supports IBM i natively, virtualized, or as a combination of both.

When the Power S1122 (MTM 9824-22A) is configured with two 4-core processors (Feature Code ERGR), full IBM i native support is available. IBM i partitions with more than four cores can be used, and VIOS is not required. When these systems are configured with any other processor feature, the 4-core-per-partition limit applies.

Here are some more requirements:

- ▶ Feature Code ERGR cannot be upgraded to any other processor feature.
- ▶ Two Feature Code ERGR 4-core processors place the system in the IBM i Power10 processor group or software tier.
- ▶ The Feature Code ERGR processor is identified as QPRCFEAT ERGR.
- ▶ Systems with this configuration must also include Feature Code EEPS (IBM i Enablement for 9824-22A with two Feature Code ERGR processors).

1.6.3 Linux operating system

The Linux distributions that are listed here are fully supported on the Power S1122 and Power S1124 server models. Other distributions, including open-source releases, can run on these servers but do not include formal enterprise-grade support.

Red Hat Enterprise Linux

The latest version of the RHEL distribution from Red Hat is supported in native Power11 mode, which enables it to access all features of the Power11 processor and platform.

At announcement, the Power S1122 and Power S1124 servers support the following minimum RHEL OS levels:

- ▶ RHEL 8.10 for Power LE or later (Power10 compatibility)
- ▶ RHEL 9.4 for Power LE or later (Power10 compatibility)
- ▶ RHEL 9.6 for Power LE or later (Power11 native)

Red Hat OpenShift Container Platform

The Red Hat OpenShift Container Platform is supported on Power servers, including the Power11 scale-out models.

Red Hat OpenShift Container Platform 4.18 runs in Power10 compatibility mode, and Red Hat OpenShift 4.19 or later runs in Power11 native mode on Power11 scale-out servers.

The Power S1122 and Power S1124 servers support the following minimum OS levels for Red Hat OpenShift Container Platform:

- ▶ Red Hat Enterprise Linux CoreOS (RHCOS) 4.18 (Power10 compatibility)
- ▶ RHCOS 4.19 or later (Power11 native)

SUSE Linux Enterprise Server

The latest version of the SUSE Linux Enterprise Server distribution is supported in native Power11 mode, which enables it to access all features of the Power11 processor and platform.

At announcement, the Power S1122 and Power S1124 servers support SUSE Linux Enterprise Server 15 SP6 or later (Power11 native).

1.6.4 VIOS

The minimum required VIOS level for Power11 scale-out servers is VIOS 4.1.1.10 when running in Power11 mode. The following levels are supported:

- ▶ VIOS 4.1.1.10 (SP)
- ▶ VIOS 4.1.0.40 (SP) (Power10 compatibility mode)
- ▶ VIOS 3.1.4.60 (SP) (Power9 compatibility mode)

IBM regularly updates the VIOS code. For more information, see [Fix Central](#).

1.7 Firmware and Hardware Management Console

The firmware on an Power server is the foundational software layer that initializes hardware components, manages system resources, and provides the interface between the hardware and the OS.

Known as the Power Firmware, it includes components such as the Flexible Service Processor (FSP), IBM PowerVM Hypervisor (PHYP), and system boot code. This firmware enables advanced features such as logical partitioning, dynamic resource allocation, and Secure Boot. It also supports system diagnostics, error reporting, and remote management, especially when integrated with tools such as the HMC. Regular firmware updates from IBM provide performance improvements, security fixes, and compatibility with new hardware and software features.

An HMC is a dedicated appliance for configuring and managing system resources on Power servers. It can be delivered as a physical server or as a virtual appliance running in either PowerVM or an x86 virtual machine. A GUI, a CLI, and Representational State Transfer (REST) application programming interfaces (APIs) are available.

1.7.1 HMC requirements

The HMC provides basic virtualization management for configuring LPARs. Most operations, including processor and memory settings, can be completed dynamically.

The HMC also supports advanced service functions such as guided repair and verification, concurrent firmware updates for managed systems, and continuous error reporting through IBM Electronic Service Agent (ESA) for faster support. HMC management features help improve server usage, simplify system management, and accelerate resource provisioning by using IBM PowerVM virtualization technology.

Supported HMCs for Power11 processor-based servers:

- ▶ 7063-CR2
- ▶ Virtual Hardware Management Console (vHMC) running on PowerVM
- ▶ vHMC running on an x86-based VM

The minimum required HMC version for Power scale-out servers is V11R1 M1110. This version is supported on the 7063-CR2 and vHMC appliances only.

Note:

- ▶ The 7063-CR1 and 7042 HMCs are not supported for Power11 processor-based servers.
- ▶ HMC code at V11R1 M1110 cannot manage Power8 or earlier processor-based servers.

1.8 Rack support

All currently available scale-out servers fit a standard 19-inch rack. The servers are certified and tested in IBM Enterprise racks (7965-S42 and 7014-T42). Customers can install the servers in other racks if those racks meet the required strength, rigidity, depth, and hole pattern specifications. Contact IBM Support to confirm whether other racks are suitable.

Note: It is a best practice to order Power11 scale-out servers with an IBM 42U Enterprise Rack (Feature Code ECR0, 7965-S42). This rack provides a high-quality environment for IBM manufacturing, assembly, and testing, and offers a complete package.

If a system is installed in a non-IBM rack or cabinet, ensure that the rack meets all requirements.

Important: The customer 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.

1.8.1 New rack considerations

Consider the following points when ordering racks:

- ▶ The new IBM Enterprise 42U Slim Rack (7965-S42) offers 42 EIA units (U) of space in a slim footprint.
- ▶ The 7014-T42 rack is no longer available for purchase with Power11 scale-out servers. Installing a Power11 scale-out server in this rack is still supported.



IBM Power11 scale-out servers

IBM Power11 scale-out servers introduce a powerful suite of capabilities that are built on IBM's hallmark strength: a fully integrated, end-to-end stack. From the Power11 processor and system architecture to the firmware, operating systems (OSs), and cloud integration, IBM Power delivers innovation through deep synergy across every layer of the infrastructure. This holistic design is designed to support autonomous IT operations (ITOps) and drive tangible business outcomes across three key pillars.

Business continuity is at the core of Power11 scale-out servers, offering a resilient and dependable platform for mission-critical workloads. It minimizes risk exposure to threats and regulatory challenges, whether the deployment is on-premises, in the cloud, or in a hybrid environment. The system enhances productivity and efficiency by maximizing uptime and streamlining operations. This approach reduces complexity and reduces operational costs, helping organizations achieve more with fewer resources.

The Power11 scale-out servers are designed for the future, enabling accelerated growth and scalability in the artificial intelligence (AI) era. They support the rapid deployment of diverse AI workloads and new applications, helping ensure that mission-critical processes scale securely and consistently.

This chapter contains the following topics:

- ▶ Power11 scale-out servers
- ▶ Power S1122 (9824-22A) server overview
- ▶ Power S1124 (9824-42A) server overview
- ▶ Processor and memory activation on Power S1122 and Power S1124 systems
- ▶ Environmental and physical specifications

2.1 Power11 scale-out servers

IBM uses the Power11 processor in multiple form factors to meet different client requirements. The Power11 scale-out servers are optimized for horizontal scalability. These systems are ideal for distributed workloads, cloud-native applications, and environments where flexibility and cost efficiency are essential. They support up to two sockets and are designed to deliver high performance in compact form factors, making them suitable for businesses that want to scale incrementally or run multiple smaller workloads in parallel.

One of the stand-out features of the Power11 scale-out models is support for Peripheral Component Interconnect Express (PCIe) Gen5, which increases data throughput and connectivity for high-speed adapters. In addition, the servers use Open Memory Interface (OMI) technology with differential DIMM (DDIMM) memory cards to enhance memory performance, resilience, and security. Security includes Transparent Memory Encryption (TME), which is essential for protecting sensitive data. The internal storage capabilities are robust and support up to 240 TB of NVMe storage in a two-socket system, helping ensure low-latency access to large datasets.

These servers are also optimized for hybrid cloud environments, offering flexible consumption-based pricing through IBM Power Private Cloud with Shared Utility Capacity. This model enables businesses to scale resources dynamically based on workload demands, making it ideal for modern applications, including AI and containerized workloads on platforms such as Red Hat OpenShift. Overall, Power11 scale-out servers provide a future-ready infrastructure that balances performance, security, and cost efficiency for enterprise IT needs.

2.1.1 Family overview

The Power11 scale-out server family consists of four new servers that are available in two form factors:

- ▶ Two-socket, 2U form factor servers (2S2U)

In the 2S2U form factor, IBM provides the IBM Power S1122 model and the IBM Power L1122 model. These two servers are similar, with the primary difference being that the Power L1122 is designed for the Linux market segment. At least 75% of the cores in the system must run Linux partitions, and up to 25% of activated cores can run AIX and IBM i.

- ▶ Two-socket, 4U form factor servers (2S4U)

In the 2S4U form factor, IBM provides the IBM Power S1124 model and the IBM Power L1124 model. As with the Power L1122, the Power L1124 is designed for the Linux market segment. At least 75% of the cores in the system must run Linux partitions, and up to 25% of activated cores can run AIX and IBM i.

All servers are built on the advanced Power11 processor architecture and use the performance benefits of many multithreaded cores, high-performance DDR5 memory, and PCIe Gen5 adapter connectivity.

Table 2-1 lists the machine types and models for the Power11 scale-out servers.

Table 2-1 Machine types and models of Power S1122 and Power S1124 server models

Server name	Machine type and model
Power S1122	9824-22A
Power L1122	9856-22H

Server name	Machine type and model
Power S1124	9824-42A
Power L1124	9856-42H

2.2 Power S1122 (9824-22A) server overview

The Power S1122 system introduces a comprehensive set of advanced capabilities, which are rooted in IBM's core strength: an end-to-end, full-stack design. By integrating every layer of the infrastructure, from the POWER processor and systems to the firmware, OSs, and the cloud, IBM Power delivers unique and innovative solutions. This integrated approach is based on autonomous IT principles and consistently translates into tangible business outcomes across three core pillars.

Figure 2-1 shows the Power S2122 from the front.



Figure 2-1 Front view of the Power S1122

2.2.1 Processors

The Power11 processor is the compute engine for the next generation of Power11 scale-out servers. It is the successor to the current Power10 processor. It offers superior performance for applications that use the Matrix Math Accelerator (MMA) facility to accelerate computation-intensive kernels, matrix multiplication, convolution, and discrete Fourier transform. To efficiently accelerate MMA operations, the Power11 processor core implements a dense math engine (DME) microarchitecture that provides an accelerator for cognitive computing, machine learning, and AI inferencing workloads.

Two Power11 processors of the same type are supported. The following processor module configurations are allowed:

- ▶ Two 4-core, typical 3.6 - 4.0 GHz (max), Power11 entry Single-Chip Module (eSCM) processors (Feature Code ERGR)
- ▶ Two 10-core, typical 3.05 - 4.0 GHz (max), Power11 eSCM processors (Feature Code ERGQ)
- ▶ Two 16-core, typical 3.0 - 4.2 GHz (max), Power11 Dual-Chip Module (DCM) processors (Feature Code EBG8)
- ▶ Two 24-core, typical 2.65 - 4.15 GHz (max), Power11 DCM processors (Feature Code EBG9)
- ▶ Two 30-core, typical 2.4 - 3.95 GHz (max), Power11 DCM processors (Feature Code EBGA)

The Power S1122 offers enhanced Workload Optimized Frequency for optimum performance. This mode dynamically optimizes the processor frequency based on CPU usage and operating environmental conditions. For a description of this feature and other power management options that are available for this server, see [IBM EnergyScale for Power10 processor-based servers](#).

Processor activations

Processors require an activation feature to be used. For eSCM-based processors, 100% of the processor cores require an activation code. For DCM-based processors, 50% of the processor cores must be activated unless the system is running with Power Enterprise Pools (PEP) 2.0, in which case the minimum number of activated cores is one. This configuration also requires the PEP 2.0 feature code.

For more information, see 2.4, “Processor and memory activation on Power S1122 and Power S1124 systems” on page 69.

2.2.2 Memory

The Power S1122 server uses next-generation DDIMMs, which are high-performance, high-reliability, and high-function memory cards that contain a buffer chip, intelligence, and 4000 MHz or 4800 MHz DRAM memory. DDIMMs are installed in DDIMM slots in the server system.

- ▶ A minimum of 32 GB of memory is required with one processor module. All memory DIMMs must be ordered in pairs.
- ▶ A minimum of 64 GB of memory is required with two processor modules. All memory DIMMs must be ordered in quads.
- ▶ Each DIMM Feature Code delivers two physical memory DIMMs.
- ▶ Plans for future memory upgrades should be considered when selecting the memory feature size during the initial system order. For the best possible performance, populate all DIMM slots, especially for OLAP and other high-bandwidth workloads.

To help with the installation rules, two DDIMMs are ordered by using one memory Feature Code. Select from the following options:

- ▶ 64 GB (2 × 32 GB) (Feature Code EM54)
- ▶ 128 GB (2 × 64 GB) (Feature Code EM5B)
- ▶ 256 GB (2 × 128 GB) (Feature Code EM5G)

Note: DDR5 memory DDIMMs require FW1110.00 or later.

2.2.3 PCIe slots

The Power S1122 server supports up to eight U.2 NVMe devices and up to 10 PCIe hot-plug slots, providing excellent configuration flexibility and expandability. All PCIe slots and NVMe devices are concurrently maintainable.

With both Power11 processor slots populated, ten PCIe slots are available:

- ▶ Four x16 Gen4 or x8 Gen5 half-height, half-length slots
- ▶ Four x8 Gen5 half-height, half-length slots (with x16 connectors)
- ▶ Two x8 Gen4 half-height, half-length slots (with x16 connectors)

The x16 slots provide up to twice the bandwidth of x8 slots because they offer twice as many PCIe lanes. PCIe Gen5 slots can support up to twice the bandwidth of PCIe Gen4 slots, and PCIe Gen4 slots can support up to twice the bandwidth of PCIe Gen3 slots, assuming an equivalent number of PCIe lanes.

At least one PCIe Ethernet adapter is required on the server to help ensure proper manufacturing, testing, and support. One of the x8 PCIe slots is used for this required adapter.

These servers optimize energy efficiency when cooling the PCIe adapter environment. They detect which IBM PCIe adapters are installed and, if an adapter requires higher cooling levels, automatically increase fan speed to improve airflow across the PCIe adapters. Faster fans increase the sound level of the server. Higher-wattage PCIe adapters include PCIe3 SAS adapters and SSD or flash PCIe adapters (Feature Codes EJ10, EJ14, and EJ0J).

For more information about PCIe slots and supported adapters, see Chapter 3, “I/O subsystem” on page 75.

2.2.4 NVMe drives

NVMe SSDs in the 15-millimeter carrier U.2 2.5-inch form factor are used for internal storage in the Power S1122 system. The Power S1122 supports up to eight NVMe U.2 devices when two storage backplanes with four NVMe U.2 drive slots (Feature Code EJ1X) are ordered. Both 7-millimeter and 15-millimeter NVMe drives are supported in the 15-millimeter carrier.

For more information about the supported NVMe drives, see Chapter 3, “I/O subsystem” on page 75.

2.2.5 I/O expansion drawers

To provide more PCIe slots and NVMe capacity, the Power11 scale-out servers support I/O expansion drawers.

NED24 NVMe Expansion Drawer

The NED24 NVMe Expansion Drawer (Feature Code ESR0) is a storage expansion enclosure with 24 U.2 NVMe bays. It supports up to 24 U.2 NVMe devices in 15 mm Gen3 carriers. The 15 mm carriers can accommodate either 7 mm or 15 mm NVMe devices.

Each NED24 NVMe Expansion Drawer includes two redundant AC power supplies, which are part of the enclosure base.

The NED24 NVMe Expansion Drawer connects to a Power server through dual CXP Converter adapters (Feature Code EJ24 or Feature Code EJ2A). A cable pair attaches the PCIe Cable Adapter (Feature Code EJ24) to the fan-out module (FOM). Feature Code ECLS provides a pair of 3-meter copper cables. Only copper cables are supported on the Power S1122. Two cables of identical length, or one feature, are required for each FOM.

The following U.2 NVMe devices are supported in the NED24 NVMe Expansion Drawer:

- ▶ Feature Code ES5B. Enterprise 800 GB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES5A. Enterprise 800 GB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5C. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5D. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES5E. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5F. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES5G. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5H. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES4B. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES4C. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES4D. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES4E. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES4F. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES4G. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ECT9. 15.3 TB Mainstream NVMe U.2 SSD 4k for AIX and Linux
- ▶ Feature Code ECTB. 15.3 TB Mainstream NVMe U.2 SSD 4k for IBM i

PCIe Gen4 I/O Expansion Drawer

The PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) is a 4U-high, 19-inch-wide, PCIe Gen4-based rack-mountable I/O drawer that is available as a feature of Power11 processor-based servers.

The PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) replaces the PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0). There is no upgrade path from the PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0) to the PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0).

The PCIe Gen4 Fan-out Module (Feature Code ENZF) is the FOM for placement in the PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0). Each FOM provides six PCIe Gen4 slots (four x16 and two x8). All six slots, including the x8 slots, use x16 connectors. A fully populated PCIe Gen4 I/O Expansion Drawer with two PCIe Gen4 Fan-out Modules provides twelve PCIe Gen4 slots (eight x16 and four x8). Each PCIe Gen4 FOM connects only to PCIe x16 to CXP Converter Card features (Feature Code EJ24 or Feature Code EJ2A).

For the Power S1122 server, a cable pair attaches the PCIe Cable Adapter (Feature Code EJ24) to the FOM. Feature Code ECLS provides a pair of 3-meter copper cables. Two cables of identical length, or one feature, are required for each FOM. Copper cables have the same OS software prerequisites as active optical cables (AOCs).

Limitations

- ▶ The combined maximum of NED24 NVMe Expansion Drawers (Feature Code ESR0) and PCIe Gen4 I/O Expansion Drawers (Feature Code ENZ0) is limited to half of the maximum number of controller cards (Feature Code EJ24 or Feature Code EJ2A) allowed per server.
- ▶ A 1 EIA (1U) space is required in the IBM 7965-S42 rack or any OEM rack between the top of the PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) and the bottom of a scale-out server (9824-22A/42A, 9856-22H/42H) or NED24 NVMe Expansion Drawer (Feature Code ESR0) when installed directly below one of these options. This requirement applies to both factory-integrated orders (Feature Code 4651 - Feature Code 4666) and systems that are integrated in the field (Feature Code 4650). This 1U space is necessary for proper function of the cable management bracket and for serviceability of the PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0).

For more information about the I/O expansion drawers, see 3.2, “Enhancing I/O scalability with expansion drawers” on page 85.

2.2.6 Power S1122 system configuration

The minimum initial order for the Power S1122 must include two processor modules of the same type, two 32 GB DIMMs (one Feature Code EM54 64 GB [2 × 32 GB] DDIMM), two power supplies and power cords, an OS indicator, a cover set indicator, and a Language Group Specify. It must also include one of the following storage options and one of the following network options if AIX or Linux is the primary OS:

- ▶ Storage options:
 - For boot from NVMe. One NVMe drive slot and one NVMe drive, or one PCIe NVMe add-in adapter.
 - For boot from storage area network (SAN). An internal SSD is not required if Feature Code 0837 (boot from SAN) is selected. A Fibre Channel (FC) adapter must be ordered if Feature Code 0837 is selected.
- ▶ Network options:
 - One PCIe2 4-port 1 Gb Ethernet adapter.
 - One of the supported 10 Gb Ethernet adapters.

Table 2-2 shows the minimum defined initial order configuration.

Table 2-2 Minimum configuration for AIX or Linux

System Feature Codes	Feature Code	Description	Default	Minimum quantity	Notes
Op-Panel	EU0K	Operator Panel LCD Display		1	Optional with AIX and Linux. Always Qty. 1 by default, but can be deselected for AIX and Linux.
Virtualization Engine	EPVT	PowerVM Enterprise Edition	1	1	Must select one option.
	or				
	EPV0	Deactivation of Live Partition Mobility (LPM)		1	

System Feature Codes	Feature Code	Description	Default	Minimum quantity	Notes	
Processor Modules	ERGR	4-core Typical 3.6 - 4.0 GHz (max) Power11 Processor		2	Must select the Processor Module option.	
	or					
	ERGQ	10-core Typical 3.05 - 4.0 GHz (max) Power11 Processor		2		
	or					
	EBG8	16-core Typical 3.0 - 4.2 GHz (max) Power11 Processor		2		
	or					
	EBG9	24-core Typical 2.65 - 4.15 GHz (max) Power11 Processor		2		
	or					
	EBGA	30-core Typical 2.40 - 3.95 GHz (max) Power11 Processor		2		
Processor Module Activations	ERFQ	One core Processor Activation for eSCM		4	100% of Static processor core activations must be ordered.	
	or					
	EBF8	One core Processor Activation Pools 2.0 or Static for DCM		1	A minimum of 50% of Static Capacity Upgrade on Demand (CUoD) activations must be ordered if you are not using PEP 2.0. A minimum of one activation is required if you are using Feature Code EP20 PEP 2.0 Enablement.	

System Feature Codes	Feature Code	Description	Default	Minimum quantity	Notes	
Memory	EM54	64 GB (2x32 GB) DDIMMs, 4000 MHz or 4800 MHz, 16 Gbit DDR5 Memory		1	Minimum of two DIMMs per one DIMM feature.	
	or					
	EM5B	128 GB (2x64 GB) DDIMMs, 4000 MHz or 4800 MHz, 16 Gbit DDR5 Memory		1		
	or					
	EM5G	256 GB (2x128 GB) DDIMMs, 4000 MHz or 4800 MHz, 32 Gbit DDR5 Memory		1		
Active Memory Mirroring (AMM)	EM8G	AMM	0	0	Optional feature. Maximum of Qty. 1 per system. Memory mirroring requires a minimum of eight DIMMS (four DIMMs). A minimum of two DIMMs per one DIMM feature.	
Storage Backplane	EJ1X	Storage Backplane with four NVMe U.2 drive slots		1	You must order Qty. 1 NVMe backplane feature except when Feature Code 0837 or Feature Code ESCZ (internet Small Computer Systems Interface (iSCSI) boot) is on the order or when an NVMe PCIe add-in adapter is used as the Load Source. Mixing NVMe devices is allowed on each backplane.	
Bezels	EJBS	Front IBM Bezel for 8 NVMe-bays Backplane Rack-Mount		1	When no NVMe backplane is ordered, the default is Feature Code EJBS.	
	or					
	EJUT	Front OEM Bezel for 8 NVMe-bays Backplane Rack-Mount		1		
NVMe Devices	ES5A	800 GB Mainstream NVMe U.2 SSD 4k for AIX and Linux	2	0	When no NVMe backplane is ordered, the default is Feature Code EJBS.	

System Feature Codes	Feature Code	Description	Default	Minimum quantity	Notes	
Required local area network (LAN) adapters	EC71	PCIe4 LP 2-Port 25/10/1 GbE RoCE SFP28 Adapter	1	1	Qty. 1 of these LAN features is required on all initial orders. The default adapter is Feature Code EC71.	
	or					
	EC85	PCIe5 LP 2-Port 200 GbE RoCE Adapter		1		
	or					
	EN2X	PCIe3 LP 4-port 10 GbE BaseT RJ45 Adapter		1		
	or					
	EN2Y	PCIe LP 4-Port 1 GbE Adapter		1		
Power Supply	EB3N	AC Power Supply - 2000W for Server (200-240 VAC)	2	2	Each initial order must have all its power supplies. Power supplies cannot be added later on. Only 200 - 240 V power cords can be used.	
Power Cables	6458	Power cord 4.3 m (14 ft), Drawer to IBM power distribution unit (PDU) (250V/10A)	2	2	Qty. 2 is required.	
Language Group	9300	Language Group Specify - US English	1	1	A Language Specify code is required.	
Primary OS	2146	Primary OS: AIX			Must select one option.	
	or					
	2147	Primary OS: Linux				

The racking approach for the initial order can be an MTM 7965-S42.

2.2.7 Comparing Power S1122, IBM Power S1022, and IBM Power S1022s servers

Table 2-3 provides a comparison between the Power10 processor-based and Power11 processor-based 2S2U servers.

Table 2-3 Comparing Power S1122, Power S1022, and Power S1022s servers

Attribute	Power S1022 (22A)	Power S1022s (22B)	Power S1122 DCM processors	Power S1122 eSCM processors
Processor	Two Power10 DCMs (5387 pins)	Two Power10 eSCM (5387 pins)	Two Power11 DCMs (5387 pins)	Two Power11 eSCMs (5387 pins)

Attribute	Power S1022 (22A)	Power S1022s (22B)	Power S1122 DCM processors	Power S1122 eSCM processors
Processor power (max)	325 W	240 W	430 W	285 W
SMP X-bus	Four 2B @ 32 Gbps	Four 2B @ 32 Gbps	Four 2B @ 32 Gbps	Four 2B @ 32 Gbps
Memory channels per system	32 OMI channels	16 OMI channels	32 OMI channels	16 OMI channels
Memory DDR4 data rate	2666 Mbps or 3200 Mbps per OMI channel	2666 Mbps or 3200 Mbps per OMI channel	4800 Mbps per OMI channel	4000 Mbps per OMI channel
Memory bandwidth per system (max theoretical)	<ul style="list-style-type: none"> ▶ 818 GBps with 3200 Mbps 16 GB, 32 GB, or 64 GB DDIMM ▶ 682 GBps with 2666 Mbps 128 GB DDIMM 	<ul style="list-style-type: none"> ▶ 409 GBps with 3200 Mbps 16 GB, 32 GB, or 64 GB DDIMM ▶ 341 GBps with 2666 Mbps 128 GB DDIMM 	<ul style="list-style-type: none"> ▶ 1228 GBps with 4800 Mbps 32 GB DDIMMs ▶ 2456 GBps with 4800 Mbps 64 GB or 128 GB DDIMMs 	614 MBps with 4800 Mbps DDIMMS
DIMMs per system	32	16	32	16
Memory capacity per system (max)	4 TB	2 TB	4 TB	2 TB
OpenCAPI port	6 ports @ 25 Gbps	0	0	0
OpenCAPI bandwidth per system (peak)	300 GBps	N/A	N/A	N/A
PCIe lanes per system (max)	128 PCIe G4 lanes @ 16 Gbps	128 PCIe G4 lanes @ 16 Gbps	128 PCIe G4 lanes @ 16 Gbps	128 PCIe G4 lanes @ 16 Gbps
PCIe slots per system	<ul style="list-style-type: none"> ▶ Four PCIe G4 x16 or G5 x8 slots ▶ Four PCIe G5 x8 slots ▶ Four PCIe G4 x8 slots 	<ul style="list-style-type: none"> ▶ Four PCIe G4 x16 or G5 x8 slots ▶ Four PCIe G5 x8 slots ▶ Two PCIe G4 x8 slots 	<ul style="list-style-type: none"> ▶ Four PCIe G4 x16 or G5 x8 slots ▶ Four PCIe G5 x8 slots ▶ Four PCIe G4 x8 slots 	<ul style="list-style-type: none"> ▶ Four PCIe G4 x16 or G5 x8 slots ▶ Four PCIe G4 x8 slots ▶ Two PCIe G4 x8 slots
Slots for internal storage	General	General	General	Genera
Drives (max)	Eight NVMe U.2	Eight NVMe U.2	Eight NVMe U.2	Eight NVMe U.2

Attribute	Power S1022 (22A)	Power S1022s (22B)	Power S1122 DCM processors	Power S1122 eSCM processors
I/O Expansion Drawer (max)	1	1	1	1
Service Interface	Enterprise Baseboard Management Controller (eBMC)	eBMC	eBMC	eBMC

2.3 Power S1124 (9824-42A) server overview

The Power S1124 is a high-performance, 4U rack-mounted scale-out server that is designed for enterprise workloads that require flexibility, scalability, and reliability. It supports one or two Power11 processors. Each processor includes two spare cores to enhance system availability. With both sockets populated, the server supports up to 60 active cores and 480 simultaneous threads through eight-way simultaneous multithreading (SMT8).

Figure 2-2 shows the Power S1024 front view.



Figure 2-2 Power S1024 front view

The Power S1124 supports CUoD, enabling organizations to activate more processor cores as needed. A minimum of 50% of installed cores must be activated at purchase, with the remainder available for future upgrades. These static activations are tied to the system but can be converted to a Power Private Cloud with Shared Utility Capacity, enabling shared processor activations across a pool of systems for greater flexibility and cost efficiency.

The Power S1124 server can also be purchased as part of a Power Private Cloud with Shared Utility Capacity pool. In this case, the system can be purchased with one or more base processor activations that are shared within the pool of systems. More base processor activations can be added to the pool in the future. It is also possible to convert a system with static activations to become part of a Power Private Cloud with Shared Utility Capacity pool.

The system includes 32 DDIMM slots, with 16 usable in single-socket configurations. These slots support DDR5 memory through the Power11 OMI interface, delivering up to 1228 GBps of memory bandwidth. With both sockets populated, the server supports up to 8 TB of memory, with a minimum of 32 GB per socket.

Here are some of the Power S1124 features:

- ▶ TME
Provides increased data security by encrypting memory with no management setup and no performance impact.
- ▶ Optional AMM for IBM PowerVM Hypervisor (PHYP) resilience
Enhances resilience by mirroring critical memory that is used by the PHYP.

The Power S1124 comes bundled with PowerVM Enterprise Edition to support robust virtualization and hybrid cloud deployments. It can run on the following OSs, which makes it ideal for both traditional enterprise applications and modern containerized workloads:

- ▶ AIX
- ▶ IBM i
- ▶ Linux, including Red Hat OpenShift Container Platform

Internal storage for the Power S1124 is exclusively NVMe-based, connecting directly to the system PCIe lanes to deliver high performance and efficiency. A maximum of 16 U.2 form-factor NVMe devices can be installed, providing up to 102.4 TB of storage in a single server. Extra NVMe storage can be connected through expansion drawers, such as the NED24 NVMe Expansion Drawer (Feature Code ESLS). FC connectivity to an external storage array is also available.

Other features that are integrated into the Power S1124

- ▶ System management by using an eBMC
- ▶ EnergyScale technology
- ▶ Redundant hot-swap cooling
- ▶ Redundant hot-swap AC Titanium power supplies
- ▶ Up to two Hardware Management Console (HMC) 1 GbE RJ45 ports
- ▶ One rear Universal Serial Bus (USB) 3.0 port
- ▶ One front USB 3.0 port
- ▶ 19 inch rack-mounting hardware (2U)

2.3.1 Processors

The Power11 processor is the compute engine for the next generation of Power11 processor-based servers and is the successor to the current Power10 processor. It offers superior performance for applications that use the MMA facility to accelerate computation-intensive kernels, matrix multiplication, convolution, and discrete Fourier transform. To efficiently accelerate MMA operations, the Power11 processor core implements a DME microarchitecture that provides an accelerator for cognitive computing, machine learning, and AI inferencing workloads.

A maximum of two Power11 dual-chip processor modules of the same type are supported:

- ▶ One or two 16-core, typical 3.4 - 4.2 GHz (max), Power11 processors (Feature Code EP3X)
- ▶ Two 24-core, typical 3.05 - 4.15 GHz (max), Power11 processors (Feature Code EP3H)
- ▶ Two 30-core, typical 2.8 - 3.95 GHz (max), Power11 processors (Feature Code EP3Y)

The MMA feature enables in-core AI inferencing and machine learning where the data resides.

Processor activations

Processors require an activation feature to be used. For more information about the options that are available, see 2.4, “Processor and memory activation on Power S1122 and Power S1124 systems” on page 69.

2.3.2 Memory

The Power S1124 server uses next-generation DDIMMs, which are high-performance, high-reliability, and high-function memory cards that contain a buffer chip, intelligence, and 4000 MHz or 4800 MHz DRAM memory. DDIMMs are installed in DDIMM slots in the server system.

- ▶ A minimum of 32 GB of memory is required with one processor module. All memory DIMMs must be ordered in pairs.
- ▶ A minimum of 64 GB of memory is required with two processor modules. All memory DIMMs must be ordered in quads.

Each DIMM feature code delivers two physical memory DIMMs.

Note: Plans for future memory upgrades should be considered when selecting the memory feature size during the initial system order.

For the best possible performance, populate all DIMM slots, or as many as possible, especially for high-bandwidth workloads.

To help with the installation rules, two DDIMMs are ordered by using one memory Feature Code. You can select from the following options:

- ▶ 64 GB (2 × 32 GB) (Feature Code EM54)
- ▶ 128 GB (2 × 64 GB) (Feature Code EM5B)
- ▶ 256 GB (2 × 128 GB) (Feature Code EM4M)
- ▶ 512 GB (2 × 256 GB) (Feature Code EM5J)

Note: DDR5 memory DDIMMs require FW1110.00 or later.

2.3.3 PCIe slots

The Power S1124 servers support up to 16 U.2 NVMe devices and up to 10 PCIe slots, providing excellent configuration flexibility and expandability. All PCIe slots and NVMe drives support concurrent maintenance.

With two Power11 processor DCMs, 10 PCIe slots are available:

- ▶ Four x16 Gen4 or x8 Gen5 full-height, half-length slots
- ▶ Four x8 Gen5 full-height, half-length slots (with x16 connectors)
- ▶ Two x8 Gen4 full-height, half-length slots (with x16 connectors)

With one Power11 processor DCM, five PCIe slots are available:

- ▶ One PCIe x16 Gen4 or x8 Gen5 full-height, half-length slot
- ▶ Three PCIe x8 Gen5 full-height, half-length slots (with x16 connectors)
- ▶ One PCIe x8 Gen4 full-height, half-length slot (with x16 connector)

The x16 slots provide up to twice the bandwidth of x8 slots because they offer twice as many PCIe lanes. PCIe Gen5 slots can support up to twice the bandwidth of PCIe Gen4 slots, and PCIe Gen4 slots can support up to twice the bandwidth of PCIe Gen3 slots, assuming an equivalent number of PCIe lanes.

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

These servers optimize energy efficiency when cooling the PCIe adapter environment. They detect which PCIe adapters are installed and, if an adapter requires higher cooling levels, automatically increase fan speed to improve airflow across the PCIe adapters. Faster fans increase the sound level of the server.

2.3.4 NVMe drive support

The Power S1124 supports up to 16 NVMe U.2 devices when two storage backplanes with eight NVMe U.2 drive slots (Feature Code EJ1Y) are ordered. NVMe U.2 form factor drives are installed in a 2.5 in., 15 mm carrier for internal storage in the Power S1124 system. Both 7 mm and 15 mm NVMe drives are supported in the 15-millimeter carrier.

2.3.5 I/O expansion drawers

To provide more PCIe slots and NVMe capacity, the Power11 scale-out servers support I/O expansion drawers.

NED24 NVMe Expansion Drawer

The NED24 NVMe Expansion Drawer (Feature Code ESR0) is a storage expansion enclosure with 24 U.2 NVMe bays. It supports up to 24 U.2 NVMe devices in 15 mm Gen3 carriers. The 15 mm carriers can accommodate either 7 mm or 15 mm NVMe devices.

Each NED24 NVMe Expansion Drawer includes two redundant AC power supplies, which are part of the enclosure base.

The NED24 NVMe Expansion Drawer connects to a Power server through dual CXP Converter adapters (Feature Code EJ24 or Feature Code EJ2A). A cable pair attaches the PCIe Cable Adapter (Feature Code EJ2A) to the FOM.

- ▶ Feature Code ECLX provides a pair of 3-meter optical cables with transceivers.
- ▶ Feature Code ECLY provides a pair of 10-meter optical cables with transceivers.
- ▶ Feature Code ECLS provides a pair of 3-meter copper cables.

Two cables of identical length, or one feature, are required for each FOM. Optical cables are smaller in diameter, more flexible, and can be longer than copper cables. Copper cables cost less. Both copper and optical cables offer the same performance and reliability characteristics.

The following U.2 NVMe devices are supported in the NED24 NVMe Expansion Drawer:

- ▶ Feature Code ES5B. Enterprise 800 GB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES5A. Enterprise 800 GB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5C. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5D. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES5E. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5F. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES5G. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES5H. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for IBM i

- ▶ Feature Code ES4B. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES4C. Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES4D. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES4E. Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ES4F. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for AIX and Linux
- ▶ Feature Code ES4G. Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for IBM i
- ▶ Feature Code ECT9. 15.3 TB Mainstream NVMe U.2 SSD 4k for AIX and Linux
- ▶ Feature Code ECTB. 15.3 TB Mainstream NVMe U.2 SSD 4k for IBM i

PCIe Gen4 I/O Expansion Drawer

The PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) is a 4U-high, 19-inch-wide, PCIe Gen4-based, rack-mountable I/O drawer that is available as a feature of Power11 processor-based servers.

The PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) replaces the PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0). There is no upgrade path from the PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0) to the PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0).

The PCIe Gen4 Fanout Module (Feature Code ENZF) is the FOM for placement in the PCIe Gen4 I/O Expansion Drawer. Each FOM provides six PCIe Gen4 slots (four x16 and two x8). All six slots, including the x8 slots, use x16 connectors. A fully populated PCIe Gen4 I/O Expansion Drawer with two PCIe Gen4 Fan-out Modules provides 12 PCIe Gen4 slots (eight x16 and four x8). Each PCIe Gen4 FOM connects only to PCIe x16 to CXP Converter Card features (Feature Code EJ24 or Feature Code EJ2A).

For Power S1124 servers, a cable pair attaches the PCIe Cable Adapter (Feature Code EJ2A) to the FOM:

- ▶ Feature Code ECLX provides a pair of 3-meter optical cables with transceivers.
- ▶ Feature Code ECLY provides a pair of 10-meter optical cables with transceivers.
- ▶ Feature Code ECLS provides a pair of 3-meter copper cables.

Two cables of identical length, or one feature, are required for each FOM. Optical cables are smaller in diameter, more flexible, and can be longer than copper cables. Copper cables cost less. Both copper and optical cables provide the same performance and reliability characteristics.

Consider the following items:

- ▶ The PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) can be ordered as Miscellaneous Equipment Specification (MES) and added to existing configurations that include the PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0).
- ▶ The combined maximum of NED24 NVMe Expansion Drawers (Feature Code ESR0), PCIe Gen4 I/O Expansion Drawers (Feature Code ENZ0), and PCIe Gen3 I/O Expansion Drawers (Feature Code EMX0) is limited to half of the maximum number of controller cards (Feature Code EJ24 or Feature Code EJ2A) that are allowed per server.
- ▶ A 1 EIA (1U) space is required in the IBM 7965-S42 rack or any OEM rack between the top of the PCIe Gen4 I/O Expansion Drawer and the bottom of a scale-out server (9824-22A/42A, 9786-22H/42H) or NED24 NVMe Expansion Drawer when installed directly below one of these options. This requirement applies to both factory-integrated orders (Feature Code 4651 - Feature Code 4666) and systems that are integrated in the field (Feature Code 4650). This 1U space is necessary for proper function of the cable management bracket and for serviceability of the PCIe Gen4 I/O Expansion Drawer.
- ▶ The PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) is supported by FW1050 and HMC 10.3.1050.

For more information about the I/O expansion drawers, see 3.2, “Enhancing I/O scalability with expansion drawers” on page 85.

2.3.6 Power S1124 system minimum configuration

The minimum initial order for the Power S1124 must include a processor module, two 32 GB DIMMs (one Feature Code EM54 64 GB [2 × 32 GB] DDIMM), two power supplies and power cords, an OS indicator, a cover set indicator, and a Language Group Specify. It must also include one of the following storage options and one of the following network options:

- ▶ Storage options:
 - For boot from NVMe
 - One NVMe drive slot and one NVMe drive, or one PCIe NVMe add-in adapter
 - For boot from SAN
 - An internal SSD is not required if Feature Code 0837 (boot from SAN) is selected. A FC adapter must be ordered if Feature Code 0837 is selected.
- ▶ Network options:
 - One PCIe2 4-port 1 Gb Ethernet adapter
 - One of the supported 10 Gb Ethernet adapters

Table 2-4 shows the minimum configuration when AIX or Linux are the designated OSs.

Table 2-4 Minimum configuration for Power S1124 with AIX or Linux

System feature	Feature Code	Description	Default	Minimum quantity	Notes
Op-Panel	EU0K	Operator Panel LCD display		1	Optional with AIX or Linux. The default is Qty. 1, but can be deselected for AIX or Linux.
Virtualization engine	EPVT	PowerVM Enterprise Edition	1	1	Must select one option.
	or				
	EPV0	Deactivation of LPM		1	
Processor modules	EP3X	16-core typical 3.4 - 4.2 GHz (max) Power11 processor		1	Must select the Processor Module option.
	or				
	EP3H	24-core typical 3.05 - 4.15 GHz (max) Power11 processor		2	
	or				
	EP3Y	30-core typical 2.8 - 3.95 GHz (max) Power11 processor		2	

System feature	Feature Code	Description	Default	Minimum quantity	Notes	
Processor module activations	EBF8	One core processor activation for DCM		1	A minimum of 50% of static CUoD activations must be ordered if you are not using PEP 2.0. A minimum of one activation is required if you are using Feature Code EP20 PEP 2.0 Enablement.	
Memory	EM54	64 GB (2x 32 GB) DDIMMs, 4000 MHz or 4800 MHz, and 16 Gbit DDR5 Memory		1	Minimum two DIMMs per one DIMM feature.	
	or					
	EM5B	128 GB (2x 64 GB) DDIMMs, 4000 MHz or 4800 MHz, and 16 Gbit DDR5 Memory		1		
	or					
	EM4M	256 GB (2x 128 GB) DDIMMs, 4000 MHz or 4800 MHz, and 32 Gbit DDR5 Memory		1		
	or					
EM5J	512 GB (2x 256 GB) DDIMMs, 4000 MHz or 4800 MHz, and 32 Gbit DDR5 Memory		1			
AMM	EM8G	AMM	0	0	Optional feature. Maximum of Qty. 1 per system. Memory mirroring requires a minimum of eight DIMMS (four DIMMs). Minimum of tow 2 DIMMs per one DIMM.	
Storage backplane	EJ1Y	Storage backplane with eight NVMe U.2 drive slots		1	You must order a Qty. 1 NVMe backplane feature except when Feature Code 0837 or Feature Code ESCZ (iSCSI boot) is on the order or when an NVMe PCIe add-in adapter is used as the Load Source. Mixing NVMe devices is allowed on each backplane.	

System feature	Feature Code	Description	Default	Minimum quantity	Notes
Bezels	EJBU	Front IBM bezel for 16 NVMe-bays backplane rack-mount		1	When no NVMe backplane is ordered, the default is Feature Code EJBU.
	or				
	EJUV	Front OEM bezel for 16 NVMe-bays backplane rack-mount		1	
NVMe devices	ES5A	800 GB mainstream NVMe U.2 SSD 4k for AIX or Linux	2	0	When no NVMe backplane is ordered, the default is Feature Code EJBU.
Required LAN adapters	EC72	PCIe4 LP 2-port 25/10/1 GbE RoCE SFP28 adapter	1	1	Qty. 1 of these LAN features is required on all initial orders. The default adapter is Feature Code EC72.
	or				
	EC86	PCIe5 LP 2-port 200 GbE RoCE adapter		1	
	or				
	EN2W	PCIe3 LP 4-port 10 GbE BaseT RJ45 adapter		1	
	or				
	EN2Z	PCIe LP 4-port 1 GbE adapter		1	
Power supply	EB3S	AC titanium power supply - 1600W for server (200 - 240 VAC)	2	2	Each initial order must have all power supplies present. Power supplies cannot be added later. Only 200 - 240 V power cords can be used.
Power cables	6458	Power cord 4.3 m (14 ft), drawer to IBM PDU (250 V and 10A)	2	2	Qty. 2 is required.
Language group	9300	Language Group Specify - US English	1	1	Language Specify code is required.
Primary OS	2146	Primary OS - AIX			Must select one option.
	or				
	2147	Primary OS - Linux			

2.3.7 Comparing Power S1124 and Power S1024

Table 2-5 compares the Power S1124 and Power S1024 servers.

Table 2-5 Comparing the Power S1124 and the Power S1024

Attribute	S1024 (24A)	Power S1124 DCM processors
Processor	Two Power10 DCMs (5387 pins)	Two Power11 DCMs (5387 pins)
Processor power (max)	415 W	500 W
SMP X-bus	Four 2B @ 32 Gbps	Four 2B @ 32 Gbps
Memory channels per system	32 OMI channels	32 OMI channels
Memory DDR4 data rate	2666 Mbps or 3200 Mbps per OMI channel	4000 Mbps or 4800 Mbps per OMI channel
Memory bandwidth per system (max theoretical)	<ul style="list-style-type: none"> ▶ 818 GBps with 3200 Mbps and 16 GB, 32 GB, or 64 GB 2U_DDIMM ▶ 682 GBps with 2666 Mbps 128 GB 2U_DDIMM ▶ 750 GBps with 2933 Mbps 128 GB or 256 GB 4U_DDIMM 	<ul style="list-style-type: none"> ▶ 1228 GBps with 4800 Mbps 32 GB DDIMMs ▶ 2456 GBps with 4800 Mbps 64 GB DDIMMs ▶ 2048 GBps with 4000 Mbps 128 GB or 256 GB DDIMMs
DIMMs per system	32	32
Memory capacity per system (max)	8 TB with 4U DIMM	8 TB with 4U DIMM
OpenCAPI port	Six ports @ 25 Gbps	0
OpenCAPI bandwidth per system (peak)	300 GBps	N/A
PCIe lanes per system (max)	128 PCIe G4 lanes @ 16 Gbps	128 PCIe G4 lanes @ 16 Gbps
PCIe slots per system	<ul style="list-style-type: none"> ▶ Four PCIe G4 x16 or G5 x8 slots ▶ Four PCIe G5 x8 slots ▶ Two PCIe G4 x8 slots 	<ul style="list-style-type: none"> ▶ Four PCIe G4 x16 or G5 x8 slots ▶ Four PCIe G5 x8 slots ▶ Two PCIe G4 x8 slots
Slots for internal storage	General	General
Drives (max)	16 NVMe U.2	16 NVMe U.2
I/O expansion drawer (max)	2	2
Service interface	Enterprise Baseboard Management Controller (eBMC)	eBMC

2.4 Processor and memory activation on Power S1122 and Power S1124 systems

The IBM Power S1122 and Power S1124 systems support Capacity on Demand (CoD) activations for all DCM or eSCM-based processor features. However, all installed memory is automatically activated upon installation.

Note: CUoD applies only to processor features. Memory features are fully activated when installed.

Processor activation options

IBM offers several processor activation types to support different workloads and pricing models:

- ▶ Static processor activations
 - They are permanent, and support any application environment.
 - At least 50% of installed cores must be statically activated at installation unless the system is configured for PEP 2.0 (Feature Code EP20), in which case a reduced minimum of one processor core activation is allowed.
 - More cores can be activated later.
- ▶ Base processor activations (PEP 2.0) (Feature Code EP20)
 - Ordered for a specific server but transferable within a Power pool.
 - Support any OS and application type.
 - Require a minimum of one activation and a maximum equal to the total number of processor cores in the system.
- ▶ CUoD static activations
 - Available for any processor module and OS.
 - Provide flexibility to scale core usage over time.

Shared Utility Capacity with Power Enterprise Pools 2.0

Power S1122 and Power S1124 systems support Shared Utility Capacity within PEP 2.0. This feature enables enhanced multi-system resource sharing and minute-by-minute tracking of compute resource consumption across a pool of systems. It enables organizations to tailor their infrastructure with a mix of purchased and pay-per-use processor and software resources. By consolidating infrastructure onto Power S1122 and Power S1124 systems, IT teams can simplify system management and focus on optimizing business outcomes rather than manually reallocating resources.

Shared Utility Capacity resources are tracked by virtual machine (VM) and monitored by the IBM Cloud Management Console (IBM CMC), which integrates with local HMCs. This setup eliminates the need to overprovision capacity on individual systems because all processor and memory resources across the pool are activated and available for use. When a pool is initiated, each system's purchased processor activations, memory, and OS entitlements become part of the pool's Base Capacity. Any extra installed resources beyond this base are considered Metered Capacity, which is activated at pool startup and monitored by the minute. Clients are charged only for metered usage that exceeds the pool's aggregate Base Capacity, with charges deducted from Capacity Credits (5819-CRD) or billed monthly, depending on availability.

Licensing and pool configuration guidelines

Pools that are created after July 12, 2022, require consistent IBM i license tiers across all systems. Clients with existing pools (Power S1022 and Power S1024) can add Power S1122 systems and migrate applications at their own pace.

Consider the following items for IBM i workloads:

- ▶ Power S1122 systems with up to 16 cores (P20) can join pools with P20-tier systems.
- ▶ Power S1122 systems with 48 or 60 cores (P30) require all systems in the pool to have P30-tier IBM i licenses.

2.4.1 Advanced memory functions on Power11 scale-out servers

The Power11 scale-out servers provide two features that improve the usability of the installed memory. AMM increases resilience by mirroring critical memory that is used by the PHYP. Active Memory Expansion (AME) uses compression and decompression functions to expand the usable memory capacity that is allocated to an AIX partition.

Active Memory Mirroring

AMM for the PHYP is available as an optional feature (Feature Code EM8G) to improve resilience by mirroring critical memory that is used by the PHYP. This capability enables the hypervisor to continue operating if there is a memory failure. A portion of the available memory can be proactively partitioned so that a duplicate set is available for use if a non-correctable memory error occurs. This feature can be implemented at the granularity of DIMMs or logical memory blocks.

Active Memory Expansion (Feature Code EMBP)

AME is an innovative technology that supports the AIX OS and enables the effective maximum memory capacity to exceed the physical memory limit. Compression and decompression of memory content can expand memory capacity by up to 100% or more. This capability enables a partition to perform more work or support more users with the same physical memory. Similarly, it enables a server to run more partitions and handle more workloads without more physical memory.

AME uses CPU resources to compress and decompress memory contents. The tradeoff between memory capacity and processor cycles can be advantageous, but the degree of expansion depends on how compressible the memory content is and on the availability of spare CPU capacity for compression and decompression.

Power11 processors include a hardware accelerator that improves AME efficiency and reduces the use of Power core resources. The Power11 accelerator includes minor enhancements and uses the higher bandwidth and reduced latency characteristics of Power11 processors.

You have extensive control over AME usage. You can enable or disable AME on each AIX partition. Control parameters define the wanted expansion level for each partition to manage CPU usage by the AME function. An IPL is required for the partition that enables memory expansion. After activation, monitoring capabilities are available through standard AIX performance tools, such as `lparstat`, `vmstat`, `topas`, and `svmon`.

A planning tool that is included with AIX enables you to sample workloads and estimate both the potential memory expansion and the CPU resources that are required. Any Power server model can run the planning tool. In addition, a one-time, 60-day trial of AME is available to provide accurate memory expansion and CPU usage measurements. You can request the trial through the [Power CoD website](#).

AME is enabled by a chargeable hardware feature (Feature Code EMBP), which can be ordered with the initial system node or as a MES order. A software key is provided when the enablement feature is ordered and applied to the system node. An IPL is not required to enable the system node. The key is specific to an individual system node and is permanent. It cannot be transferred to another server.

The additional CPU resources that are used for memory expansion are part of the CPU resources that are assigned to the AIX partition running AME. Normal licensing requirements apply.

2.5 Environmental and physical specifications

This section describes the environmental requirements and physical characteristics of IBM Power scale-out servers. It includes key details such as voltage requirements, thermal output, and cooling needs to help ensure proper operation in data center environments. It also provides the dimensions and weight of each server model to help with planning for rack installation, transportation, and serviceability.

Voltage

- ▶ The Power S1122 has an AC-rated voltage and frequency of 200 - 240 V AC at 50 Hz or 60 Hz plus or minus 3 Hz.
- ▶ The Power S1124 has an AC-rated voltage and frequency of 200 - 240 V AC at 50 Hz or 60 Hz plus or minus 3 Hz.

Planning for electrical power

Consider the following points when planning electrical power for a Power scale-out server:

- ▶ Power supply redundancy is supported.
- ▶ The Power S1122 supports a maximum of two power supplies. There are no specific plugging rules or sequences when connecting the power supplies to the rack PDUs. All power supplies feed a common DC bus.
- ▶ The Power S1124 supports a maximum of four power supplies. There are no specific plugging rules or sequences when connecting the power supplies to the rack PDUs. All power supplies feed a common DC bus.
- ▶ The power supplies automatically accept any voltage within the published rated-voltage range. If multiple power supplies are installed and operating, they draw approximately equal current from the utility and provide approximately equal current to the load.
- ▶ Power draw and heat load vary by configuration. When planning the electrical system, use the maximum values. For heat load planning, use the [IBM Systems Energy Estimator](#) to obtain an estimate based on a specific configuration.
- ▶ To calculate amperage, multiply the kVA by 1,000 and divide the result by the operating voltage.

Thermal output

Here are the thermal outputs of the servers:

- ▶ Power S1122 and Power L1122
 - Thermal output (maximum). 7643 BTU/hr
 - Maximum power consumption. 2240 W
 - Maximum kVA. 2.31 kVA
 - Phase. Single
- ▶ Power S1124 and Power L1124
 - Thermal output (maximum). 9383 BTU/hr
 - Maximum power consumption. 2750 W
 - Maximum kVA. 2.835 kVA
 - Phase. Single

Operating environment

The allowable temperature and humidity ranges for both operating and non-operating are described in this section.

Environment (operating)

- ▶ ASHRAE class
 - Allowable. A3 (fourth edition)
- ▶ Airflow direction
 - Recommended. Front-to-back
- ▶ Temperature
 - Recommended. 18.0°C to 27.0°C (64.4°F to 80.6°F)
 - Allowable. 5.0°C to 40.0°C (41.0°F to 104.0°F)
- ▶ Low-end moisture
 - Recommended. 9.0°C (15.8°F) dew point
 - Allowable. -12.0°C (10.4°F) dew point and 8% relative humidity
- ▶ High-end moisture
 - Recommended. 60% relative humidity and 15°C (59°F) dew point
 - Allowable. 85% relative humidity and 24.0°C (75.2°F) dew point
- ▶ Maximum altitude
 - 3,050 m (10,000 ft)

Allowable environment (non-operating)

- ▶ Temperature
 - Recommended. 5°C to 45°C (41°F to 113°F)
- ▶ Relative humidity
 - Recommended. 8% to 85%
- ▶ Maximum dew point
 - Recommended. 27.0°C (80.6°F)

Physical specifications

Table 2-6 provides the physical specifications (dimensions and weight) of the scale-out servers.

Table 2-6 Physical specifications for the scale-out servers

Specification	2S4U Power S1124 and Power L1124	2S2U Power S1122 and Power L1122
Max weight (lbs)	96 (43.5 kg)	70 (32.2 kg)
Height (mm)	173 (6.8")	86.5 (3.4")
Width, overall (mm) ^a	482 (18.97")	482 (18.97")
Width within rack (mm) ^b	446 (17.6")	446 (17.6")
Depth (mm) ^c	712 (28")	813 (32")
Rack mountable?	Yes (4U)	Yes (2U)

a. The width is measured to the outside edges of the rack-mount bezels.

b. The width is 446 mm (17.6 in.) for the main chassis, which fits in between a 482.6 mm (19 in.) rack-mounting flanges.

c. The cable management arm with the maximum cable bundle adds 248 mm (9.8 in.) to the depth.



I/O subsystem

The internal I/O system of IBM Power systems is designed for high-speed, low-latency data movement to help ensure optimal performance for enterprise workloads. At the core of this architecture is the IBM Power11 processor, which integrates advanced I/O subsystems that support high throughput and efficient communication between processors, memory, and peripheral devices. IBM Power servers use technologies such as Peripheral Component Interconnect Express (PCIe) Gen5 and Non-Volatile Memory Express (NVMe) to deliver fast data access and reduce bottlenecks.

IBM Power systems also use a modular and scalable I/O architecture that allows flexible configuration and expansion. The system supports Single Root I/O Virtualization (SR-IOV) for efficient sharing of I/O resources across virtual machines (VMs). PowerVM virtualization technology enables dynamic allocation of I/O bandwidth to match workload demands.

This internal I/O design enhances performance and improves system reliability and availability by incorporating features such as redundant paths, hot-swappable components, and advanced error detection and correction. These capabilities make the IBM Power internal I/O system a key enabler of high-throughput, resilient, and scalable enterprise computing. This chapter describes the I/O subsystem for the IBM Power S1122 and IBM Power S1124 servers.

This chapter contains the following topics:

- ▶ Internal I/O
- ▶ Enhancing I/O scalability with expansion drawers
- ▶ Available adapter list
- ▶ NVMe U.2 devices
- ▶ Other device support

3.1 Internal I/O

The IBM Power scale-out servers (Power S1122, IBM Power L1122, Power S1124, and IBM Power L1124) are built on the Power11 architecture and feature an advanced internal I/O subsystem that is designed to support high-throughput, low-latency data movement within the server. At the core of the internal I/O is the Power11 processor's Open Memory Interface (OMI), which decouples memory from the processor and allows for flexible, high-bandwidth memory configurations. This architecture supports DDR5 memory through memory buffer chips, enabling faster data access and improved scalability. The scale-out servers also integrate PCIe Gen5 technology, which doubles the bandwidth of previous generations and enables faster communication between internal components such as storage controllers and network adapters.

The internal I/O design of the Power11 scale-out servers emphasizes reliability and serviceability. The system supports NVMe drives natively, which enhances internal storage performance, and features a modular design that allows for expansion through I/O drawers. These drawers connect through high-speed links and can house more PCIe slots, enabling the system to scale with growing workload demands. Overall, the internal I/O architecture of the Power11 scale-out servers is designed to deliver exceptional performance, flexibility, and resilience for mission-critical applications.

The internal I/O subsystem of the Power11 scale-out servers connects to PCIe Express controllers on the Power11 processors in the system. Each processor features two PCI Express controllers (PECs), PEC0 and PEC1, which support up to three PCI host bridges that directly connect to a PCIe slot or device.

In the Power11 scale-out servers, the Power11 processor is packaged as either a Dual-Chip Module (DCM) or an entry Single-Chip Module (eSCM). These modules contain two Power11 processors that are integrated into a single component that plugs into a socket on the system board.

The DCM and the eSCM are described in 1.5.2, "Comparing Power11, Power10, and Power9 processors" on page 38. The DCM is a module in which both Power11 processors are fully functional and provide cores, memory, and I/O. In an eSCM, only the first chip (P0) is fully functional with cores, memory, and I/O; the second chip (P1) supports I/O only.

All modules that are used in the 2S4U versions of the servers (Power S1124 and Power L1124) are DCMs. The 2S2U versions (Power S1122 and Power L1122) use either DCMs or eSCMs, depending on the processor feature chosen.

Table 3-1 shows the module type based on the processor feature that is chosen.

Table 3-1 Module type by feature code

Feature Code Power S1122 and Power L1122	Number of cores	Module type	Frequency range
EP3Y and EP4J	30 + 2 spare	DCM	2.4 GHz - 3.95 GHz
EP3H and EP4B	24 + 2 spare	DCM	2.65 GHz - 4.15 GHz
EP3 and EP4C	16 + 2 spare	DCM	3.00 GHz - 4.2 GHz
ERGQ/ N/A	10	eSCM	3.05 GHz - 4.00 GHz
ERGR/ an	4	eSCM	3.60 GHz - 4.00 GHz

3.1.1 Internal PCIe Gen 5 subsystem

All IBM Power11 2S2U and 4S4U scale-out server models include 10 physical PCIe adapter slots in the chassis. The number of slots that are available for use depends on the number of processor sockets that are populated in the system. The number of available slots is the same for both DCM- and eSCM-based processor features:

- ▶ With one processor socket populated, five PCIe adapter slots are available.
- ▶ With two processor sockets populated, all 10 PCIe adapter slots are available.

These internal PCIe adapter slots support a range of adapters. For more information, see 3.3, “Available adapter list” on page 98.

The adapter slots are a mix of PCIe Gen5 and PCIe Gen4 slots, with some operating at x8 speed and others at x16 speed. All PCIe adapter slots support hot-plug capability when used with Hardware Management Console (HMC) or enterprise Baseboard Management Controller (eBMC)-based maintenance procedures.

One extra slot is available at the rear of each server. This slot is dedicated to the eBMC management controller and cannot be used for any other PCIe adapter type.

Each system requires at least one local area network (LAN) adapter to support a connection to local networks. This requirement enables initial system testing and configuration and the preinstallation of any operating systems (OSs), if required. The required network adapter is installed by default in slot C10.

All PCIe slots support hot-plug adapter installation and maintenance when service procedures are activated through the eBMC or HMC interfaces and enhanced error handling (EEH). PCIe EEH-enabled adapters respond to a special data packet that is generated by the affected PCIe slot hardware by calling the system firmware, which examines the affected bus, enables the device driver to reset it, and continues without a system restart. For Linux, EEH support extends to most frequently used devices, although some third-party PCI devices might not provide native EEH support.

All PCIe adapter slots support hardware-backed network virtualization through SR-IOV technology. Configuring an SR-IOV adapter into SR-IOV shared mode might require more hypervisor memory. If sufficient hypervisor memory is not available, the request to move to SR-IOV shared mode fails. The user is then instructed to free extra memory and attempt the operation again.

The servers optimize energy efficiency when cooling the PCIe adapter environment. They detect which IBM PCIe adapters are installed in their PCIe slots and, if an adapter requires higher levels of cooling, automatically increase fan speed to improve airflow across the PCIe adapters. Faster fans increase the sound level of the server.

PCIe slot capabilities

If both sockets in a server are populated, the system supports a maximum of 10 PCIe adapters. If only one socket is populated, the maximum number of PCIe slots is 5.

Table 3-2 lists the adapter slots that are available in the Power11 scale-out servers that use DCM modules with either one or two processor modules installed.

Table 3-2 PCIe slot details for Power 2S2U and 2S4U servers with DCM-based processors

Adapter slot	Type	Sockets that are populated
P0-C0	PCIe4 x16 or PCIe5 x8 slots	2
P0-C1	PCIe4 x8 with x16 connector	2
P0-C2	PCIe5 x8 with x16 connector	2
P0-C3	PCIe4 x16 or PCIe5 x8 slots	2
P0-C4	PCIe4 x16 or PCIe5 x8 slots	2
P0-C5 ^a	eBMC	All systems
P0-C7	PCIe5 x8 with x16 connector	1 or 2
P0-C8	PCIe4 x8 with x16 connector	1 or 2
P0-C9	PCIe5 x8 with x16 connector	1 or 2
P0-C10	PCIe4 x16 or PCIe5 x8 slots	1 or 2
P0-C11	PCIe5 x8 with x16 connector	1 or 2

a. Used for eBMC only.

Table 3-3 shows the PCIe slot configuration when a Power S1122 system is populated with eSCM processor modules.

Table 3-3 PCIe slot details for Power S1122 servers with eSCM based processors

Adapter slot	Type	Sockets that are populated
P0-C0	PCIe4 x16	2
P0-C1	PCIe4 x8 with x16 connector	2
P0-C2	PCIe4 x8 with x16 connector	2
P0-C3	PCIe4 x16	2
P0-C4	PCIe4 x16	2
P0-C5 ^a	eBMC	All systems
P0-C7	PCIe4 x8 with x16 connector	1 or 2
P0-C8	PCIe4 x8 with x16 connector	1 or 2
P0-C9	PCIe4 x8 with x16 connector	1 or 2
P0-C10	PCIe4 x16	1 or 2
P0-C11	PCIe4 x8 with x16 connector	1 or 2

a. Used for eBMC only.

x16 slots provide up to twice the bandwidth of x8 slots because they offer twice as many PCIe lanes. PCIe Gen5 slots support up to twice the bandwidth per lane of PCIe Gen4 slots, and PCIe Gen4 slots support up to twice the bandwidth per lane of PCIe Gen3 slots.

Slot locations

The Power S1124 and Power L1124 servers are 4U (EIA units), and support the installation of full-height PCIe adapters.

Figure 3-1 shows the PCIe adapter slot locations for these 2S4U server models.




Figure 3-1 PCIe adapter slot locations on the Power S1124 and Power L1124 servers

The Power S1122 and Power L1122 servers are 2U (EIA units), and support the installation of low-profile PCIe adapters.

Figure 3-2 shows the PCIe adapter slot locations for these 2S2U server models.



Figure 3-2 PCIe adapter slot locations on the Power S1122 and Power L1122 server models

More PCIe slots

The total number of PCIe adapter slots can be increased by adding PCIe I/O expansion drawers. Each I/O expansion drawer supports one or two fan-out modules (FOMs). Each FOM provides six more PCIe slots. Connecting a FOM in a PCIe expansion drawer requires installing a PCIe optical cable adapter in one of the internal PCIe x16 adapter slots (C0, C3, C4, or C10).

For more information, see 3.2, “Enhancing I/O scalability with expansion drawers” on page 85.

Direct memory access space allocation

The server PCIe slots are allocated direct memory access (DMA) space by using the following algorithm:

- ▶ All slots are allocated a 2 GB default DMA window.
- ▶ All I/O adapter slots, except the embedded Universal Serial Bus (USB), are allocated Dynamic DMA Window (DDW) capability based on installed platform memory.

DDW capability is calculated by assuming 4 KB I/O mappings. Consider the following points:

- ▶ For systems with less than 64 GB of memory, slots are allocated 16 GB of DDW capability.
- ▶ For systems with at least 64 GB but less than 128 GB of memory, slots are allocated 32 GB of DDW capability.
- ▶ For systems with 128 GB or more of memory, slots are allocated 64 GB of DDW capability.
- ▶ Slots can be enabled with Huge Dynamic DMA Window (HDDW) capability by using the I/O Adapter Enlarged Capacity setting in the Advanced System Management Interface (ASMI).
- ▶ HDDW-enabled slots are allocated enough DDW capability to map all installed platform memory by using 64 KB I/O mappings.
- ▶ The minimum DMA window size for HDDW-enabled slots is 32 GB.
- ▶ Slots that are HDDW-enabled are allocated the larger of the calculated DDW and HDDW capability.

3.1.2 Internal NVMe storage subsystem

In both the IBM Power S1122 and Power S1124 systems, NVMe drives are the exclusive internal storage option, delivering high-performance, low-latency capabilities that are crucial for demanding workloads. These drives use a 15 mm carrier in a U.2 2.5-inch form factor and are designed for direct integration into the system's front slots.

The connectivity of these NVMe drives is native PCIe. This direct connection to the Power11 processor architecture, which uses PCIe Gen5, helps ensure maximum bandwidth and minimal latency. This capability is essential for applications that require rapid data access, such as databases, analytics, and artificial intelligence (AI) inferencing. The Power S1122 and Power S1124 systems are designed to fully use the performance advantages of NVMe by connecting these flash devices directly to high-speed PCIe lanes.

For advanced configurations, especially in virtualized environments that use PowerVM, the NVMe drives can be managed and allocated through the Virtual I/O Server (VIOS). This approach allows flexible sharing of NVMe resources across multiple logical partitions (LPARs) while maintaining high performance. For enhanced redundancy and availability, the Power S1122 and Power S1124 systems can support multiple NVMe backplanes and controllers (for example, through features such as the NVMe Just a Bunch of Flash (JBOF) adapter). These configurations enable NVMe drives to be split between redundant VIOS instances to eliminate single points of failure for critical internal storage. This comprehensive approach to NVMe connectivity helps ensure that these Power servers deliver both extreme performance and enterprise-grade reliability for internal storage.

Connectivity for internal NVMe drives

The NVMe drives connect to the system through one or two backplanes that connect four or eight NVMe devices to one of the system's PCIe slots. Each installed backplane attaches to an NVMe JBOF adapter, which contains a PCIe Gen4 switch. Each NVMe JBOF adapter requires a PCIe slot in the system.

The NVMe JBOF card in the Power S1122 supports up to four NVMe slots by using two cables. Each NVMe JBOF card in the Power S1124 supports up to eight NVMe slots by using four cables. This difference in cabling explains why the Power S1122 supports a maximum of eight drives and the Power S1124 supports a maximum of 16 drives.

Table 3-4 summarizes the available internal storage options.

Table 3-4 Internal storage capabilities

Feature	2S2U	2S4U
NVMe 4-device backplane	1 or 2 Up to 8 devices	N/A
NVMe 8-device backplane	N/A	1 or 2 Up to 16 devices
NVMe U.2 7 mm device (max of 4)	800 GB	
NVMe U.2 15 mm devices	0.8, 1.6, 3.2, 6.4, or 15.3 TB	
Concurrently maintainable NVMe	Yes	

3.1.3 NVMe backplane

Different backplanes are provided for the 2S2U and 2S4U servers:

► 2S2U

The storage backplane that supports four NVMe U.2 drive slots (Feature Code EJ1X) is the base storage backplane. You can install one or two backplanes to support up to eight devices in the system.

► 2S4U

The storage backplane that supports eight NVMe U.2 drive slots (Feature Code EJ1Y) is the base storage backplane. You can install one or two backplanes to support up to 16 devices.

3.1.4 NVMe JBOF card

The internal NVMe storage connects to the processor through a plug-in PCIe NVMe JBOF card. The card is installed in designated PCIe slots. Up to two cards can be installed in either the Power S1122 or Power S1124 system. There is a one-to-one correspondence between each JBOF card and an NVMe backplane.

- In the 2S4U system, each JBOF card contains four connectors that are cabled to four connectors on a single eight-drive backplane. Each cable provides signaling for two NVMe drives.
- In the 2S2U system, each JBOF card contains two connectors that are cabled to two connectors on a single four-drive backplane. Each cable provides signaling for two NVMe drives.

The NVMe JBOF card contains a 52-lane PCIe Gen4 switch. The connected NVMe devices are individually addressable and can be allocated individually to LPARs running on the system.

Figure 3-3 shows the JBOF card in the 2S4U system with four connectors to the eight-port backplane. Two JBOF cards are supported in the 2S4U server.

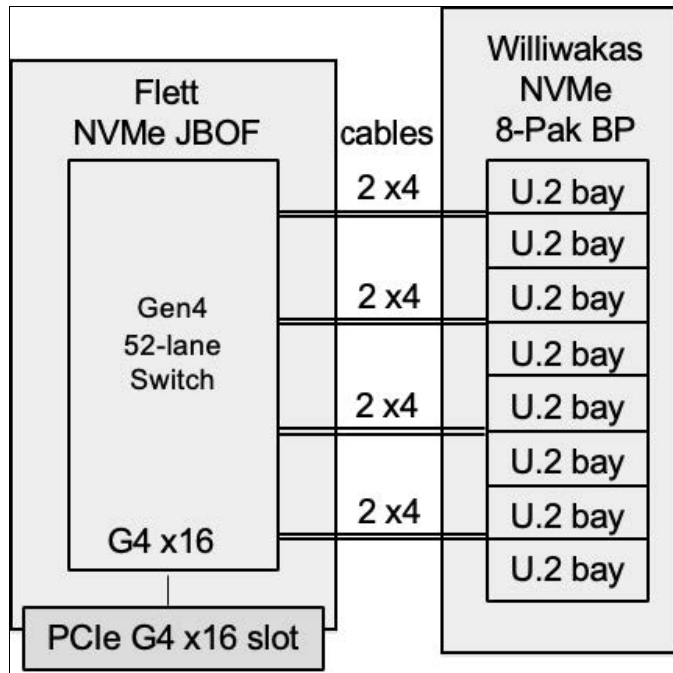


Figure 3-3 4U version of the NVMe JBOF card

Similarly, Figure 3-4 shows the JBOF card in the 2S2U. For this system, only two cables are connected to the four-drive backplane. Two of the JBOF cards are supported on the 2S2U system.

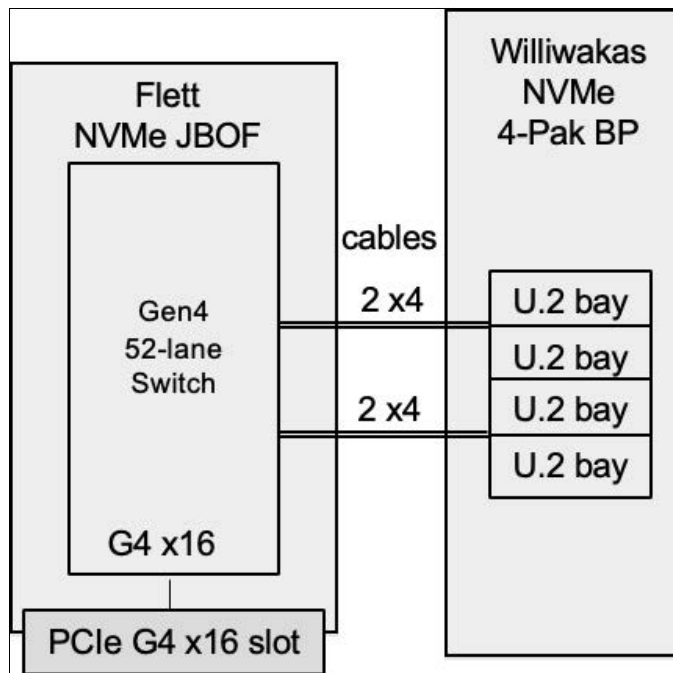


Figure 3-4 2U version of the NVMe JBOF card

The NVMe JBOF card is treated as a regular cable card, with error handling support that is similar to the support of a system board switch. The card is not concurrently maintainable because of the cabling that is required to the NVMe backplane.

Table 3-5 shows the available features by model.

Table 3-5 NVMe features

Feature	CCIN	FRU	Config	Feature Code for 2S4U	Feature Code for 2S2U
NVMe JBOF Card (4U)	6B87	Yes	Feature	EJ1Y	N/A
NVMe U.2 8-Pak BP (4U)	6B89	Yes	Feature		
NVMe JBOF Card (2U)	6B87	Yes	Feature	N/A	EJ1X
NVMe U.2 4-Pak BP (2U)	6B89	Yes	Feature		

3.1.5 NVMe JBOF to backplane cabling

Three PCIe slots support NVMe JBOF cards: C8, C10, and C11. Only two of these PCIe slots can be connected to JBOF cards at a time. PCIe slots C8 and C10 can be cabled only to NVMe backplane P1, and C11 can be cabled only to NVMe backplane P2. C10 is the default location for the JBOF that is connected to backplane P1, but the JBOF can be moved to C8 if another PCIe card must be installed in C10. You cannot install a JBOF card in both slots C8 and C10.

Figure 3-5 shows the cabling options.

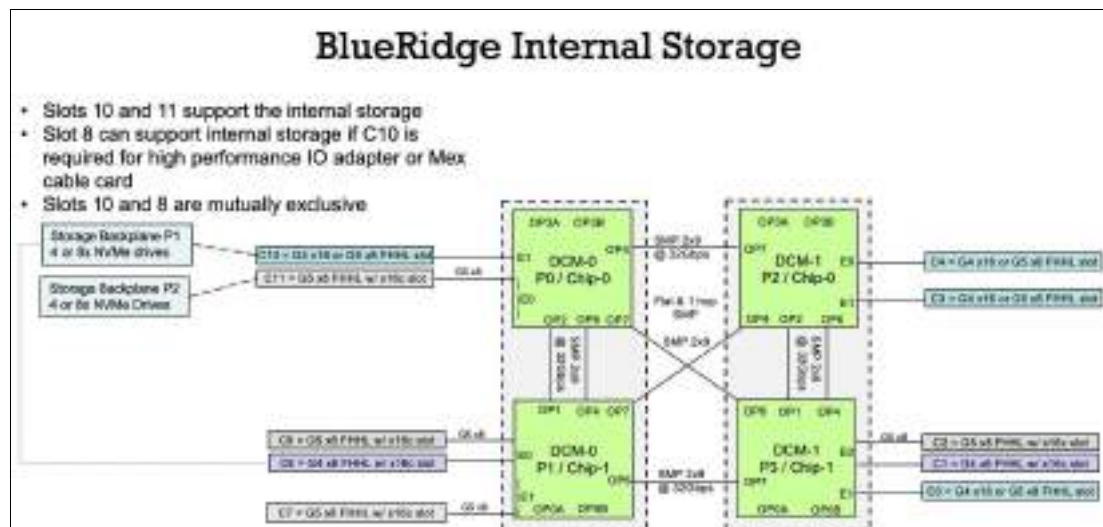


Figure 3-5 Cabling diagram for the NVMe backplane

Table 3-6 shows how the NVMe JBOF card is placed in the PCIe slots under various configurations.

Table 3-6 PCIe card slots for JBOF cards

NVMe backplanes			
	P1 (left)	P2 (middle)	Notes
Default plugging	C10	C11	The NVMe cards are plugged in to the x16 slot first by default plug order: C10, C11, and C8.
A x16 adapter is required in C10.	C8	C11	The plug order is C8 and C11.

Figure 3-6 shows the connector numbering on the NVMe JBOF card (left) and the NVMe backplane (right). Each connector on the JBOF connects to the corresponding connector on the backplane.

- ▶ C0 provides signaling for NVMe drives 0 and 1.
- ▶ C1 provides signaling for drives 2 and 3.
- ▶ C2 provides signaling for drives 4 and 5.
- ▶ C3 provides signaling for drives 6 and 7.

In the Power S1122 systems, only C1 and C2 are connected. The other connections on both the JBOF and the backplane remain unconnected.




Figure 3-6 Connector locations for JBOF card (left) and NVMe backplane (right)

3.2 Enhancing I/O scalability with expansion drawers

Adding I/O drawers to IBM Power11 2S2U and IBM Power11 2S4U servers enhances scalability, flexibility, and overall performance in enterprise environments. The Power11 2S2U and Power11 2S4U servers support the following options:

- ▶ PCIe Gen4 I/O Expansion Drawers

These drawers provide increased throughput and bandwidth for connected devices, enabling the system to handle more demanding workloads and a broader range of peripheral devices. They expand the available PCIe slots beyond what is natively supported on the system board, enabling more network adapters, storage controllers, and accelerators to be integrated without compromising existing configurations. This capability is beneficial for data-intensive applications such as AI inferencing, high-speed networking, and large-scale database operations.

- ▶ NED24 NVMe Expansion Drawers

These drawers provide more NVMe capacity for expanding the direct-attached storage of the system.

The modular nature of I/O expansion drawers supports a pay-as-you-grow model, aligning with dynamic business needs and reducing upfront capital expenditure. By expanding the number and type of adapters and the amount of internal storage that can be installed, these drawers enable more flexible system configurations. This flexibility is especially beneficial in virtualized environments where multiple VMs or containers require dedicated I/O resources.

More I/O expansion supports higher VM density and more granular resource allocation, improving overall system usage. IBM expansion drawers are designed with enterprise-grade reliability features, including hot-plug capabilities, redundant paths, and integration with PowerVM and the HMC. These features help ensure that I/O resources can be added or serviced without system downtime, supporting continuous operations and minimizing the risk of service disruptions.

The I/O expansion capability aligns with the broader design goals of the Power S1122 and Power S1124 systems: resilience, security, and performance. The system architecture, including its I/O subsystem, is built to support mission-critical workloads with features such as redundant paths, error recovery, and Secure Boot. The drawers integrate seamlessly with the HMC and PowerVM virtualization, helping ensure centralized control and simplified management. This integration makes the Power11 2S2U and Power11 2S4U servers with I/O drawers an ideal platform for enterprises that want to modernize their infrastructure while maintaining high availability (HA) and robust performance.

3.2.1 PCIe expansion card

The PCIe expansion function is enabled by using the PCIe Gen4 cable adapter (Feature Code EJ24). Each adapter occupies one PCIe Gen5 slot in the system node. A maximum of eight adapters can be installed in a system node and up to 32 in a system with four system nodes. Each adapter connects to either one Feature Code ENZ0 FOM or one NED24 Enclosure Services Manager (ESM). The cable adapter can be installed in any slot in the system node. Table 3-9 on page 90 shows for the adapter locations based on the expansion drawer being attached.

Figure 3-7 shows the PCIe Gen4 cable adapter (Feature Code EJ24).




Figure 3-7 PCIe Gen4 cable adapter with Feature Code EJ24

Feature Code EJ2A (CCIN 6B99) is a PCIe Gen4 x16 cable adapter that provides high-speed connectivity between the system and supported I/O expansion drawers. It is a full-height, half-length card with two ports for connecting expansion drawer cables. The adapter can be installed in either PCIe Gen4 x16 or x8 slots. Using x8 slots offers greater flexibility for connecting more drawers but delivers reduced performance compared to installation in full x16 slots.

This adapter supports the attachment of the following components:

- ▶ One ESM in an NED24 NVMe Expansion Drawer
- ▶ One PCIe4 six-slot FOM in ENZ0 PCIe4 Expansion Drawer

The EJ2A adapter is a key component for extending the PCIe bus from IBM Power scale-out servers to external I/O resources, enabling scalable configurations for high-performance workloads. It includes two status LEDs: a green LED that indicates link status and an amber LED used for identifying the adapter during maintenance activities.

Each EJ2A adapter supports two CXP interface ports, which can use either optical or copper cables to connect to PCIe Gen4 expansion drawers. A single adapter connects to one FOM within an expansion drawer. Therefore, to use both FOMs in a drawer, two EJ2A adapters are required. When connecting an NED24 NVMe Expansion Drawer, two adapters are also necessary for full connectivity.

Internally, the EJ2A includes a built-in PCIe switch that extends the IBM Power E1080 internal PCIe Gen4 bus to external I/O drawers while maintaining bandwidth and performance integrity.

Figure 3-8 on page 87 shows the adapter's physical layout and connectivity.

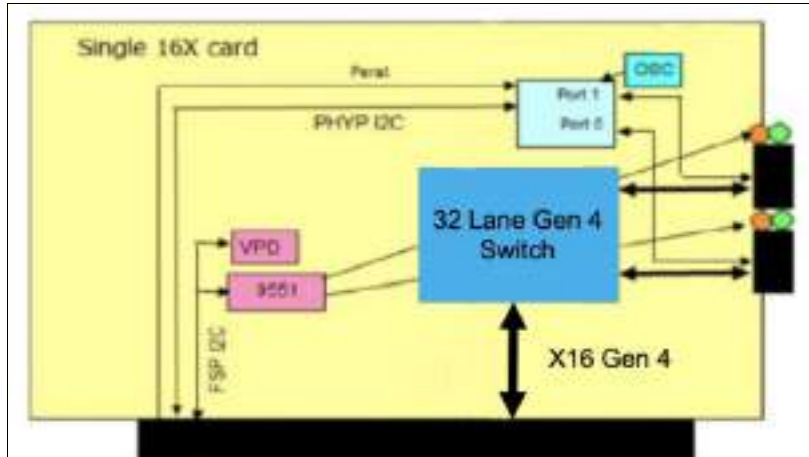


Figure 3-8 PCIe expansion card

Supported PCIe4 cable adapters

Table 3-7 shows the supported cable card adapters for Power11 processor-based servers.

Table 3-7 Supported cable card adapters by system type

System name	Adapter Feature Code and CCIN
Power S1124 and Power L1124	PCIe4 cable adapter (Feature Code EJ2A CCIN 6B99)
Power S1122 and Power L1122	PCIe4 cable adapter (Feature Code EJ24 CCIN 6B92)

Copper cable versus active optical cable

The low-profile cable adapter (Feature Code EJ24 CCIN 6B98) is designed for 2U systems and supports copper cables only because of thermal constraints. The full-height cable adapter (Feature Code EJ2A CCIN 6B99) is designed for 4U systems, supports both copper and active optical cables (AOCs), and is fully supported on the Power S1124 and Power L1124 servers.

Table 3-8 lists the supported cable adapter feature codes for Power11 processor-based servers.

Table 3-8 Supported PCIe cable adapters on Power11 processor-based servers

Feature Code	CCIN	Form factor	Cable support	Supported systems
EJ24	6B98	Low-profile (2U)	Copper only	Power S1122 and Power L1124
EJ2A	6B99	Full-height (4U)	Copper and AOCs	Power S1124 and Power L1124

The choice of copper or optical cables depends on the specific configuration and the required cable length. Use AOC cables for thinner and longer cable requirements, such as Feature Code ECLX (3 m optical) and Feature Code ECLY (10 m optical). Each cable Feature Code includes two cables.

Active optical cables

An AOC consists of a fiber cable and two active electrical-to-optical converter modules that are combined into one assembly (a CXP converter). The AOC has an allowed bend radius of 1 inch (25 mm).

Note: AOC cables are not supported on the 2S2U models because of thermal restrictions with the EJ24 adapter.

Figure 3-9 shows the AOC connector.



Figure 3-9 AOC connector

The following cables are available:

- ▶ 3.0 m CXP x16 Copper Cable Pair for PCIe4 Expansion Drawer (Feature Code ECLS)
This 3.0 m cable pair connects a PCIe4 FOM in the PCIe Gen4 I/O Expansion Drawer to a PCIe4 optical converter adapter in the system unit. The pair includes two identical copper cables, each with two CXP connectors. One cable attaches to the top CXP port of the PCIe4 FOM and the top CXP port of the PCIe4 optical converter adapter. The other cable attaches to the bottom CXP ports.
- ▶ 3.0 m Active Optical Cable x16 Pair for PCIe4 Expansion Drawer (Feature Code ECLX)
This 3.0 m AOC x16 pair connects a PCIe4 module in the PCIe Gen4 I/O Expansion Drawer to a PCIe4 optical converter adapter in the system unit. The pair includes two identical cables, each with two CXP connectors. One cable attaches to the top CXP ports, and the other attaches to the bottom CXP ports.

- ▶ 10 m Active Optical Cable x16 Pair for PCIe4 Expansion Drawer (Feature Code ECLY)

This 10 m AOC x16 pair connects a PCIe4 module in the PCIe Gen4 I/O Expansion Drawer to a PCIe4 optical converter adapter in the system unit. The pair includes two identical cables, each with two CXP connectors. One cable attaches to the top CXP ports, and the other attaches to the bottom CXP ports.
- ▶ 20 m Active Optical Cable x16 Pair for PCIe4 Expansion Drawer (Feature Code ECLZ)

This 20 m AOC x16 pair connects a PCIe4 module in the PCIe Gen4 I/O Expansion Drawer to a PCIe4 optical converter adapter in the system unit. The pair includes two identical cables, each with two CXP connectors. One cable attaches to the top CXP ports, and the other attaches to the bottom CXP ports.

Note: Consider the following points:

- ▶ Use the 3 m cables for intra-rack installations and the 10 m cables for inter-rack installations.
- ▶ You cannot mix copper and optical cables on the same PCIe Gen4 I/O drawer. Both FOMs must use copper cables or both must use optical cables.

A PCIe Gen4 I/O Expansion Drawer with two I/O FOMs connects to one host system node through two PCIe4 cable adapters and four FOM-out modules.

Connecting to the PCIe expansion drawer

Figure 3-10 illustrates the connection of two expansion drawer cable pairs for two PCIe4 6-slot FOMs.




Figure 3-10 Connecting to the PCIe Gen4 Expansion Drawer

Connecting to the NED24 NVMe Expansion Drawer

Figure 3-11 shows the connectivity to the NED24 NVMe Expansion Drawer. Both connections to the drawer must be populated and from the same system.




Figure 3-11 Connecting to the NED24 drawers

3.2.2 Supported I/O drawers

The Power11 2S2U and Power11 2S4U models support the following I/O expansion drawers:

- ▶ PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0)

The PCIe Gen4 I/O Expansion Drawer is a 4U-high, 19-inch-wide, PCIe Gen4-based rack-mountable I/O drawer that is available as a feature of Power11 processor-based servers. It replaces the PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0). There is no upgrade path from the PCIe Gen3 I/O Expansion Drawer to the PCIe Gen4 I/O Expansion Drawer.

- ▶ NED24 NVMe Expansion Drawer (Feature Code ESR0)

The NED24 NVMe Expansion Drawer is a storage expansion enclosure with 24 U.2 NVMe bays. It supports up to 24 U.2 NVMe devices in 15 mm Gen3 carriers. The 15 mm carriers can accommodate either 7 mm or 15 mm NVMe devices.

Within the Power11 scale-out servers, the PCIe slots support the PCIe x16 to CXP converter card (Feature Code EJ2A), which is used to attach expansion drawers. Two PCIe x16 to CXP converter cards are required to attach each expansion drawer. Depending on the server type, the maximum number of supported drawers is shown in Table 3-9.

Table 3-9 Maximum I/O drawer configuration

Configuration	Max I/O drawers	Max NED24	Max ENZ0	Comments
2S4U	2	1	2	N/A
2S2U	1	1	1	AOC cabling is not supported.

3.2.3 Non-supported drawers

The Power11 scale-out servers no longer support the following drawers:

- ▶ PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0)

The PCIe Gen3 I/O Expansion Drawer is not supported by the Power11 platform. This drawer was replaced by the PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0). None of the FOMS that were available with I/O expansion drawers (Feature Code EMXF, Feature Code EMXG, and Feature Code EMXH) are supported in the new PCIe Gen4 I/O Expansion Drawer.

- ▶ IBM EXP24SX SAS Storage Enclosure (Feature Code ESLS and Feature Code ESLL)

The IBM EXP24SX SAS Storage Enclosure is not supported by the Power11 platform. For more internal storage, the NED24 NVMe Expansion Drawer (Feature code ESR0) is available and can be populated with up to 24 NVMe drives.

3.2.4 PCIe Gen 4 I/O Expansion Drawer

The 19 inch, 4 EIA (4U) PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0) and two PCIe Fan-out Modules (Feature Code ENZF) provide 12 PCIe Gen4 full-length, full-height slots. Each FOM provides six PCIe slots C0 - C5. Slots C0 - C3 are x16, and slots C4 and C5 are x8. PCIe Gen1, Gen2, and Gen3 full-height adapters are also supported.

Important: The PCIe Gen3 I/O Expansion Drawer (Feature Code EMX0) is not supported by the Power11 platform. None of the FOMs (Feature Code EMXF, Feature Code EMXG, and Feature Code EMXH) are supported in the new PCIe Gen4 I/O Expansion Drawer (Feature Code ENZ0).

A blind swap cassette (BSC) houses the full-height adapters that are installed in these slots. The BSC is the same as the one that is used with the previous-generation Feature Code EMX0 drawer. The drawer is shipped with a full set of BSCs.

A PCIe CXP converter adapter (Feature Code EJ24) occupies one of the PCIe Gen5 slots in the system node, and a pair of AOCs or copper cables is used for the system node-to-FOM connection. Both FOMs are independent PCIe domains and can be serviced independently. A minimum of one FOM is required in the ENZ0 drawer, in location P0, which is on the left side of the drawer when viewed from the rear. Each PCIe Gen4 I/O Expansion Drawer has two power supplies.

Drawers can be added to a server dynamically. Concurrent repair and adding or removing expansion drawers and PCIe adapters is done through HMC-guided menus or by OS support utilities.

Careful balancing of I/O, assigning adapters through redundant ENZ0 expansion drawers, and connecting to different system nodes can help ensure HA for I/O resources that are assigned to LPARs.

Figure 3-12 shows the PCIe Gen 4 I/O Expansion Drawer front view.



Figure 3-12 PCIe Gen 4 I/O Expansion Drawer front view

Figure 3-13 shows the PCIe Gen 4 I/O Expansion Drawer rear view.




Figure 3-13 PCIe Gen 4 I/O Expansion Drawer rear view

Fan-out modules supported by system

Each PCIe Gen4 I/O Expansion Drawer supports one or two FOMs, which must connect to a PCIe expansion adapter in the system to extend the PCIe bus to the drawer. The number of supported FOMs varies by system type and configuration, which directly affects how many PCIe Gen4 Expansion Drawers can be connected.

For example, systems such as the Power11 2S4U can support up to four FOMs when equipped with two DCMs. However, if only one socket is installed, the system supports only one FOM. In such cases, a filler is required in place of the second module.

Table 3-10 shows a detailed breakdown of the maximum number of supported FOMs and the total cable requirements for each scale-out server model.

Table 3-10 Number of supported fan-out modules per system

System name	Maximum FOMs	Maximum number of G4 copper cables	Maximum number of G4 AOC cables
Power S1124	4	4	8
Power S1122	2	4	0 ^a

a. The 2S2U servers do not support optical cables.

PCI slots that are available in the PCIe Gen4 Expansion Drawer

Table 3-11 lists the PCI slots in the PCIe Gen4 I/O Expansion Drawer that is equipped with two PCIe4 6-slot FOMs.

Table 3-11 PCIe slot configurations in the PCIe Gen4 I/O Expansion Drawer

Slot	Location code	Description
Left I/O module Slot 0	P0-C0	PCIe x16 adapter
Left I/O module Slot 1	P0-C1	PCIe x16 adapter
Left I/O module Slot 2	P0-C2	PCIe x16 adapter
Left I/O module Slot 3	P0-C3	PCIe x16 adapter
Left I/O module Slot 4	P0-C4	PCIe x8 adapter
Left I/O module Slot 5	P0-C5	PCIe x8 adapter
Right I/O module Slot 0	P1-C0	PCIe x16 adapter
Right I/O module Slot 1	P1-C1	PCIe x16 adapter
Right I/O module Slot 2	P1-C2	PCIe x16 adapter
Right I/O module Slot 3	P1-C3	PCIe x16 adapter
Right I/O module Slot 4	P1-C4	PCIe x8 adapter
Right I/O module Slot 5	P1-C5	PCIe x8 adapter

Note:

- ▶ All slots are PCIe Gen4 slots.
- ▶ All slots support full-length, full-height adapters or short form-factor adapters with a full-height tail stock in single-wide, Gen3, BSCs.
- ▶ Slots C0 - C3 in each PCIe Gen4 six-slot FOM are PCIe Gen4 x16 buses, and slots C4 and C5 are PCIe Gen4 x8 buses.
- ▶ All slots support EEH.
- ▶ All PCIe slots can be serviced with the power on.
- ▶ All six slots in a PCIe Gen4 six-slot FOM support SR-IOV shared mode.
- ▶ Only four adapters with Feature Code EC2S, Feature Code EC2U, or Feature Code EC72 can operate in SR-IOV mode simultaneously per six-slot FOM.

NED24 NVMe Expansion Drawer

IBM continues to provide industry-leading I/O capabilities with a PCIe direct-attached expansion drawer that supports NVMe drive attachment. The NED24 NVMe Expansion Drawer (Feature Code ESR0) is a storage expansion enclosure with 24 U.2 NVMe bays.

Important: The EXP24SX SAS Storage Enclosure is not supported by the Power11 platform.

Figure 3-14 shows the front view of a NED24 drawer with four NVMe drives in slots C8 - C11.



Figure 3-14 NED24 NVMe drawer front view

Figure 3-15 shows the rear view of the NED24 drawer with two power supplies and two ESMs.



Figure 3-15 NED24 NVMe drawer rear view

Each of the 24 NVMe bays in the NED24 drawer is individually addressable and can be assigned to a specific LPAR or VIOS to provide native boot support for up to 24 partitions. At the time of writing, each drawer supports up to 367 TB when using the 15.3 TB flash drive.

Note: Both ESMs must connect to the same server. Single connections and multiple server connections are not supported.

Each U.2 NVMe device is installed in the NED24 drawer by using a 15-mm carrier. The 15-mm carriers can accommodate either 7 mm or 15 mm NVMe devices.

The NED24 drawer is supported in the Power11 2S2U and Power11 2S4U configurations by using the same interconnect card that is used for PCIe Gen 4. A maximum of one NED24 NVMe Expansion Drawer is supported per system. When mixing different expansion drawers, the maximum number of drawers that are supported depends on the number of EJ24 FOMs that are supported.

NVMe devices that are supported in the NED24 NVMe Expansion Drawer

Table 3-12 on page 95 shows the NVMe devices that are supported in the NED24 NVMe Expansion Drawer.

Table 3-12 NVMe devices that are supported in the NED24 NVMe Expansion Drawer

Feature Code	CCIN	Description	OSs supported
ES5A	5B53	800 GB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5B	5B53	800 GB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES5C	5B52	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5D	5B52	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.
ES5E	5B51	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5F	5B51	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.
ES5G	5B50	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5H	5B50	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.
ES4B	N/A	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES4C	N/A	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.
ES4D	N/A	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES4E	N/A	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.
ES4F	N/A	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES4G	N/A	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.
ECT9	5941	15 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ECTB	5941	15 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i VIOS is not supported.

a. IBM i through VIOS only

Multipath support

At initial GA, the NED24 NVMe Expansion Drawer operated in Mode 1 (single connect). In Mode 1, the NVMe drives are configured as single-path devices, with only one ESM controlling each device. The switch in each ESM is configured to logically control only 12 of the 24 NVMe drives. No device failover capability is available.

Figure 3-16 shows the NED24 NVMe Expansion Drawer in single path mode.

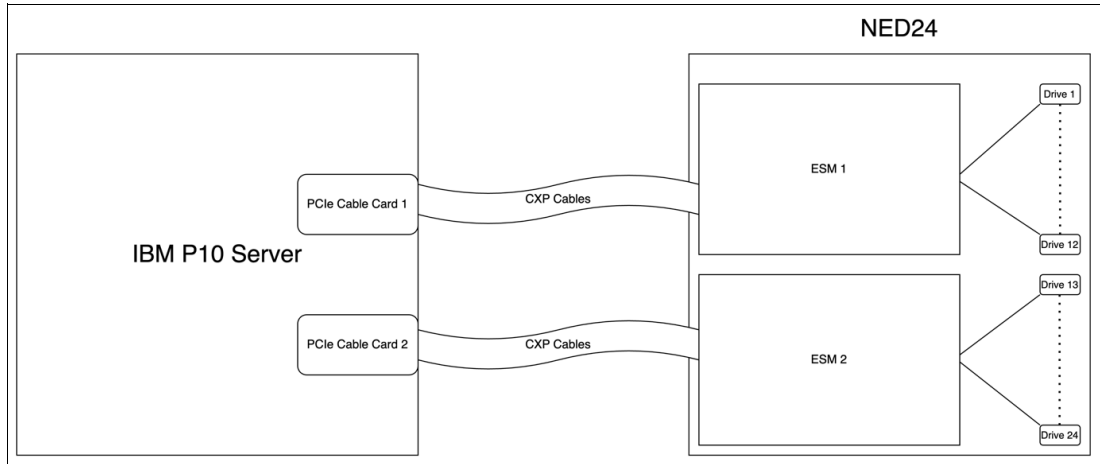


Figure 3-16 NED24 NVMe Expansion Drawer in single path mode

When running in Mode 1, use OS-level mirroring to avoid a single point of failure in the connection to the drives in the NED24 NVMe Expansion Drawers. For more information about the drive locations within the drawer for availability and reliability, see “Drive installation order” on page 97. Both ESMs must connect to the same server. Single connections and multiple server connections are not supported.

Starting with FW1060, the NVMe NED24 NVMe Expansion Drawer supports multipath. The multipath function provides two connections for each drive because each port on the multi-port drives connects through both ESMs. This configuration offers more reliability, availability, and serviceability (RAS) and improved performance.

Figure 3-17 shows the multipath function of the NED24 NVMe Expansion Drawer.

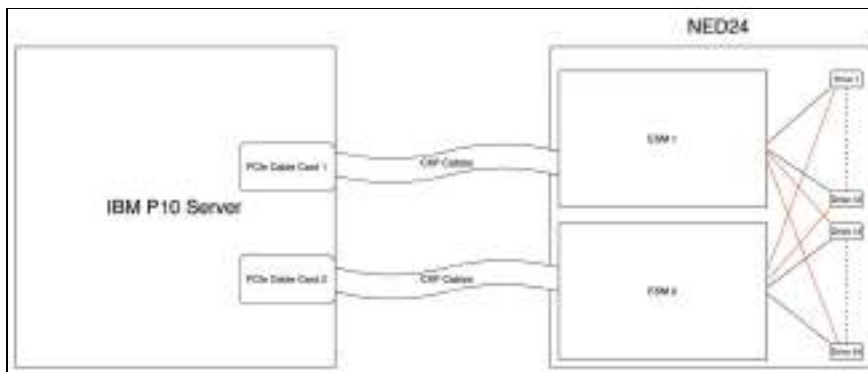


Figure 3-17 Multipath on the NED24 NVMe Expansion Drawer

Multipath is automatically enabled with the installation of FW1060 or later and is enabled when the appropriate OS level is installed.

Multipath support is provided in the OSs as shown in Table 3-13.

Table 3-13 Multipath support by operating system

OS	Supported releases
AIX	<ul style="list-style-type: none"> ▶ AIX 7.3 with the 7300-02 Technology Level (TL) and Service Pack (SP) 7300-02-02-2420 or later ▶ AIX 7.2 with the 7200-05 TL and SP 7200-05-08-2420, or later ▶ AIX 7.3 with the 7300-01 TL and SP 7300-01-04-2420, or later
IBM i	IBM i 7.6, IBM i 7.5 Technology Refresh (TR) 4, and IBM i 7.4 TR 10
Linux	<ul style="list-style-type: none"> ▶ SUSE Linux Enterprise Server 15.5 or later ▶ Red Hat Enterprise Linux (RHEL) 9.2 or later

Drive installation order

Although there is no performance difference for drives in any of the NED24 NVMe Expansion Drawer slots, there is a best practice order for installing drives within the enclosure. This order provides good separation for the mirroring between drives and helps ensure optimal cooling and airflow within the enclosure.

Figure 3-18 shows the placement for the first four drives.



Figure 3-18 Drive installation order

Table 3-14 shows the placement of all drives.

Table 3-14 Drive installation order

Drive pair	First drive slot	Second drive slot
1	1	13
2	7	19
3	2	14
4	8	20
5	3	15
6	9	21
7	4	16
8	10	22
9	5	17
10	11	23
11	6	18
12	12	24

3.3 Available adapter list

This section describes the various types and functions of the PCIe adapters that are available on the following servers:

- ▶ Power S1122 (9824-22A)
- ▶ Power S1124 (9824-42A)

This list is subject to change as more PCIe adapters are tested and certified or as listed adapters become unavailable. The following sections describe the available adapters by feature number. The tables indicate OS support (AIX, IBM i, and Linux) for each adapter.

3.3.1 Storage area network adapters

The Power11 scale-out servers support connections to devices that use Fibre Channel (FC) connectivity, either directly or through a storage area network (SAN). A range of PCIe-connected FC adapters are available in low-profile and full-height form factors.

All supported FC adapters feature LC connections. If you are attaching a switch or a device with an SC-type fiber connector, an LC-SC 50-micron Fibre Converter Cable or an LC-SC 62.5-micron Fibre Converter Cable is required.

Table 3-15 lists the low-profile FC adapters that are available for the Power S1122.

Table 3-15 Available SAN adapters for the Power S1122

Feature Code	CCIN	Description	OSs supported ^a
EN1K	579C	PCIe4 LP 32 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux
EN1M	2CFC	PCIe4 32 Gb 4-port Optical FC Adapter	AIX, IBM i ^b , and Linux
EN2M	2CFC	PCIe4 32 Gb 4-port Optical FC Adapter	AIX, IBM i ^b , and Linux
EN1B	578F	PCIe3 32 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux
EN1P	2CFD	PCIe4 64 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux
EN2P	2F05	PCIe4 64 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux

a. Validate the specific OS requirements.

b. IBM native support only with Feature Code EEEPS. Otherwise, IBM i is supported through VIOS with partition size restrictions.

Table 3-16 on page 99 lists the full-height FC adapters that are supported for the Power S1124 server models and within the PCIe expansion drawer that is connected to any of the Power11 scale-out server models.

Table 3-16 Full-height SAN adapters that are available in Power S1124 and PCIe Gen4 Expansion Drawer

Feature Code	CCIN	Description	OSs supported ^a
EN1J	579C	PCIe4 32 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux
EN1L	2CFC	PCIe4 32 Gb 4-port Optical FC Adapter	AIX, IBM i ^b , and Linux
EN1A	578F	PCIe3 32 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux
EN1L	2CFC	PCIe4 32 Gb 4-port FC Adapter	AIX, IBM i ^b , and Linux
EN1N	2CFD	PCIe4 64 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux
EN2L	2F06	PCIe4 32 Gb 4-port FC Adapter	AIX, IBM i ^b , and Linux
EN2N	2F05	PCIe4 64 Gb 2-port FC Adapter	AIX, IBM i ^b , and Linux

- a. Validate the specific OS requirements.
- b. If installed in a PCIe Gen4 Expansion Drawer that is attached to a Power S1122, IBM native support only with Feature Code EEPS. Otherwise, IBM i is supported through VIOS with partition size restrictions. For Power S1124, IBM i is supported natively.

3.3.2 Network adapters

To connect the Power S1122 and Power S1124 server models to a LAN, you can use the LAN adapters that are supported in the PCIe slots of the system. Various connection speeds and physical interfaces are supported.

Table 3-17 lists the low-profile LAN adapters that are supported for the Power S1122.

Table 3-17 Network adapters that are available in the Power S1122

Feature Code	CCIN	Description	OSs supported ^a
EN2X	2F04	PCIe2 LP 4-port 10 GbE Adapter	AIX, IBM i ^b , and Linux
EN2Y	5775	PCIe LP 4-port 1 GbE Adapter	AIX, IBM i ^b , and Linux
EC71	2CF9	2-PORT 25 Gb EN CONNECTX-6 Lx SFP28 NO CRYPTO PCIe4 x8 LP ADAPTER	AIX, IBM i ^b , and Linux
EC73	2CF8	2-PORT 25 Gb EN CONNECTX-6 Lx SFP28 CRYPTO PCIe4 x8 LP ADAPTER	IBM i ^c
EC85 ^d	EC2C	2-PORT 200 GbE NIC&RoCE SR/Cu PCIe5x16 NO CRYPTO LP ADAPTER	AIX, IBM i ^b , and Linux

- a. Validate the specific OS requirements.
- b. IBM native support only with Feature Code EEPS. Otherwise, IBM i is supported through VIOS with partition size restrictions.
- c. Supported only by IBM i native (requires Feature Code EEPS) and runs in dedicated mode only (no PowerVM virtualization). RoCE and IP Security (IPsec) are supported only by IBM i Db2 Mirror.
- d. Requires an x16 slot.

Table 3-18 lists the full-height LAN adapters that are supported for the Power S1124 server models and within the PCIe expansion drawer (Nimitz) that is connected to any of the Power11 scale-out server models.

Table 3-18 Full-height LAN adapters that are available in Power S1124 and the PCIe Gen4 Expansion Drawer

Feature Code	CCIN	Description	OSs supported ^a
EN2W	2F04	PCIe2 4-port 10 GbE Adapter	AIX, IBM i ^b , and Linux
EN2Z	5775	PCIe 4-port 1 GbE Adapter	AIX, IBM i ^b , and Linux
EC72	2CF9	2-PORT 25 Gb EN CONNECTX-6 Lx SFP28 NO CRYPTO PCIe4 x8 LP CAPABLE ADAPTER	AIX, IBM i ^b , and Linux
EC74 ^c	2CF8	2-PORT 25 Gb EN CONNECTX-6 Lx SFP28 CRYPTO PCIe4 x8 LP CAPABLE ADAPTER	IBM i ^d
EN26	EC2A	4-PORT 25/10/1 GbE NIC&RoCE SR/Cu PCIe4 x16 LP CAPABLE ADAPTER	AIX, IBM i ^b , and Linux
EC86	EC2C	2-PORT 200 GbE NIC&RoCE SR/Cu PCIe5x16 NO CRYPTO LP CAPABLE ADAPTER	AIX, IBM i ^b , and Linux

- a. Validate the specific OS requirements.
- b. IBM native support only with Feature Code EEPS. Otherwise, IBM i is supported through VIOS with partition size restrictions.
- c. Requires an x16 slot.
- d. Supported only by IBM i native (requires Feature Code EEPS) and runs in dedicated mode only (no PowerVM virtualization). RoCE and IP Security (IPsec) are supported only by IBM i Db2 Mirror.

3.3.3 Transceiver (small form-factor pluggable) replacement support for FC and Ethernet Adapter Features

IBM supports the replacement of transceivers (small form-factor pluggable [SFP]) for certain FC and Ethernet adapter features. However, replacement is not supported for all features. In particular, Ethernet adapters that are designed specifically for copper media do not support conversion to optical transceivers.

Also, some FC adapter features do not support SFP replacement, even if the transceivers appear physically removable. For these adapters, a separate SFP part is not available and cannot be ordered.

For more information about which adapters support SFP replacement, see the following resources:

- ▶ [Transceiver component \(SFP\) not replaceable for Power Fibre Channel and Ethernet adapter features.](#)
- ▶ [Transceiver component \(SFP\) is replaceable for following Power Fibre Channel and Ethernet adapter features.](#)

3.3.4 SAS and tape adapters

The internal storage in the Power11 scale-out servers is based entirely on NVMe devices that are connected directly over PCIe. If you need more storage, consider using the NED24 NVMe Expansion Drawer. Other options for extra storage include SAN-attached devices (either Ethernet-based or FC-based).

For SAS tape support, a four-port SAS tape adapter is available in both low-profile and full-height form factors.

Table 3-19 lists the low-profile SAS tape adapter that is supported for the Power S1122.

Table 3-19 Low-profile SAS tape adapter that is available for the Power S1122

Feature Code	CCIN	Description	OSs supported ^a
EJ2C	57F2	SAS TAPE HBA with 4X HD MINISAS, LOW-PROFILE PCIe3 12 Gb x8	IBM i ^b VIOS is not supported.

a. Validate the specific OS requirements.

b. Native support only with Feature Code EEPS. Otherwise, IBM i is supported through VIOS with partition size restrictions.

Table 3-20 lists the full-height USB adapters that are supported for the Power S1124 server models and within the PCIe expansion drawer (Nimitz) that is connected to any of the Power11 scale-out server models.

Table 3-20 Full-height SAN adapters that are available in Power S1124 and PCIe Gen4 Expansion Drawer

Feature Code	CCIN	Description	OSs supported ^a
EJ2B	57F2	SAS TAPE HBA with 4X HD MINISAS, LOW-PROFILE CAPABLE PCIe3 12 Gb x8	IBM i ^b VIOS is not supported.

- a. Validate the specific OS requirements.
- b. When connected to the Power S1122, IBM native support only with Feature Code EEPS. Otherwise, IBM i is supported through VIOS with partition size restrictions.

Note: The 12G SAS Tape Adapter was introduced in 2023. AIX support for this adapter on the Power11 platform is planned. This adapter supports up to four LTO tape drives.

3.4 NVMe U.2 devices

Table 3-21 shows the supported NVMe U.2 devices for Power11 scale-out servers. For more information, see “Differences between mainstream and enterprise NVMe SSDs” on page 104.

Table 3-21 NVMe U.2 Devices supported in the scale-out servers

Feature Code	CCIN	Description	OSs supported
EC7T	59B7	800 GB NVMe U.2 15 mm Carrier SSD PCIe4 (AIX and Linux)	AIX and Linux
EC5V	59BA	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
EC5W	59BA	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES1K	5947	800 GB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES1E	59B8	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES1F	59B8	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES1G	59B9	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES1H	59B9	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.

Feature Code	CCIN	Description	OSs supported
ES5A	5B53	800 GB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5B	5B53	800 GB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES5C	5B52	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5D	5B52	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES5E	5B51	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5F	5B51	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES5G	5B50	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES5H	5B50	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS not supported
ES4B	N/A	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES4C	N/A	1.6 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES4D	N/A	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES4E	N/A	3.2 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ES4F	N/A	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ES4G	N/A	6.4 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.
ECT9	5941	15 TB 4 K NVMe U.2 15 mm SSD PCIe4 (AIX and Linux)	AIX, IBM i ^a , and Linux
ECTB	5941	15 TB 4 K NVMe U.2 15 mm SSD PCIe4 (IBM i)	IBM i ^b VIOS is not supported.

a. IBM i is supported through VIOS only.

b. When connected to the Power S1122, there is IBM native support only with Feature Code EEPS.

Differences between mainstream and enterprise NVMe SSDs

This section describes the differences between enterprise NVMe SSDs and mainstream NVMe SSDs.

- ▶ Endurance and use cases

Enterprise NVMe SSDs are designed for high-endurance workloads, supporting up to 10 drive writes per day (DWPD) and typically using durable MLC NAND (formerly called eMLC). Mainstream NVMe SSDs usually support approximately 1 DWPD, making them suitable for read-heavy or less write-intensive applications.

- ▶ Over-provisioning and performance

Enterprise NVMe SSDs have more over-provisioned NAND capacity, which helps manage garbage collection, reduces write amplification, and improves random write performance and endurance. Mainstream NVMe SSDs have less over-provisioning, which reduces cost but affects performance and lifespan during write-heavy operations.

- ▶ Cost

Mainstream NVMe SSDs are more cost-effective because of lower over-provisioning and endurance specifications. They offer a reduced cost per GB but reduced performance in write-intensive tasks.

- ▶ Monitoring and lifecycle management

Mainstream NVMe SSDs must be monitored for end-of-life (EOL) indicators. When nearing EOL, a predictive failure analysis (PFA) alert is triggered. The drive should be replaced promptly to avoid degraded performance or write failures.

- ▶ Warranty considerations

Mainstream NVMe SSDs are not covered under warranty if they exceed their rated write lifecycle. Usage beyond rated endurance can result in degraded or failed write operations.

3.4.1 USB adapters

USB adapters are available to support the connection of devices, such as external DVD drives to the Power11 scale-out server models.

Table 3-22 lists the low-profile USB adapter that is supported within the Power S1122.

Table 3-22 USB adapter that is supported in the Power S1122

Feature Code	CCIN	Description	OSs supported
EC6J	590F	2-PORT USB 3.0 ADPTR, PCIe2 x1/SHORT LP	AIX, Linux, and IBM i Assignment to VIOS is supported.

Table 3-23 on page 105 lists the full-height USB adapter that is supported within the Power S1124 server models, and within the PCIe expansion drawer (Nimitz) that is connected to any of the Power11 scale-out server models.

Table 3-23 Full-height SAN adapter that is supported in Power S1124 and PCIe Gen4 Expansion Drawer

Feature Code	CCIN	Description	OSs supported
EC6K	590F	2-PORT USB 3.0 ADPTR, PCIe2 x1/SHORT LP CAPABLE	AIX, Linux, and IBM i Assignment to VIOS is supported

The embedded TI TUSB7340 USB3.0 host controller is used to provide four USB 3.0 ports:

- ▶ 1x USB 3.0 port in the rear
- ▶ 1x USB 3.0 port in the front
- ▶ 1x USB 3.0 connection to the Baseboard Management Controller (BMC) for a virtual keyboard, video, and mouse USB 3.0 port

The host USB controller can be unconfigured by the customer through the eBMC System Management web interface, which disables all host USB ports. Individual ports cannot be disabled.

3.4.2 Cryptographic adapters

Two different cryptographic coprocessors or accelerators are supported by the Power11 scale-out server models. These adapters work with the IBM Common Cryptographic Architecture (CCA) to deliver acceleration to cryptographic workloads.

For more information about the cryptographic coprocessors, the available associated software, and the available CCA, see the [IBM Security IBM PCIe Cryptographic Coprocessor](#).

PCIe Gen3 Cryptographic Coprocessor 4769

The PCIe Gen3 Cryptographic Coprocessor 4769 features a PCIe local bus-compatible interface. The coprocessor includes a security-enabled subsystem module and batteries for backup power.

The hardened, encapsulated subsystem contains redundant IBM PowerPC 476 processors and custom symmetric key and hashing engines to perform Advanced Encryption Standard (AES), DES, TDES, SHA-1, SHA-2, MD5, and HMAC, and public key cryptographic algorithm support for RSA and Elliptic Curve Cryptography (ECC).

Other hardware features include a secure real-time clock, a hardware random number generator, and a prime number generator. It also includes a separate service processor that manages self-tests and firmware updates. The secure module is protected by a tamper-responding design that defends against various system attacks.

The coprocessor provides acceleration for AES, DES, Triple DES, HMAC, CMAC, MD5, multiple SHA hashing methods, modular-exponentiation hardware (such as RSA and ECC), and full-duplex DMA communications.

A security-enabled code-loading arrangement enables control program and application program loading and refresh after coprocessor installation in your server. IBM provides an embedded subsystem control program and a cryptographic application programming interface (API) that implements IBM CCA.

The IBM 4769 is verified by NIST at FIPS 140-2 Level 4, the highest level of certification that is achievable at the time of writing for commercial cryptographic devices.

This adapter is available only in a full-height form factor and in two variations with different feature codes:

- ▶ Feature Code EJ35
Does not include a BSC and can be installed only within the chassis of a Power S1124.
- ▶ Feature Code EJ37
Includes a BSC housing and can be installed only in a PCIe I/O expansion drawer enclosure. This option is supported for the Power S1122 server.

PCIe Gen3 Cryptographic Coprocessor 4770

Here are some cryptographic applications that benefit from the strong security characteristics of the PCIe Gen3 Cryptographic Coprocessor 4770:

- ▶ Financial applications such as PIN generation and verification in automated teller and point-of-sale transaction servers
- ▶ Key management systems
- ▶ Internet business and web-serving
- ▶ Public key infrastructure applications
- ▶ Smart card applications
- ▶ Custom proprietary solutions
- ▶ Secure key support for quantum-safe cryptography for both signature (ML-DSA, FIPS 204) and key encapsulation method (ML-KEM), and FIPS 203)

The PCIe Gen3 Cryptographic Coprocessor 4770 includes enhanced hardware to perform more than 300 cryptographic algorithms, including quantum-safe algorithms. It has strong logical and physical security characteristics and is designed to meet the highest level of security certifications.

Here are the PCIe Gen3 Cryptographic Coprocessor 4770 attributes:

- ▶ Enhanced hardware to perform asymmetric, symmetric, and hashing algorithms (300+)
- ▶ Hardware support for quantum-safe algorithms (CRYSTALS-Dilithium)
- ▶ CCA (PKCS #11 supported through OpenCryptoki)
- ▶ Quantum-safe protected firmware by using parallel signatures (ECDSA + CRYSTALS-Dilithium)
- ▶ Tamper-responding secure hardware that is designed to meet the highest level of security for PCI HSMv4, FIPS 140-2 Level 4, Common Criteria, GBIC, AusPayNet, and other certifications
- ▶ X9 TR-31 key block native support, X9 TR-34 remote key load, and AES and TDES PIN processing, including AES DUKPT and TDES DUKPT

Table 3-24 on page 107 lists the low-profile cryptographic adapter that is supported within the Power S1122.

Table 3-24 Low-profile cryptographic adapter

Feature Code	CCIN	Description	OSs supported
EPG4 ^a	C138	4770 CRYPTO COPROC PCI3 x4 LP	AIX, Linux, and IBM i Assignment to VIOS is supported.

a. This feature is not available in China.

Table 3-25 lists the full-height cryptographic adapters that are supported within the Power S1124 server and the PCIe expansion drawers.

Table 3-25 Full -height cryptographic adapters

Feature Code	CCIN	Description	OSs supported
EJ35	4769	PCIe3 Crypto Coprocessor no BSC 4769 (Power S1124 chassis only)	AIX, Linux, and IBM i Assignment to VIOS is supported.
EJ37	4769	CRYPTO COPROC PCI3 x4 LP CAPABLE IN GEN3' BSC For I/O Drawers	AIX, Linux, and IBM i Assignment to VIOS is supported.
EPG5	C138	4770 CRYPTO COPROC PCI3 x4 LP CAPABLE	AIX, Linux, and IBM i Assignment to VIOS is supported.
EPG6	C138	4770 CRYPTO COPROC PCI3 x4 LP CAPABLE IN GEN3' CASSETTE FOR I/O DRAWER	AIX, Linux, and IBM i Assignment to VIOS is supported.

3.4.3 Graphical adapters

At the time of writing, there is not a graphics adapter that is available for the Power11 processor-based servers. For applications where a local graphics terminal is required, IBM is planning on supporting a new graphics adapter for scale-out servers. The adapter is planned to provide both AIX and Linux support.

3.5 Other device support

This section describes other devices that attach to the Power11 scale-out servers.

RDX removable hard disk cartridge

The RDX removable hard disk cartridge is not supported by Power11 processor-based servers. It cannot be purchased and used with these servers.

There are alternative solutions to an RDX based backup strategy that are offered:

- ▶ Cloud-based backup services
- ▶ On-premises entry tape drives and libraries

IBM System Storage 7226 Model 1U3 Multi-Media Enclosure

The IBM System Storage 7226 Model 1U3 Multi-Media Enclosure can accommodate up to two LTO tape drives, two RDX removable disk drive docking stations, or up to four DVD-RAM drives. The IBM System Storage 7226 Model 1U3 Multi-Media Enclosure offers SAS, USB, and FC electronic interface drive options for attachment to the Power S1122 and Power S1224 servers.

The IBM System Storage 7226 Model 1U3 Multi-Media Enclosure is a 1U rack-mountable dual bay enclosure with storage device options for LTO 7, LTO 8, and LTO 9 tape drives with SAS-6 or SAS-12 interface or an FC 8-Gbps interface. and offers up to 18 TB single cartridge native capacity.

The RDX feature is not supported on Power11 processor-based servers. Up to two devices can be installed in any combination in the IBM System Storage 7226 Model 1U3 Multi-Media Enclosure.

The IBM System Storage 7226 Model 1U3 Multi-Media Enclosure offers the following drive features:

- ▶ Half-High LTO Ultrium 8 SAS Tape Drive (Feature Code 8541)
- ▶ Half-High LTO Ultrium 8 FC Tape Drive (Feature Code 8546)
- ▶ Half-High LTO Ultrium 9 SAS Tape Drive (Feature Code 8641)
- ▶ Half-High LTO Ultrium 9 FC Tape Drive (Feature Code 8646)

3.5.1 No longer supported

Starting with the Power11 platform, IBM no longer supports the following features:

- ▶ SAS storage adapters and drawers. SAS tape support is still available (Feature Code EJ2B and Feature Code EC2C)
- ▶ Gen3 I/O Expansion Drawer (Feature Code EMX0) and the associated FOMs
- ▶ 8 Gb FC adapters
- ▶ All RDX features
- ▶ NVMe slot adapters

This list is not exhaustive. To validate whether your device is supported, see the [IBM Sales Manual](#).

3.5.2 PCIe adapters that require increased cooling

PCIe adapters can generate higher levels of heat and often have strict temperature limits. IBM Power scale-out servers support various adapters. To ensure adequate cooling, the fan speed floor increases when these adapters are installed. The minimum fan RPM varies depending on the specific adapter and the ambient temperature.

Note: Higher fan speed floors can result in a noticeable increase in overall system noise. Fan activity is often triggered based on temperature readings from sensors on the operator panel.

IBM Power scale-out servers include three sensors on the operator panel to measure ambient temperature. You can view the current readings from these sensors through the ASMI (eBMC) interface by selecting **Hardware Status** → **Sensors**, as shown in Figure 3-19.

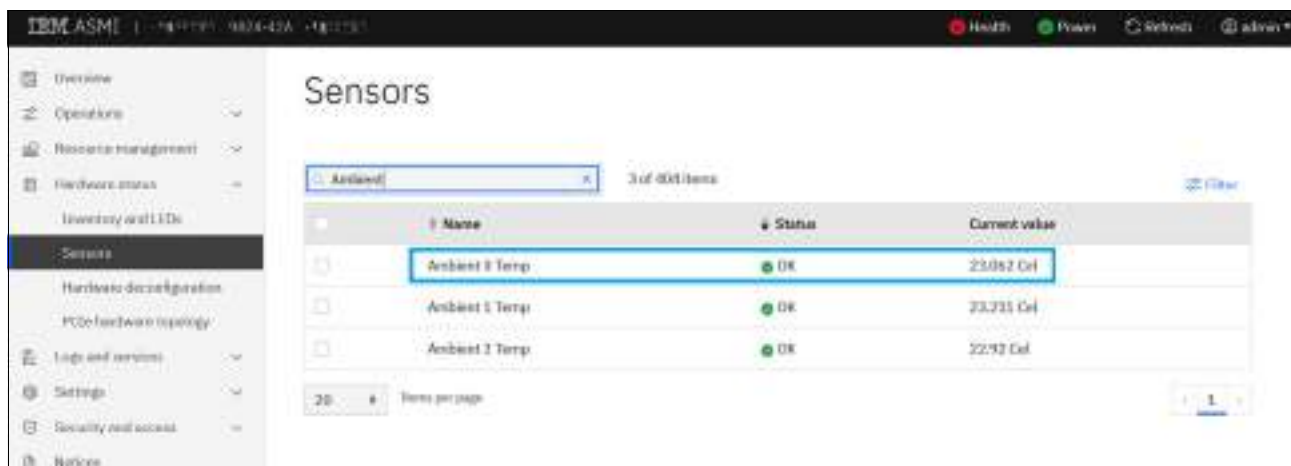


Figure 3-19 Ambient temperature readings from three sensors on the operator panel

The firmware uses an internal algorithm to determine which ambient sensor value to use for fan control. For example, if the readings are as follows (see Figure 3-19), the firmware selects Ambient 0 Temp (23.062°C) as the controlling value because it represents the median reading among the three sensors.

- ▶ Ambient 0 Temp = 23.062 °C
- ▶ Ambient 1 Temp = 23.211 °C
- ▶ Ambient 2 Temp = 22.92 °C

On the Power S1124 (eBMC) system model 9824-42A, if the intake air temperature that is measured by the front-mounted sensors is under 27°C (as shown in Figure 3-19 on page 109), the default fan speed target is 5000 RPM, as illustrated in Figure 3-20, assuming that no hot adapters are installed.

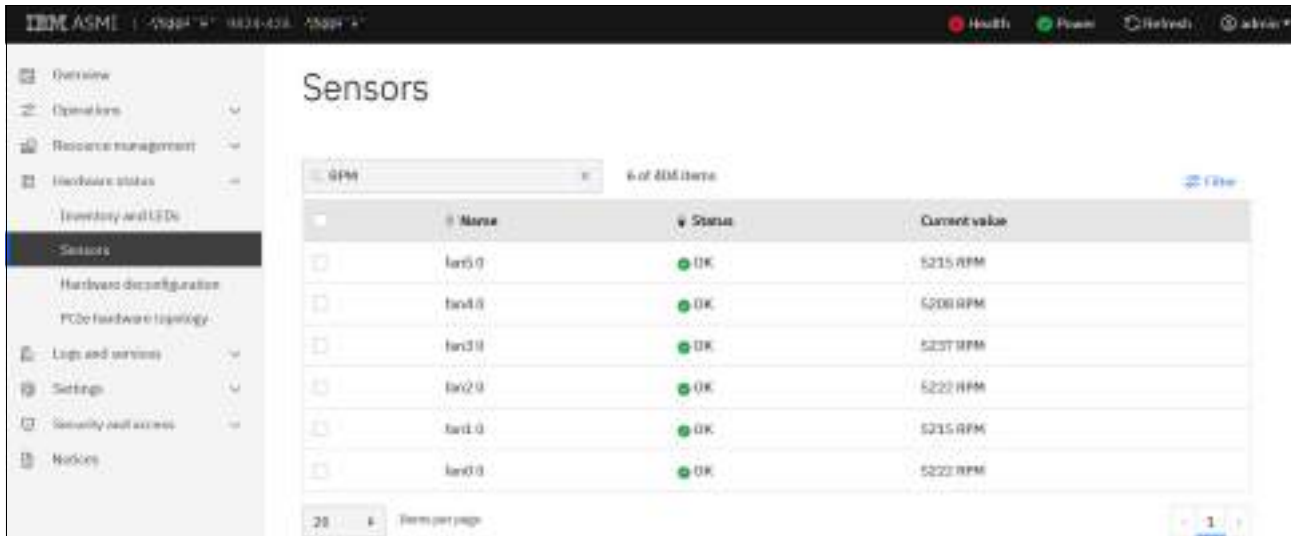


Figure 3-20 Fan speed readings in RPM as viewed through the ASMI GUI (eBMC) under the Sensors tab

Note: The fan speed target and the actual speed readings (in RPM), as displayed in the ASMI GUI (eBMC) under the **Sensors** tab, might vary slightly from the expected values. This variation is normal and results from system-specific fan control policies. For example, 2U (Power S1122) and 4U (Power S1124) systems follow different rules and ambient temperature zones, which influence their fan speed behavior.

If hot adapters are present, the fan speed behavior adjusts based on the [Hot PCIe Adapters List](#) and the corresponding ambient temperature range.

Note: PCIe4 cable adapters (Feature Code EJ2A and Feature Code EJ24) are not treated as standard PCIe hot adapters. Their presence in the system does not automatically trigger higher fan floor speeds. Instead, these adapters include temperature sensors, and the system firmware uses predefined thresholds to determine when to increase fan speeds to help ensure adequate cooling.

Note: The Hot Adapter List is continuously updated. To access the latest version, see this [public GitHub repository](#). Find your system-specific `com.ibm.Hardware.Chassis.Model.xxxxxx` directory and view the `pcie_cards.json` file for the most up-to-date list.

Table 3-26 lists the PCIe adapters that require increased cooling and result in higher fan speeds on the IBM Power scale-out servers.

Table 3-26 PCIe adapters that require increased cooling

Feature Code	Description
EC2S and EC2R	PCIe3 2-port 10 Gb NIC RoCE SR/Cu adapter
EC2U and EC2T	PCIe3 2-port 25/10 Gb NIC RoCE SFP28 adapter

Feature Code	Description
EC66 and EC67	PCIe4 2-port 100 GbE RoCE x16 adapter
EJ14	PCIe3 12 GB Cache RAID SAS 4 Adapter Quad-port with Advanced Features
EJ0L	PCIe3 12 GB Cache RAID SAS Adapter Quad-port 6 Gb
EJ0J and EJ0M	PCIe3 RAID SAS Adapter Quad-port 6 Gb
EJ10 and EJ11	Media PCIe3 4 x8 SAS Port Adapter (Tape and DVD)
EC5G and EC5B	PCIe3 NVMe x8 1.6 TB SSD Adapter
EC5C and EC5D	NVMe PCIe3 NVMe x8 3.2 TB SSD Adapter
EC5E and EC5F	PCIe3 NVMe x8 6.4 TB SSD Adapter
EC6U and EC6V	PCIe3 NVMe x8 1.6 TB SSD Adapter IBM i
EC6W and EC6X	PCIe3 NVMe x8 3.2 TB SSD Adapter IBM i
EC6Y and EC6Z	PCIe3 NVMe x8 6.4 TB SSD Adapter IBM i
EJ32	PCIe Crypto Coprocessor
EJ1X	PCIe JBOF for NVMe U.2 4-Pack Backplane
EJ1Y	PCIe JBOF for NVMe U.2 8-Pack Backplane
EJ24	PCIe x16 to CXP Converter Card 2U Gen4 for MEX I/O drawer
EJ2A	PCIe x16 to CXP Converter Card 4U Gen4 for MEX I/O drawer
EC76 and EC75	PCIe4 2-port 100 Gb RoCE EN ConnectX-6Dx Adapter
EC78 and EC77	PCIe4 2-port 100 Gb RoCE EN ConnectX-6Dx Encryption Offload Adapter
EC7A and EC7B	NVMe PCIe4 NVMe x8 1.6 TB SSD Adapter (AIX and Linux)
EC7C and EC7D	NVMe PCIe4 NVMe x8 3.2 TB SSD Adapter (AIX and Linux)
EC7E and EC7F	NVMe PCIe4 NVMe x8 6.4 TB SSD Adapter (AIX and Linux)
EC7J and EC7K	NVMe PCIe4 NVMe x8 1.6 TB SSD Adapter (IBM i)
EC7L and EC7M	NVMe PCIe4 NVMe x8 3.2 TB SSD Adapter (IBM i)
EC7N and EC7P	NVMe PCIe4 NVMe x8 6.4 TB SSD Adapter (IBM i)
EJ35	PCIe3 x7 4769 Crypto Coproc
EN1E and EN1F	PCIe3 x8 4-Port Fibre Channel (16 Gbps)
EN1C and EN1D	PCIe3 16 Gb 4-port FC Adapter
EN2L	PCIe4 32 Gb 4-port FC Adapter
EC64	2-Port EDR 100 Gb InfiniBand ConnectX-5 GEN4 PCIe LP ADAPTER – RPQ only
EN26	PCIe4x16 4-port 25/10/1 GbE NICRoCE SR/Cu
EN2D	PCIe4x16 4-port 25/10/1 GbE NICRoCE SR/Cu with Crypto

Note: The availability of the PCIe adapters that are listed in Table 3-26 on page 110 might vary, and not all adapters are offered for IBM Power scale-out servers.



Artificial intelligence support

Artificial intelligence (AI) is a cornerstone of digital transformation across industries. It enables organizations to automate processes, gain deeper insights, and deliver personalized experiences. From predictive analytics and natural language processing to computer vision and generative AI, the demand for AI-driven solutions is growing rapidly. To support these workloads effectively, enterprises need an infrastructure that can handle the computational intensity and data throughput that AI requires. IBM Power11 processor-based servers play a critical role in meeting these needs.

IBM Power is purpose-built for AI workloads with its high-performance architecture, including the latest Power11 processors that feature integrated AI acceleration. These processors handle low-precision arithmetic operations that are commonly used in AI models, boosting inferencing speed without compromising accuracy. IBM Power also supports advanced technologies such as the IBM Spyre Accelerator, a Peripheral Component Interconnect Express (PCIe)-attached AI card that is optimized for enterprise AI workloads, and Power11 on-chip AI inferencing, which enables real-time decision-making at scale. These capabilities make Power servers ideal for deploying AI models in production environments where performance, reliability, and scalability are essential.

Beyond hardware, IBM Power integrates with Red Hat OpenShift AI, enabling organizations to build, train, and deploy AI models in a containerized, hybrid cloud environment. This combination supports consistent DevOps practices, efficient resource usage, and simplified management of AI workflows. With support for AI frameworks such as PyTorch, TensorFlow, and vLLM, IBM Power provides a complete, enterprise-ready platform for running AI. Whether running AI at the edge, in the data center, or across hybrid cloud environments, IBM Power delivers the performance and flexibility that are needed to turn AI potential into real-world impact.

This chapter contains the following topics:

- ▶ On-chip support
- ▶ AI acceleration with IBM Spyre Accelerator
- ▶ Solutions

4.1 On-chip support

With the upcoming availability of the IBM Spyre Accelerator on IBM Power servers, starting with systems that are powered by the new Power11 processor, organizations can use a new class of hardware acceleration that is designed to support a wide spectrum of enterprise workloads. The IBM Spyre Accelerator delivers performance benefits, particularly for compute-intensive and AI-driven use cases, while aligning with the architectural strengths of the IBM Power platform. The IBM Spyre Accelerator is not a GPU but a dedicated hardware device that is developed by IBM Research® to provide superior acceleration.

A common misconception among IT decision-makers is that one or more GPUs are always required for any workload involving AI. This perception often leads to infrastructure decisions that prioritize GPU integration by default, regardless of the nature or scope of the AI tasks. However, GPUs are not a prerequisite for all AI use cases. Their inclusion in every server configuration can result in disproportionate acquisition costs, elevated power consumption, and increased complexity in thermal and workload management, without guaranteed benefits for the workload.

Beginning with IBM Power10 and enhanced with the Power11 platform, each core integrates four Matrix Math Accelerators (MMAs) that support a wide range of AI inference workloads directly on the CPU. This innovation enables customers to run AI models natively on IBM Power cores without requiring a discrete GPU. The architecture is suitable for traditional AI use cases such as fraud detection, text extraction, document analysis, domain adaptation through retrieval augmented generation (RAG), pattern recognition, forecasting, and image, video, or audio processing.

The Power11 processor further enhances these capabilities. Although it does not always match the raw throughput of high-end GPUs in certain generative AI scenarios, Power11 technology provides excellent performance for tasks such as entity extraction, translation, summarization, and classification, while offering reduced energy consumption and improved data protection. By enabling AI workloads to run closer to where the data is, Power11 processor-based servers support secure, efficient, and scalable deployment of AI without unnecessary data movement.

When the IBM Spyre™ Accelerator becomes available on Power11 processor-based servers, it will introduce more acceleration capabilities that complement the on-chip features. Power11 processor-based servers, with integrated AI accelerators and future support for the IBM Spyre Accelerator, present a highly optimized, energy-aware, and secure platform for deploying AI across a broad range of use cases without default dependency on discrete GPUs.

4.2 AI acceleration with IBM Spyre Accelerator

Improvements to Power11 processor core strength and system capacity boost the performance of the MMA for inference workloads. Adding IBM Spyre Accelerator adapters to the Power11 platform provides more AI inference capabilities. Working together, IBM POWER processors and the IBM Spyre Accelerator enable next-generation infrastructure to scale demanding AI workloads for businesses.

The IBM Spyre Accelerator extends the capabilities of Power11 processor-based servers by offering a low-power, high-efficiency acceleration solution for workloads that demand frequent memory access and streamlined data movement.

Because of the advanced virtualization and workload consolidation capabilities of IBM Power servers, IBM Spyre Accelerator based applications can be colocated with other mission-critical services within the same physical server so that AI inference engines, real-time analytics, or data preprocessing workloads that use the IBM Spyre Accelerator can run close to databases or transactional systems that are hosted in logical partitions (LPARs) or containers on the same server. This architectural proximity reduces latency, improves throughput, and eliminates the impact that is typically associated with cross-node or cross-platform communication.

IBM Spyre Accelerator is uniquely positioned as one of the only AI solutions that combines the following benefits:

- ▶ Data privacy
 - AI data sovereignty on a reliable, trusted, and on-premises infrastructure
- ▶ Skills
 - Ready-to-use enterprise AI services
- ▶ Complexity
 - Accelerated plug-and-play AI for business workflows

Important: Support for the IBM Spyre Accelerator on IBM Power is expected to be announced and available in late 2025. Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice and represent goals only.

4.2.1 Deploying AI in the enterprise

AI is transforming enterprise operations across industries by enabling organizations to optimize decision-making, streamline processes, and deliver personalized experiences at scale. However, the path to successful AI integration in enterprise environments is complex. It requires the right algorithms and models and a robust, secure, and scalable infrastructure that is tailored to diverse AI workloads. IBM Power, an AI-optimized platform, provides a unique approach to meet this challenge by offering both on-chip and off-chip acceleration capabilities for different levels of AI maturity.

Most enterprise AI adoption begins with experimentation and initial use cases such as proof-of-concept models for customer segmentation, log analysis, or anomaly detection. These workloads typically require limited computational resources and can be efficiently handled by the on-chip AI acceleration that is embedded in Power11 processors. The acceleration includes MMA engines and single instruction, multiple data (SIMD) vector instructions that are combined with high memory bandwidth, which is a critical component for feeding data into AI models rapidly. On-chip acceleration enables real-time inference close to the data source, minimizing latency and reducing the need for extra hardware. It is effective for traditional machine learning models, data warehouse analytics, and vector database operations such as RAG queries.

As enterprise AI use cases evolve to include more complex deep learning models, such as image classification, video processing, and time-series forecasting, the demands on compute performance and memory bandwidth increase. At this stage, enterprises need a more flexible and powerful acceleration option, particularly as workloads move from experimentation to production.

4.2.2 Fit-for-purpose AI acceleration

IBM introduced the IBM Spyre Accelerator, which is a dedicated off-chip accelerator that is designed specifically for AI-intensive workloads. Compared to the on-chip accelerators in the Power11 platform, the IBM Spyre Accelerator offers higher throughput, parallelism, and model capacity because of its dedicated memory architecture, optimized data paths, and ability to offload processing from the CPU. The IBM Spyre Accelerator is ideal for large-scale transformer models that are used in generative AI applications such as translation, summarization, sentiment analysis, and more. This architecture supports massive parallel compute operations with higher efficiency than general-purpose CPUs.

A key advantage of the IBM Power platform is its ability to combine both acceleration strategies (on-chip and off-chip) within a unified, enterprise-grade ecosystem. Initial AI workloads can begin on existing IBM Power servers with no added hardware by using the built-in accelerators for cost efficiency. As demands grow, the IBM Spyre Accelerator can be seamlessly integrated to boost performance without requiring a platform change or software rewrite. This flexibility protects existing investments while enabling future scaling, as shown in Figure 4-1.

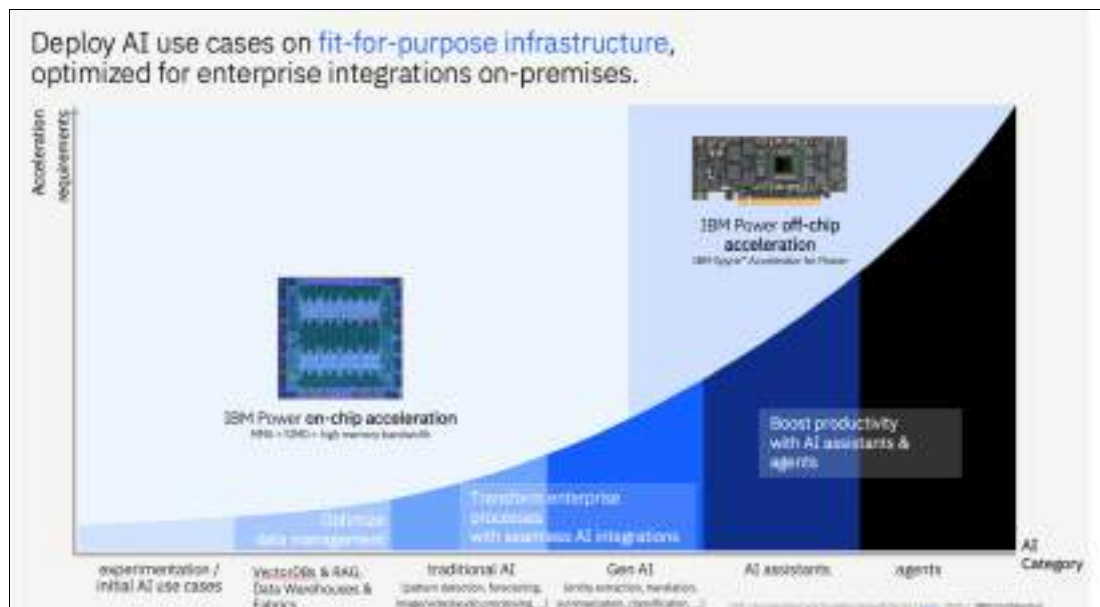


Figure 4-1 The IBM fit-for-purpose AI architecture for IBM Power

In high-value use cases such as AI assistants and autonomous agents, where real-time interaction, low-latency inference, and contextual awareness are essential, off-chip acceleration becomes a requirement. These applications often rely on large language models (LLMs) and must manage vast knowledge graphs or context windows that exceed the practical limits of CPU-based inference alone. The IBM Spyre Accelerator is built for these scenarios, and when paired with the Power11 platform, it provides a hybrid architecture that is secure, performant, and manageable within enterprise IT constraints.

This fit-for-purpose infrastructure aligns well with existing on-premises environments, especially ones where data sovereignty, latency, and compliance are non-negotiable. Unlike cloud-only solutions, the Power11 platform and IBM Spyre Accelerator combination gives enterprises full control over their data and compute stack while supporting AI workloads that rival hyperscale offerings in performance.

This approach to enterprise AI infrastructure stands out for its modularity and adaptability. Whether starting with lightweight models or scaling to multimodal AI agents, Power11 processor-based servers enable organizations to deploy AI on their terms. With on-chip acceleration for efficient starting points and IBM Spyre Accelerator off-chip acceleration for high-performance growth, businesses gain the ability to move from AI exploration to enterprise-wide transformation, and all on a platform that is built for the future of AI.

4.2.3 IBM Spyre Accelerator

The IBM Spyre Accelerator is a purpose-built, enterprise-grade accelerator that offers scalable capabilities for complex AI models and generative AI use cases. The new accelerator features 32 individual accelerator cores that are mounted on a PCIe card. Jointly designed by IBM Research and IBM Infrastructure, the IBM Spyre Accelerator architecture enables more efficient AI computation. The chip sends data directly from one compute engine to the next, helping ensure efficient energy usage. This family of processors also uses a range of lower-precision numeric formats, such as INT4 and INT8, to make running AI models more energy efficient and less memory intensive.

The IBM Spyre Accelerator is designed with a system-on-a-chip architecture that is optimized for enterprise AI workloads. It is implemented in 5 nm technology with a high-performance, low-power design that uses 75 W and provides the following capabilities:

- ▶ 32 low-power AI cores
- ▶ Multi-precision support for inference and training (FP16/8 and INT8/4)
- ▶ Support for foundation models that is enabled in the Red Hat software stack (Red Hat OpenShift Container Platform AI)
- ▶ Support for AI frameworks and libraries such as PyTorch and vLLM

Figure 4-2 shows the IBM Spyre Accelerator.

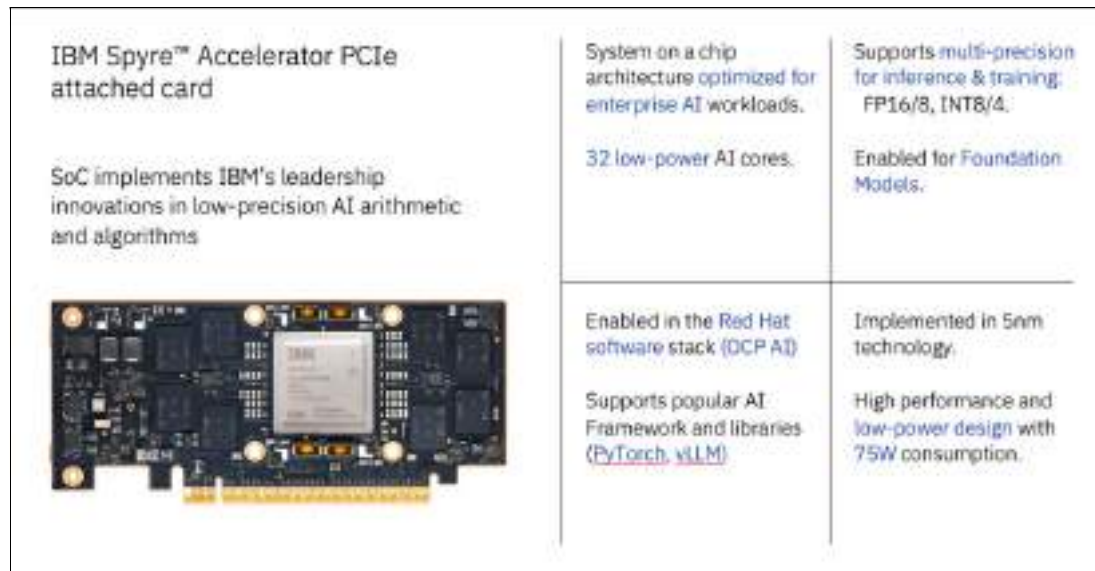


Figure 4-2 The IBM Spyre Accelerator

This dedicated, enterprise-grade AI accelerator sits on a 75 W PCIe adapter and is paired with 128 GB of LPDDR5 memory to support many LLMs for the heterogeneous workloads that are typically seen on IBM Power. Because a single IBM Spyre Accelerator does not provide enough compute capacity for most use cases, the IBM solution uses current I/O expansion technology to attach a cluster of eight IBM Spyre Accelerators in a single I/O expansion drawer, creating a logical cluster. The firmware on these eight cards coordinates compute distribution and data transfer among the cards. This architecture makes the cluster appear to the software as one high-performance compute engine with 1 TB of memory and 1.6 TBps of memory bandwidth.

Figure 4-3 shows an IBM Spyre Accelerator cluster in an expansion drawer.



Figure 4-3 IBM Spyre Accelerator cluster in an expansion drawer

This expansion drawer with IBM Spyre Accelerator cards can be attached to any announced Power11 processor-based server. The number of supported drawers depends on the Power11 processor-based server model.

At announcement, the IBM Spyre Accelerator is planned to be available in a fixed configuration that includes eight IBM Spyre Accelerator cards that are installed in a PCIe Gen4 I/O Expansion Drawer. The IBM Spyre Accelerator expansion drawer can be attached to any Power11 processor-based server: IBM Power S1122, IBM Power S1124, IBM Power E1150, and IBM Power E1180. One expansion drawer is supported on all models except the Power E1180, which supports up to two. New I/O drawer components, such as power supplies and fan-out modules (FOMs), are used to support the additional power requirements of the IBM Spyre Accelerator configuration.

Here are the initial supported software and hardware:

- ▶ FW1110.10
- ▶ Hardware Management Console (HMC) 1111
- ▶ Red Hat Enterprise Linux (RHEL) 9.6
- ▶ IBM Spyre Accelerator software stack container
- ▶ Red Hat OpenShift AI (technology preview 4Q 2025, with GA 1Q 2026)

4.3 Solutions

AI is no longer an emerging trend; it is a foundational pillar of competitive advantage. According to the IBM Institute for Business Value, about 72% of top-performing CEOs identify advanced generative AI as essential to their future success¹. However, only a fraction of enterprises has successfully moved beyond pilot projects and proofs of concept (PoCs) into full production deployments.

The Power11 processor-based server provides the AI-optimized infrastructure and integrated software stack that are required to scale AI workloads, secure enterprise data, and modernize mission-critical processes.

Organizations across industries are realizing tangible benefits from deploying AI on IBM Power:

- ▶ Logistics service providers reduce order processing time by embedding generative AI into their enterprise resource planning (ERP) systems.
- ▶ Financial institutions accelerate anomaly detection by four times while reducing total cost of ownership (TCO).
- ▶ Utility providers reduce equipment downtime and energy consumption by using AI-powered visual inspection.

These outcomes highlight the value of deploying generative AI in real business workflows by using the trusted IBM Power platform.

IBM Power is optimized for a wide range of AI workloads across key sectors:

- ▶ Finance and ERP
Fraud detection, anti-money laundering (AML), order processing, invoice compliance
- ▶ Healthcare
Medical transcription, claims, electronic health record (EHR) matching, diagnostic imaging, and assistant bots
- ▶ IT and development
Predictive IT operations (ITOps), code assistants (RPG and Ansible), transcription, documentation
- ▶ Cross-industry
Document digitization, summarization, know-your-customer (KYC), visual quality inspection, customer churn prediction, and business intelligence

Each use case uses the inherent strengths of IBM Power, such as resilient compute, secure data access, and scalable architecture.

¹ Source: <https://www.ibm.com/thought-leadership/institute-business-value/en-us/think-2025>

4.3.1 Barriers to scaling AI and the IBM Power technology advantage

Although the promise of AI is great, common enterprise challenges persist, as described in Table 4-1.

Table 4-1 Challenges to scale AI

Barrier	Description
Data complexity	Fragmented formats, distributed sources, and lack of standardized connectors
AI integration difficulty	Earlier workflows, lack of AI skills, and software maturity
Security and privacy	Concerns around data leakage, adversarial threats, and model protection
Infrastructure and cost concerns	High cost of AI development and fears of redesigning core systems
Talent and skill gaps	Shortage of data scientists and AI engineers

The Power11 platform addresses these challenges by offering a complete AI platform that is integrated into existing systems, with hardware acceleration, data fabric, and AI assistants.

4.3.2 IBM Power AI-ready portfolio

To scale AI in production, IBM Power provides an AI-ready portfolio:

- ▶ AI-ready infrastructure
 - Built-in accelerators and support for hybrid and cloud-native deployment (for example, IBM Power Virtual Server (PowerVS))
- ▶ Optimized software stack
 - Red Hat, IBM, open source, and certified independent software vendor (ISV) AI solutions
- ▶ Data fabric
 - Secure, unified access to enterprise data across silos
- ▶ AI assistants and agents
 - Tools to enhance developer productivity, automate ITOps, and streamline business workflows

This stack empowers organizations to infuse AI into core business processes while maintaining governance, performance, and security.

4.3.3 Strategic outcomes that are enabled by IBM Power

Table 4-2 describes example outcomes that are enabled by the implementation of AI-solutions on IBM Power servers.

Table 4-2 Example outcomes that are enabled by IBM Power

Objective	Outcome example
Transform processes.	Accelerate ERP or financial workflows that are integrated with AI services.
Boost productivity.	Roll out new application features faster with generative AI code assistants.
Optimize data use.	Reduce TCO and improve resource efficiency through unified data governance.

IBM Power enables clients to put AI to work today securely, efficiently, and at enterprise scale.

4.3.4 Uniqueness of the platform

With its AI-optimized infrastructure, IBM Power is uniquely positioned to help organizations move from experimentation to real AI-driven transformation. Its scalable ecosystem, which spans hardware, software, data, and automation, makes it a reliable platform for deploying production-grade AI at scale.

A key differentiator of the Power11 platform is its ability to deliver AI capabilities without requiring external GPUs. Unlike the common assumption that all AI workloads demand dedicated GPU acceleration, the Power11 platform uses advanced on-core accelerators that are integrated within each processor core and combined with a system-wide infrastructure that is optimized for AI and data-intensive tasks.

This architectural approach enables the deployment of a wide range of AI use cases, such as inference, predictive analytics, and cognitive automation while maintaining enterprise-grade attributes that include resilience, reliability, performance, security, and sustainability. By eliminating or reducing a dependency on external GPUs, organizations benefit from simplified infrastructure, reduced energy consumption, and reduced TCO without compromising AI readiness. The Power11 platform is designed to support scalable and secure AI solutions across diverse industry environments.



Automation and management

Managing IBM Power-based processor servers effectively requires a combination of tools that provide automation, orchestration, and visibility across the infrastructure. At the core of this management stack is the Hardware Management Console (HMC), which acts as the central point for configuring and monitoring Power servers. The HMC enables administrators to manage logical partitions (LPARs), perform firmware updates, and monitor system health. It provides a CLI, a GUI, and a Representational State Transfer (REST) application programming interface (API), enabling integration with automation tools for scalable operations.

Ansible and Terraform bring infrastructure as code (IaC) capabilities to IBM Power environments. Ansible is widely used for configuration management and automation of tasks such as fixes, user management, and software deployment. It supports Power servers through modules and collections that are tailored for AIX, IBM i, and Linux on Power. Terraform complements Ansible by enabling declarative provisioning of infrastructure components, including IBM Power Virtual Server (PowerVS) and Power Virtualization Center (PowerVC) managed resources. Together, they enable teams to automate the full lifecycle of infrastructure from provisioning to configuration, helping ensure consistency and reducing manual effort.

PowerVC plays a crucial role in managing virtualized environments on IBM Power. Built on OpenStack, PowerVC provides advanced virtualization management capabilities such as image management, dynamic resource allocation, and integration with cloud platforms. It supports automation through REST APIs and integrates with both Ansible and Terraform, enabling seamless orchestration of virtual machines (VMs) and workloads. By combining HMC, Ansible, Terraform, and PowerVC, organizations can build a robust, automated, and scalable management framework for their IBM Power infrastructure.

This chapter contains the following topics:

- ▶ Hardware Management Console overview
- ▶ Ansible
- ▶ Terraform
- ▶ PowerVC

5.1 Hardware Management Console overview

The HMC is a hardware or virtual appliance that is used to configure and manage your systems. The HMC connects to one or more managed systems and provides capabilities for the following primary functions:

- ▶ Provide systems management functions, including the following examples:
 - Power off
 - Power on
 - System settings
 - Capacity on Demand (CoD)
 - Enterprise pools
 - Shared processor pools (SPPs)
 - Performance and Capacity Monitoring
- ▶ Starting Advanced System Management Interface (ASMI) for managed systems
- ▶ Deliver virtualization management through by creating, managing, and deleting LPARs, Live Partition Mobility (LPM), Remote Restart, configuring Single Root I/O Virtualization (SR-IOV), managing Virtual I/O Servers (VIOs), dynamic resource allocation, and operating system (OS) terminals.
- ▶ Acts as the service focal point for systems and supports service functions, including call home, dump management, guided repair and verify, concurrent firmware updates for managed systems, and around-the-clock error reporting with IBM Electronic Service Agent (ESA) for faster support.
- ▶ Provides appliance management capabilities for configuring network users on the HMC, and updating and upgrading the HMC.

5.1.1 HMC options

Power11 processor-based servers can connect to either the 7063-CR2 HMC or a Virtual Hardware Management Console (vHMC).

Note: Power11 processor-based servers do not support the 7063-CR1.

7063-CR2 HMC

The 7063-CR2 HMC (see Figure 5-1 on page 125) is a second-generation IBM Power HMC.

The 7063-CR2 HMC includes the following features:

- ▶ A 6-core IBM Power9 130 W processor chip
- ▶ 64 GB (4 × 16 GB) or 128 GB (4 × 32 GB) of memory
- ▶ 1.8 TB of internal disk capacity with RAID 1 protection
- ▶ Four 1 Gbps Ethernet ports (RJ-45), two 10 Gbps Ethernet ports (RJ-45), two Universal Serial Bus (USB) 3.0 ports (front), two USB 3.0 ports (rear), and one 1 Gbps Intelligent Platform Management Interface (IPMI) Ethernet port (RJ-45)
- ▶ Two 900 W power supply units
- ▶ Remote Management Service. IPMI port (OpenBMC) and Redfish API

The base warranty is 1 year, 9x5, with optional upgrades available. A USB Smart Drive is not included.

The 7063-CR2 is compatible with flat panel console kits 7316-TF3, TF4, and TF5.

Note: The 7316-TF3 and TF4 are withdrawn from marketing



Figure 5-1 7063-CR2 HMC

Virtual Hardware Management Console

Initially, the HMC was available only as a hardware appliance with the HMC firmware installed. IBM later extended this offering to include either a hardware appliance or a virtual appliance that can run on ppc64le architectures or x86 platforms.

Any customer with a valid contract can download the HMC from [IBM Entitled Systems Support \(ESS\)](#), or it can be included in an initial IBM Power S1122 or IBM Power S1124 order.

The virtual HMC (vHMC) supports the following hypervisors:

- ▶ On x86 processor-based servers
 - Kernel-based Virtual Machine (KVM)
 - Xen
 - VMware
- ▶ On Power servers
 - IBM PowerVM

The following minimum requirements are needed to install the vHMC:

- ▶ 16 GB of memory
- ▶ Four virtual processors
- ▶ Two network interfaces (maximum of four allowed)
- ▶ One disk drive with 500 GB of available capacity

For an initial Power S1122 or Power S1124 order that uses the IBM Configurator (e-config), the HMC virtual appliance can be found by selecting **Add software** → **Other System Offerings** and then choosing one of the following options:

- ▶ 5765-HMD. IBM HMC Virtual Appliance for Power 11
- ▶ 5765-HMU. IBM HMC Virtual Appliance x86 11

For more information and an overview of the vHMC, see [Virtual Hardware Management Console appliance \(vHMC\) overview](#). For more information about how to install the vHMC appliance and all requirements, see [Installing the HMC virtual appliance](#).

5.1.2 BMC network connectivity rules for the 7063-CR2 HMC

The 7063-CR2 HMC includes a Baseboard Management Controller (BMC), which is a specialized service processor that monitors the physical state of the system by using sensors. The OpenBMC implementation on the 7063-CR2 provides a GUI that can be accessed from a workstation with network connectivity to the BMC. This connection requires an Ethernet port that is configured for BMC use.

The 7063-CR2 provides two network interfaces (eth0 and eth1) for configuring BMC network connectivity on the appliance. Each interface maps to a different physical port on the system. Different management tools might use different names for these interfaces. The HMC task **Console Management** → **Console Settings** → **Change BMC/IPMI Network Settings** modifies only the Dedicated interface.

Table 5-1 lists the BMC ports.

Table 5-1 BMC ports

Management tool	Logical port	Shared or dedicated	CR2 physical port
OpenBMC UI	eth0	Shared	eth0
OpenBMC UI	eth1	Dedicated	Management port only
ipmitool	lan1	Shared	eth0
ipmitool	lan2	Dedicated	Management port only
HMC task (change the BMC/IPMI Network settings)	lan2	Dedicated	Management port only

Figure 5-2 shows the BMC interfaces of the HMC.




Figure 5-2 BMC interfaces

The main difference is that the shared and dedicated interfaces to the BMC can coexist. Each interface has its own local area network (LAN) number and physical port. Ideally, you configure one port, but both can be configured. The rules for connecting Power servers to the HMC remain the same as in previous versions.

5.1.3 High availability HMC configuration

For manageability and redundancy, a dual HMC configuration is a best practice. This configuration can include two hardware appliances, one hardware appliance and one virtual appliance, or two virtual appliances.

The following requirements must be met:

- ▶ Both HMCs must be at the same version.
- ▶ The HMCs must use different subnets to connect to the BMCs.
- ▶ The HMCs must be able to communicate with the servers' partitions over a public network to allow full synchronization and function.

5.1.4 HMC code level requirements

The minimum required HMC version for the IBM Power S1122 and Power S1124 is V11R1M1110. This version is supported only on the 7063-CR2 and vHMC appliances. It is not supported on the 7063-CR1 or the 7042 machine types. An HMC running V11R1M1110 cannot manage POWER7 processor-based servers.

An HMC running V11R1M1110 includes the following features:

- ▶ Support for managing Power11 processor-based servers
- ▶ Support for new I/O adapters
- ▶ VIOS management enhancements
 - Resource Groups
 - IBM i Secure Boot
 - Increased platform keystore size
 - Removal of support for vTPM 1.2
 - Quantum-safe LPM
 - Minimum Affinity Score and actions
- ▶ Console management and user experience improvements
 - User experience
 - Network topology enhancements
 - Trusted keystore improvements
 - Import certificate function
 - Multi-factor authentication allowlist
 - Ability to advertise device information through LLDP
- ▶ Power infrastructure maintenance and automated updates
 - Run Power platform updates with minimal touchpoints to reduce the risk of human error
 - Seamless platform updates with one-touch solutions
 - Automatic operational recovery and resiliency
- ▶ Autonomous Error Resolution (AER)
 - Ability to collect platform logs (firmware, hypervisor, HMC, and VIOS) from a single interface
 - Ability to create a case and upload logs directly from the HMC

- ▶ Sustainability
 - New Energy efficiency mode for higher performance per watt
 - Partition-level energy monitoring
 - Real-time monitoring and reporting of energy and carbon emissions at the VM or partition level
 - Scheduling of energy modes

5.1.5 Keeping your HMC up to date

In recent years, cybersecurity has emerged as a national security issue and an increasingly critical concern for CIOs and enterprise IT managers.

The IBM Power Architecture has consistently ranked highly for end-to-end security, which is why it remains a platform of choice for mission-critical enterprise workloads.

A key aspect of maintaining a secure Power environment is helping ensure that the HMC or vHMC is current and fully supported, including hardware, software, and Power firmware updates.

Outdated or unsupported HMCs represent a technology risk that can be mitigated quickly and easily by upgrading to a current release.

5.1.6 New features

The minimum HMC level that is required to support the Power11 platform is V11R1M1110. HMC V11 is not supported on the 7063-CR1. It runs on the 7063-CR2 or the vHMC. Also, HMC V11 does not support Power8 or earlier servers.

HMC V11R1M1110 became generally available (GA) in July 2025. Key new features include support for managing Power11 processor-based servers and new I/O adapters.

- ▶ Virtualization management enhancements
 - Resource Groups
 - IBM i Secure Boot
 - Increased platform keystore size
 - Removal of support for vTPM 1.2
 - Quantum-safe LPM
 - Minimum Affinity Score and actions
- ▶ Console management and user experience improvements
 - Network topology enhancements
 - Trusted keystore improvements
 - Import certificate function
 - Multi-factor authentication allowlist
 - Ability to advertise device information through LLDP
- ▶ Power infrastructure maintenance and automated updates
 - Run Power platform updates with minimal touchpoints to reduce human error
 - Update system firmware, VIOS, and I/O adapters from a single update flow

- Seamless platform updates with one-touch solutions for evacuation and return or in-place updates
 - Automatic operational recovery and resilience
 - Validation for LPM and VIOS redundancy (VIOS maintenance readiness check)
 - Automatic partition migration and return during updates
- ▶ AER
 - Collect platform logs (firmware, hypervisor, HMC, and VIOS) from a single interface
 - Create a case and upload logs directly from the HMC
 - ▶ Sustainability
 - New Energy efficiency mode for higher performance per watt
 - Partition-level energy monitoring
 - Real-time monitoring and reporting of energy and carbon emissions at the VM or partition level
 - Scheduling of energy modes

5.1.7 Using the automated maintenance tool

This section shows how to use the new automated maintenance tool that is provided by the Power11 HMC.

Launchpoint

Before you start automated maintenance, ensure that the system has the latest 11.10 firmware installed.

Complete the following steps:

1. To start the process, select **Update system** → **VIOS** → **Adapter levels**, as shown in Figure 5-3.

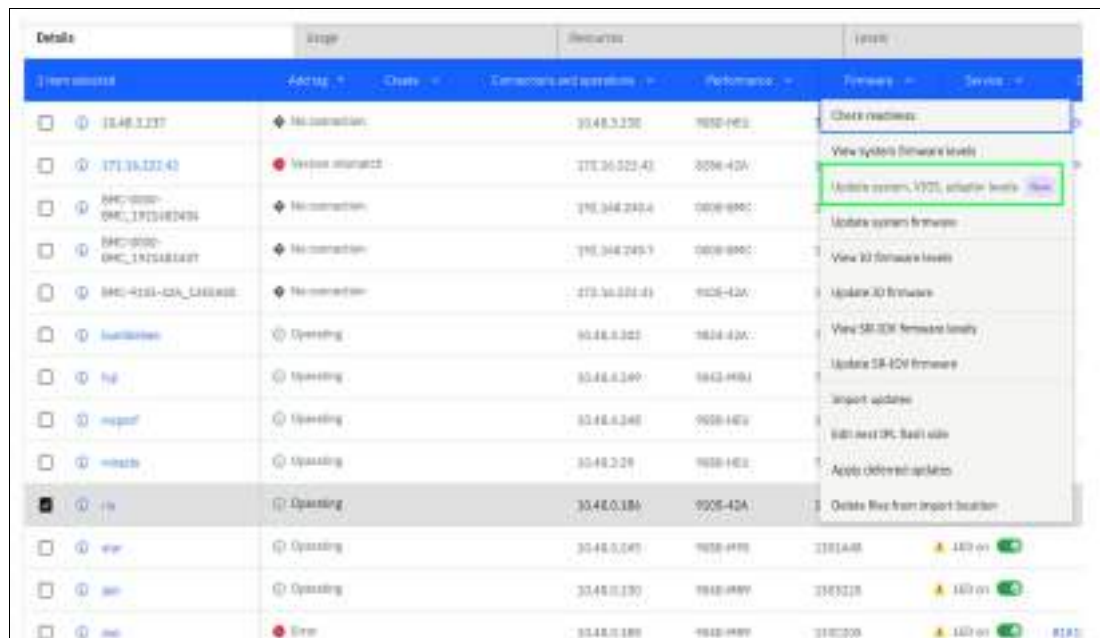


Figure 5-3 Starting the update

4. The following options are available:

- Remote HMC

Enables migration of partitions to a system that is managed by a remote HMC. When this option is selected, partitions are not automatically migrated back to the original system after the update.

- Partition Selection

By default, all partitions are selected for migration. You can select individual partitions by toggling the Server evacuation switch.

- Migration Sequence

When specific partitions are selected for migration, you can define the migration sequence for a set of partitions that must be migrated first.

The update process is automatically done. You can see the results in Figure 5-6.

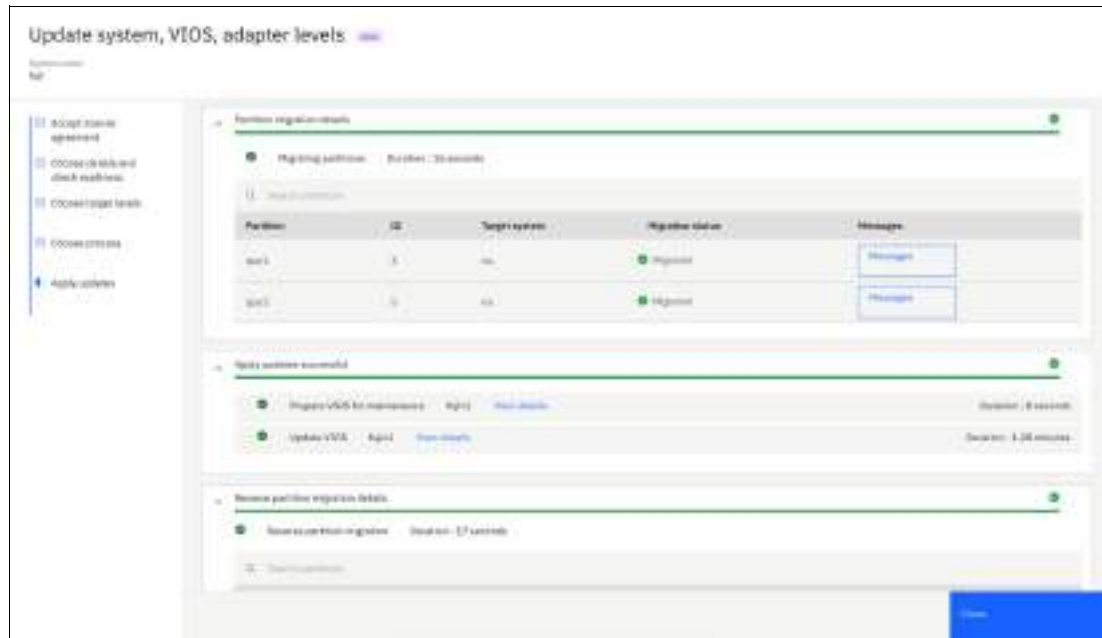


Figure 5-6 Update status

This solution provides unmatched business resiliency with zero planned downtime (ZPD) through platform automation capabilities that are integrated with IBM Concert AI-infused workloads. It enables faster, more frequent maintenance updates to keep systems secure, stable, and compatible with evolving software and hardware requirements, reducing the risk of performance degradation, security breaches, or unplanned downtime.

5.2 Ansible

Ansible is an open-source, cross-platform tool for resource provisioning automation. DevOps professionals use it for continuous software deployment by using an IaC approach. The Ansible automation platform evolved to deliver advanced automation solutions for operators, administrators, and IT decision-makers across multiple technical disciplines. It is an enterprise automation solution that is built on a thriving open-source ecosystem.

Ansible runs on several UNIX like platforms and can manage systems based on UNIX and Microsoft architectures. It uses a descriptive language to define system configurations.

Because of its broad adoption, open-source design, and extensive support for devices and platforms, Ansible is becoming a dominant tool in the market. However, it is common to use other automation tools alongside Ansible for more complex automation. For example, many organizations use Ansible with Terraform to automate infrastructure provisioning.

Ansible architecture

Figure 5-7 shows the Ansible architecture. It includes an Ansible controller and one or more Ansible client hosts. The controller runs automation tasks and stores Ansible collections. These collections contain modules, plug-ins, and roles that define the actions that Ansible performs on client nodes.

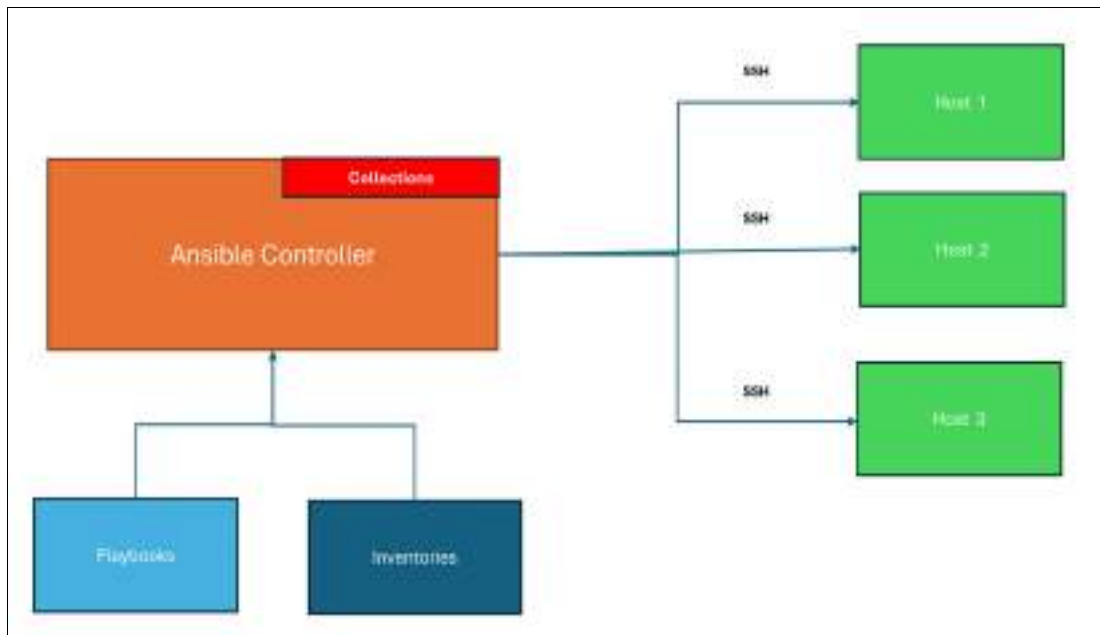


Figure 5-7 Simplified Ansible architecture

Playbooks

Ansible playbooks are YAML files that define sequences of tasks to run on remote hosts. These tasks include installing packages, configuring services, and copying files. Playbooks enable IT teams to automate infrastructure provisioning, configuration management, and application deployment.

Why choose Ansible

Ansible offers several benefits for IT professionals who want to improve efficiency, scalability, and consistency in their infrastructure. Here are some key advantages:

- ▶ Versatility
Ansible supports many devices and scales to meet growing automation needs.
- ▶ Agentless architecture.
Ansible manages devices by using Secure Shell (SSH), which removes the need for agents on target systems.
- ▶ Flexibility
Ansible handles simple CLI tasks and complex work flows that are defined in playbooks.

- ▶ Extensive module library
Ansible provides a large collection of modules for managing systems, cloud infrastructures, and OpenStack.
- ▶ Declarative approach
With declarative syntax, you define the system state, and Ansible takes the steps to achieve it.
- ▶ Minimal learning curve
The YAML syntax makes Ansible accessible to IT professionals at all levels.

Ansible is a powerful automation tool that helps organizations improve efficiency, scalability, and reliability in their IT infrastructure. By using Ansible playbooks, IT teams can streamline routine tasks, automate complex work flows, and help ensure consistent configurations across environments.

Options for implementing Ansible

When you implement Ansible for IT management, select the correct product and support level to meet your organization's needs. This section describes the available options.

Ansible Community

Here are the main community versions of Ansible:

- ▶ Ansible Core
The core automation engine. It is an open-source tool that provides basic functions for configuration management, application deployment, and task automation. Ansible Core includes modules, plug-ins, and the CLI that is needed to run playbooks and manage configurations.
- ▶ AWX
The upstream, open-source project that serves as the community version of Red Hat Ansible Tower. AWX provides a web-based UI, REST API, and task engine for managing Ansible automation at scale. It offers role-based access control (RBAC), job scheduling, graphical inventory management, and more.
- ▶ Ansible Collections
Pre-packaged modules, roles, and plug-ins that are created and shared by the community. Collections extend Ansible functions with more content that is maintained by the community or specific organizations. Collections can be downloaded from Ansible Galaxy, which is a hub for sharing and discovering Ansible content.
- ▶ Ansible Galaxy
A repository for sharing and discovering Ansible roles and collections. It is a community-driven platform that provides reusable content to simplify automation tasks. Users can search for roles and collections that are created by the community and integrate them into automation work flows.

These community versions are suitable for individual users, small teams, and development environments. However, they lack the formal support and advanced features that are provided by Red Hat Ansible Automation Platform.

Ansible Automation Platform

Ansible Automation Platform is a subscription-based enterprise solution that combines more than 20 community projects into a fully supported automation platform. It provides curated, certified, and validated Ansible Collections and roles from partners such as IBM, Juniper, Cisco, and public cloud providers.

Here are the key considerations for choosing Ansible Automation Platform:

- ▶ **Support level**
Enterprise-grade support with service-level agreements (SLAs) for security, compatibility, and upgrades. Community options offer limited support.
- ▶ **Features**
Includes capabilities beyond Ansible Core, such as a web interface and integration with other tools.
- ▶ **Cost**
Subscription-based. Community options are available at no charge.
- ▶ **Scale and complexity**
For large organizations with complex automation needs, Ansible Automation Platform is often the better choice because of its enterprise-grade features and support.

Evaluate these factors to select the Ansible offering that aligns with your organization's goals, budget, and support requirements.

Using Ansible to automate your IBM Power infrastructure

Ansible is a powerful automation tool that adds value to managing IBM Power servers. Here are some key use cases that improve efficiency, consistency, and scalability in Power environments:

- ▶ **System configuration and provisioning**
Ansible automates the setup and configuration of AIX, IBM i, and Linux on Power servers. Tasks include user and group management, network configuration, software installation, and system tuning. By using Ansible playbooks, administrators help ensure consistent configurations across multiple systems, reduce manual errors, and accelerate provisioning.
- ▶ **Fix management and compliance**
Keeping systems up to date is critical for security and stability. Ansible automates fixes for AIX and IBM i, including downloading fixes, applying them, and verifying system health after updates. It also supports compliance checks by validating configurations against predefined baselines, helping organizations meet regulatory and security standards.
- ▶ **Integration with HMC and PowerVC**
Ansible collections for IBM Power include modules that interact with HMC and PowerVC. This approach enables automation of tasks such as LPAR creation, VIOS configuration, and VM lifecycle management. Combined with dynamic inventory capabilities, Ansible discovers and manages Power infrastructure components in real time.

For more information about implementing automation with Ansible in an IBM Power environment, see *Using Ansible for Automation in IBM Power Environments*, SG24-8551.

5.3 Terraform

Terraform is an open-source tool that was originally developed by IBM HashiCorp® and is now owned by IBM. It is written in the Go programming language and compiles down into an executable file that is named Terraform. Terraform is an IaC tool that you can use to build, change, and produce versions of cloud and on-premises resources safely and efficiently. Terraform provides a mechanism to access any API for any cloud provider to manage infrastructure as a service (IaaS).

Figure 5-8 shows the process for calling the API. The configuration files define which APIs to call. These files represent the code that is referenced in IaC.



Figure 5-8 Terraform functions

Terraform uses a high-level configuration language called HashiCorp Configuration Language (HCL) to define and provision infrastructure. With Terraform, you can describe infrastructure components such as servers, databases, networking, and storage in code, enabling version control, collaboration, and repeatability. This approach eliminates manual setup and reduces the risk of configuration drift across environments.

Terraform operates through a work flow that includes writing configuration files, initializing the working directory, planning changes, and applying them. The **terraform init** command sets up the environment by downloading provider plug-ins. The **terraform plan** command creates an execution plan that shows the actions Terraform takes to reach the wanted state. Finally, **terraform apply** runs the plan and changes the infrastructure. Terraform maintains a state file that tracks the infrastructure state, which is essential for determining changes during future runs.

One of Terraform's key strengths is its provider ecosystem, which enables resource management across platforms such as PowerVS, Amazon Web Services (AWS), Azure, Google Cloud, Kubernetes, and others. This ecosystem makes Terraform a powerful tool for managing hybrid and multi-cloud environments. Terraform also supports modules, which are reusable configurations that promote consistency and reduce duplication. By codifying infrastructure, Terraform enables DevOps practices such as continuous integration and delivery (CI/CD), infrastructure testing, and automated deployments.

Figure 5-9 shows how Terraform works by defining the intended end state through configuration files, creating a plan, and applying it through infrastructure providers.

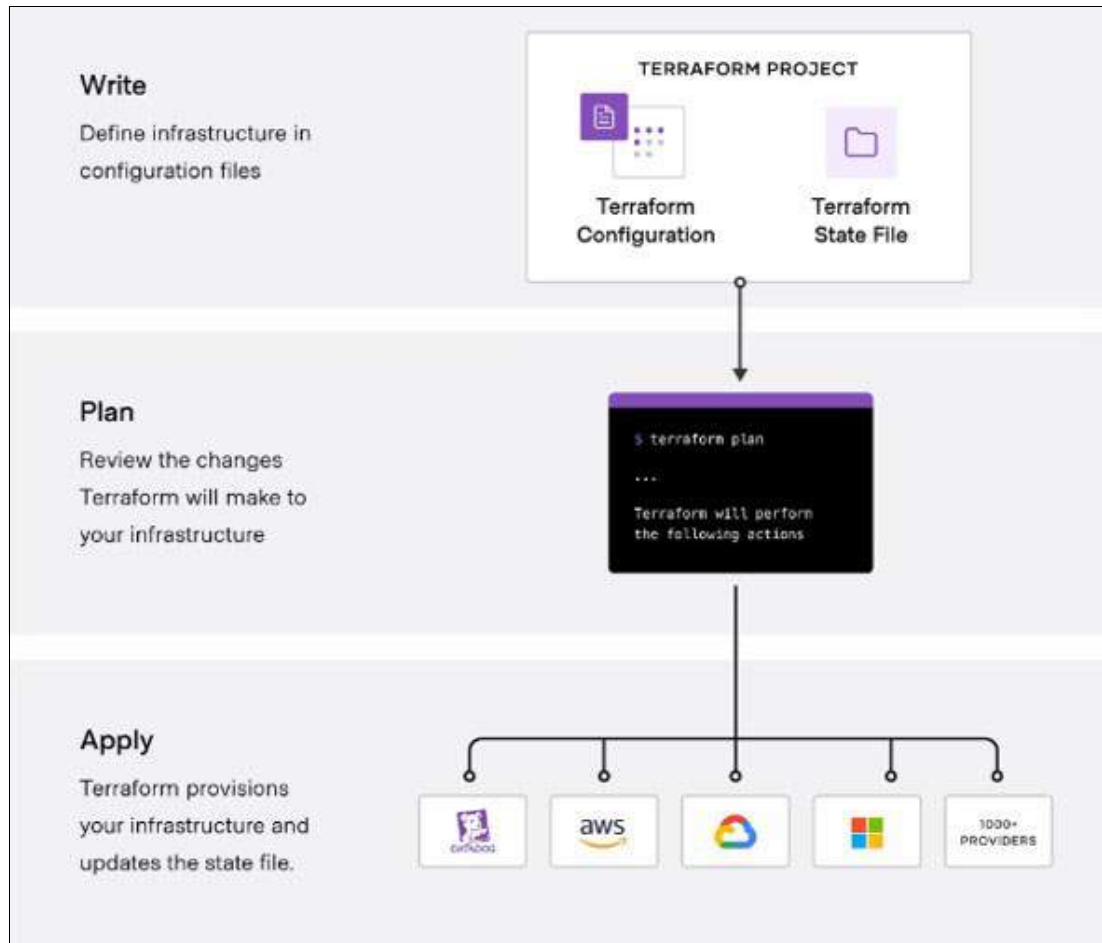


Figure 5-9 Terraform process

Here are some of the functions in the Terraform process:

- ▶ **Tracking your infrastructure**

Terraform generates a plan and prompts for approval before modifying infrastructure. It stores the infrastructure state in the `terraform.tfstate` file, which can be kept in Git, GitLab, or HCP Terraform to make versions, encrypt, and securely share with your team. This file acts as the single source of truth for your environment.
- ▶ **Automating changes**

Terraform configuration files are declarative and describe the wanted end state, so you can automate by using tools such as Ansible.
- ▶ **Standardizing configurations**

Terraform provides standardization through modules. A module is a collection of `.tf` and `.tf.json` files that are stored in a directory. Modules are the primary way to package and reuse resource configurations in Terraform.
- ▶ **Collaboration**

Terraform configurations can be saved as versions and shared through platforms such as Git, GitHub, and HCP Terraform, which are ideal for collaboration.

Combining Ansible and Terraform

Terraform connects to any provider, including Ansible, to manage infrastructure. You can browse the [Terraform Registry](#) for available providers. The Terraform provider for Ansible offers a more straightforward and robust way to run Ansible automation from Terraform than by using `local-exec`.¹

Figure 5-10 shows the Ansible provider entry in the registry.

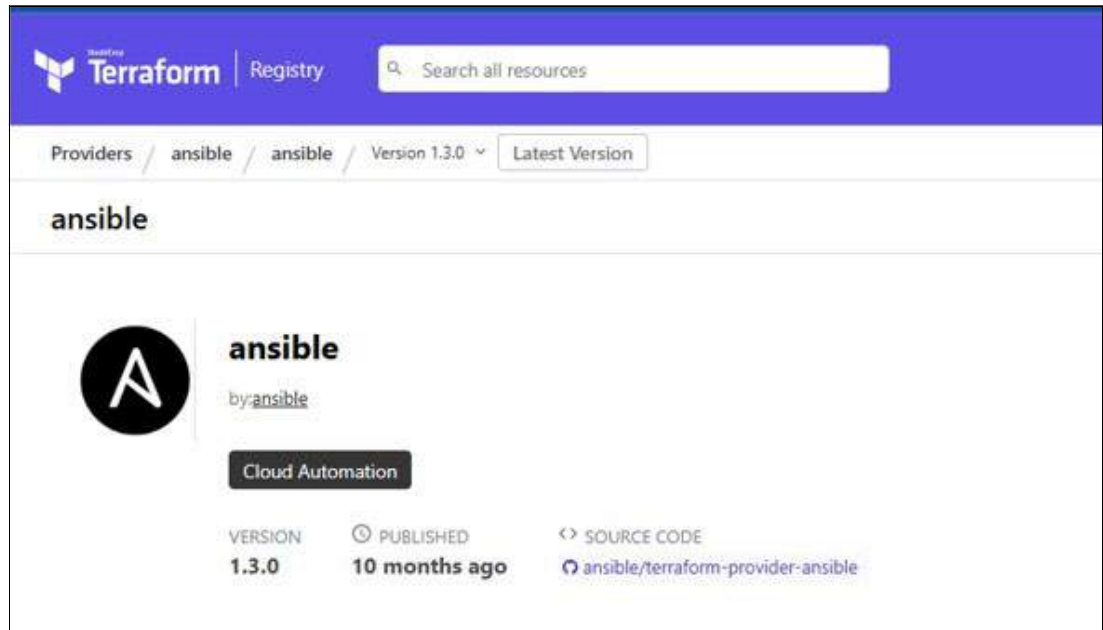


Figure 5-10 Terraform Ansible provider

Here are the prerequisites for using the Ansible provider:

1. Install Go. For installation instructions, see the [official installation guide](#).
2. Install Terraform. Install the ppc64le version for operation on IBM Power. You can find the installation instructions on the [GitHub registry](#).
3. Install Ansible. To install Ansible, see the [Ansible official installation guide](#)

For more information about Terraform in the IBM Power environment, see *Modernization Techniques for IBM Power*, SG24-8582.

5.4 PowerVC

IBM PowerVC is an advanced virtualization and cloud management platform that is built on OpenStack and designed for managing virtualized workloads on Power servers. PowerVC supports AIX, IBM i, and Linux OS and provides a unified interface for deploying, managing, and automating VMs across Power environments. PowerVC simplifies private cloud creation and integrates with higher-level cloud orchestrators such as Ansible, Terraform, and Red Hat OpenShift, enabling hybrid cloud strategies.

¹ Source: <https://developer.hashicorp.com/terraform/language/resources/provisioners/local-exec>

PowerVC offers two editions:

► **Standard Edition**

Includes core virtualization management features:

- VM image capture, import and export, and deployment
- Policy-based VM placement for optimized resource usage
- Snapshots and cloning for backup and testing
- Live VM mobility and remote restart for high availability (HA)
- RBAC and automated I/O configuration

► **Private Cloud Edition**

Builds on the Standard Edition by adding the following functions:

- A self-service portal for user VM provisioning
- Approval work flows for provisioning requests
- Prebuilt deployment templates
- Cloud management policies and metering for chargeback

PowerVC also supports Dynamic Resource Optimization (DRO), which automatically balances workloads based on CPU and memory usage, and Simplified Remote Restart, which helps ensure that VMs can restart on alternative hosts if there is a failure. These features reduce administrative impact and improve system resilience.

NovaLink is a lightweight, Linux-based virtualization management interface that acts as a bridge between PowerVC and IBM PowerVM Hypervisor (PHYP). It enables direct, scalable, and efficient control of VMs and system resources without relying solely on the traditional HMC.

NovaLink runs as a service on a dedicated LPAR on the same Power server it manages. This LPAR communicates directly with the PHYP and has direct access to hypervisor APIs. This approach enables faster and more granular control of virtualization tasks such as VM creation, deletion, migration, and resource allocation, making PowerVC more responsive and scalable, especially in environments with many VMs or frequent provisioning changes.

Table 5-2 shows a comparison of PowerVC with HMC management.

Table 5-2 Comparing PowerVC with HMC management:

Feature	PowerVC	HMC
Policy-based placement	Automatically places VMs based on resource availability and policies.	Manual placement.
Integration with cloud tools	Integrates with OpenStack, Red Hat OpenShift, and hybrid cloud platforms.	No native integration.
Live VM migration	Fully supported and automated.	Supported but requires manual steps.
DRO	Automatically balances workloads across hosts.	Not available.
Simplified networking storage setup	Automated configuration of virtual I/O, storage, and networking.	Manual configuration is required.

Feature	PowerVC	HMC
NovaLink support	Works directly with NovaLink for faster, scalable VM operations.	HMC is external and less scalable.
Scalability	Designed for large-scale environments with many VMs.	Suitable for smaller or static environments.
User interface	Modern, web-based UI with dashboards and templates.	Traditional interface that is more administrative.



Operating systems

IBM Power servers support three robust operating systems (OSs): IBM i, AIX, and Linux. Each OS is tailored to meet distinct enterprise computing needs. IBM i is an integrated operating environment that is known for reliability and security. It includes a built-in database (IBM Db2 for i), middleware, and development tools, making it ideal for running mission-critical business applications with minimal administrative impact. Its tight integration and compatibility with earlier versions make it widespread in industries such as finance, manufacturing, and retail.

AIX, IBM's UNIX-like OS, is designed for high-performance, scalable enterprise workloads. Built on a UNIX System V base and enhanced with IBM innovations, AIX offers advanced features such as dynamic system tuning, workload partitioning, and robust security. It is widely used in environments that demand stability and performance, including enterprise resource planning (ERP) systems, databases, and analytics platforms. AIX is optimized for IBM Power hardware, helping ensure seamless integration and efficient resource usage.

IBM Power also supports enterprise-grade Linux distributions, including Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server. These Linux environments provide flexibility and open-source innovation, making them ideal for cloud-native applications, containerized workloads, and artificial intelligence (AI) and machine learning (ML) development. Running Linux on Power enables organizations to use the performance advantages of the IBM Power Architecture while maintaining compatibility with the broader Linux ecosystem. Together, these OSs make IBM Power a versatile platform capable of supporting a wide range of modern and earlier workloads.

Complementing these OSs is Red Hat OpenShift, a Kubernetes-based container platform that runs natively on Power servers. Red Hat OpenShift enables organizations to build, deploy, and manage containerized applications across hybrid cloud environments with consistency and scalability. It provides a unified DevOps experience and supports modern application architectures, including microservices and AI workloads. This integration helps enterprises modernize their IT infrastructure while using the performance, security, and reliability of IBM Power.

This chapter provides the following topics:

- ▶ Subscription licensing
- ▶ AIX
- ▶ IBM i
- ▶ Linux on Power
- ▶ Red Hat OpenShift
- ▶ Introducing Kernel-based Virtual Machine support
- ▶ PowerVM Virtual I/O Server
- ▶ Setting your LPAR compatibility mode

6.1 Subscription licensing

IBM offers subscription-based licensing models for both AIX and IBM i. These models provide customers with greater flexibility, predictability, and cost efficiency. Organizations can license the OS and associated software on a monthly or annual basis instead of using perpetual licenses. This approach aligns with modern IT consumption trends, enabling businesses to scale usage up or down based on workload demands while simplifying budgeting and procurement processes. Subscription licensing also includes access to updates, fixes, and IBM Support to help ensure that systems remain secure and up to date.

For IBM i, subscription licensing includes the core OS along with integrated middleware such as Db2, security features, and application development tools. Similarly, AIX subscription licensing covers the base OS and can be bundled with other IBM software components. These subscriptions are available for deployment on Power servers, including the Power11 platform, and are suitable for hybrid cloud environments where agility and operational efficiency are essential. By adopting subscription licensing, organizations can modernize their licensing strategy while continuing to use the reliability and performance of AIX and IBM i.

6.1.1 AIX subscription licensing

AIX Standard Edition, AIX Enterprise Edition, IBM Private Cloud Edition, and IBM Private Cloud Edition with AIX are available under a subscription licensing model. This model provides access to an IBM program and IBM software maintenance for a specified subscription term of 1 year or 3 years. The subscription term begins on the start date and ends on the expiration date, as reflected on the IBM Entitled Systems Support (ESS) website.

Customers are licensed to run the product through the expiration date of the subscription term and can renew at the end of the term to continue using the product. This model offers flexible and predictable pricing over a specific period with a reduced upfront cost.

Another benefit of this model is that the licenses are *customer-number entitled*, meaning that they are not tied to a specific hardware serial number, which is the case with perpetual licenses. Therefore, the licenses can move between on-premises and cloud environments as needed, which is important for hybrid workloads.

Table 6-1 shows the product IDs for the subscription licenses.

Table 6-1 Subscription license product IDs (one or three year terms)

Product ID	Description
5765-2B1	IBM AIX 7 Standard Edition Subscription 7.3.0
5765-2E1	IBM AIX Enterprise Edition Subscription 1.10.0

Product ID	Description
5765-2C1	IBM Private Cloud Edition with AIX 7 Subscription 1.10.0
5765-6C1	IBM Private Cloud Edition Subscription 1.10.0 (without AIX)

Subscription licenses are orderable through IBM configuration tools. AIX perpetual and monthly term licenses for Standard Edition are still available.

6.1.2 IBM i subscription licensing

IBM offers several subscription licensing options for IBM i to provide flexibility and align with modern IT consumption models. These subscriptions are available on Power servers, including the Power11 platform, and are ideal for clients seeking to simplify licensing, reduce upfront costs, and scale.

The core offering is the IBM i Operating System Subscription, which includes the IBM i OS along with integrated middleware such as Db2 for i, security features, and development tools. This subscription is licensed on a per-core basis and is available in monthly or annual terms, so organizations may align software costs with usage. Support and updates are included to help ensure that systems remain secure and up to date without separate maintenance agreements.

In addition to the base OS subscription, IBM provides IBM i Value Packages, which bundle the OS with other IBM software such as IBM Rational® Development Studio for i, IBM Db2 Mirror for i, and IBM i Access Client Solutions. These packages are tailored to support specific workloads or modernization efforts.

For cloud-based deployments, the IBM i Cloud Subscription offers a flexible, usage-based model that includes the OS, support, and infrastructure when running IBM i on platforms such as IBM Power Virtual Server (PowerVS). This option is beneficial for organizations adopting hybrid or cloud-first strategies because it enables scalable, enterprise-grade IBM i environments without the complexity of traditional licensing.

6.2 AIX

AIX is a secure, scalable, and robust UNIX like OS that is based on open standards. For more than 30 years, AIX is the cornerstone of mission-critical computing for enterprise organizations in complex industries. It evolved to introduce hybrid cloud and open-source capabilities.

AIX 7.3 is the latest release. It builds on a solid foundation by offering new functions and capabilities that enhance performance, scalability, availability, and security while preserving application-binary compatibility to protect IT investments in AIX.

Coupled with IBM Power11 processor-based servers, AIX 7.3 provides an optimized and resilient computing platform that adapts to changing business demands, including new cloud use cases and improved economics.

6.2.1 AIX 7.3 key features

Here is a list of the key features in the latest AIX release:

- ▶ Workload scalability and automation

AIX 7.3 offers enhanced workload scalability, improved cloud automation through Ansible, and more than 300 open-source packages that are available through the AIX Toolbox for Open Source Software. These capabilities enable modern application development.
- ▶ Live kernel update

Introduced in AIX 7.2 and enhanced in Version 7.3, this feature enables interim fixes, Service Packs (SPs), and Technology Level (TL) updates without requiring restarts. It supports IBM Power Virtualization Center (PowerVC) managed environments and IBM Power Enterprise Pool (PEP) systems for resource optimization.
- ▶ High availability and disaster recovery (HADR)

IBM PowerHA and VM Recovery Manager (VMRM) provide automated recovery and multi-site replication to minimize downtime and help ensure business continuity in hybrid or public cloud environments.
- ▶ AI integration

AIX workloads are a natural source for AI because they host large volumes of high-quality data on customer behavior and transactions. This data can be used for AI development, enabling machine learning and deep learning for actionable insights on a unified platform.
- ▶ Cloud flexibility

AIX enables private, on-premises cloud transformation with PowerVC, offering hybrid cloud functions for seamless AIX virtual machine (VM) import and export. It also supports software-defined infrastructure for storage area network (SAN)-less DevOps environments. Available through PowerVS, AIX supports mission-critical databases with enhanced scalability, cloud automation, robust security, and flexible licensing. Workloads can run in hybrid or public clouds without refactoring, helping ensure reliability and efficiency.

6.2.2 Supported levels

At the time of writing, Power11 processor-based servers support the following minimum AIX OS levels when they are installed by using direct I/O connectivity:

- ▶ AIX 7.3 with TL 7300-03 and SP 1 or later
- ▶ AIX 7.3 with TL 7300-02 and SP 4 or later
- ▶ AIX 7.2 with TL 7200-05 and SP 10 or later

At the time of writing, Power11 processor-based servers support the following minimum AIX OS levels when they are installed by using virtual I/O:

- ▶ AIX 7.3 with TL 7300-03 and SP 0 or later
- ▶ AIX 7.3 with TL 7300-02 and SP 2 or later
- ▶ AIX 7.2 with TL 7200-05 and SP 8 or later

Important:

- ▶ Starting with AIX 7.3 TL 3 SP 1, AIX logical partitions (LPARs) operate in native Power11 processor compatibility mode, fully using Power11 capabilities.
- ▶ LPARs running versions earlier than AIX 7.3 TL 3 SP 1 operate in Power9 or Power10 processor compatibility mode on Power11 processor-based servers.
- ▶ LPARs running AIX 7.2 can operate in Power9 or Power10 processor compatibility mode.

6.2.3 AIX maintenance levels

IBM periodically releases maintenance packages (SPs and TLs) for the AIX OS. For more information about these packages, and downloading and obtaining installation files, see [IBM Fix Central](#). For more information about hardware features compatibility and the corresponding AIX TLs, see [IBM Support](#).

Service Update Management Assistant (SUMA), which automates checking and downloading OS updates, is part of the base OS. For more information about the `suma` command, see [Service Update Management Assistant \(SUMA\)](#).

The [Fix Level Recommendation Tool \(FLRT\)](#) provides cross-product compatibility information and fix recommendations for IBM products. Use FLRT to plan upgrades of key components or verify the current health of a system.

The [IBM AIX Operating System Service Strategy and Best Practices](#) resource is available at no cost to AIX clients. It provides insight into the AIX service strategy and helpful lifecycle information to maintain your version of AIX. For more information, see [AIX on IBM Power](#).

For more information, see the following resources:

- ▶ [AIX support lifecycle information](#)
- ▶ [System Software Maps](#)
- ▶ [System to AIX maps](#)

6.2.4 Licensing

The AIX OS is available as the following editions:

- ▶ AIX Standard Edition
- ▶ AIX Enterprise Edition
- ▶ AIX Cloud Edition

The Enterprise Edition and Cloud Edition include software that typically is required to manage larger IBM Power environments, including hybrid clouds.

There are two licensing models for AIX:

- ▶ CPU-based licensing
- ▶ Subscription licensing

For CPU-based licensing, AIX licenses can be ordered with the server or purchased later as a Miscellaneous Equipment Specification (MES) upgrade. The license grants you the right to use AIX on a specific server. To receive support for AIX, you must have a valid Software Maintenance Agreement (SWMA).

Subscription licensing is a newer model that offers greater flexibility in acquiring AIX. The subscription license includes access to an SWMA for a specified term of 1 or 3 years. After this term, you can renew the subscription to continue using AIX. This model provides predictable pricing over a specific period with reduced upfront costs.

Another benefit is that the licenses are customer-number entitled, meaning that you can use them on any Power server in your environment. You can also move them between on-premises, private cloud, and public cloud environments if needed.

6.2.5 New AIX editions

The IBM AIX OS is an open, standards-based UNIX like platform that is the foundation of mission-critical workloads and databases for tens of thousands of customers for more than 35 years. AIX provides an enterprise-class IT infrastructure that delivers reliability, availability, performance, and security.

At the time of writing, IBM offers the following updates and enhancements to AIX and related products for Power8, Power9, Power10, and Power11 processor-based servers:

- ▶ IBM AIX 7 Enterprise Edition 1.13
- ▶ IBM Private Cloud Edition 1.13
- ▶ IBM Private Cloud Edition with AIX 1.13
- ▶ VMRM 1.9

IBM AIX 7 Enterprise Edition 1.13

IBM updated AIX 7 Enterprise Edition and its corresponding subscription offering to Version 1.13. Here are the bundled software components that are included with AIX 7 Enterprise Edition 1.13 (5765-CD3 and 5765-2E1):

- ▶ IBM AIX 7.3 TL3 or IBM AIX 7.2 TL5
- ▶ IBM PowerSC 2.3
- ▶ IBM PowerVC for Private Cloud 2.3.1
- ▶ IBM VMRM HA 1.9
- ▶ IBM Tivoli® Monitoring 6.3

Here are the bundle components that were updated:

- ▶ AIX 7.3 TL3 updated with SP 1.
- ▶ AIX 7.2 TL5 updated with SP 10.
- ▶ IBM PowerSC updated from Version 2.2 to Version 2.3.
- ▶ IBM PowerVC for Private Cloud updated from Version 2.3.0 to Version 2.3.1.
- ▶ IBM VMRM HA updated from Version 1.8 to Version 1.9.

Additional information for IBM AIX 7 Enterprise Edition

Clients with an active SWMA or subscription for earlier versions of AIX Standard Edition or AIX Enterprise Edition are entitled to upgrade at no charge. To update, download, or install, see the [IBM ESS](#) website. Clients can choose either AIX 7.3 TL 3 or AIX 7.2 TL 5. Clients selecting AIX 7.2 TL 5 can upgrade to AIX 7.3 TL 3 at any time if their SWMA or subscription is current. Clients with AIX Enterprise Edition can trade up to IBM Private Cloud Edition with AIX.

IBM Private Cloud Edition 1.13

IBM updated the Private Cloud Edition and its corresponding subscription offering to Version 1.13. Here are the bundled software components that are included with Private Cloud Edition 1.13 (5765-ECB and 5765-6C1):

- ▶ IBM PowerSC 2.3
- ▶ IBM PowerVC for Private Cloud 2.3.1
- ▶ IBM VMRM DR 1.9
- ▶ IBM Tivoli Monitoring 6.3

Here are the bundle components that were updated:

- ▶ IBM PowerSC updated from Version 2.2 to Version 2.3.
- ▶ IBM PowerVC for Private Cloud updated from Version 2.3.0 to Version 2.3.1.
- ▶ IBM VMRM DR updated from Version 1.8 to Version 1.9.
- ▶ IBM Cloud Management Console (IBM CMC) updated from Version 1.22 to Version 1.23.

IBM Private Cloud Edition with AIX 1.13

IBM updated Private Cloud Edition with AIX and its corresponding subscription offering to Version 1.13. Here are the bundled software components that are included with Private Cloud Edition with AIX 1.13 (5765-CBA and 5765-2C1):

- ▶ IBM AIX 7.3 TL 3 or IBM AIX 7.2 TL 5
- ▶ IBM PowerSC 2.3
- ▶ IBM PowerVC for Private Cloud 2.3.1
- ▶ IBM VMRM DR 1.9
- ▶ IBM Tivoli Monitoring 6.3

Here are the bundle components that were updated:

- ▶ AIX 7.3 TL 3 updated with SP 1.
- ▶ AIX 7.2 TL 5 updated with SP 10.
- ▶ IBM PowerSC updated from Version 2.2 to Version 2.3.
- ▶ IBM PowerVC for Private Cloud updated from Version 2.3.0 to Version 2.3.1.
- ▶ IBM VMRM DR updated from Version 1.8 to Version 1.9.
- ▶ IBM CMC updated from Version 1.22 to Version 1.23.

Additional information for IBM Private Cloud Edition

Private Cloud Edition 1.13 and Private Cloud Edition 1.13 with AIX 7 include an entitlement for a subscription to IBM CMC (5765-CMT) for the same term as their SWMA. Clients with an active SWMA or subscription for earlier versions of Private Cloud Edition or Private Cloud Edition with AIX are entitled to upgrade at no charge. To update, download, or install, see the [IBM ESS](#) website. Clients can choose either AIX 7.3 TL 3 or AIX 7.2 TL 5. Clients selecting AIX 7.2 TL 5 can upgrade to AIX 7.3 TL 3 at any time if their SWMA or subscription is current. Clients with AIX Enterprise Edition can change to Private Cloud Edition with AIX.

VM Recovery Manager 1.9

VMRM provides automated VM management in the data center and disaster recovery (DR) management between sites. In a PowerVS public cloud, the product is called DR Automation. DR Automation provides automated DR management between geographically dispersed PowerVS data centers.

There are other capabilities:

- ▶ Workgroup support for IBM PowerHA
- ▶ Failover rehearsal for PowerVS

6.3 IBM i

IBM i is a powerful, integrated OS that is designed to run on IBM Power hardware. It is known for robustness and security. It changed names several times, reflecting technological advancements and IBM's strategic branding changes.

IBM i remains a premier enterprise computing platform that is recognized for strong security, scalability, and seamless integration. It supports modern workloads while preserving application compatibility, which makes it useful in industries such as finance, retail, and manufacturing, where reliability and performance are critical. IBM's ongoing innovation helps ensure that IBM i, a key component of IBM Power, continues to help enterprises efficiently manage and deploy their most critical applications, building on its stability and business continuity.

6.3.1 Introducing IBM i

IBM i is a highly integrated and dependable platform that streamlines IT infrastructure by combining the OS, Db2 for i database, security, and middleware into a single environment. This integration reduces the need for complex configurations and third-party components. As a result, organizations benefit from a reduced total cost of ownership (TCO) and reduced administrative impact, which is valuable in environments with limited IT resources.

One of the IBM i key strengths is its reliability, availability, and serviceability (RAS). Designed to support mission-critical workloads, it offers features such as automatic error detection, self-healing capabilities, and robust HADR options. These capabilities make IBM i a trusted platform across industries such as finance, healthcare, and manufacturing, where system uptime and data integrity are critical. IBM i also supports earlier and modern application development so that businesses can run RPG and COBOL applications alongside modern languages such as Java, Python, and Node.js while integrating open-source tools and cloud-native services.

Since 2008, IBM released new versions of IBM i approximately every 3 years, with IBM i 7.6 as the latest. Each release introduces new features, workload support, and language enhancements while maintaining compatibility with earlier versions so that customers can preserve and extend existing investments. IBM i is recognized as one of the most secure, self-optimizing, and scalable OSs in the world. IBM continues to deliver updates and enhancements throughout each release lifecycle through Program Temporary Fixes (PTFs) and Technology Refreshes (TRs). A key differentiator in the enterprise OS market is IBM's long-term commitment to IBM i, demonstrated by its publicly available roadmap, which at the time of writing extends to 2035 and is updated annually.

Figure 6-1 on page 149 shows the roadmap.




Figure 6-1 IBM i support matrix

For more information about IBM i and its s functions, see *IBM i 7.6 features and function*, SG24-8588.

6.3.2 Supported levels

For the IBM Power S1122, Virtual I/O Server (VIOS) is required for IBM i support. IBM i partitions must be set to restricted I/O mode, and there are limitations on the maximum partition size (up to four cores per IBM i partition are supported). Multiple IBM i partitions can be created and run concurrently, and each partition can have up to four cores.

Here are the minimum supported levels:

- ▶ IBM i 7.6 with Power11 GA PTF Group or later
- ▶ IBM i 7.5 TR 6 or later
- ▶ IBM i 7.4 TR 12 or later

The IBM Power S1124 supports the following minimum levels of IBM i, both natively and as a client of VIOS:

- ▶ IBM i 7.4 TR 12
- ▶ IBM i 7.5 TR 6
- ▶ IBM i 7.6

6.3.3 New IBM i features that are supported by the Power11 platform

IBM i 7.6 Power11 GA PTF Group, IBM i 7.5 Hardware TR 6, and IBM i 7.4 Hardware TR 12 all support the new IBM i Power servers with Power11 technology and their new I/O offerings.

Several new features are available only with Power11 technology and newer IBM i code levels:

- ▶ NVMe Hot Spare (IBM i 7.6)
Enables immediate replacement of a mirror-protected disk unit under a failing NVMe device. For more information, see [NVMe Hot Spare](#).
- ▶ Secure Boot (IBM i 7.6)
Verifies the authenticity of the boot process. Security administrators can extend Secure Boot technology that is used for PowerVM firmware, such as **hostboot**, IBM PowerVM Hypervisor (PHYP), and partition firmware (PFW) to include IBM i Licensed Internal Code (LIC) through digital signatures. The IBM i OS verifies the integrity of the boot chain up through IBM i LIC before running by using a digital signature and a hardware-protected asymmetric key.
- ▶ Common Cryptographic Architecture Cryptographic Service Provider (CCA CSP)
Delivered as IBM i Option 35 with IBM i 7.6, 7.5, and 7.4, supporting the IBM 4770 Cryptographic Coprocessor with release 8.4.x for CCA.
- ▶ NX-accelerated compression (IBM i 7.5 and 7.6)
Enables accelerated DEFLATE compression for PASE programs on Power11 processor-based servers by using the zlibNX library, as supported in AIX. The pigz program automatically benefits, reducing CPU usage by up to 95% and speeding up compression tasks.
- ▶ Tape Library Virtualization (IBM i 7.6)
Extended to include IBM 3584 tape libraries when running on Power10 and Power11 processor-based servers. This enhancement builds on existing virtualization capabilities for libraries such as the IBM 3572, IBM 3573, and IBM 3555. Tape Library Virtualization enables multiple IBM i partitions to share a single physical tape library, reducing infrastructure costs by eliminating the need for a dedicated tape drive per LPAR. Support also extends to Virtual Tape Libraries that emulate the 3584 model. For more information, see [IBM i Removable Media: Tape Library Virtualization](#).

6.3.4 IBM i software tiers

The IBM i software tiers for the Power S1122 and Power S1124 range from Power10 to P30, depending on the configuration of your server.

Power S1122 software tiers for IBM i

All IBM Power S1122 configurations are designated as IBM i software tier Power10. IBM i support is provided at a cost-effective Power10 software tier even though the Power S1122 has two sockets.

There are limitations on partition size, and all I/O must be virtualized through VIOS. VIOS is required, and IBM i partitions must be set to restricted I/O mode. Up to four cores (real or virtual) per IBM i partition are supported. Multiple IBM i partitions can be created and run concurrently, and each partition can have up to four cores.

Native IBM i support is available on the Power S1122 (MTM 9824-22A) only when populated with two 4-core processors (Feature Code ERGR), with a maximum of eight cores active. This configuration is available at a Power10 IBM i software tier and supports IBM i natively, virtualized, or as a combination of both.

When configuring the Power S1122 (MTM 9824-22A) with two 4-core processors (Feature Code ERGR), IBM i partitions with more than four cores can be used. When configuring these systems with any other processor feature, the four-core-per-partition limit is enforced.

Consider the following information:

- ▶ Feature Code ERGR cannot be upgraded to any other processor feature.
- ▶ Two Feature Code ERGR 4-core processors constitute an IBM i Power10 processor group/ or software tier. Feature Code ERGR processor is QPRCFEAT ERGR.
- ▶ Systems with this configuration also require IBM i Enablement for 9824-22A (Feature Code EEPS) with two Feature Code ERGR processors.

Power S1124 software tiers for IBM i

All Power S1124 servers are either P20 or P30 software tiers. Table 6-2 shows the details.

Table 6-2 Power S1124 IBM i software tiers

Number of cores	Feature Code	Number of processor modules	QPRCFEAT	IBM i software tier
16	EP3X	1	EP3X	P20
32	EP3X	2	EP3X	P20
48	EP3H	2	EP3H	P30
64	EP3Y	2	EP3Y	P30

Note: At the time of writing, there is not a Power11 processor-based server with P05 support. If you need P05 support, consider one of these options:

- ▶ Power S1012 (9028-21B) 1-core or 4-core processors
- ▶ Power S1014 (9105-41B) 4-core processor

It is expected that IBM will produce a Power11 replacement for the Power S1012 that will support the P05 tier sometime in the future.

6.3.5 IBM i licensing

On the Power11 platform, IBM i licenses are available only through subscription, except for customers performing MES or serial number upgrades from IBM Power E1080 to IBM Power E1180, where licensing remains unchanged.

For new Power10 purchases, IBM announced the withdrawal of perpetual IBM i licenses for P20 and P30 software tiers in July 2025, effective January 1, 2026.

The following actions may be done with new Power10 purchases before January 2026:

- ▶ Purchase new subscription licenses.
- ▶ Purchase new perpetual licenses.
- ▶ Transfer existing perpetual licenses.

IBM previously withdrew new perpetual licenses for P05 and Power10 tiers.

When you acquire an IBM i license, IBM generates license keys that are unique to the server where the license is installed. These licenses fall into two categories based on the license keys that are provided:

- ▶ Non-expiring (perpetual). License keys have no expiration date.
- ▶ Subscription. License keys expire at the end of the subscription term.

IBM i licenses are also grouped into *Software Tiers*, also called *Processor Groups*. These tiers correspond to the size of the hardware platform that runs the OS. The Software Tiers are P05, P10, P20, and P30. For a list of Power machines and their software tiers, see 6.3.4, “IBM i software tiers” on page 150.

The IBM i licensing metric differs across tiers:

- ▶ P05. Based on number of cores and number of users
- ▶ P10. Based on number of cores and number of users
- ▶ P20. Based on number of cores
- ▶ P30. Based on number of cores

For many customers, IBM i 7.6 is their first experience with subscription-based licensing. This change occurred on May 7, 2024 when IBM moved to subscription-only licenses for customers purchasing specific P05 and P10 tier machines.

Note: IBM plans to migrate all software tiers to this subscription model in the future.

User licensing

For IBM i licenses for the P05 and Power10 tiers, clients can choose how many users to license for IBM i. IBM i users are ordered in bundles of 10, or an unlimited user license is available. As stated in the list in 6.3.5, “IBM i licensing” on page 151, the P20 and P30 tiers are licensed solely core. Customers at the P20 and P30 tiers can have any number of IBM i users.

For more information on IBM i subscription licensing, including a comprehensive FAQ section, see [IBM i Subscription Transformation FAQ](#).

6.3.6 IBM i subscription bundles for P20 and P30

Subscription bundles combine multiple products under a single product ID. All products share start and end dates for the subscription, and each product retains its own terms. This approach simplifies licensing by streamlining purchase decisions and license compliance, and it increases the flexibility to deploy capabilities. Bundles accelerate innovation because you gain more capabilities from the start, helping you modernize your infrastructure. They also reduce total investment because bundles cost less than individual licenses, lowering both upfront and ongoing expenses.

P20 and P30 customers can choose either stand-alone subscriptions or subscription bundles:

- ▶ IBM i P20 Standard Edition
- ▶ IBM i P30 Enterprise Edition

IBM i P20 Standard Edition

The IBM i P20 Standard Edition bundle is a simplified bundle with core capabilities. IBM created the subscription bundles to simplify the acquisition of IBM i and IBM i enterprise capabilities.

Figure 6-2 shows the contents of the P20 Standard Edition bundle.



Figure 6-2 P20 Bundle contents

Note: The * character in Figure 6-2 on page 153 designates a license for assets only.

IBM i P30 Enterprise Edition

The IBM i P30 Enterprise Edition is a strategic bundle that offers comprehensive IBM i enterprise capabilities to enhance availability and cyber resilience, improve productivity through automation, modernize application development, and use AI.

Figure 6-3 shows the contents of the P30 bundle.

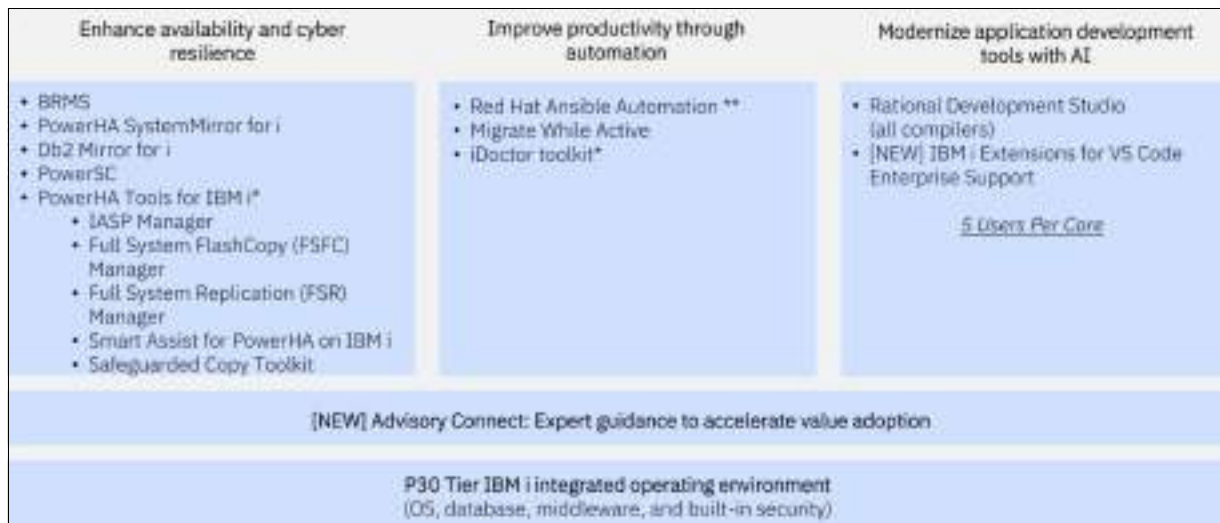


Figure 6-3 P30 bundle contents

Note:

- ▶ The * character in Figure 6-3 on page 154 designates a license for assets only.
- ▶ The ** characters in Figure 6-3 on page 154 designate that for Red Hat Ansible Automation, there is 1 year of support for 100 managed nodes with a minimum purchase of 10 P30 Enterprise Edition cores

Subscription bundles and Power Enterprise Pools 2.0

Power Enterprise Pools (PEP) 2.0 enables sharing of IBM i OS and IBM PowerHA SystemMirror® across systems. For more information about PEP, see 10.1, “IBM Power Private Cloud with Shared Utility Capacity” on page 226.

In a pool, you can mix systems with the following licenses:

- ▶ Subscription licenses
 - Subscription bundles for IBM i P30 Enterprise Edition or IBM i P20 Standard Edition
 - Stand-alone subscriptions
- ▶ Perpetual licenses

Each system must use a single license type, either subscription or perpetual. Licenses are shared as base capacity or available as metered capacity by the minute in a pool.

Other IBM i licensed program products (LPPs) must be licensed per system based on peak usage. They are not shared as base capacity or available as metered capacity by the minute in a pool.

Note: For each serial number, a system can have either an OS subscription or an OS perpetual license, but not both.

Rational Development Studio for i named user license

Rational Development Studio for i named user licenses that are acquired through IBM i P30 Enterprise Edition can be shared across all systems with Enterprise Edition licenses. For example, if you have three systems with 10 cores each, for a total of 30 cores, you are entitled to 150 named users (five users per core). These 150 named users can be assigned to any or all three systems.

IBM i license terms

IBM i license terms do not change based on subscriptions, subscription bundles, or perpetual licenses.

Determine the number of processors that are used by the program in a machine by using the smaller value from one of the following methods:

- ▶ The total number of activated processors in the machine.
- ▶ The sum of a and b is as follows (round any remaining fraction of a processor up to a full processor in the final calculation):
 - When IBM i runs in partitions with dedicated processors, use the sum of the processors that are assigned to those partitions.
 - When IBM i runs in partitions that are members of a shared processor pool (SPP), for each SPP, use the smaller of the following values:
 - The maximum capacity of the SPP
 - The sum of the virtual processors of each uncapped partition plus the processing units that are assigned to each capped partition running the program

6.3.7 IBM i Software Maintenance Agreement

SWMA is included with subscription licenses. For non-expiring licenses, 1 year of SWMA is included with the license purchase, and extra years are billed separately.

IBM i SWMA is critical for your business. You use it to engage with IBM Support centers worldwide, and it entitles you to software updates, called PTFs on IBM i, and new OS releases. These PTFs deliver new functions and enhancements across the IBM i portfolio, including security, database, high availability (HA), performance improvements, and more.

6.3.8 Capacity BackUp and dedicated backup machine licensing models for IBM i

IBM offers more licensing models for IBM i customers:

- ▶ Capacity BackUp (CBU)

The CBU feature enables temporary transfer of IBM i license entitlements between a production server and a CBU-designated server when you order a new IBM i server with the CBU option. Both servers must be owned by the same company. CBU provides a cost-effective backup solution.

With the Power S1122 and Power S1124 CBU designation, you can temporarily transfer IBM i processor and user license entitlements that are purchased for a primary machine to a secondary CBU-designated system for HADR operations. This temporary transfer can reduce costs compared to purchasing licenses for the secondary system. Processor activations cannot be transferred.

The CBU specify feature (Feature Code 0444 or Feature Code 4891) is available only with a new server purchase. System prerequisites, registration, and approval are required before applying the CBU feature. Standard IBM i terms and conditions do not allow permanent or temporary transfer of processor or user license entitlements outside the approved CBU process. These entitlements remain with the machine that they were ordered for. When you register the association between your primary and CBU system, you must agree to terms and conditions for temporary transfers.

After registration and approval, you can temporarily move optional IBM i processor and user license entitlements from the primary system to the CBU system when the primary system is down or inactive. The CBU system then supports failover and role swapping for test, HA, and DR scenarios. *Temporary entitlement transfer* means that the entitlement is temporarily assigned to the CBU system during HA or DR operations.

Before transferring IBM i processor entitlements, the primary machine must have more than one IBM i processor license, and the CBU must have at least one IBM i processor license. The CBU must be configured to avoid compliance warnings before failover. An activated processor must be available on the CBU server to use the transferred entitlement. You can transfer any processor entitlements above the minimum if the primary system workload does not require them during transfer or replication reversal.

During temporary transfers, the CBU system does not update its internal entitlement records, and you might see noncompliance warnings. These warnings do not indicate actual noncompliance. For PowerHA, you need at least one IBM i and one PowerHA license, and temporary keys that are equal to the number of PowerHA licenses on the primary system. PowerHA for IBM i also supports compliance message suppression keys.

The minimum permanent entitlement on the CBU is one. You must license all permanent workloads, such as replication. For example, if replication consumes four cores at peak, you must permanently license four cores on the CBU.

Servers with P20 or higher software tiers do not allow user entitlement transfers; only processor entitlements can be transferred.

Here are the eligible primary systems for Power S1122 and Power S1124 CBUs:

- Power S1122 CBU (Power10 tier):
 - 4-core and 8-core servers (Feature Code ERGR)
 - 10-core and 20-core servers (Feature Code ERGQ)
 - 32-core server (Feature Code EBG8)
 - 48-core server (Feature Code EBG9)
 - 60-core server (Feature Code EBGA)
- Power S1124 CBU (P20 tier):
 - Power S1124 (9824-42A) with 12 or 32 cores
 - Power E1080 (9080-HEX)
 - Power S1024 (9105-42A) with 12, 24, 32, or 48 cores
- Power S1124 CBU (P30 tier):
 - Power S1124 (9824-42A) with 48 or 60 cores
 - Power E1080 (9080-HEX)
 - Power E980 (9080-M9S)
 - Power S1024 (9105-42A) with 32 or 48 cores

If your primary or CBU machine is sold or retired, temporary entitlement transfers must return to the original machine. For registration details and terms, see the [IBM Power Systems: Capacity BackUp website](#).

- ▶ **Dedicated Backup Machine (DBM)**

DBM licensing supports backup environments. Unlike CBU, a DBM machine does not need to belong to the same company as the production server. For DBM, IBM i license entitlements must cover all cores running IBM i on the server.

For more information about temporary keys for IBM i LPPs for CBU and DBM, see [IBM Power: Capacity BackUp](#).

6.3.9 Upgrading an existing IBM i license to Version 7.6

If you have an active IBM i SWMA, you can upgrade your IBM i license to Version 7.6 at no additional charge.

If you have IBM i subscription licenses at Version 7.5 or earlier, go to the [IBM i Subscription Transformation FAQ](#) and search for “upgrade”.

If you have IBM i non-expiring licenses at Version 7.5 or earlier, upgrade IBM i through the [IBM ESS website](#). Select **My Entitled Software** → **Software Updates**. Select your IBM system by customer number, server model, and serial number, and then choose the option to upgrade from the current release to Version 7.6. After completing this process, you are directed to the Software Download section to obtain installation media for installing or upgrading your system.

After upgrading your IBM i non-expiring or subscription licenses to Version 7.6, use the IBM ESS website to download IBM i license keys for Version 7.6. Until license keys are installed, a 70-day evaluation period is available. Failure to install license keys makes IBM i and LPPs unusable until the appropriate keys are installed.

Your licensing model does not change during the upgrade process. If you had a non-expiring or subscription license before the upgrade, you retain the same license type afterward.

Upgrades to IBM i 7.6 are supported from IBM i 7.4 and IBM i 7.5.

6.3.10 IBM i simplification

IBM i and many products in the IBM i portfolio require license keys to help ensure proper software licensing. Recent simplification reduced the number of products that require keys. For more information about IBM i simplification, see [IBM i Portfolio Simplification](#)

6.4 Linux on Power

Linux is a powerful, open-source, cross-platform OS that runs on a wide range of hardware to provide a consistent, UNIX like environment across many architectures. Its compatibility with IBM Power offers a flexible and cost-effective option for running many applications while using the platform’s performance, availability, and reliability.

As a no-charge and open-source solution, Linux reduces TCO by eliminating licensing fees and enabling deep customization. It supports modern technologies such as containers, Kubernetes, and cloud-native applications, making it ideal for building scalable, agile IT environments. Its portability means that it can run seamlessly on architectures such as x86, ARM, and IBM Power. Linux is supported by a global community that provides extensive documentation, tools, and peer assistance.

IBM advocated for Linux for more than two decades, integrating it into its enterprise ecosystem, including Power servers, IBM Z mainframes, and IBM Cloud. IBM collaborates with major Linux distributions such as Red Hat, SUSE, and Canonical to help ensure optimized performance and integration. Through initiatives such as LinuxONE and contributions to the Linux kernel and open-source projects, IBM positions Linux as a cornerstone of secure, modern infrastructure.

Security is a key strength of Linux. Its user privilege model combined with community-driven updates and tools such as SELinux and AppArmor helps protect systems from vulnerabilities and malware. Linux is efficient and lightweight, and it can run many platforms, such as IoT devices to high-performance enterprise servers. With its modular design, you can tailor it to your needs, whether that means a minimal installation or a full desktop environment.

Linux also provides access to a large ecosystem of tools and software, some of which are exclusive or better supported on Linux than on other OSs. Its open-source nature offers flexibility and adaptability to align with business or technical requirements. Linux serves as a strong foundation for hybrid cloud infrastructure, supporting application modernization and deployment at scale. Its scalability, reliability, and cost-effectiveness make it suitable for enterprise environments and mission-critical workloads.

6.4.1 Supported distributions

The Linux distributions that are describe in this section are supported on Power scale-out servers. Other distributions, including open-source releases, can run on these servers but do not include formal enterprise-grade support.

Red Hat Enterprise Linux

The latest version of the RHEL distribution is supported in native Power11 mode, which enables full access to all Power11 processor and platform features.

At announcement, Power S1122 and Power S1124 servers support the following minimum RHEL levels:

- ▶ RHEL 10: Native Power11 support, fully using advanced architecture features
- ▶ RHEL 9.6: For Power LE or later (Power11 native support)
- ▶ RHEL 9.4: For Power LE or later (Power10 compatibility)
- ▶ RHEL 8.10: For Power LE or later (Power10 compatibility)

Red Hat Enterprise Linux roadmap

Power servers support Linux workloads while using the performance, reliability, and availability of Power hardware. This section outlines supported enterprise Linux distributions. Other distributions may run on Power, but only select enterprise versions provide formal support.

RHEL is a leading open-source platform for Linux, hybrid cloud, containers, and Kubernetes. It offers secure, transparent, and proactive lifecycle management to support autonomous operations. Lifecycle planning is essential for customers, partners, independent software vendors (ISVs), and the broader ecosystem.

Starting with RHEL 8, the lifecycle includes three phases: Full Support, Maintenance Support, and Extended Life, as shown in Figure 6-4 on page 159. Red Hat also publishes projected release timelines and extended support details for minor versions. RHEL 8, 9, and 10 each provide a 10-year lifecycle across the first two phases, followed by Extended Life.



Figure 6-4 Red Hat Enterprise Linux support lifecycle

Note: The RHEL lifecycle phases are designed to minimize changes within each major release over time and help ensure predictable availability and content.

Here are the lifecycle phases:

► Full Support phase

During this phase, Red Hat provides the following updates:

- Security updates for Common Vulnerabilities and Exposures (CVEs) with a CVSS score of 7 or higher, delivered through Red Hat Security Advisories (RHSAs)
- Urgent and select high-priority bug fixes through Red Hat Bug Advisories (RHBA)
- Extra errata as needed

New or improved hardware support and select software enhancements, typically in minor releases. Minor releases are cumulative and focus on resolving medium or higher-priority issues. Updated installation images are also provided.

► Maintenance Support phase

For RHEL 8, 9, and 10, Red Hat continues to release the following items:

- Security updates for CVEs with a CVSS score of 7 or higher
- Urgent and select high-priority bug fixes

New features and hardware support are not included during this phase.

► Extended Life phase

Customers retain access to existing content through the Red Hat Customer Portal, including documentation and migration guidance. Only limited technical support is available. No new fixes, hardware support, or root-cause analysis is provided. Support applies only to existing installations, and Red Hat may end support at its discretion.

Release cadence

To provide predictability, minor RHEL releases are scheduled every 6 months during the Full Support phase. Each release includes details on Extended Update Support (EUS) and SAP-specific services.

Figure 6-5 shows the planned release schedule for RHEL 8.

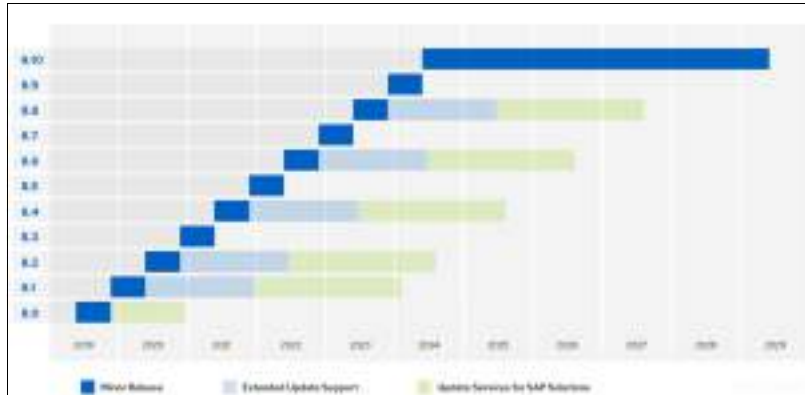


Figure 6-5 Release schedule for RHEL8

Figure 6-6 shows the planned release schedule for RHEL 9.




Figure 6-6 Release schedule for RHEL9

Figure 6-7 shows the release schedule for RHEL 10.




Figure 6-7 Release schedule for RHEL10

For more information regarding the RHEL roadmap, see the [Red Hat Enterprise Linux Life Cycle](#). To find the latest Red Hat certified Power servers, see the [Red Hat Certified Hardware catalog](#) and [System to Red Hat Enterprise Linux maps](#).

SUSE Linux Enterprise Server

The latest version of the SUSE Linux Enterprise Server distribution of Linux from SUSE is supported in native Power11 mode. It can access all features of the Power11 processor and platform. At announcement, the Power E1050 servers support SUSE Linux Enterprise Server 15 SP 6 or later (Power11 native)

SUSE roadmap

SUSE Linux Enterprise Server is designed for long-term enterprise stability and support. Its product lifecycle provides predictable maintenance phases, enabling organizations to plan deployments, updates, and migrations effectively.

► Lifecycle duration

SUSE Linux Enterprise Server offers a 13-year lifecycle per major release, which is divided into two phases:

- General Support (10 years)

Includes full maintenance, security fixes, bug fixes, hardware enablement, and new certified third-party software. This phase is ideal for production environments requiring continuous updates and active support.

- Extended Support (3 years)

After General Support ends, Extended Support provides critical security updates and selected bug fixes, giving customers more time to migrate to newer versions while maintaining operational security.

► Release cadence

- Major releases

Every 4 years, introducing new features, platform support, and architectural improvements.

- Service Packs (SPs)

Released every 12 - 14 months, including incremental updates, feature enhancements, and hardware enablement.

Each SP is supported for 6 months after the release of the subsequent SP to provide customers time to validate and upgrade within a predictable window.

Figure 6-8 shows the SP schedule.




Figure 6-8 SUSE Linux Enterprise Server major releases and Service Packs

Figure 6-9 shows the release lifecycle, including long-term SP support.




Figure 6-9 SUSE Linux Enterprise Server long-term Service Pack support

For more information regarding the SUSE roadmap, see [SUSE Linux Enterprise Server Documentation](#) and [Product Support Lifecycle](#). To find the latest SUSE certified Power servers, see the [System to SUSE Linux Enterprise Server maps](#) and [SUSE YES Certified Hardware - Bulletin Search](#).

Linux and Power11 technology

The Power11 specific toolchain is available in IBM Advance Toolchain for Linux 15.0. This toolchain enables customers and developers to use all new Power11 technology instructions when programming. The cross-module function call impact is reduced through a new PC-relative addressing mode.

One key benefit of Power11 technology is a 10x to 20x performance advantage over Power9 for AI inferencing workloads, which is driven by increased memory bandwidth and new instructions. A notable example is the Matrix Math Accelerator (MMA), which is a purpose-built feature that is designed for machine learning and deep learning inference. MMA also supports multiple AI data types.

6.4.2 Licensing

Linux licensing on IBM Power follows the same principles as on other platforms, with considerations that are tailored for enterprise environments. The Linux kernel is licensed under the GNU General Public License version 2 (GPLv2), which helps ensure that the software remains a no-charge and open source so that you can run, study, modify, and distribute it.

On IBM Power, major supported Linux distributions include RHEL, SUSE Linux Enterprise Server, and Ubuntu. Each distribution adheres to open-source licensing models but differs in how software is packaged and delivered.

For example, RHEL and SUSE Linux Enterprise Server provide access to source code under open-source licenses but require a paid subscription for precompiled binary files, updates, and enterprise support. This model combines open-source flexibility with the stability and support that are needed for mission-critical workloads.

On Power servers, these distributions are optimized to use the architecture's performance, scalability, and reliability features. Licensing terms typically include support for virtualization technologies such as PowerVM.

Red Hat licensing

RHEL is sold on a subscription basis, with initial subscriptions and support that are available for 1, 3, or 5 years. Support is provided directly by Red Hat or through IBM Technical Support Services.

RHEL 8 for Power LE subscriptions cover up to four cores and up to four LPARs. Subscriptions can be stacked to cover more cores or LPARs.

When you order RHEL from IBM, a subscription activation code is automatically published in Enterprise Storage Server. After retrieving this code from Enterprise Storage Server, use it to establish proof of entitlement and download the software from Red Hat.

For more information about Red Hat licensing, see the [Red Hat Enterprise Linux subscription guide](#).

SUSE Linux Enterprise Server licensing

SUSE Linux Enterprise Server is sold on a subscription basis, with initial subscriptions and support available for 1, 3, or 5 years. Support is provided directly by SUSE or through IBM Technical Support Services.

SUSE Linux Enterprise Server 15 subscriptions cover one socket or one LPAR and can be stacked to cover more sockets or LPARs.

When you order SUSE Linux Enterprise Server from IBM, a subscription activation code is automatically published in Enterprise Storage Server. After retrieving this code from Enterprise Storage Server, use it to establish proof of entitlement and download the software from SUSE.

You must register the purchased Linux offering on the distributor's website by using the activation code. After registration, you can download the software packages electronically and obtain the latest upgrades that are available for the product.

For more information about SUSE Linux Enterprise Server licensing, see [SUSE Licensing](#).

Linux and Power11 technology

The Power11 specific toolchain is available in IBM Advance Toolchain for Linux 15.0. This toolchain enables customers and developers to use all new Power11 technology instructions when programming. The cross-module function call impact is reduced through a new PC-relative addressing mode.

One key benefit of Power11 technology is a 10x to 20x performance advantage over Power9 for AI inferencing workloads, which is driven by increased memory bandwidth and new instructions. A notable example is the MMA, which is designed for machine learning and deep learning inference. MMA also supports multiple AI data types.

6.5 Red Hat OpenShift

Red Hat OpenShift on IBM Power provides a platform for building, deploying, and managing containerized applications across hybrid cloud environments. By using the performance and reliability of IBM Power Architecture, Red Hat OpenShift enables enterprises to modernize IT infrastructure with cloud-native technologies such as Kubernetes, containers, and microservices. With support for AIX, IBM i, and Linux, you can colocate containerized workloads with applications to reduce latency and improve data access. Integration with IBM Cloud services and automation tools simplify cluster provisioning, scaling, and lifecycle management, so developers can focus on innovation instead of infrastructure.

The latest release, Red Hat OpenShift 4.18, introduces enhanced capabilities for IBM Power, including improved networking with user-defined networks, advanced Operator Lifecycle Management, and deeper GitOps integration. These features support secure, scalable, and flexible deployments on-premises or in the PowerVS cloud. Red Hat OpenShift on Power also supports IBM Cloud Pak and AI workloads, enabling enterprises to accelerate digital transformation while maintaining enterprise-grade security and performance. This combination of technologies helps businesses innovate faster, optimize resource usage, and build a resilient hybrid cloud foundation.

For more information about how Red Hat OpenShift can improve your hybrid cloud experience, see 10.2, “Red Hat OpenShift” on page 228.

Licensing

Red Hat OpenShift Container Platform is available by subscription with terms of 1, 3, or 5 years. Each subscription covers two processor cores and can be stacked to meet workload requirements. Support is available from Red Hat or through IBM Technical Support Services.

When you order Red Hat OpenShift Container Platform for Power from IBM, you receive a subscription activation code through the IBM ESS portal. This code serves as proof of entitlement and is used to download the software from the Red Hat Customer Portal.

For more information and documentation, see the official [Red Hat OpenShift Documentation](#).

Supported Levels

IBM Power scale-out servers support Red Hat OpenShift Container Platform 4.18 and later. Red Hat OpenShift 4.18 supports the Power11 platform when running in Power10 compatibility mode at general availability. Full support for Power11 native mode is planned for Red Hat OpenShift 4.19.

Red Hat OpenShift Container Platform roadmap

Red Hat provides a defined product lifecycle for Red Hat OpenShift Container Platform to help customers and partners plan, deploy, and support infrastructure. This lifecycle is published for transparency, although exceptions can occur.

Red Hat OpenShift 4 follows a phased, time-based lifecycle with at least four minor versions supported concurrently, as shown in Figure 6-10 on page 165. Each minor version has a fixed support period that offers varying levels of maintenance. Red Hat targets a release cadence of every 4 months to support customer planning. All errata remain available to active subscribers throughout the lifecycle.



Figure 6-10 Red Hat OpenShift Container Platform lifecycle phases

Full Support

This phase begins at the general availability of a minor version and ends 6 months later or 90 days after the next minor version reaches general availability, whichever is later.

The following events occur during this phase:

- ▶ Full and development support are provided according to the scope of coverage and the service-level agreement (SLA).
- ▶ Critical and important security fixes (RHSAs) are released as needed.
- ▶ Urgent and high-priority bug fixes (RHBAs) are released promptly; other fixes might be included in periodic updates.
- ▶ You must stay on a supported microversion (for example, 4.x.z) to receive updates.

Maintenance Support

This phase starts after Full Support and ends 18 months after the minor version reaches general availability.

The following events occur during this phase:

- ▶ Critical and important security advisories (RHSAs) continue.
- ▶ Select urgent bug fixes (RHBAs) might be released.
- ▶ Other fixes and enhancements (Red Hat Enhancement Advisories (RHEAs)) can be issued at Red Hat's discretion.
- ▶ No technical support is provided after this phase, except assistance with upgrading.
- ▶ Access to hosted services for unsupported versions is not guaranteed.

Extended Update Support

Even-numbered minor versions (for example, 4.8 and 4.10) are designated as EUS releases. These releases offer extended support phases to simplify upgrades and reduce node restarts.

- ▶ EUS Add-On – Term 1 (6 Months)
 - Optional support after maintenance support
 - Includes critical and important security updates and urgent bug fixes.
 - You may stay on a minor release for up to 24 months.
 - Included with Premium subscriptions (x86_64) and available as an add-on to Standard subscriptions. Contact your Red Hat sales representative for access or guidance.
- ▶ EUS Add-On – Term 2 (12 Months)
 - Optional support after Term 1.
 - Includes critical and important updates and urgent fixes for platform-aligned Operators and selected Red Hat OpenShift Container Platform components.

With both Term 1 and Term 2, support extends up to 36 months for stable, mission-critical environments.

Table 6-3 shows the dates for versions of Red Hat OpenShift Container Platform at the time of writing.

Table 6-3 Red Hat OpenShift Container Platform lifecycle dates

Version	General availability	Full Support ends	Maintenance Support ends	Extended Update Support Add-On - Term 1 ends	Extended Update Support Add-On - Term 2 ends	Extended life phase ends
Full Support Red Hat OpenShift Container Platform 4						
4.19	17-Jun-25	GA of 4.20 + 3 Months	17-Dec-26	N/A	N/A	N/A
4.18	25-Feb-25	17-Sep-25	25-Aug-26	25-Feb-27	25-Feb-28	N/A
Maintenance Support Red Hat OpenShift Container Platform 4						
4.17	1-Oct-24	25-May-25	1-Apr-26	N/A	N/A	N/A
4.16	27-Jun-24	1-Jan-25	27-Dec-25	27-Jun-26	27-Jun-27	N/A
4.15	27-Feb-24	27-Sep-24	27-Aug-25	N/A	N/A	N/A
Extended Update Support Red Hat OpenShift Container Platform 4						
4.14	31-Oct-23	27-May-24	1-May-25	31-Oct-25	31-Oct-26	N/A
4.12	17-Jan-23	17-Aug-23	17-Jul-24	17-Jan-25	17-Jan-26	N/A

Note: Red Hat OpenShift Container Platform 4 uses Red Hat Enterprise Linux CoreOS (RHCOS) as its managed node OS. RHCOS is updated during cluster upgrades, which can include changes between Red Hat Enterprise Linux minor versions.

For more information about the Red Hat OpenShift Container Platform roadmap, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

6.6 Introducing Kernel-based Virtual Machine support

Kernel-based Virtual Machine (KVM) is a widely used virtualization technology on x86_64 systems that can also run on Power servers. IBM remains committed to PowerVM as the premier enterprise virtualization software in the industry.

With KVM on Power, IBM targets x86 customers on entry servers but offers both KVM and PowerVM to meet the varying virtualization needs of PowerLinux customers. KVM technology provides an opportunity to simplify virtualization infrastructure with a single hypervisor and management software across multiple platforms.

KVM guests can run within a PowerVM LPAR. These guests are VMs that run on the LPAR and use its existing resources.

One benefit of this approach is that it combines the advantages of PowerVM virtualization capabilities with the flexibility of KVM. KVM guests running in a PowerVM LPAR have a unique runtime architecture that differs from other virtualization mechanisms on Power servers.

For more information about KVM support on the Power11 platform, see 9.2, “KVM support” on page 222.

6.7 PowerVM Virtual I/O Server

IBM PowerVM software is a virtualization environment that runs AIX, IBM i, and Linux VMs on Power servers. Businesses use server virtualization to consolidate multiple workloads onto fewer systems, increase server usage, and reduce costs. PowerVM provides a secure and scalable virtualization environment for your applications, which is built on the advanced RAS features and leading performance of the IBM Power platform.

PowerVM is designed to protect and isolate critical workloads through a highly secure, enterprise-grade hypervisor. It enforces strong workload isolation and I/O integrity to help ensure the reliability of mission-critical applications. With robust automation capabilities, PowerVM accelerates service delivery by streamlining the provisioning and management of VMs and storage resources, making it suitable for cloud-based infrastructures. It also enhances operational efficiency and maximizes return on investment through features such as Live Partition Mobility (LPM), which enables zero-downtime workload migration, and resource optimization tools that improve the usage of compute and storage resources.

The VIOS is part of the PowerVM Editions hardware feature. VIOS runs in an LPAR and facilitates the sharing of physical I/O resources between client LPARs within the server.

At the time of writing, Power11 processor-based servers support the following minimum levels of the PowerVM VIOS:

- ▶ VIOS 4.1.1.10
- ▶ VIOS 4.1.0.40
- ▶ VIOS 3.1.4.60

6.8 Setting your LPAR compatibility mode

IBM Power Processor Compatibility Mode enables newer Power servers to run applications and OSs that were originally compiled for earlier generations of POWER processors. This feature is important when migrating to the latest hardware, such as the Power11 platform, where some OSs and applications might not yet support native run times. In these cases, compatibility mode enables workloads to run in a Power10 (or earlier) environment, helping ensure continuity and minimizing disruption during upgrades.

Compatibility mode is defined at the LPAR level during LPAR creation. Administrators can specify which processor generation the partition should emulate, which enables multiple workloads to operate in different compatibility modes on the same physical server. This flexibility is useful in mixed environments where not all software is updated to support the latest architecture.

Running in a previous-generation compatibility mode helps maintain application portability and facilitates migration. However, it might restrict access to newer processor features and performance enhancements. For this reason, use compatibility mode as a temporary solution. Your long-term goal should be to recompile applications and upgrade OSs to fully use the advanced capabilities and performance improvements of the latest POWER processor architecture.

Figure 6-11 shows how to set compatibility mode in your LPAR definition.

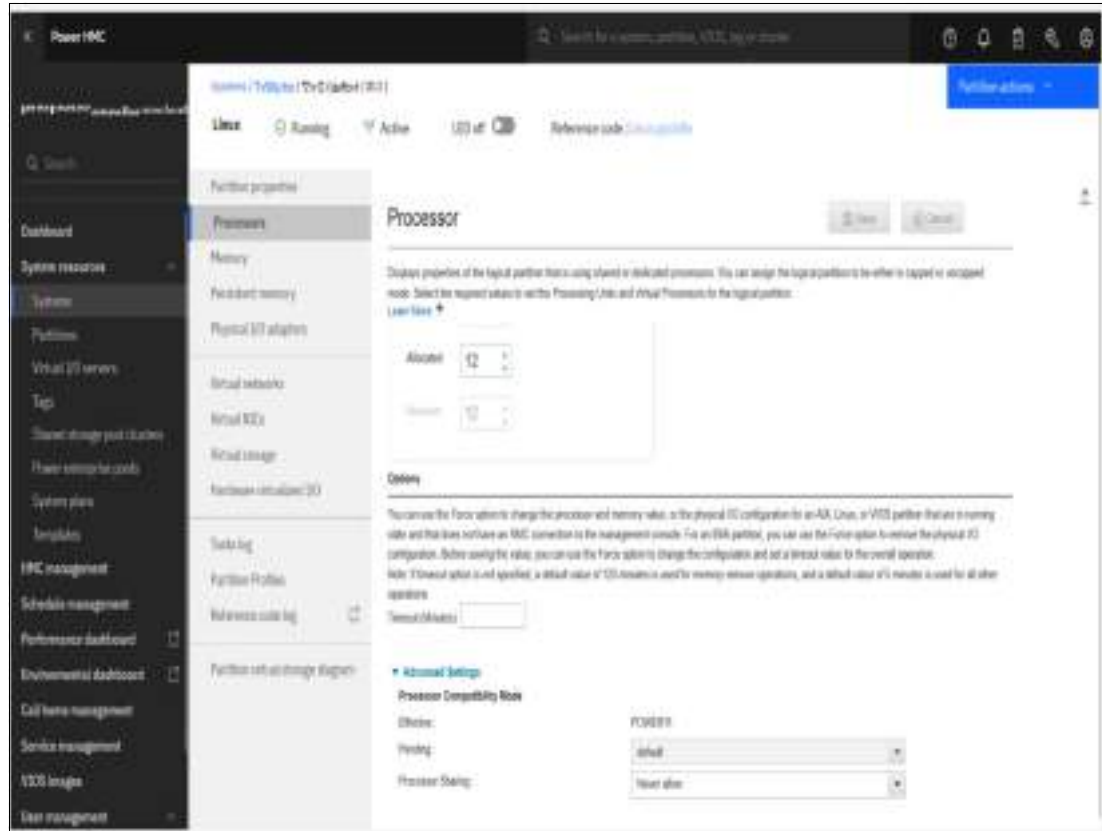


Figure 6-11 Setting LPAR processor compatibility mode



Enterprise solutions

IBM Power servers support a wide range of enterprise workloads that require high availability (HA), performance, and scalability. These systems are ideal for mission-critical applications such as SAP HANA, Oracle, and IBM Db2, where consistent uptime and fast transaction processing are essential. With advanced reliability, availability, and serviceability (RAS) features and capabilities such as Live Partition Mobility (LPM), Power servers minimize downtime and enable seamless workload management. They are also suitable for enterprise resource planning (ERP) and customer relationship management systems, providing the processing power and memory capacity that are needed to handle large-scale, real-time data operations.

In addition to traditional workloads, Power servers are optimized for modern enterprise requirements, including artificial intelligence (AI), machine learning, and hybrid cloud deployments. Built-in accelerators and support for containerized environments such as Red Hat OpenShift enable organizations to run AI models and cloud-native applications efficiently. These servers also perform well in high-performance computing (HPC) scenarios such as financial modeling, scientific research, and healthcare analytics because of their high core counts and memory bandwidth. With robust security features such as Transparent Memory Encryption (TME) and Secure Boot, IBM Power provides a trusted platform for industries with strict compliance requirements. Overall, Power servers deliver a flexible, secure, and cost-effective foundation for evolving enterprise IT landscapes.

This chapter contains the following topics:

- ▶ High availability and disaster recovery solutions
- ▶ IBM Db2
- ▶ Oracle
- ▶ SAP HANA
- ▶ Banking
- ▶ Healthcare

7.1 High availability and disaster recovery solutions

IBM offers a suite of high availability and disaster recovery (HADR) solutions for its Power servers. Each solution addresses different operational needs and workloads. PowerHA SystemMirror for AIX is a clustering solution that helps ensure application uptime through an active-standby model. It enables fast failover and minimizes disruption during outages or maintenance. For IBM i environments, PowerHA SystemMirror for i provides integrated HADR capabilities by using features such as independent auxiliary storage pools (IASPs) and real-time data replication to protect mission-critical workloads. VM Recovery Manager (VMRM) complements these solutions as a flexible option that supports AIX, IBM i, and Linux workloads. It manages full logical partition (LPAR) recovery across systems. VMRM is ideal for environments where full system recovery is acceptable. It offers automated failover and disaster recovery (DR) without requiring application-level configuration. Together, these tools deliver a comprehensive HADR strategy across the IBM Power ecosystem.

7.1.1 PowerHA SystemMirror for AIX

PowerHA SystemMirror for AIX, commonly known as PowerHA, is a HA clustering solution that helps ensure application uptime and fast recovery. It uses an active-standby or active-passive model for application availability. This model is also called an application restart model, where an application and its resources, such as a file system with data, an IP address, application scripts, and a volume group are logically combined and operate as a single entity on a virtual machine (VM) or LPAR, referred to as a node in a cluster. The Resource Group moves as a unit to a standby VM or LPAR, and the application restarts.

During a planned outage, PowerHA can gracefully quiesce the application on the primary node. The outage lasts only for the period that it takes to acquire the resources and restart the application, which is typically measured in minutes. PowerHA uses a warm standby method of HA because the target node already runs AIX.

PowerHA for AIX can also operate across data centers, where data is replicated through the storage subsystem or over IP by using the AIX Geographic Logical Volume Manager (GLVM) subsystem. In both cases, within or across data centers, the failover process can be automated to reduce application outages. Applications can move to another running node for maintenance, enabling near-zero downtime. PowerHA itself can be maintained with a live kernel update.

PowerHA for AIX is tightly integrated with the AIX operating system's (OS) inherent reliability and IBM Power hardware features such as virtual I/O, concurrent hardware maintenance, and dynamic resizing of LPARs. PowerHA supports critical enterprise applications such as Oracle, SAP, Db2, and Epic systems, which are common workloads. Licensing uses an $n+1$ model per cluster, where n is the number of cores in the production online copy of the application, with one extra core for the standby. For example, on a 24-core Power server, if two clusters are created (one supporting an LPAR with four cores and another with five), a total of 11 cores must be licensed: $4+1=5$ and $5+1=6$.

7.1.2 IBM PowerHA SystemMirror for i

PowerHA SystemMirror for i is a HADR solution that is designed for the IBM i OS. It provides automated failover and data replication to help ensure business continuity during planned or unplanned outages. PowerHA integrates with IBM i native features, such as IASPs to enable fast and efficient switching between systems with minimal downtime.

The core capability of PowerHA is real-time data replication between primary and backup systems by using technologies such as Geographic Mirroring or IBM storage-based replication. This approach helps ensure that critical applications and data remain available even if the primary system becomes unavailable. PowerHA supports both local HA within a single data center and remote DR across geographically dispersed locations, offering flexibility for various business continuity strategies.

PowerHA also includes tools for monitoring, automation, and management, which simplify configuration and maintenance of HADR environments. With features such as role-based access, cluster Resource Groups, and automated failover policies, PowerHA reduces the complexity of managing HA systems. It is a key component for organizations that rely on IBM i to run mission-critical workloads and require continuous access to data and applications.

7.1.3 VM Recovery Manager

VMRM operates on an active-inactive model. It re-creates the entire definition of a VM or LPAR on a target Power server and then restarts it. VMRM supports any workload running on a Power server, including AIX, IBM i, and Linux. Recovery takes longer than an application restart model because the entire LPAR must be restarted after being redefined, rather than just restarting the application.

In HA mode, VMRM uses LPM to dynamically move an LPAR to a system that shares a Hardware Management Console (HMC). In DR mode, it manages the data replication process from the storage provider to a set of disks at the target system. The LPAR exists in only one location at a time, except in limited circumstances where a DR exercise can be tested on the target system.

Like PowerHA, VMRM automates DR. It can detect when a VM or an entire Power server is down and initiate a policy-based failover to a single target or a group of targets. VMRM can coexist with PowerHA, enabling a cluster of LPARs to fail over as a group to Power servers at an alternative site. It is often used to provide simple DR support for Linux VMs without requiring software installation or OS configuration (the Linux VM restarts automatically).

VMRM is ideal for development, test, and sandbox environments that do not require the same RTO as production systems, and for Linux on Power. It integrates with solutions such as SAP HANA and Oracle. No configuration work is required on the individual OSs being managed, although optional agents can be installed to check for application failures and trigger an LPAR failover.

Similar to PowerHA, VMRM is licensed based on the number of production cores in use, not on the entire Power server. However, unlike PowerHA, no standby licenses are required. For example, on a 24-core Power server, if the VMs requiring VMRM support use nine cores, only nine cores must be licensed.

7.2 IBM Db2

In today's data-driven world, organizations across industries, including manufacturing, healthcare, and the public sector must extract maximum value from rapidly growing volumes of information. Whether the goal is to uncover actionable insights from customer data or accelerate the processing of high-volume online transactions, businesses need solutions that are optimized for performance, scalability, and cost efficiency. These solutions must support mission-critical workloads with HA while controlling operational costs.

By combining IBM Db2 with Power servers built on the advanced Power11 architecture, organizations can meet these demands confidently. This integrated, workload-optimized stack delivers exceptional performance for both transactional and analytical workloads. Db2 uses advanced database innovations to maximize efficiency and throughput, earning top rankings in industry benchmarks such as Transaction Processing Performance Council Benchmark C (TPC-C), Transaction Processing Performance Council Benchmark H (TPC-H), and SAP SD 3-Tier. Power servers provide a robust hardware foundation, offering superior price and performance and the ability to handle more than one million transactions per minute at a cost of less than \$1 per transaction¹. Db2 automatically uses Power11 parallelism and large page sizes, simplifying deployment and enabling cost-effective scaling for web applications, messaging backbones, and workload consolidation.

Beyond performance, the combination of Db2 and Power servers helps ensure HA and operational efficiency. Power servers are designed with built-in redundancy, error handling, and reliability features, and IBM PowerHA SystemMirror adds automated monitoring and recovery to minimize downtime. Cost control is further enhanced through Db2 Deep Compression, which reduces storage requirements, and IBM PowerVM virtualization, which enables efficient resource usage and workload consolidation. These capabilities reduce hardware, energy, and management costs, creating a more agile and cost-effective IT infrastructure. Together, Db2 and Power servers deliver a powerful, scalable, and resilient platform for modern data-driven enterprises.

For more information about IBM Db2 HA, see *IBM PowerHA SystemMirror for AIX Cookbook*, SG24-7739

7.3 Oracle

For more than 35 years, clients have relied on IBM Power to deploy Oracle database and application workloads. Organizations of all sizes benefit from class-leading reliability and security of Power, and its advanced recovery, self-healing, and diagnostic capabilities that are designed to reduce application downtime.

IBM Power11 processor-based servers, fully certified by Oracle software, enable organizations to consolidate multiple workloads on fewer servers, increasing overall system usage and reducing costs. This higher usage enables clients to maximize Oracle license use while reducing IT costs by managing fewer physical machines.

The cost of security breaches continues to grow, averaging USD 4 million per breach.² IBM Power technology helps protect businesses against cyberthreats with end-to-end security protocols, including new Trusted Memory Encryption (TME) that does not impact performance. Power servers are considered more secure than other servers based on reported security incidents. Combined with IBM AIX Trusted Execution, Power servers enable you to verify system integrity and implement advanced security policies to enhance trust across the entire environment.

As data volumes expand rapidly, organizations often struggle to modernize IT infrastructure. Power servers and the IBM AIX OS provide a solid foundation for modernizing Oracle database workloads, developing new applications, and consolidating workloads. With the Red Hat Ansible Automation Platform, clients can manage Oracle workloads on Power servers as part of a broader enterprise automation strategy.

¹ Source: https://www.tpc.org/tpcc/results/tpcc_results5.asp?print=false&orderby=priceperf&sortby=asc

² Source: <https://www.ibm.com/think/insights/cost-of-a-data-breach-2024-financial-industry>

Mission-critical workloads require a server and OS with reliability, HA, and scalability without compromising performance. Power servers deliver up to 2.5 times better core performance compared to x86 servers³, enabling Oracle workloads to scale linearly without bottlenecks.

Oracle certifies its products on Power servers, delivering benefits such as comprehensive end-to-end support, portability, and efficiency. Power servers provide 99.999% reliability to maintain maximum availability. The combined design of AIX on Power servers with IBM PowerHA technology delivers exceptional uptime and enables proactive management and monitoring to prevent planned and unplanned outages.

Benefits of running Oracle on the Power11 platform

Here are some benefits of running your Oracle database workloads on the Power11 platform:

- ▶ **Industry-leading security**
IBM Power and AIX keep critical Oracle workloads protected and available while reducing costs.
- ▶ **Simplified management**
Automatically deploy a cloud-ready OS that meets private cloud requirements with IBM Power Virtualization Center (PowerVC).
- ▶ **Unmatched uptime**
IBM Power supports demanding workloads and provides 99.999% reliability to maintain maximum availability.
- ▶ **Improved workload cost-effectiveness**
Use Power LPARs and dynamic logical partition (DLPAR) eligibility for Oracle hard partitioning. Clients can license only the Power cores that are available to Oracle software.
- ▶ **Standards-based automation**
Improve manageability and scalability by helping ensure consistent and repeatable outcomes through an enhanced automation portfolio for Oracle workloads that is built on Ansible.
- ▶ **Data protection**
Benefit from end-to-end security in Power with trusted boot, main memory encryption, runtime verification of OS files, and role-based access and runtime control with PowerSC.

For more information about running Oracle on Power, see *Oracle on IBM Power Systems*, SG24-8485.

7.3.1 Running Oracle Standard Edition 2 on IBM Power

Oracle Standard Edition 2 (SE2) is a database option that is suitable for many business needs and can run on Power servers. SE2 offers a full-featured database with performance and flexibility. It includes features such as relational, JSON, and XML data handling, and Real Application Clusters for clustering services. SE2 is licensed on servers with up to two sockets, and licensing costs remain the same regardless of the number of cores within each socket.

³ Source:

<https://newsroom.ibm.com/2021-09-08-IBM-unveils-new-generation-of-IBM-Power-servers-for-frictionless-scalable-hybrid-cloud>

Here are some key benefits of running Oracle SE2 on IBM Power:

- ▶ **Licensing**
SE2 can be licensed by using either the processor metric or the Named User Plus metric. Licensing rules include a maximum of two occupied sockets and a maximum of 16 CPU threads per socket, effectively limiting the number of cores.
- ▶ **Power server compatibility**
Oracle Database SE2 is compatible with Power servers. IBM designed specific Power10 processor-based servers (Power S1012 and Power S1014) to compete with x86 systems in the two-socket space, addressing Oracle's need for SE2 to remain cost-competitive with SQL Server. IBM is expected to provide a Power11 processor-based server similar to the Power S1012 in the future to support Oracle SE2.

Benefits of IBM Power for Oracle Workloads

IBM Power offers several advantages for Oracle databases, including main memory encryption, runtime verification of OS files, role-based access control (RBAC), and advanced recovery, self-healing, and diagnostic capabilities.

- ▶ **HA**
Starting with Oracle Database 19c Release Update 19.7, Standard Edition High Availability is supported on IBM AIX running on Power servers. This feature provides cluster-based database failover for Oracle Database 19c Standard Edition.
- ▶ **Other features**
SE2 includes capabilities for building business applications, such as support for relational, JSON, XML, spatial, graph, and unstructured data, and Oracle Multitenant Architecture.

Licensing and implementation

Consider the following items when planning your Oracle SE2 implementation:

- ▶ Although the number of cores in each socket does not affect the license price, you must adhere to the maximum number of CPUs that are allowed for the cluster, which applies to the entire cluster, not per node.
- ▶ Standard Edition High Availability provides cluster-based failover for single-instance Standard Edition Oracle Databases by using Oracle Clusterware.
- ▶ Oracle Standard Edition High Availability benefits from cluster capabilities and storage solutions that are part of Oracle Grid Infrastructure, such as Oracle Clusterware, Oracle Automatic Storage Management (Oracle ASM), and Oracle ASM Cluster File System (Oracle ACFS).
- ▶ Using integrated, shared, and concurrently mounted storage, such as Oracle ASM and Oracle ACFS for database files and unstructured data enables Oracle Grid Infrastructure to restart an Oracle Database on a failover node faster than cluster solutions that rely on failing over and remounting volumes and file systems.
- ▶ Starting with Oracle Database 19c Release Update (19.13), Standard Edition High Availability is supported on IBM AIX on Power servers (64-bit).

7.4 SAP HANA

IBM Power offers many options to meet clients wherever they are on their journey to SAP S/4HANA. Whether you are considering SAP HANA or next-generation S/4HANA, Power servers provide on-premises, off-premises, or fully managed (RISE with SAP) solutions. These solutions are built to run mission-critical applications such as SAP, accelerate ERP and application deployments, and maximize the impact on data management, integration, automation, and business processes.

Here are some IBM and SAP partnership highlights:

- ▶ Over 4,800 clients run SAP HANA on Power servers
- ▶ More than 120 external client references for SAP HANA on IBM Power
- ▶ 39 SAP Pinnacle Awards won by IBM
- ▶ 30,000 organizations run essential workloads such as SAP on IBM Power

The key to IT efficiency and business continuity is a platform that integrates with your infrastructure while supporting digital transformation. Power servers are purpose-built for data-intensive applications such as SAP HANA and S/4HANA. These applications require large-scale in-memory computing while maintaining HA and flexibility for hybrid cloud environments.

Benefits of running SAP HANA on IBM Power

With global data volumes projected to exceed 180 ZB in 2025, organizations across all sectors face increasing pressure to manage, process, store, and extract insights from critical data. Running SAP HANA on Power servers enables businesses to achieve the following goals:

- ▶ Provision faster.
Simplify system management and boost agility. Create environments flexibly by allocating resources incrementally, starting as low as 0.01 cores and 1 GB of memory.
- ▶ Maximize uptime.
Minimize disruption to business operations. IBM Power delivers best-in-class reliability for more than 15 years.⁴
- ▶ Cut energy usage.
Reduce data center costs and improve sustainability. Power11 scale-out servers provide comparable performance while using 50% less energy than x86-based servers. At maximum input power, Power11 scale-out servers consume 15% less energy and deliver 54% more performance than comparable x86 servers.⁵
- ▶ Scale in an affordable way.
Avoid over-provisioning with flexible hybrid cloud solutions, instant scaling, and pay-per-use options. IBM Power supports up to 40 TB of scale-up capacity, which is the largest supported for SAP S/4HANA and SAP BW.

⁴ Source:
<https://www.itjungle.com/2024/10/21/ibm-nears-the-end-of-the-road-for-server-reliability-improvements/>

⁵ Source:
<https://www.techspot.com/news/108599-ibm-introduces-power11-servers-boosts-uptime-security-energy.html>

- ▶ Strengthen security.
Protect critical data and applications with end-to-end security, including TME with no performance impact. Power servers are considered more secure than other servers based on the number of security incidents that are reported.
- ▶ Gain faster insights.
Accelerate decision-making to maximize efficiency. IBM Power delivers up to 2.5 times better per-core performance compared to x86 servers.⁶

7.5 Banking

As financial institutions accelerate digital transformation, IT architects must design platforms that are secure, scalable, and AI-ready. Power servers, particularly the Power10 processor-based and Power11 processor-based servers, provide a robust foundation for deploying AI workloads in banking environments. With built-in AI acceleration, enterprise-grade RAS features, and seamless integration with hybrid cloud and container platforms, Power servers enable architects to modernize infrastructures while embedding intelligence into core banking work flows.

Hybrid cloud deployment models for AI in banking

Hybrid cloud is the preferred architecture for modern banking platforms because it enables agility, regulatory compliance, and cost optimization. Power servers support hybrid cloud deployments with the following products:

- ▶ IBM Power Virtual Server (PowerVS) provides public cloud scalability.
- ▶ Red Hat OpenShift delivers container orchestration across environments.
- ▶ IBM Cloud Pak for Data and watsonx duports AI lifecycle management.

Banking use cases that use AI on IBM Power

Here is a list of some of the use cases for AI for banks and other financial institutions:

- ▶ On-premises AI training with cloud-based inference
Banks often train AI models on-premises to comply with strict data residency and regulatory requirements. Using IBM Power10 processor-based or Power11 processor-based servers, financial institutions can securely train models on sensitive customer or transactional data. Once trained, these models are deployed to PowerVS for inference, enabling scalable, real-time decision-making in the cloud. This hybrid approach is ideal for applications such as real-time credit scoring or fraud detection, where data privacy is critical and rapid response times are essential.
- ▶ End-to-end AI on PowerVS with secure on-premises data access
In this model, banks run the full AI lifecycle (training, tuning, and inference) on PowerVS while maintaining secure access to on-premises data sources. This setup enables institutions to use cloud scalability without moving sensitive data offsite. For example, a bank might deploy a generative AI model on PowerVS to automate the summarization of complex financial reports, pulling data securely from internal systems. This approach helps ensure compliance with data governance policies while accelerating insights and reporting.

⁶ Source:

<https://newsroom.ibm.com/2021-09-08-IBM-unveils-new-generation-of-IBM-Power-servers-for-frictionless-scalable-hybrid-cloud>

- ▶ **Distributed AI microservices across hybrid cloud**
Some banks adopt a microservices architecture, distributing AI components across on-premises Power servers and PowerVS in the cloud. This model supports modular, scalable AI applications that operate across environments. For example, asset valuation engines or personalized client engagement tools can run inference in the cloud while accessing real-time data from on-premises systems. This approach enables agility and performance, especially for institutions managing diverse workloads worldwide or across business units.
- ▶ **AI for risk and compliance on IBM Power**
Global Tier-1 banks increasingly use AI to automate risk management and compliance processes. By deploying generative AI and machine learning models on Power servers, these institutions can continuously monitor transactions, flag anomalies, and generate regulatory reports with minimal human intervention. This model supports high-throughput, low-latency processing, making it ideal for meeting stringent regulatory requirements such as anti-money laundering (AML), know-your-customer (KYC), and Basel III compliance.
- ▶ **AI-enhanced core banking applications**
Power servers also support the integration of AI directly into core banking platforms. This integration enables real-time intelligence within mission-critical applications such as loan underwriting, treasury operations, and fraud analytics. For example, a bank might use AI models running on the Power11 platform to detect unusual transaction patterns or predict liquidity needs, enhancing operational efficiency and decision-making. This model helps ensure that AI capabilities are embedded where they deliver the most immediate business value.

Integrating IBM watsonx with Power servers

IBM watsonx is a modular AI and data platform that is designed to accelerate the deployment of enterprise-grade AI. When integrated with IBM Power10 and Power11 processor-based servers, watsonx enables financial institutions to build, deploy, and govern AI models across hybrid cloud environments, including PowerVS.

Here are the components of watsonx:

- ▶ **IBM watsonx.ai®**
Model training and inference
- ▶ **watsonx.data**
Data lakehouse integration
- ▶ **watsonx.governance**
Model risk and compliance
- ▶ **watsonx.orchestrate**
AI agents and workflow automation

Power servers are high-performance systems that are optimized for data-intensive workloads. The integration between watsonx and IBM Power delivers the following benefits:

- ▶ **Accelerated AI workloads**
Power11 processors are optimized for AI inference and training, making them ideal for running watsonx models efficiently.
- ▶ **Hybrid cloud flexibility**
watsonx can run on IBM Power in hybrid cloud environments, enabling businesses to keep sensitive data on-premises while using cloud-native AI capabilities.

- ▶ Enterprise-grade security and reliability
Power servers provide robust security and uptime, aligning with the watsonx focus on trustworthy AI.
- ▶ Tight integration with Red Hat OpenShift
Both watsonx and IBM Power support Red Hat OpenShift, enabling containerized AI workloads to run seamlessly across environments.

7.6 Healthcare

Power servers provide a secure, high-performance, and resilient infrastructure that enables healthcare organizations to manage and protect digital health data, drive data-informed decisions, and meet regulatory requirements with confidence. Designed for mission-critical workloads, Power servers help healthcare providers improve patient outcomes, enhance operational efficiency, and reduce risk.

- ▶ Security and compliance
Healthcare organizations handle sensitive patient data that must comply with regulations such as HIPAA. Power servers deliver advanced security features, including Secure Boot, RBAC, and encryption at rest and in transit. With quantum-safe firmware signing and encrypted LPM, Power servers help ensure data integrity and privacy. According to Information Technology Intelligence Consulting (ITIC), organizations that use Power servers experience an average of only 3.3 minutes of unplanned downtime per year due to security issues, which demonstrates industry-leading resilience.⁷
- ▶ Unmatched reliability for clinical systems
Power servers are designed for continuous availability of mission-critical applications such as electronic health records (EHRs), PACS, and healthcare information systems (HISs). Features such as Predictive Failure Analysis (PFA), dynamic resource allocation, and system redundancy contribute to its 99.9999% availability rating, as reported by 1,900 C-level executives in ITIC's global reliability survey.⁸ Power servers have been the most reliable non-mainframe server platform for more than 15 years.
- ▶ High-performance computing for healthcare analytics
Healthcare workloads from medical imaging to genomics and real-time analytics demand exceptional compute power. Power11 processor-based servers deliver up to 2.5 times more performance per core than previous generations⁹, enabling faster insights for research, drug discovery, and personalized medicine.
- ▶ Scalable and flexible infrastructure
Healthcare environments often face unpredictable demand. Power servers support dynamic scaling to accommodate workload fluctuations, whether driven by seasonal surges or new clinical initiatives. Its support for multiple OSs and virtualization technologies enable healthcare IT teams to run diverse workloads efficiently on a single platform.

⁷ Source: <https://itic-corp.com/itic-2022-global-server-reliability-results/f>

⁸ Source:

<https://www.itjungle.com/2024/10/21/ibm-nears-the-end-of-the-road-for-server-reliability-improvements/>

⁹ Source: <https://newsroom.ibm.com/2025-07-08-ibm-power11-raises-the-bar-for-enterprise-it>

- ▶ AI-ready for clinical innovation

IBM Power is optimized for AI and machine learning workloads, with on-chip Matrix Math Accelerators (MMAs) that enable real-time inferencing at the point of care.

IBM Power S1122 servers can process up to 42% more batch queries per second than comparable x86 systems under peak load, with subsecond inferencing latency¹⁰, which is ideal for applications such as diagnostic imaging, predictive analytics, and treatment optimization.

- ▶ Seamless integration with healthcare ecosystems

Power servers integrate with many healthcare applications, including electronic medical records (EMRs), imaging platforms, analytics tools, and telemedicine systems. Power servers support interoperability and data exchange across systems and work closely with leading independent software vendors (ISVs) such as Epic to certify hardware for industry-standard healthcare software.

7.6.1 Epic

Epic is a HIS provider that develops and delivers a comprehensive EMR platform that covers all aspects of healthcare. The Epic solution includes applications for medical billing, emergency care, radiology, outpatient and inpatient services, and ambulatory care.

IBM and Epic Systems maintain a strategic partnership that enables healthcare organizations to run Epic's EHR platform on IBM infrastructure, including Power servers and IBM Storage solutions. This collaboration supports the performance, security, and compliance requirements of healthcare providers managing mission-critical clinical workloads.

Power servers are certified to run Epic Operational Database (ODB) workloads, with best practices that are published for tuning AIX environments to meet Epic's stringent performance requirements. These configurations are optimized for high throughput, low latency, and HA, which are critical for real-time clinical operations. At the time of writing, IBM Power10 processor-based servers are certified by Epic Systems, supporting key infrastructure components on both AIX and Linux for large healthcare deployments, primarily in the database and back-end services tier. These platforms host InterSystems Caché or IRIS for Health, which serves as the core databases behind Epic's EHR system.

IBM and Epic are expected to certify Power11 configurations after Power11 processor-based servers become available.

Supported products

Here is a list of supported Epic products at the time of writing:

- ▶ InterSystems Caché or IRIS for Health

Officially supported by Epic on both AIX and Linux (RHEL) on Power10 processor-based servers. Deployed in production for the Epic ODB, also known as Chronicles.

- ▶ Epic Production and Reporting Database Tiers

Most often run on AIX or RHEL on Power10 processor-based servers for performance and reliability. Other back-end services, such as database copies that are used for Clarity extracts, may also run on AIX or Linux, depending on the customer's architecture.

¹⁰ Source:

<https://www.ibm.com/new/announcements/combining-ai-with-a-trusted-data-approach-on-ibm-power-to-fuel-business-outcomes>

Note: The Epic application layer (Hyperspace, Interconnect, and related components) typically runs on Windows or Linux on x86. Power based AIX or Linux deployments focus on the infrastructure and database tiers.

Here are the required OS and firmware levels for Power10 processor-based servers:

▶ AIX

- AIX 7.2 Technology Level (TL) 5 Service Pack (SP) 4 or later.
- AIX 7.3 TL1 or later.

Certified by both Epic and InterSystems for Power10 compatibility

▶ Linux

- Red Hat Enterprise Linux (RHEL) 8.6 or later on Power10 (ppc64le).
- Certified by InterSystems for IRIS for Health on Power.
- Required for newer installations that use containerized services or performance-optimized Linux environments.

▶ Virtualization

- Both AIX and Linux deployments use IBM PowerVM for LPAR management.
- Linux LPARs can also be deployed by using Red Hat OpenShift on Power for modern container-based workloads.

▶ Storage

A high-performance storage area network (SAN), such as IBM FlashSystem, is commonly used with tuning for Caché or IRIS.

▶ HADR

- PowerHA for AIX.
- Linux HA tools, including IBM VMRM).
- InterSystems mirroring for resilience.



Servicing IBM Power11 scale-out servers

The goal of *serviceability* is to enable efficient system repair while minimizing disruption to operations. It includes system installation, upgrades or rollbacks (Miscellaneous Equipment Specification (MES)), and ongoing maintenance or repair activities. Depending on the system configuration and warranty agreement, service tasks can be performed by the client, an IBM technician, or an authorized service provider.

IBM Power servers include advanced serviceability features that support an efficient maintenance environment. These features streamline service operations and reduce downtime by incorporating the following key attributes:

- ▶ Simplified installation and upgrade processes
- ▶ Support for concurrent maintenance and guided repair
- ▶ Automated diagnostics and error reporting
- ▶ End-to-end service workflows, from issue detection to resolution

The following topics are covered in this chapter:

- ▶ IBM maintenance
- ▶ IBM Expert Care
- ▶ IBM tools and interfaces
- ▶ Baseboard Management Controller card
- ▶ ASMI
- ▶ Entitled System Support
- ▶ System firmware

8.1 IBM maintenance

IBM offers a comprehensive global maintenance and support framework for Power servers. This framework helps ensure high availability (HA), rapid issue resolution, and minimal disruption to business operations. Maintenance services are available under various service-level agreements (SLAs) that meet the needs of different environments, from standard business hours to mission-critical 24x7 operations. IBM Support includes remote diagnostics, onsite service, and proactive monitoring, backed by a global network of skilled service professionals and parts depots.

You can choose between 8x5 and 24x7 service coverage. The 8x5 option provides support during standard business hours (typically 8 AM - 5 PM., Monday through Friday, excluding holidays). This option is ideal for non-critical systems or environments with internal IT support. For systems that require continuous uptime, the 24x7 option helps ensure around-the-clock support, including weekends and holidays. This level of coverage is essential for industries such as finance, healthcare, and manufacturing, where downtime can cause operational or financial consequences.

IBM defines response time targets based on the selected service level. For example, under a 24x7 agreement, IBM may commit to a 2-hour or 4-hour onsite response for critical hardware issues, depending on the location and contract terms. These response times are supported by the IBM global logistics and service infrastructure, which includes strategically located parts centers and field engineers in more than 170 countries.

This worldwide coverage helps ensure that Power server customers receive consistent, high-quality support regardless of location. IBM maintenance services are enhanced by features such as Call Home diagnostics, automated error reporting, and remote problem determination. These features accelerate issue resolution and reduce manual intervention. Combined, these capabilities provide a reliable maintenance ecosystem that supports both traditional on-premises deployments and modern hybrid cloud environments.

8.2 IBM Expert Care

Managing Power servers should be seamless and efficient. IBM Power Expert Care provides immediate access to a curated bundle of high-value services that are trusted by the IBM Power community. This approach eliminates delays and reduces the complexity of traditional procurement processes. With this preselected service package, you benefit from streamlined support that is enhanced by artificial intelligence (AI)-driven tools that accelerate response times, improve case resolution, and increase overall satisfaction.

Power Expert Care Premium focuses on mission-critical environments where downtime is not an option. Backed by the experts who designed and built the Power11 platform, this service tier helps ensure that your infrastructure is supported by technical depth and operational excellence.

Here are the key features of the Power Expert Care Premium:

- ▶ A 30-minute response time for hardware and software support cases
- ▶ Four-hour, onsite response target for urgent issues
- ▶ System and microcode compatibility guidance to help ensure optimal performance
- ▶ Dedicated Technical Account Manager (TAM) and mission-critical support resources
- ▶ Automated case creation and log analysis for faster issue resolution

- ▶ Health checks for hardware, operating system (OS), and applications (available as add-ons)
- ▶ Zero planned downtime (ZPD) support with proactive planning and TAM assistance

For non-mission-critical systems, IBM offers Power Advanced Expert Care, which provides 24x7 standard response times for repair and maintenance cases.

Technical Account Manager

The TAM plays a vital role in delivering proactive, product-based support for Power servers. Acting as the primary point of contact, the TAM provides strategic guidance and direct engagement for hardware and software within the scope of your support agreement. TAM services are delivered in English during your business hours. Support in other languages is available on request and mutual agreement, and are subject to availability.

TAMs streamline operations and reduce downtime through the following responsibilities:

- ▶ Enabling Call Home for proactive error reporting, Autonomous Error Resolution (AER), and ZPD on the Power11 platform
- ▶ Activating Support Insights for predictive analytics and delivering monthly reports
- ▶ Providing firmware and microcode compatibility analysis (exclusive to Power11 Premium clients)
- ▶ Sharing software lifecycle and roadmap updates
- ▶ Delivering HIPER alerts to help avoid high-impact issues
- ▶ Offering best practices documentation and priority handling of Severity 1 and 2 cases
- ▶ Leading complex case resolution and managing case progression
- ▶ Coordinating welcome calls, support planning, monthly reporting, and quarterly reviews
- ▶ Engaging IBM resources for Remote Code Load on Power10 and Power11 processor-based servers (onsite code load available separately for the Power11 platform)
- ▶ Supporting change management by communicating planned events to relevant teams

For clients running SAP HANA on Power10 or the Power11 platform, the TAM also provides the following services:

- ▶ Advises on SAP HANA best practices.
- ▶ Coordinates troubleshooting across the full software and hardware stack.
- ▶ Provides technical recommendations for error identification, environment optimization, and known defect mitigation.

8.3 IBM tools and interfaces

Servicing Power servers relies on a combination of proactive maintenance, intelligent diagnostics, and modern remote management tools. These capabilities help ensure HA and system reliability. At the center of this ecosystem is the Hardware Management Console (HMC), which acts as the central hub for monitoring system health, managing logical partitions (LPARs), and coordinating service activities. Complementing the HMC is the service processor, a dedicated hardware component that operates independently of the main system. It continuously monitors hardware status, logs errors, and enables early detection of potential issues. Features such as Call Home support and automated error reporting streamline the support process by alerting IBM Support teams to critical events, often before they affect operations.

You can access and manage the service processor through the HMC or directly through the Advanced System Management Interface (ASMI), which provides low-level control and more troubleshooting capabilities. For automated environments, the service processor also supports DMTF Redfish application programming interfaces (APIs), enabling secure, standardized remote management and integration with broader data center orchestration tools. Together, these technologies form a resilient and intelligent service infrastructure that supports both on-premises and hybrid cloud deployments, helping ensure that Power servers remain secure, efficient, and continuously operational.

8.4 Baseboard Management Controller card

Power11 scale-out systems use an enterprise Baseboard Management Controller (eBMC) for system service management, monitoring, maintenance, and control. The eBMC also provides access to system event log files. The eBMC controller card is installed in a dedicated slot (C5).

The eBMC is a specialized service processor that monitors the physical state of the system by using sensors. A system administrator or service representative can communicate with the Baseboard Management Controller (BMC) through an independent connection.

8.5 ASMI

The ASMI is the GUI for the eBMC. It provides functions similar to ASMI on Flexible Service Processor (FSP)-managed servers but features a redesigned interface based on customer feedback from a Design Thinking workshop.

To access the ASMI GUI, use the HMC to select the server, and then select **Operations** → **Launch Advanced System Management**. A window opens and shows the system name, model, type, serial number, and the IP address of the service processor (eBMC). Click **OK** to open the ASMI window.

If the eBMC is connected to a network that is accessible from your workstation, you can connect directly by entering `https://<eBMC IP>` in your web browser.

Figure 8-1 on page 185 shows the ASMI login window.



Figure 8-1 ASMI login window

After you log in, the Overview window opens. It includes server, firmware, network, power, and status information. When you log in for the first time, the default username and password is admin, but this credential is invalidated after the first login. Immediately change the admin password. This change is also required after a factory reset of the system. This policy helps ensure that the eBMC is not left with a well-known password, which improves system security. The password must meet specific criteria (for example, abcd1234 is invalid).

The new ASMI for eBMC-managed servers introduces important differences from the ASMI version that is used by FSP-based systems. It also delivers valuable new features:

- ▶ Update system firmware.

You can install a firmware update for the server by using the ASMI GUI, even if the system is managed by an HMC. In this case, the firmware update is always disruptive. To install a concurrent firmware update, you must use the HMC. This option is not available in the ASMI GUI.

- ▶ Download memory dumps.

Memory dumps can be downloaded by using the HMC or from the ASMI menu if necessary. You can also start a memory dump from the ASMI. Select **Logs** → **Dumps**, select the memory dump type, and click **Initiate memory dump**. The following memory dump types are available:

- BMC memory dump (nondisruptive)
- Resource memory dump
- System memory dump (disruptive)
- Network Time Protocol (NTP) server support

- Lightweight directory access protocol (LDAP) for user management
- Host console

Use the host console to monitor the server start process. It can also be used to access the OS when a single LPAR uses all resources.

▶ User management.

You can create users in the eBMC. This feature can also create an individual user for HMC access to the server. A user can have one of the following privileges:

- Administrator
- Read Only (You cannot modify anything except the user's password, so this privilege cannot be used for HMC access.)

▶ IBM security with Access Control Files (ACFs).

On FSP-managed servers, IBM Support generates a password by using the serial number and date for root access through the `celogin` user. On eBMC-managed systems, the support team generates an ACF. This file must be uploaded to the server to gain access. This procedure is required, for example, if the admin password must be reset. Physical access to the system is necessary for this process.

▶ Jumper reset.

A physical jumper can reset everything on the server. This factory reset process clears all configurations, including LPAR definitions, eBMC settings, and NVRAM.

8.5.1 More ASMI functions

This section describes more functions of the ASMI.

▶ Real-time progress indicator

The ASMI of an eBMC server provides a real-time progress indicator for viewing operator panel codes. To open the window that shows the codes, select **Logs** → **Progress logs**, and then click **View code in real time**.

▶ Inventory and LEDs

When you select **Hardware status** → **Inventory and LEDs**, you can view most hardware components with their current state. You can also control identification LEDs for each component, the system identification LED, and the system attention LED. You can turn on or off all identification LEDs individually. The system attention LED can only be turned off.

You can also display component details, such as the size of a differential DIMM (DDIMM) or the part number of a component that must be replaced.

▶ Sensors

The ASMI displays data from various sensors within the server and its components. To view this data, select **Hardware status** → **Sensors**. Loading sensor data takes time, during which a progress bar appears at the top of the window.

▶ Network settings

The default network settings for the two eBMC ports use DHCP. When you connect a port to a private HMC network with the HMC as a DHCP server, the new system receives its IP address from the HMC during firmware startup. The system then automatically appears in the HMC and can be configured.

Note: DHCP is the best practice method for connecting the eBMC of a server to the HMC.

If you do not use DHCP and want to configure a static IP address, you can set the IP address in the ASMI GUI. Before you make this change, you must connect to the ASMI. Because default IP addresses are not the same for every server, you must first determine the configured IP.

To determine the configured IP address, use the operator window. This optional component is one per rack of Power11 processor-based scale-out servers.

In the control window, complete the following steps:

1. Use the Increment or Decrement options to scroll to function 02.
2. Click Enter until the value changes from N (normal) to M (manual). This process activates access to function 30.
3. Scroll to function 30 and click Enter until 30 appears.
4. Scroll to 3000 and click Enter to display the IP of the ETH0 port. Scroll to 3001 and click Enter to display the IP address of ETH1.
5. After you determine the IP address, scroll again to function 02 and set the value back from M to N.

For more information about function 30 in the operator window, see [Function 30: Service processor IP address and port location](#).

After you determine the IP address, configure any computer with a web browser to an IP address in the same subnet (Class C) and connect the computer to the correct Ethernet port of the server. After connecting the cable, use a web browser to access the ASMI at `https://<IP address>` and configure the network port address settings.

To configure the network ports, select **Settings** → **Network** and select the correct adapter to configure. Before you configure a static IP address, turn off DHCP. Several static IP addresses can be configured on one physical Ethernet port.

In the ASMI network settings window, you cannot configure the VMI address. The VMI address is another IP address that is configured on the physical eBMC Ethernet port of the server to manage virtualization. The VMI address can only be configured in the HMC.

Using an Access Control File

If you lose the access password for the ASMI service user, you can access the ASMI by using an ACF. The ACF is a digital certificate that IBM Support provides when you open a support case. To use the ACF, the system must be enabled at the server by using the operator panel.

Complete the following steps:

1. On the operator panel, use Increment or Decrement to scroll to function 74.
2. Click Enter and then select 00 to accept the function (FF rejects it). The ACF function is now active for 30 minutes. To use it, complete the following steps:
 - a. Open the ASMI login window.
 - b. Click **Upload service login certificate** to upload the ACF into the system and allow the user to enter the ASMI with the associated password that IBM Support also supplies.

For more information, see [Function 74: Authentication Override for ACF upload](#).

Policies

If you select **Security and access** → **Policies**, you can turn on and off security-related functions:

- ▶ You can enable or disable management over the Intelligent Platform Management Interface (IPMI).
- ▶ You can turn off Host Universal Serial Bus (USB) enablement. Some customers require disabling the server USB ports. You can make this change in the Policies window.

8.6 Entitled System Support

IBM Enterprise Storage Server is available to view and manage IBM Power and IBM Storage software and hardware. In general, most IBM Systems products that customers purchase through IBM Digital Sales representatives or IBM Business Partners can be accessed on this site when the IBM Configurator is used.

The site features the following three main sections:

- ▶ My entitled software

This section lists activities that are related to IBM Power and IBM Storage software, including the ability to download licensed, no-charge, and trial software media, place software update orders, and manage software keys.

- ▶ My entitled hardware

This section lists activities that are related to IBM Power and IBM Storage hardware, including the ability to renew Update Access Keys (UAKs), buy and use Elastic Capacity on Demand (CoD), assign or buy credits for new and existing pools in a Power Private Cloud environment (Power Enterprise Pools 2.0 (PEP 2.0)), download Storage Capacity on Demand codes, and manage Hybrid Capacity credits.

- ▶ My inventory

This section lists activities that are related to IBM Power and IBM Storage inventory, including the ability to browse software licenses, software maintenance, and hardware inventory, manage inventory retrievals through Base Composer, or generate several types of reports.

8.7 System firmware

System firmware provides low-level control for the system hardware. New firmware release levels introduce new features and fixes. Service Packs (SPs) often bundle fixes. An SP is called an *update level*. A new release is called an *upgrade level*. All system firmware is available for download from [IBM Fix Central](#).

Terminology

Release Level: A major new function, such as the introduction of new hardware models and functions or features that are enabled through firmware. This firmware upgrade is disruptive.

Service Pack: Primarily firmware fixes and minor function changes that are applicable to a specific Release Level. These firmware updates usually occur concurrently.

Concurrent:	A code update that enables the OSs running on the Power server to continue running while the update is installed and activated.
Deferred:	A code fix that is installed concurrently but activates only after the system restarts.
Partition Deferred:	A code fix that is installed concurrently but activates only after the partition restarts.
Disruptive:	A code fix that requires a system restart during the code update process.

Service Pack severity

The severity classification is specific to each SP that becomes available in IBM FixCentral. All types are listed here.

NEW:	Introduces new features and functions. This classification is considered as a new release level for a product.
PE (Program Temporary Fix in error):	Addresses minor issues. This SP can be installed when convenient.
ATT (attention):	Addresses low-impact and low-potential issues. This SP should be installed at the customer's earliest convenience.
SPE (special attention):	Addresses high-impact but low-potential issues. This SP should be installed at the customer's earliest convenience.
HIPER (high impact or pervasive):	Addresses high-impact or pervasive issues with significant customer impact. This SP should be installed as soon as possible.

For more information, see the [IBM Support Preventive Service Planning glossary](#).

System firmware update strategy for Power11 processor-based servers

The system management model determines the appropriate firmware update strategy for Power11 processor-based servers.

- ▶ For managed systems, a best practice method for updating system firmware is through the HMC.
- ▶ For unmanaged systems, the best practice method is to update the system firmware directly through the OS.

You can also update the system firmware by using the embedded Baseboard Management Controller (eBMC).

The update method impacts the type of update:

- ▶ The HMC can perform firmware updates that are either concurrent or disruptive, depending on the firmware version.
- ▶ Firmware updates from the OS or eBMC are always disruptive.

A concurrent firmware update does not require a system restart. The firmware description file indicates this detail in the Service Pack Summary, which categorizes the update as either Disruptive Service Pack, Deferred Service Pack, or Concurrent Service Pack.

Any fixes that are marked as DEFERRED in the Service Pack Summary do not take effect until the next system initial program load (IPL). Fixes that are labeled as DEFERRED: PARTITION_DEFERRED require a partition restart.

Depending on the system management model, the following firmware update options are available:

- ▶ HMC-managed systems
Perform firmware updates by using the HMC or the eBMC ASMI.
- ▶ Co-managed systems with both PowerVM NovaLink and the HMC
Perform firmware updates by using the HMC, the NovaLink partition, or the eBMC or ASMI interface. To perform the update through the HMC, first migrate the system from PowerVM NovaLink management to HMC management.
- ▶ Unmanaged systems running IBM i
Apply updates by using PTFs or the eBMC ASMI interface.
- ▶ Unmanaged systems running AIX
Perform firmware updates through AIX system diagnostics or the eBMC ASMI interface.

All Power11 processor-based servers also support system firmware updates through the eBMC USB port.

Note: The customer is responsible for performing system firmware updates. If the customer requests that an IBM Support Service Representative do the firmware update, the request is considered as a billable service unless the customer has a valid support agreement that explicitly includes onsite firmware update coverage.

The firmware description file that is available on IBM Fix Central provides essential information about dependencies between HMC versions, AIX APARs, IBM i PTFs, and system firmware levels. Before installing a new system firmware release or SP, review the firmware description package.

For cross-version compatibility details, see the Power servers Fix Level Recommendation Tool (FLRT) in the [POWER Code Matrix](#).

Updating system firmware from the HMC

To update or upgrade the system firmware by using the Hardware Management Console (HMC), complete the following steps:

1. In the System View, select the system to update.
2. Click **Firmware**.
3. Click **Update System Firmware**.

The Update System Firmware wizard guides you through the necessary steps to complete the firmware update process, as shown in Figure 8-2 on page 191.

Figure 8-4 shows that the System Readiness Check completed successfully for the Power11 processor-based server firmware update process.



Figure 8-4 System Readiness Check completed successfully

Figure 8-5 shows the **Choose source files location** menu, which you use to update the system firmware. Several options are available for loading the required update files:

- ▶ IBM Website
- ▶ FTP Server
- ▶ SFTP Server
- ▶ Mount Point on the HMC
- ▶ CD/DVD
- ▶ USB Drive

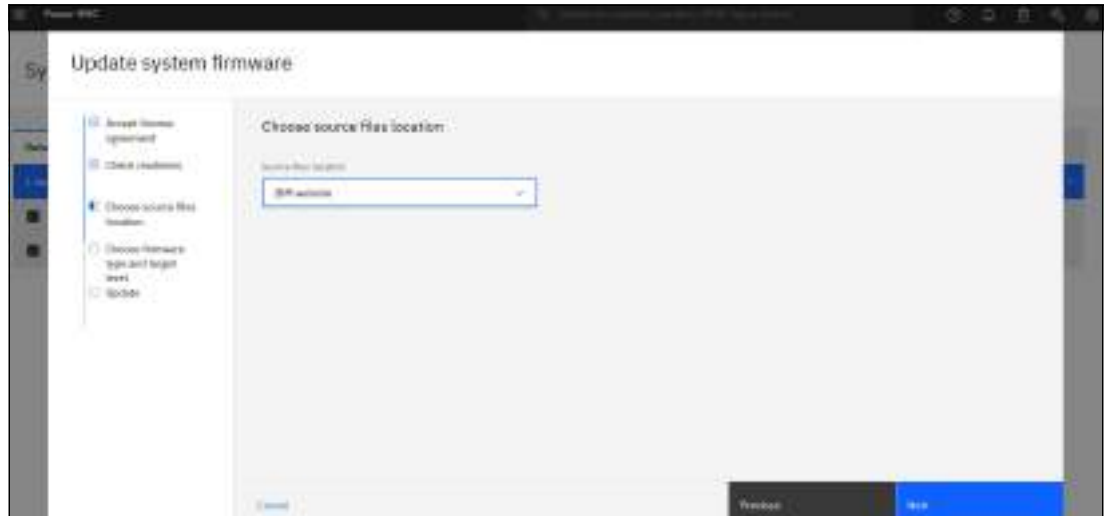


Figure 8-5 Selecting "Choose Source File Location" to update the system firmware

Click **Search Available Levels** to find the new system firmware on the installation media, as shown in Figure 8-6. Use the drop-down menu under the “Target eBMC Level” column to select the system firmware level.

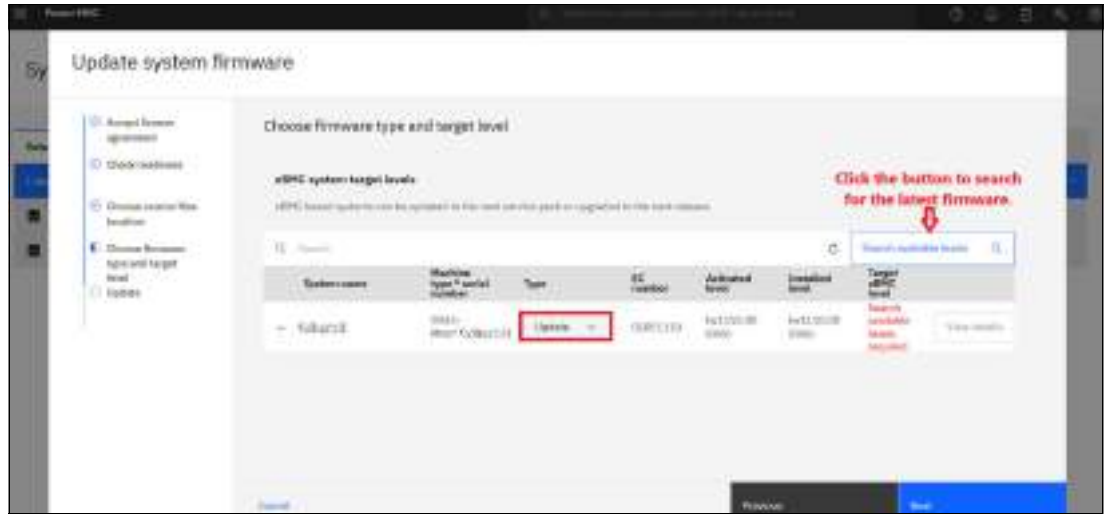


Figure 8-6 System Firmware Target Level Selection window

Import Update Files from HMC wizard

The Import Update Files wizard in the HMC GUI, as shown in Figure 8-7, guides you through the steps that are required to import firmware update or upgrade files. You can then reuse these files to update the firmware on other Power11 processor-based servers.



Figure 8-7 The “Import Update Files” option in the HMC GUI

Delete Files from Import Location wizard

The Delete Files from Import Location wizard in the HMC GUI, as shown in Figure 8-8, guides you through the steps that are required to delete imported firmware files. This process enables the management and cleanup of firmware files that are stored on the HMC.

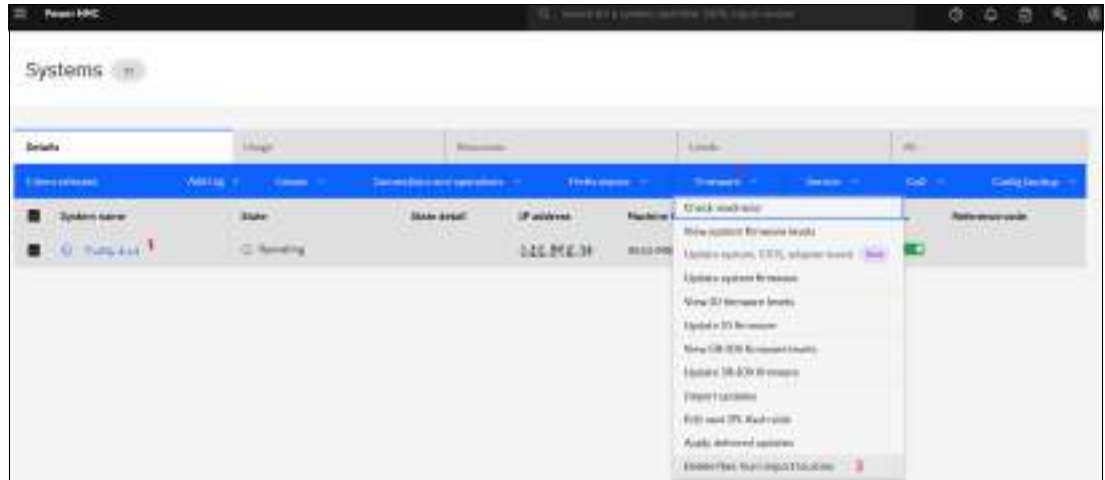


Figure 8-8 The “Delete Files from Import Location” option in the HMC GUI

SR-IOV Shared mode adapter firmware management

Single Root I/O Virtualization (SR-IOV)-capable adapters running in Shared mode use a different firmware versioning mechanism than adapters operating in Dedicated mode.

Here are the SR-IOV configuration requirements:

- ▶ Ensure that the managed system includes Peripheral Component Interconnect Express (PCIe) adapters that support SR-IOV functions.
- ▶ Each SR-IOV-capable adapter must be installed in a PCIe slot that supports SR-IOV. For more information, see the adapter placement guidelines in [IBM Documentation](#).
- ▶ Verify that the system is running supported levels for the following items:
 - System firmware
 - HMC or PowerVM NovaLink
 - An OS with SR-IOV driver support
- ▶ Configure the adapters in SR-IOV Shared Mode by using either the HMC or NovaLink interface.

Here are the requirements in SR-IOV (Shared mode):

- ▶ Adapter firmware and drivers are integrated with the system firmware.
- ▶ Firmware updates and their accompanying readme files are available from [IBM Fix Central](#).

- ▶ When an adapter is configured in SR-IOV Shared mode, its firmware and driver are automatically updated to the latest versions that are included with the system firmware. This update occurs during standard maintenance operations such as a system IPL, adapter replacement, or when switching the adapter mode between Shared and Dedicated.
- ▶ If system firmware updates are installed concurrently, you can perform selective manual updates of SR-IOV adapter firmware by using the HMC. For more information, see [IBM Documentation: Update SR-IOV firmware](#).

Note: You can update either the adapter firmware alone or both the adapter driver and firmware. During the firmware update process, network traffic on the configured logical ports of the SR-IOV adapter can be temporarily disrupted. Updating each SR-IOV adapter typically takes 2 - 5 minutes. The update is performed sequentially across all SR-IOV adapters in the system.

Figure 8-9 shows the **View SR-IOV Firmware Levels** and **Update SR-IOV Firmware** options in the HMC GUI.



Figure 8-9 “View SR-IOV Firmware Levels” and “Update SR-IOV Firmware” options

Updating the system firmware by using eBMC

You can use the eBMC ASMI to update the system firmware on both managed and unmanaged systems. To perform a firmware update by using the eBMC ASMI, power off the system. The firmware package includes multiple files, but only the firmware image with the .tar file extension is required for the update process, as shown in Figure 8-10.

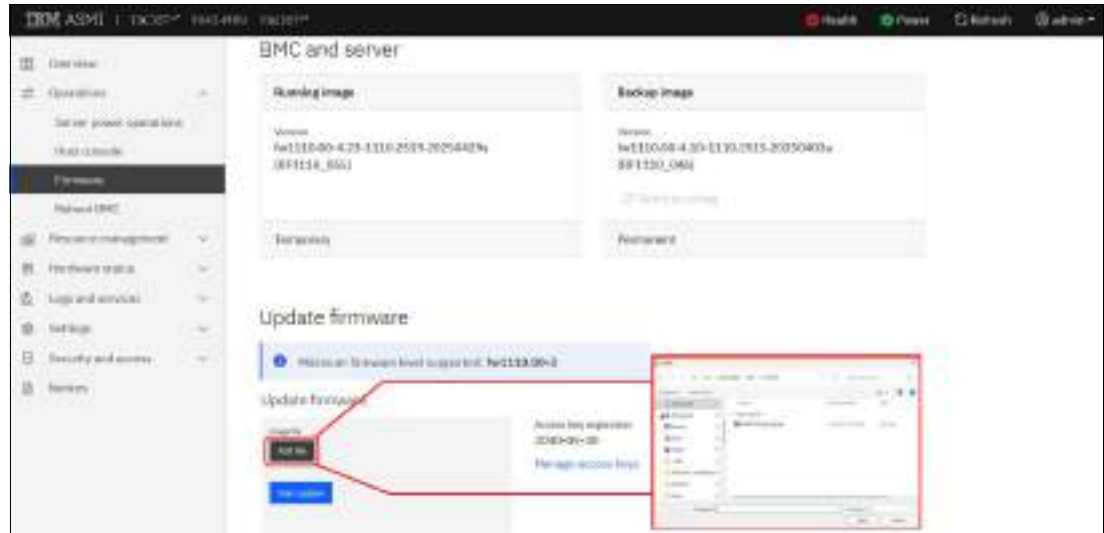


Figure 8-10 Firmware update by using the eBMC ASMI

I/O microcode update strategy

The customer is responsible for microcode updates unless the services are provided under a Microcode Support Services contract.

Device microcode can be updated by using the following methods:

- ▶ HMC
 - Use the **Update I/O Firmware** feature in the HMC GUI to apply microcode updates to supported I/O devices.
- ▶ Diagnostic menus (AIX and Virtual I/O Server (VIOS))
 - Perform microcode updates by using either the system diagnostics or [stand-alone diagnostics utilities](#), from a USB key or a network location. This method is applicable to AIX and VIOS environments.
- ▶ Linux
 - For more information about updating microcode on Linux systems, see the microcode readme files that are available from [IBM FixCentral](#). These instructions can include the use of vendor-provided tools, depending on the I/O devices.
- ▶ IBM i (PTFs)
 - This method is applicable only to IBM i systems. Microcode updates are delivered through PTFs, which can be downloaded from IBM FixCentral.

Viewing or updating I/O firmware from the HMC

You can use the HMC to view the current I/O firmware levels for a system and update those levels by using a firmware repository. The HMC does not have direct access to the partitions to perform I/O firmware updates. Instead, it relies on the Resource Monitoring and Control (RMC) infrastructure, which is also used for features such as dynamic logical partitioning and Service Focal Point, to facilitate communication between the HMC and the partitions.

Figure 8-13 illustrates the **Choose Source Files Location** menu, which you can use to specify the source location for updating the microcode on I/O devices. Several options are available for loading the required update files:

- ▶ IBM Website
- ▶ FTP Server
- ▶ SFTP Server
- ▶ Mount Point on the HMC
- ▶ CD/DVD
- ▶ USB Drive

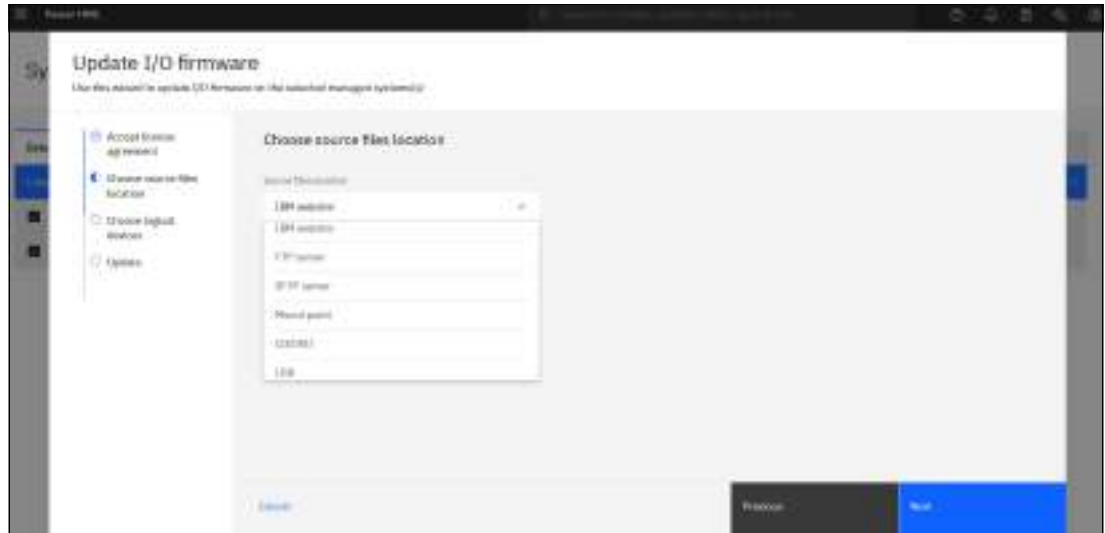


Figure 8-13 Update I/O firmware window

For more information about how to view or update the I/O firmware from the HMC, see [View or Update the I/O Firmware from the HMC](#).

8.7.1 Update Access Keys

Since the introduction of IBM Power8 processor-based servers, IBM uses the concept of a UAK for each server.

When the system applies firmware updates, it checks the UAK and its expiration date. System firmware updates include a release date. If the release date for the firmware updates is later than the UAK expiration date, the system does not process the updates.

Managing Update Access Keys

The system uses a UAK to control the application of system firmware updates. Each UAK includes an expiration date, and firmware updates include a release date. When the system applies a firmware update, it checks whether the update's release date is later than the UAK expiration date. If the release date is later, the system does not process the firmware update.

When a UAK expires, it must be replaced by using either the HMC or the eBMC ASMI. The expiration of a UAK does not affect I/O microcode updates; these updates can still be applied even if the UAK expired. Customers can obtain a new UAK by opening a case with IBM Support and requesting a renewal key, or by downloading it directly from the [IBM Entitled Systems Support \(ESS\) website](#).

By default, newly delivered systems include a UAK that expires after 3 years. Thereafter, the UAK can be extended every 6 months, but only if a current hardware maintenance contract exists for that server. The contract can be verified on the Enterprise Storage Server web page.

You can check the validity and expiration date of the current UAK by using either the HMC or eBMC GUIs or their CLIs. Also, you can retrieve the expiration date from the OS level.

Verifying the expiration date of UAKs by using the HMC

Figure 8-14 shows the current UAK expiration date on the View Current System Firmware Levels window of the HMC.



Figure 8-14 UAK expiration date displayed on the HMC View current system firmware levels window

You can use the HMC or the Virtual Hardware Management Console (vHMC) to configure the Automatic Firmware UAK. For more information, see [How to Automatically Update the Access Key](#).

Verifying the expiration date of UAKs by using the eBMC ASMI

Figure 8-15 and Figure 8-16 show the current UAK expiration date on the Firmware and Overview windows of the eBMC ASMI.



Figure 8-15 UAK expiration date displayed on the eBMC Firmware window

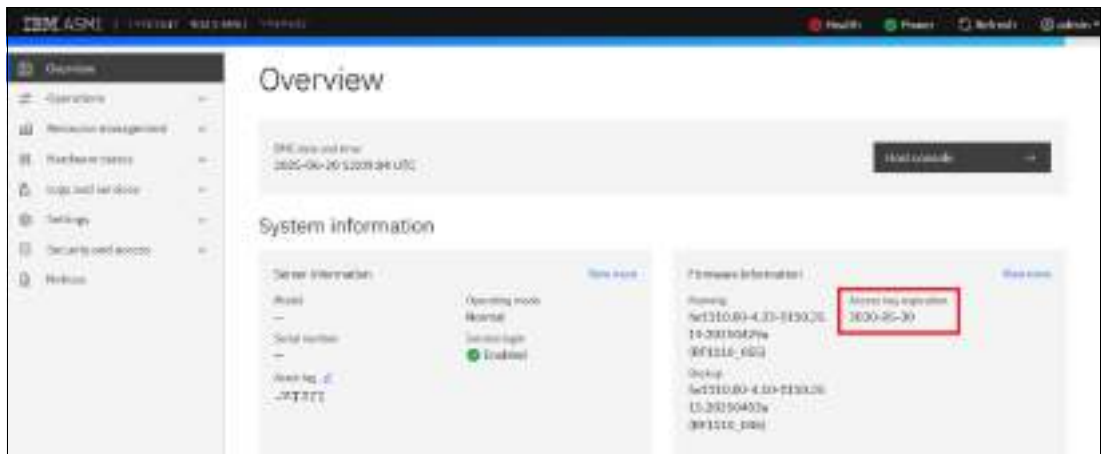


Figure 8-16 UAK expiration date displayed on the eBMC Overview window

Verifying the expiration date of UAKs by using AIX

There are multiple methods of checking the UAK expiration date within AIX:

- ▶ Use the `lscfg` command:

```
lscfg -vpl sysplanar0 |grep -p "System Firmware"
```

The output is similar to the output that is shown in Example 8-1 on page 201.

Example 8-1 Output of the lscfg command to check the UAK expiration date

```
lscfg -vpl sysplanar0 |grep -p "System Firmware"
System Firmware:
Code Level, LID Keyword.....PHYP_1 21372025061280A00701
Code Level, LID Keyword.....PFW 21212025060681CF0681
Code Level, LID Keyword.....FSP_Fil 16112025061681E00109
Code Level, LID Keyword.....FipS_BU 16112025061681E00208
Microcode Image.....RK1110_058 RK1110_057 RK1110_058
Microcode Level.....FW1110.00 FW1110.00 FW1110.00
Microcode Build Date.....20250725 20250725 20250725
Update Access Key Exp Date..20301014
Hardware Location Code.....U9080.HEU.DEN0013-Y1
Physical Location: U9080.HEU.DEN0043-Y1
```

- ▶ Run a **grep** command on Access Key, as shown in Example 8-2.

Example 8-2 Output of the lscfg command to check the UAK expiration date

```
# lscfg -vpl sysplanar0 | grep "Access Key"
Update Access Key Exp Date..20301014
```

- ▶ On AIX 7.3, run the **lparstat** command, as shown in Example 8-3.

Example 8-3 Output of the lparstat command to check the UAK expiration date

```
# lparstat -u
FW Update Access Key Expiration (YYYYMMDD): 20301014
AIX Update Access Key Expiration (YYYYMMDD): 20301014
AIX Image Date (YYYYMMDD): 20250725
```

Verifying the expiration date of UAKs by using Linux

There is not a single Linux command to view a UAK directly. Instead, UAKs are managed through IBM tools and interfaces, such as IBM Electronic Service Agent (ESA), ASMI, and HMC.

Managing the UAK with IBM i

ESA on IBM i introduces a new feature to help clients manage the refresh of the UAK for IBM i stand-alone partitions. There are two ways of using this UAK management feature:

- ▶ Manual

You can manually start the process to check the UAK expiration date, download the UAK, and apply the UAK by running the **WRKSRVAGT TYPE(*UAK)** command. You can also run the **GO SERVICE** command to access the ESA on IBM i main menu, and then select Option 20. Check and refresh Update Access Key, as shown in Figure 8-17 and Figure 8-18.



Figure 8-17 Electronic Service Agent on IBM i main menu option 20

The help instructions for Option 20 can help you understand this feature.

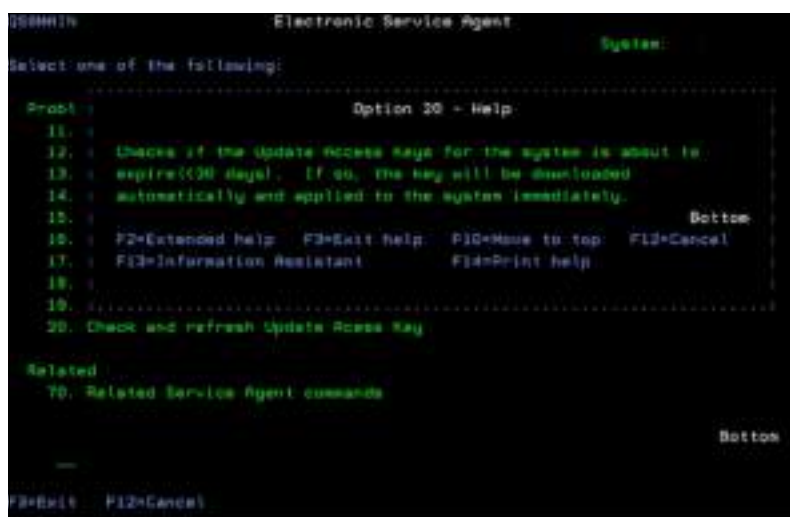


Figure 8-18 Help information on Electronic Service Agent on IBM i main menu option 20

► Automatic

You can use the Change Service Agent Attribute (**CHGSRVAGTA**) command to enable automatic checking. You can use the **CHGSRVAGTA** command parameter **REFRESHUAK** to establish automatic UAK management. By default, the job that is named QS9UAK runs every Sunday, but you can customize it.

To enable the automatic function, complete the following steps:

- a. Enter **CHGSRVAGTA**, and then press F4.
- b. Enter Yes for the Enable field.
- c. Press Page Down to continue.

The Change Service Agent Attribute panel appears, as shown in Figure 8-19. Specify the weekday and time for UAK refresh processing.

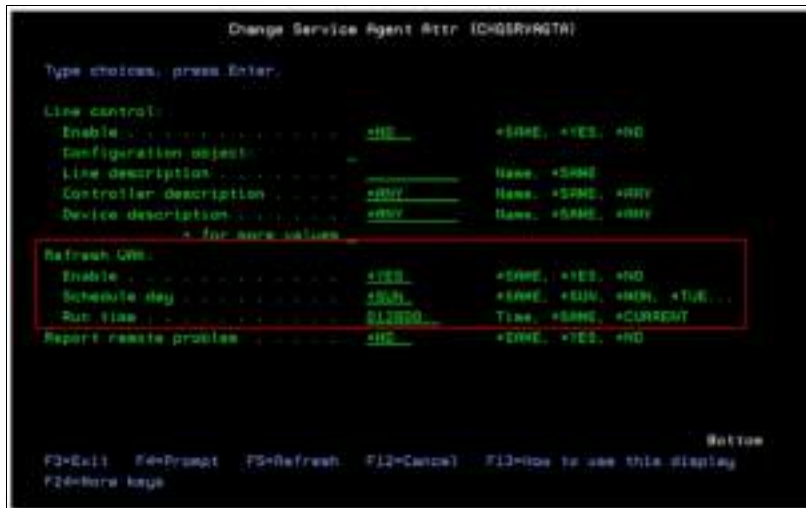


Figure 8-19 Change Service Agent Attribute refresh UAK configuration panel



Virtualization

Virtualization on Power servers forms a cornerstone of IBM enterprise computing strategy, providing robust, scalable, and secure environments for running multiple workloads on a single physical server. At the heart of this capability is IBM PowerVM, an IBM enterprise-grade virtualization technology. PowerVM enables the creation of logical partitions (LPARs) enabling multiple operating systems (OSs) such as AIX, IBM i, and Linux to run concurrently on the same hardware. It supports advanced features such as Live Partition Mobility (LPM), dynamic resource allocation, and IBM Micro-Partitioning®, which help maximize hardware usage and reduce operational costs.

In addition to PowerVM, Power servers also support Kernel-based Virtual Machine (KVM), an open-source virtualization option. KVM on Power provides a flexible and cost-effective alternative for Linux-based workloads, particularly in cloud-native and containerized environments. It integrates with modern orchestration tools and supports features such as nested virtualization and Single Root I/O Virtualization (SR-IOV) for high-performance networking. KVM is ideal for organizations that want to use open-source technologies while benefiting from the performance and reliability of IBM Power hardware.

A key component of virtualization efficiency on Power servers is the usage of shared processor pools (SPPs). SPPs enable multiple LPARs to share a pool of physical processors, enabling dynamic allocation of CPU resources that is based on workload demand. With the IBM Power11 platform, PowerVM adds support for Resource Groups, which enhance processor sharing capabilities and improve efficiency. Processor Pools extend the capabilities of SPPs to improve processor usage and enhance isolation and control by capping the maximum CPU resources that are available to each pool. When combined with PowerVM or KVM, SPPs help optimize performance, reduce licensing costs, and help ensure consistent service levels across virtual environments.

This chapter contains the following topics:

- ▶ PowerVM
- ▶ KVM support

9.1 PowerVM

The PowerVM platform is composed of a family of technologies, capabilities, and offerings that deliver industry-leading virtualization for enterprises. It is the overall branding term for Power server virtualization, including IBM PowerVM Hypervisor (PHYP), logical partitioning, IBM Micro-Partitioning, Virtual I/O Server (VIOS), LPM, and other related technologies. PowerVM combines hardware enablement and software.

9.1.1 IBM PowerVM Hypervisor

Power servers, when combined with PowerVM technology, offer key capabilities that help consolidate and simplify IT environments:

- ▶ Increase server usage and share I/O resources to reduce 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 match changing business requirements or handle unexpected changes in demand.
- ▶ Simplify IT infrastructure management by making workloads independent of hardware resources so that business-driven policies can deliver resources based on time, cost, and service-level requirements.

When combined with features in Power11 scale-out servers, the PHYP delivers functions that enable other system technologies, including LPAR technology, virtualized processors, IEEE virtual local area network (VLAN)-compatible virtual switches, virtual SCSI adapters, virtual Fibre Channel (FC) adapters, and virtual consoles.

The PHYP is a basic component of the system 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 VLAN channels between LPARs to reduce the need for physical Ethernet adapters for inter-partition communication.
- ▶ Monitors the Flexible Service Processor (FSP) and performs a reset or reload if the system detects the loss of an FSP, notifying the OS if the problem is not corrected.

The PHYP is always active, regardless of the system configuration or whether it is connected to the managed console. It requires memory to support the resource assignment of the LPARs on the server. The amount of memory that the PHYP firmware requires varies according to several factors:

- ▶ Memory usage for hardware page tables (HPTs)
- ▶ Memory usage to support I/O devices
- ▶ Memory usage for virtualization

Memory usage for hardware page tables

Each partition on the system includes its own HPT, which contributes to hypervisor memory usage. The OS uses the HPT to translate effective addresses to real physical addresses in the hardware. This translation enables multiple OSs to run simultaneously in their own logical address space. When the system dispatches a virtual processor for a partition on a physical processor, the hypervisor specifies to the hardware the location of the partition HPT to use 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 for AIX, VIOS, and Linux partitions is 1:128 of the maximum memory size of the partition. AIX, VIOS, and Linux use larger page sizes (16 KB and 64 KB) instead of 4 KB pages. Using larger page sizes reduces the overall number of pages that must be tracked; therefore, the overall size of the HPT can be reduced. For example, the HPT is 2 GB for an AIX partition with a maximum memory size of 256 GB.

When defining a partition, the specified maximum memory size is based on the amount of memory that can be dynamically added to the dynamic logical partition (DLPAR) without changing the configuration and restarting the partition.

In addition to setting the maximum memory size, the HPT ratio can be configured. You can issue the `chsyscfg` Hardware Management Console (HMC) command with the `hpt_ratio` parameter to define the HPT ratio that is used for a partition profile. 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 increase CPU consumption if the OS must reload entries in the HPT more frequently. Most customers choose to use the IBM-provided default values for HPT ratios.

Memory usage for I/O devices

To support I/O operations, the hypervisor maintains structures that are called *translation control entities* (TCEs), which provide an information path between I/O devices and partitions. TCEs contain the address of the I/O buffer, indications of read versus write requests, and other I/O-related attributes. Each I/O device uses many TCEs, enabling multiple requests to be active simultaneously for the same physical device. To provide better affinity, the system distributes TCEs across multiple processor chips or drawers to improve performance when accessing them.

For physical I/O devices, the hypervisor defines the base amount of space for TCEs based on the number of supported I/O devices. Systems that support high-speed adapters can be configured to allocate more memory to improve I/O performance. Linux is the only OS that uses these extra TCEs, enabling the system to free memory for use by partitions if only AIX is running.

Memory usage for virtualization features

Virtualization requires the PHYP to allocate more memory for hardware statesave areas and various virtualization technologies. For example, on Power11 processor-based servers, each processor core supports up to eight simultaneous multithreading (SMT) threads (SMT8) of execution, and each thread contains more than 80 different registers.

The PHYP must reserve save areas for the register contents for the maximum number of virtual processors that are configured. The greater the number of physical hardware devices, virtual devices, and virtualization features, the more hypervisor memory that is required. For efficient memory consumption, wanted and maximum values for various attributes (processors, memory, and virtual adapters) must be based on business needs and should not be set higher than actual requirements.

Predicting memory that is used by the PowerVM Hypervisor

The IBM System Planning Tool (SPT) is a resource that estimates the amount of hypervisor memory that is required for a specific server configuration. After you download and install the SPT executable file, define a configuration by selecting the correct hardware platform, installed processors and memory, and partition attributes. SPT can estimate the amount of memory that is assigned to the hypervisor, which helps with changing a configuration or deploying servers.

The PHYP provides the following types of virtual I/O adapters:

- ▶ Virtual SCSI

The PHYP provides a virtual SCSI mechanism for the virtualization of storage devices. Storage virtualization is accomplished by using two paired adapters: a virtual SCSI server adapter and a virtual SCSI client adapter.

- ▶ Virtual Ethernet

The PHYP provides a virtual Ethernet switch function that enables partitions to communicate quickly and securely on the same server without requiring physical interconnection or external connectivity. If a Layer 2 bridge to a physical Ethernet adapter is configured in a VIOS partition, it is known as a Shared Ethernet Adapter (SEA).

- ▶ Virtual FC

A virtual FC adapter provides customer LPARs with a Fibre Channel connection (IBM FICON®) to a storage area network (SAN) through the VIOS partition. The VIOS partition connects the virtual FC adapters on the VIOS partition to the physical FC adapters on the managed system.

- ▶ Virtual (tty) Console

Each partition must have access to a system console. Tasks such as OS installation, network setup, and problem analysis activities require a dedicated system console. The PHYP provides the virtual console by using a virtual tty or serial adapter and a set of hypervisor calls to operate them. Virtual tty does not require the purchase of more features or software, such as the PowerVM Edition features.

Logical partitions

LPARs and virtualization increase system resource usage and add configuration flexibility. *Logical partitioning* enables a server to operate as though it were two or more independent servers. When you logically partition a server, you divide its resources into subsets that are called LPARs. You can install software on an LPAR, and the LPAR operates as an independent logical server with the resources that are allocated to it.

An LPAR is also referred to in some documentation as a *virtual machine* (VM), which makes it appear similar to what other hypervisors offer. However, LPARs provide a higher level of security, isolation, and other features.

You can assign processors, memory, and I/O devices to LPARs. AIX, IBM i, Linux, and VIOS can run on LPARs. VIOS provides virtual I/O resources to other LPARs that run general-purpose OSs.

LPARs share a few system attributes, such as the system serial number, system model, and processor Feature Codes. All other system attributes can vary from one LPAR to another.

Micro-Partitioning

With IBM 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 (SPP), and virtual processors enable the OS to manage the fractions of processing power that are assigned to the LPAR.

From an OS perspective, a virtual processor is indistinguishable from a physical processor unless the OS is enhanced to detect the difference. Physical processors are abstracted into virtual processors that are available to partitions.

On an Power11 processor-based server, you can define a partition with a processor capacity as small as 0.05 processing units. This number represents 0.05 of a physical core. Each physical core can be shared by up to 20 shared processor partitions, and the partition's entitlement can be incremented fractionally by as little as 0.05 of the processor. The PHYP dispatches and time-slices the shared processor partitions on the physical processors. The HMC creates and manages the shared processor partitions.

Processing mode

When you create an LPAR, you can assign entire processors for dedicated use or partial processing units from an SPP. This setting defines the processing mode of the LPAR.

Dedicated mode

In dedicated mode, the system assigns physical processors as a whole to partitions. The SMT feature in the Power11 processor core enables the core to run instructions from two, four, or eight independent software threads simultaneously.

Shared dedicated mode

On Power11 processor-based servers, you can configure dedicated partitions to act as processor donors for idle processors that they own, enabling the donation of spare CPU cycles from dedicated processor partitions to a shared processor pool (SPP). The dedicated partition retains absolute priority for its dedicated CPU cycles. Enabling this feature can increase system usage without compromising computing power for critical workloads in a dedicated processor mode LPAR.

Shared mode

In shared mode, LPARs use virtual processors to access fractions of physical processors. Shared partitions can define any number of virtual processors, with a maximum equal to 20 times the number of processing units that are assigned to the partition. The PHYP dispatches virtual processors to physical processors according to the partition's processing unit entitlement. One processing unit represents the processing capacity of one physical processor. All partitions receive total CPU time that is equal to their processing unit entitlement.

Logical processors are defined on top of virtual processors. Therefore, even with a virtual processor, the concept of a logical processor applies, and the number of logical processors depends on whether SMT is enabled or disabled.

9.1.2 Multiple shared processor pools

Power11 processor-based servers support multiple shared processor pools (MSPPs). This capability enables a system administrator to create a set of micropartitions to control the processor capacity that can be used from the physical SPP.

Micropartitions are created and identified as members of either the default processor pool or a user-defined SPP. The PHYP monitors the virtual processors within the set of micropartitions, and processor capacity is managed according to user-defined attributes.

If the IBM Power server is under heavy load, each micropartition within an SPP is assured of its processor entitlement, plus any capacity that might be allocated from the reserved pool if the micropartition is uncapped. If specific micropartitions in an SPP do not use their processing capacity entitlement, the unused capacity is ceded, and other uncapped micropartitions within the same SPP can use the extra capacity according to their uncapped weighting. In this way, the entitled pool capacity of an SPP is distributed among the set of micropartitions within that SPP.

All IBM Power servers that support the MSPP capability have a minimum of one SPP (the default) and up to a maximum of 64 SPPs. This capability can help customers reduce TCO when the cost of software or database licenses depends on the number of assigned processor cores.

SPPs: SPPs can be used with Resource Groups, and such configurations can improve performance and isolation.

9.1.3 Virtual I/O Server

The VIOS is part of IBM PowerVM. It is the specific appliance that enables the sharing of physical resources among LPARs to improve efficiency, for example, through consolidation. In this configuration, the VIOS owns the physical I/O resources (SCSI, FC, network adapters, and optical devices) and enables customer partitions to share access to them, minimizing and optimizing the number of physical adapters in the system.

The VIOS eliminates the requirement for every partition to have a dedicated network adapter, disk adapter, and disk drive. It supports OpenSSH for secure remote logins and provides a firewall to limit access by ports, network services, and IP addresses.

Figure 9-1 on page 211 shows an overview of a VIOS configuration.

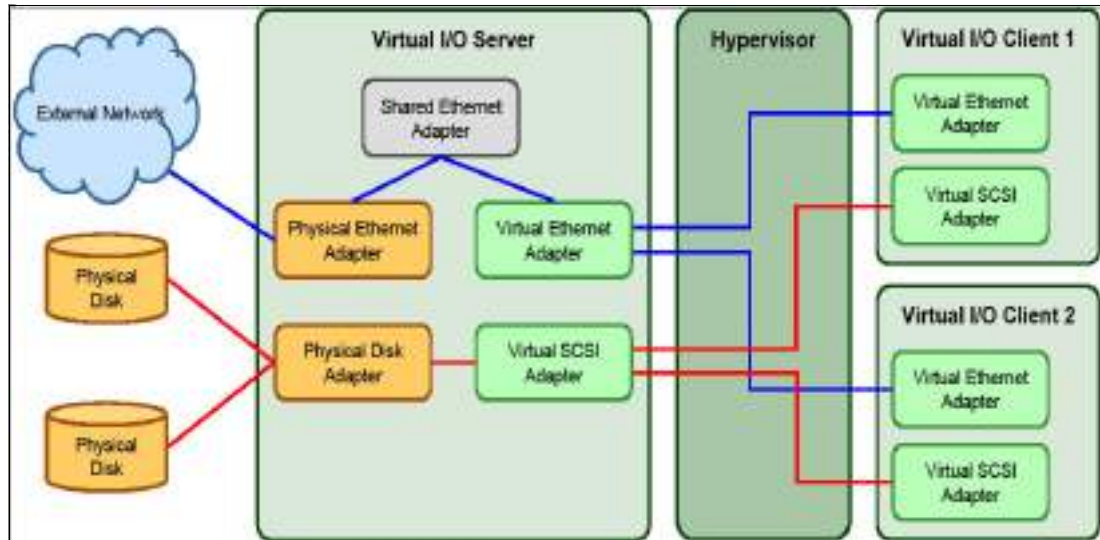


Figure 9-1 Architectural view of the VIOS

It is a best practice to run dual VIOSs per physical server.

Shared Ethernet Adapter

A SEA connects a physical Ethernet network to a virtual Ethernet network. The SEA provides access by connecting PHYP VLANs to VLANs on external switches. Because the SEA processes packets at Layer 2, the original MAC address and VLAN tags remain visible to other systems on the physical network. IEEE 802.1 VLAN tagging is supported.

By using the SEA, several customer LPARs can share one physical adapter. You can also 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 internal virtual Ethernet networks and one or more external physical network adapters that are configured with Etherchannel. The PHYP defines these virtual Ethernet network adapters on the VIOS.

Virtual SCSI

Virtual SCSI provides a virtualized implementation of the Small Computer System Interface (SCSI) protocol. It is based on a client/server relationship. The VIOS LPAR owns the physical I/O resources and acts as a server, or in SCSI terms, a target device. Client LPARs access the virtual SCSI backing storage devices that are provided by the VIOS.

The virtual I/O adapters (a virtual SCSI server adapter and a virtual SCSI client adapter) are configured by using the HMC. The virtual SCSI server (target) adapter is responsible for running any SCSI commands that it receives and is owned by the VIOS partition. The virtual SCSI client adapter enables a client partition to access physical SCSI and SAN-attached devices and LUNs that are mapped for use by the client partitions. The VIOS provisions virtual disk resources.

Internet Small Computer Systems Interface

The Internet Small Computer Systems Interface (iSCSI) disk is supported in VIOS 3.1.0 and later. An iSCSI disk provides block-level access to storage devices by carrying SCSI commands over an IP network. The iSCSI disk facilitates data transfers over the internet by using TCP. The iSCSI disk can manage storage over long distances.

iSCSI support in the VIOS enables iSCSI disks to be exported to client LPARs as vSCSI disks. This support is available in VIOS 3.1 and later. VIOS 3.1 enables multipath I/O (MPIO) support for the iSCSI initiator. With MPIO support, configure and create multiple paths to an iSCSI disk, similar to other protocols. The client LPAR can run either an AIX or Linux OS.

VIOS 3.1.1 enables support for multiple iSCSI initiators on the VIOS. This support also includes performance enhancements for the iSCSI driver. With multiple iSCSI initiator support, create multiple iSCSI software initiator devices on a single AIX OS instance.

N_Port ID Virtualization

N_Port ID Virtualization (NPIV) is a technology that enables multiple LPARs to access one or more external physical storage devices through the same physical FC adapter. This adapter is attached to a VIOS partition that acts only as a pass-through, managing the data transfer through the PHYP.

Each partition features one or more virtual FC adapters, each with a unique pair of worldwide port names. This configuration enables you to connect each partition to independent physical storage on a SAN. Unlike virtual SCSI, only the client partitions can access the disk.

For more information and requirements for NPIV, see *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

9.1.4 Live Partition Mobility

LPM enables you to move a running LPAR from one system to another without disruption. Inactive partition mobility enables you to move a powered-off LPAR from one system to another.

LPM provides systems management flexibility and improves system availability by avoiding the following situations:

- ▶ Planned outages for hardware upgrades or firmware maintenance.
- ▶ Unplanned downtime. With preventive failure management, if a server indicates a potential failure, move its LPARs to another server before the failure occurs.

HMC 10.1 (HMC V10R1) and VIOS 3.1.3 or later provide the following enhancements to the LPM feature:

- ▶ Automatically choose the fastest network for LPM memory transfer.
- ▶ Allow LPM when a virtual optical device is assigned to a partition.

9.1.5 Active Memory Expansion

Active Memory Expansion (AME) is an optional Feature Code (#EMBP) under the IBM PowerVM purview that enables memory expansion on the system. AME is an innovative technology that supports the AIX OS. It enables the effective maximum memory capacity to be larger than the true physical memory maximum. Compression and decompression of memory content can expand memory capacity by up to 100% or more. This expansion can enable a partition to complete more work or support more users with the same physical amount of memory. Similarly, AME can enable a server to run more partitions and perform more work with the same physical amount of memory.

AME uses CPU resources to compress and decompress memory contents. The tradeoff of memory capacity for processor cycles can be an excellent choice, but the degree of expansion depends on how compressible the memory content is. It also depends on having adequate spare CPU capacity that is available for compression and decompression.

The IBM Power E1080 includes a hardware accelerator that is designed to boost AME efficiency and use fewer processor core resources. Each AIX partition can enable or disable AME. Control parameters set the wanted amount of expansion in each partition to help manage the CPU resources that are used by the AME function.

An initial program load (IPL) is required for the specific partition that is enabling memory expansion. When enabled, monitoring capabilities are available in standard AIX performance tools such as `lparstat`, `vmstat`, `topas`, and `svmon`.

A planning tool that is included with AIX enables you to sample workloads and estimate how expandable the partition's memory is and how many CPU resources are needed. The feature can be ordered with the initial order of the Power E1080 or as a Miscellaneous Equipment Specification (MES) order. A software key is provided when the enablement feature is ordered, which is applied to the system node. An IPL is not required to enable the system node. The key is specific to an individual system, is permanent, and cannot be moved to a different server.

IBM i does not support AME.

9.1.6 Remote Restart

Remote Restart is a high availability (HA) option for LPARs. If an error causes a server outage, a partition that is configured for Remote Restart can be restarted on a different physical server. If starting the server takes longer, the Remote Restart function can be used for faster reprovisioning of the partition. Typically, Remote Restart is faster than restarting the stopped server and then restarting its partitions.

The Remote Restart function relies on technology similar to LPM, where a partition is configured with storage on a SAN that is shared and accessible by the server hosting the partition.

HMC V10R1 provides an enhancement to the Remote Restart feature that enables remote restart when a virtual optical device is assigned to a partition.

9.1.7 POWER processor modes

Although they are not virtualization features, IBM POWER processor modes are described here because they affect various virtualization capabilities.

On IBM Power servers, partitions can be configured to run in several processor modes, including the following ones:

- ▶ The Power9 platform
The native mode for IBM Power9 processors implements IBM Power Instruction Set Architecture (ISA) 3.0.
- ▶ The Power10 platform
The native mode for IBM Power10 processors implements ISA 3.1.
- ▶ The Power11 platform
The native mode for Power11 processors implements ISA 3.1.

For more information, see the [Power11 documentation](#).

Figure 9-2 shows the available processor modes on a Power11 processor-based server.

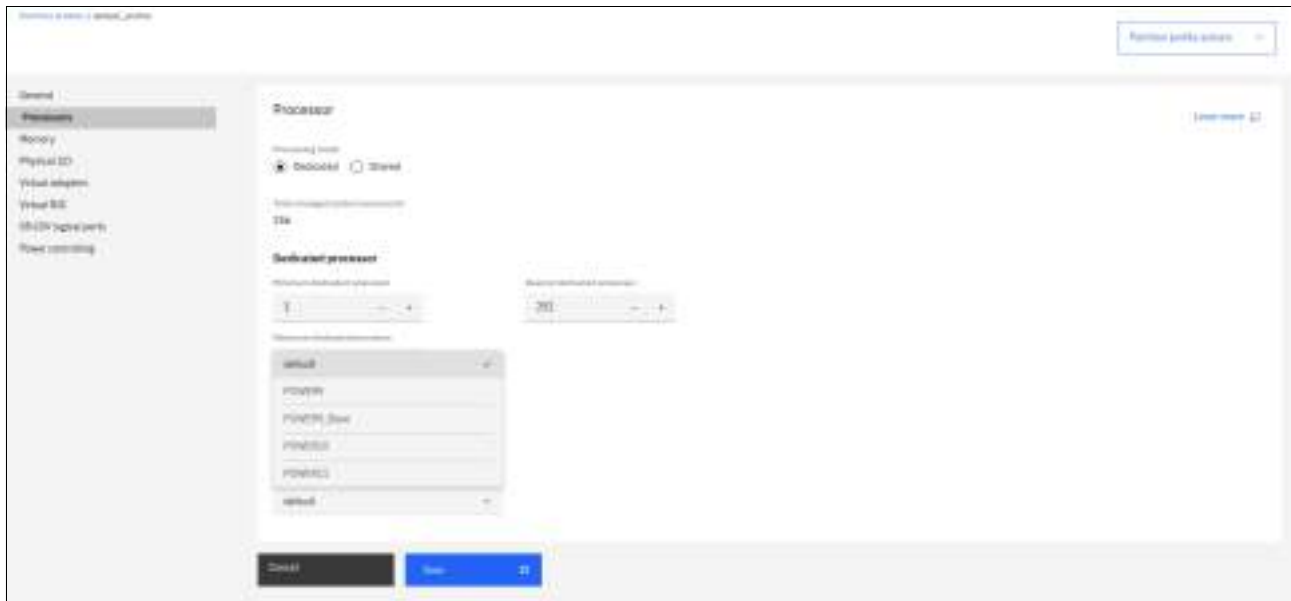


Figure 9-2 Processor modes

Processor compatibility mode is important when LPM migration is planned between different generations of servers. An LPAR that might be migrated to a server that is managed by a processor from another generation must be activated in a specific compatibility mode.

9.1.8 Single Root I/O Virtualization

SR-IOV is an extension to the Peripheral Component Interconnect Express (PCIe) specification that enables multiple OSs to share a PCIe adapter simultaneously with little or no runtime involvement from a hypervisor or other virtualization intermediary. SR-IOV is a PCI standard architecture that enables PCIe adapters to become self-virtualizing. It enables adapter consolidation through sharing, similar to how logical partitioning enables server consolidation.

With an adapter capable of SR-IOV, assign virtual *slices* of a single physical adapter to multiple partitions through logical ports without using a VIOS.

9.1.9 Virtualization features resources

For more information about virtualization features, see the following IBM Redbooks publications:

- ▶ *IBM PowerVM Best Practices*, SG24-8062
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- ▶ *IBM Power Systems SR-IOV: Technical Overview and Introduction*, REDP-5065

9.1.10 Resource Groups

Resource Groups are a new feature that is introduced with Power11 processor-based systems and designed to enhance Resource Groups are a new feature that was introduced with the Power11 platform. It is designed to enhance system performance by up to 25% through improved workload optimization and resource affinity.

Although SPPs already provide compute capacity isolation by capping the maximum resources that are available to each pool, Resource Groups extend this capability by introducing advanced affinity-based optimizations for more efficient workload dispatching.

Early performance evaluations indicate that when Resource Groups are configured effectively, they can deliver shared processor performance that closely matches dedicated processors, particularly in large, partitioned environments.

Resource Groups properties

With the introduction of Resource Groups in Power11 processor-based systems, system administrators now have a powerful tool for organizing and optimizing compute resources. Designed for flexibility and performance, Resource Groups enable the grouping of both dedicated and shared processor partitions, which enables more granular control over resource allocation and workload management. This feature enhances system efficiency, supports dynamic reconfiguration, and integrates seamlessly with existing technologies such as SPPs, LPM, and Power Enterprise Pools (PEP) 2.0.

Here are the core features of Resource Groups:

- ▶ Performance optimization

Resource Groups improve workload affinity, enabling shared processor performance that is nearly equivalent to dedicated processors.

- ▶ Flexible configuration

Resource Groups can include both dedicated processor partitions and shared processor partitions. A Resource Group configuration specifies the number of general-purpose cores (AIX, IBM i, Linux, or VIOS) and IFL cores (Linux or VIOS). Administrators can reassign powered-off partitions between groups. The system places resources that are not assigned to a user-defined group in the default Resource Group. Cores can be dynamically reallocated among groups.

- ▶ Affinity and isolation

Each Resource Group has its own set of SPPs. This configuration enhances compute capacity isolation and affinity-based dispatching.

- ▶ Monitoring management

Usage metrics are available for Resource Groups, SPPs, and the overall system. Dynamic Platform Optimizer can run at the system or group level.

- ▶ Mobility compatibility

When using LPM, administrators can select the Resource Group on the target system. Resource Groups are fully compatible with PEP 2.0.

Resource Groups use cases

There are multiple scenarios in which Resource Groups can deliver significant benefits to customers. This section presents several example use cases. Beyond these examples, organizations can combine use cases to maximize functions and tailor performance to specific business needs.

Consolidating across multiple lines of business

Figure 9-3 illustrates resource allocation and sharing across multiple lines of business on Power11 processor-based servers. With the introduction of Resource Groups in Power11 processor-based servers, the system allocates resources more efficiently, which enables improved workload isolation and performance across the system.

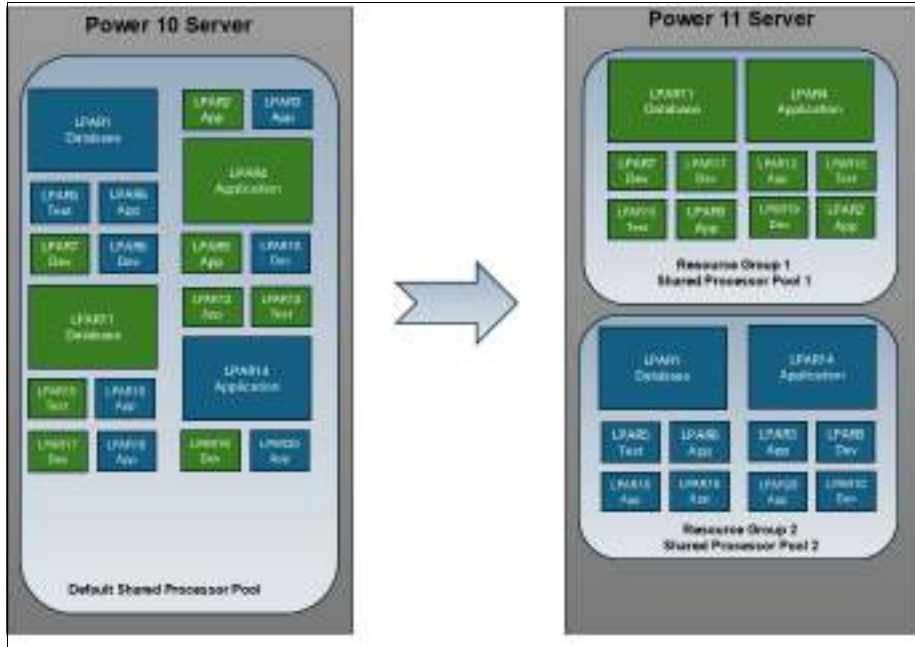


Figure 9-3 Example where Resource Groups consolidate multiple lines of business

Isolating production workloads from test/dev workloads

In this example, the system restricts the development and test LPARs to using resources only from the default Resource Group, which limits their access to shared system resources and isolates them from production workloads.

Figure 9-4 on page 217 illustrates how Resource Groups provide isolation for development and test workloads.

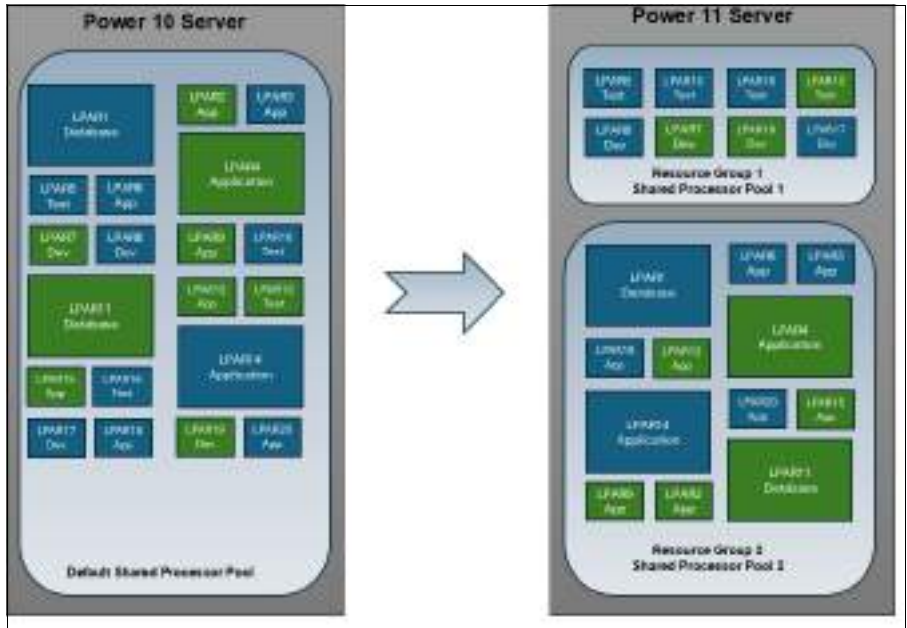


Figure 9-4 Example where Resource Groups isolate workloads

Improving application performance by grouping workload tiers into Resource Groups

In this example, improved affinity between the database and application server enhances performance, enabling more efficient resource usage and faster communication within the same Resource Group.

Figure 9-5 illustrates how grouping workloads into different Resource Groups can improve application affinity and performance.

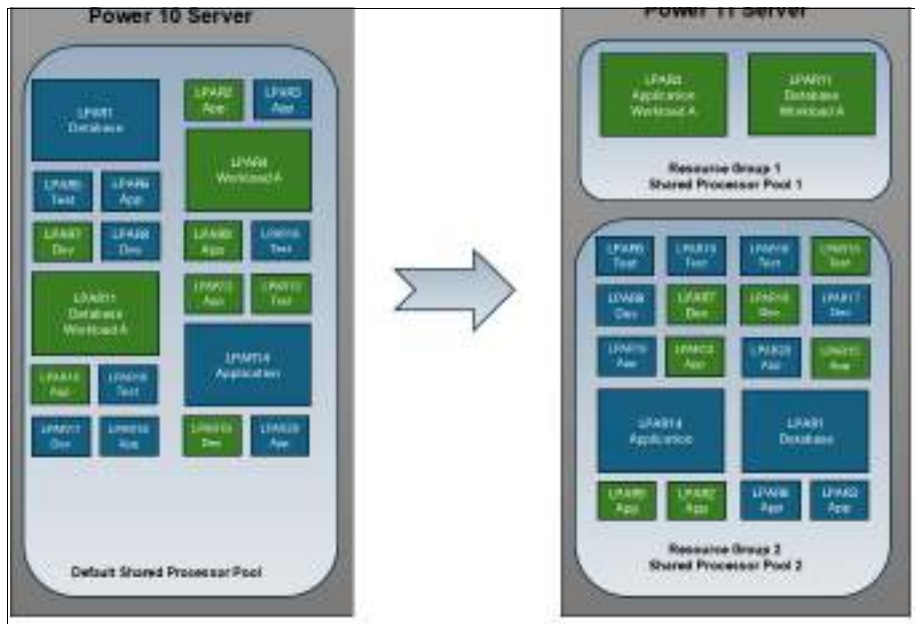


Figure 9-5 Example where Resource Groups improve performance by grouping workload tiers

Providing system-level isolation in multi-server consolidation scenarios

Resource Groups enable effective workload isolation on Power11 processor-based servers. This capability enables organizations to consolidate multiple smaller servers into a single, more powerful Power11 processor-based server while logically separating workloads into distinct Resource Groups. Each group can be managed independently, helping ensure performance, security, and resource control across different business functions or environments.

Figure 9-6 illustrates how using Resource Groups can help maintain system-level isolation during server consolidation.

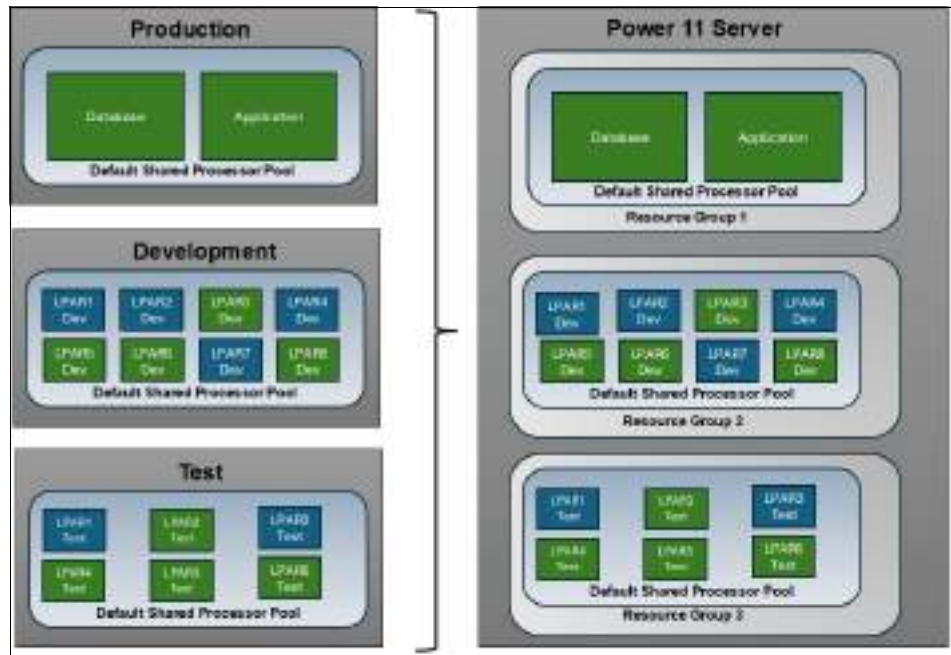


Figure 9-6 Example where Resource Groups provide an extra layer of workload isolation

Improving performance by mapping SPPs into Resource Groups

You can use SPPs and Resource Groups together on Power11 processor-based servers to optimize application performance and reduce TCO. This combination enables precise control over resource allocation and workload isolation, helping ensure efficient usage of compute capacity and maintaining performance consistency across diverse workloads.

Figure 9-7 on page 219 shows how SPPs and Resource Groups enhance performance and reduce your TCO.

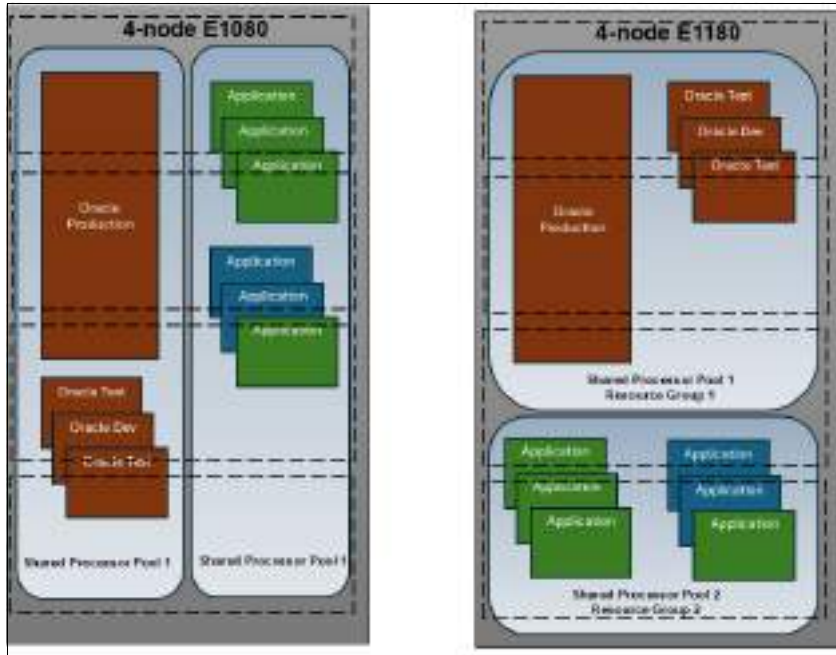


Figure 9-7 Improving performance by mapping shared processor pools into Resource Groups

Resource Groups Advisor

Resource Groups Advisor (RGA) is a web-based tool that is designed to assist with the configuration of Resource Groups on Power servers. Offered as a no-charge service, RGA analyzes a customer's server setup and provides tailored best practices for optimal Resource Group configurations. It helps model and validate configurations to help ensure efficient resource allocation and system performance.

For more information about RGA, see [Resource Groups Advisor for IBM Power](#).

Figure 9-8 shows the welcome window for RGA.



Figure 9-8 Welcome window

Click **Getting Started**. A window opens where you can create a configuration or upload a previously saved version. Click **New Configuration**, which opens the window that is shown in Figure 9-9.

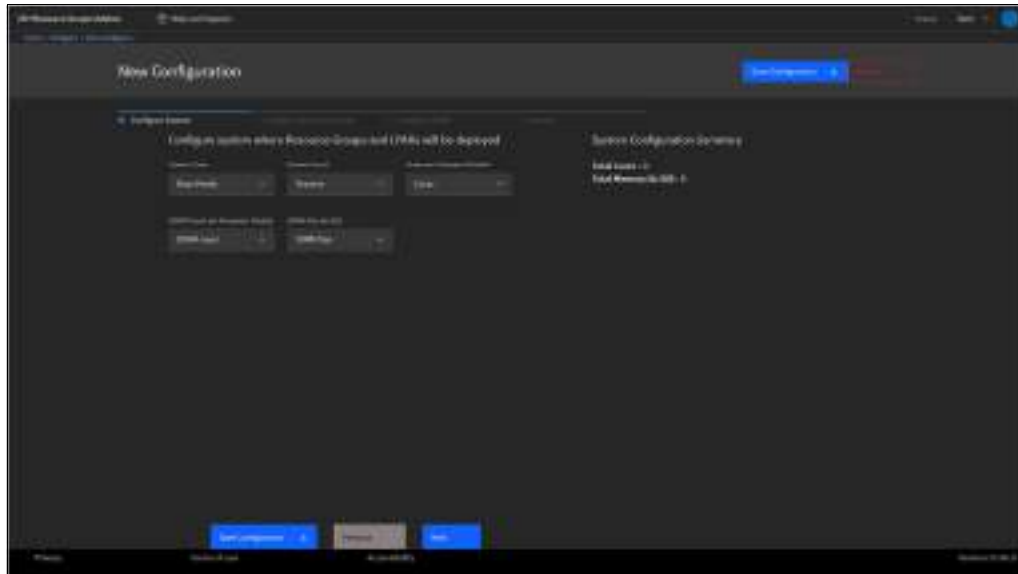


Figure 9-9 New configuration window

After you provide the system configuration parameters, click **Next**. The system displays the planned Resource Group configuration for your system, as shown in Figure 9-10. The system always configures the default Resource Group. You can add more Resource Groups to your configuration.

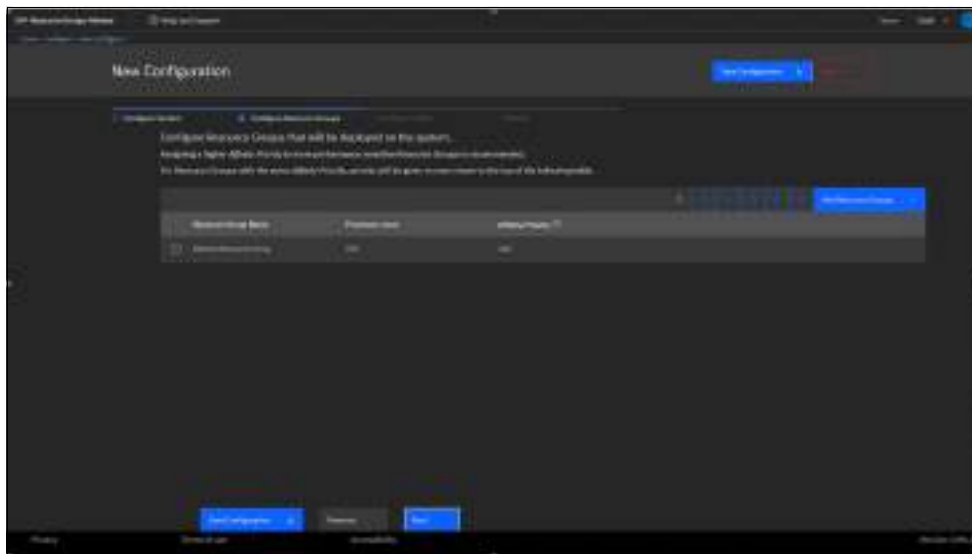


Figure 9-10 Resource Group definitions

After you define your Resource Groups, the system prompts you to define the LPAR details. Click **Add LPAR** to enter details for each LPAR, including whether it runs in Shared or Dedicated mode, the processor and memory allocations, and the Resource Group to which the LPAR is assigned. You can also add LPARs from the [IBM System Planning Tool](#).

When all your LPARs are defined, you see an LPAR summary window that is similar to Figure 9-11.

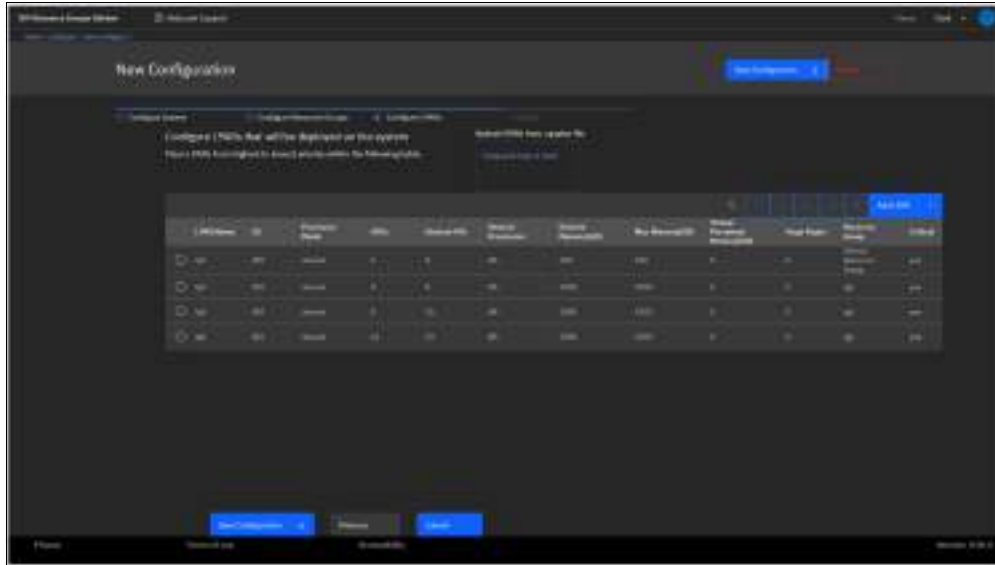


Figure 9-11 LPAR configuration

After you enter all the LPAR information, click **Submit**. The system displays the model output, similar to what is shown in Figure 9-12.




Figure 9-12 Output from the tool

Save the configuration or start a new model.

Note: RGA simulates how PowerVM creates Resource Groups and allocates resources to LPARs on a system, based on the provided input configurations. The system does not retrieve or assess any customer data from actual servers.

9.2 KVM support

Power11 processor-based servers can run KVM within a PowerVM LPAR. This configuration enables you to create and manage lightweight Linux VMs by using standard KVM tools while using the existing resources of the PowerVM LPAR. KVM on Power11 processor-based servers does not replace PowerVM. It is an extra capability that brings the power, speed, and flexibility of KVM virtualization technology to a PowerVM LPAR.

KVM-enabled LPARs can host PPC64-LE KVM guests, which are Linux VMs. The KVM guests within the LPAR use resources (CPU, memory, and I/O) that are allocated to the LPAR by the PHYP. This approach offers flexibility in deploying Linux workloads and can be more cost-effective than other virtualization solutions, especially for organizations that are already invested in the Linux ecosystem. KVM integration with the Linux kernel can deliver high performance, particularly when running Linux-based workloads.

This setup enables use cases such as running standard Linux distributions, containers, and other workloads that benefit from the KVM virtualization stack. It also can consolidate different types of workloads on one or more Power11 processor-based servers (see Figure 9-13 for an example). KVM on Power11 processor-based servers uses the strengths of both PowerVM and KVM to provide a powerful and flexible virtualization environment for Linux workloads.

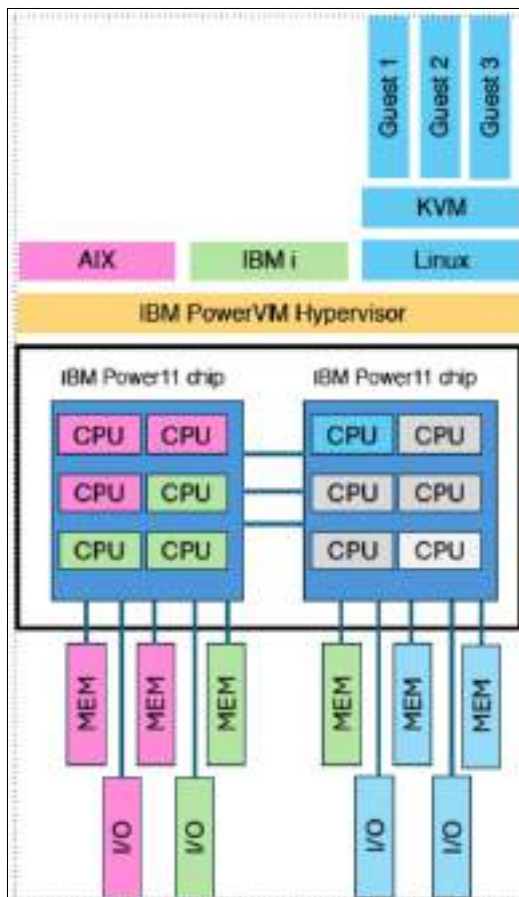


Figure 9-13 Power processor-based server stack for KVM

KVM in a PowerVM LPAR uses the industry-standard Linux KVM virtualization stack and integrates with existing Linux virtualization ecosystems (see Figure 9-14). KVM in an LPAR is enabled by the following components:

- ▶ IBM Power Architecture and Power11 implementation provide advanced virtualization capabilities to run multiple OS instances that share hardware resources while maintaining isolation.
- ▶ The Radix MMU architecture can independently manage page tables for the LPAR and the KVM guest instances on the LPAR.
- ▶ The PowerVM virtualization stack provides new functions to create and manage KVM guests, extending the Power Architecture to include new hypervisor interfaces.
- ▶ A Linux kernel with the KVM kernel module provides a core virtualization infrastructure to run multiple VMs in a Linux host LPAR. Upstream kernels and enabled downstream distributions such as Fedora and Ubuntu use the newly introduced Power Architecture extensions to create and manage KVM guests in the PowerLinux LPAR.
- ▶ QEMU is a user-space component that implements VMs on the host by using KVM functions.
- ▶ LibVirt provides a toolkit for managing VMs.

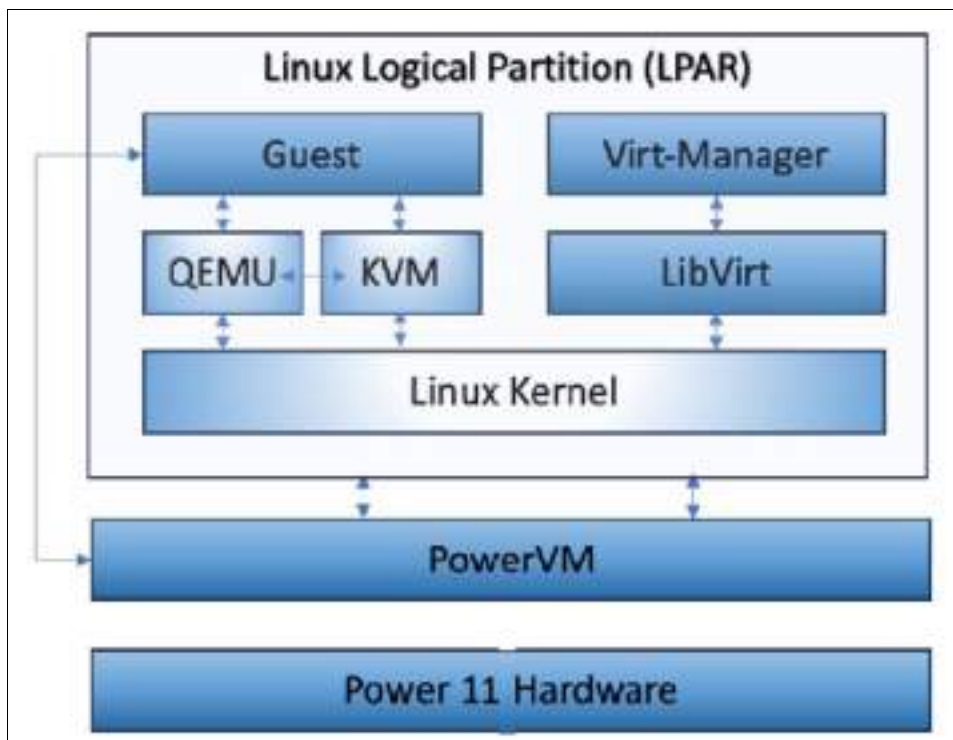


Figure 9-14 Industry-standard Linux virtualization stack

KVM support on Power11 processor-based servers requires the following items:

- ▶ The partition must be a Linux partition that runs in Power10 processor compatibility mode.
- ▶ The partition must be enabled for KVM:
 - For HMC-managed systems, set the partition to KVM Capable.
 - For unmanaged systems, set the default partition environment to Linux KVM on the Baseboard Management Controller (BMC).

- ▶ The partition must be running in Radix mode, which is the default MMU mode for Linux LPARs.
- ▶ The partition must be assigned to dedicated CPUs, with processor sharing set to Never Allow.

The following features are not supported on KVM LPARs:

- ▶ Shared processors
- ▶ vPMEM LUNs
- ▶ Platform keystore
- ▶ LPM
- ▶ Dynamic platform optimization
- ▶ Adding or removing memory, processor, and I/O DLPAR

PCI pass-through of LPAR-attached PCI devices to KVM guests is not supported. This capability is planned for future releases.

For more information, see [KVM in a PowerVM LPAR](#).



Hybrid cloud solutions

This chapter explores IBM Power servers as a foundational platform for hybrid cloud environments. It emphasizes their ability to support mission-critical workloads with exceptional performance, security, and flexibility. Power servers, which are built on the advanced IBM Power11 processor-based architecture, deliver significant improvements in compute efficiency and reduce operational costs. These systems integrate seamlessly across on-premises, private, and public cloud infrastructures. They support a wide range of enterprise operating systems (OSs), including AIX, IBM i, and Linux.

This chapter highlights the IBM strategic approach to hybrid cloud and demonstrates how Power servers enable scalable, cloud-native application development by using Kubernetes and Red Hat OpenShift. Key management tools such as IBM Power Virtualization Center (PowerVC), IBM Cloud Management Console (IBM CMC), and IBM Power Enterprise Pools (PEP) 2.0 play important roles in simplifying virtualization, optimizing resource allocation, and enhancing operational agility.

The chapter also introduces IBM Power Virtual Server (PowerVS), a cloud-based extension of Power servers that enables organizations to modernize at their own pace. PowerVS provides flexible, pay-as-you-go access to a virtualized Power infrastructure. It enables seamless workload migration, supports rapid provisioning, and delivers robust business continuity capabilities.

This chapter contains the following topics:

- ▶ IBM Power Private Cloud with Shared Utility Capacity
- ▶ Red Hat OpenShift
- ▶ IBM Power Virtual Server

10.1 IBM Power Private Cloud with Shared Utility Capacity

IBM Power Private Cloud with Shared Utility Capacity introduces a flexible and efficient approach to resource management through PEP 2.0. This innovation enables multiple Power servers to operate as a unified resource pool and share compute capacity across systems. Designed for clients deploying private cloud infrastructure, PEP 2.0 enhances operational agility by supporting dynamic, minute-by-minute consumption of compute resources.

In this model, each system within the pool is provisioned with a set number of base processor and memory activations. When configured as a PEP, these base activations and their associated OS entitlements are aggregated to define the pool baseline capacity. Concurrently, all hardware resources across the pool are fully activated and available for workload deployment, which enables usage to exceed the base capacity when needed. IBM CMC continuously monitors average resource usage across the pool and compares it to the base entitlement. Any usage beyond the base is billed as metered capacity, either through Prepaid Capacity Credits or monthly billing. PEP 2.0 eliminates the need for manual reallocation of activations between systems and offers granular resource sharing, improved cost efficiency, and simplified management for enterprise IT environments.

PEP with Shared Utility Capacity on IBM Power S1122 and IBM Power S1124 systems provides enhanced multi-system resource sharing and by-the-minute tracking and consumption of compute resources across a collection of Power S1122, Power S1124, IBM Power S1022, and IBM Power S1024 systems within a single PEP 2.0. Shared Utility Capacity offers complete flexibility to tailor initial system configurations with the right mix of purchased and pay-for-use consumption of processor, memory, and software. Clients with existing PEP 2.0 configurations of Power S1022 and S1024 systems can add one or more Power S1122 and Power S1124 systems to their pool and migrate workloads at their preferred rate and pace. Any Power S1122, Power S1124, Power S1022, and Power S1024 systems can interoperate and share compute resources within the same pool. Clients with Mobile Capacity on a Power S1122 or Power S1124 system can upgrade their system to support PEP 2.0 and use Shared Utility Capacity resources.

With IBM CMC 1.20, IBM i license entitlements are monitored and metered by each software tier based on processor usage and the tier of systems within a pool. Clients can deploy a mix of two-socket scale-out Power servers with processors requiring Power10, P20, and P30 tiers of IBM i license entitlement within the same pool without acquiring all Base Capacity IBM i license entitlements at the highest tier required for processors in the pool. Idle Base IBM i capacity of one software tier can offset minutes of IBM i consumption exceeding Base IBM i entitlement of an equal or smaller software tier (for example, idle P20 Base IBM i can offset usage exceeding P20 or Power10 Base IBM i).

A Power infrastructure that is consolidated onto Power S1122 and Power S1124 systems can greatly simplify system management so IT teams can focus on optimizing business results instead of moving resources within the data center. Shared Utility Capacity resources are tracked by virtual machine (VM) and monitored by IBM CMC, which integrates with local Hardware Management Consoles (HMCs) to manage the pool and track resource usage by system and VM, by the minute, across a pool.

Clients do not need to over-provision capacity on each individual system to support growth because all available processor and memory on all systems in a pool are activated and available for use. On Power S1122 and Power S1124 systems, Mobile and Shared Utility Capacity capabilities are offered through the purchase of a PEP Subscription (5765-P2E) for each Power S1122 or Power S1124 system that is configured for use in a PEP.

When a PEP 2.0 is started, each eligible Power S1122 and Power S1124 system's purchased processor activations, memory activations, and supported OS entitlement resources become Base Capacity resources as part of the pool and are aggregated across defined systems for consumption monitoring. Metered Capacity is the additional installed processor and memory resource above each system's Base Capacity. It is activated and available for immediate use when a pool is started, and then monitored by the minute by the IBM CMC.

Metered resource usage is charged only for minutes exceeding the pool's aggregate Base resources. Usage charges are debited in real time against a client's Capacity Credits (5819-CRD) on their account or may be billed monthly, in arrears, where available.

Important: Only two consecutive generations of Power servers are supported in the same Power Enterprise Pool. Customers who want to add a Power11 processor-based server to an existing pool with Power9 and Power10 processor-based servers should contact their IBM representative or IBM Expert Labs to discuss transition planning and best practices.

10.1.1 IBM Cloud Management Console

As private and hybrid cloud deployments continue to expand, enterprises require deeper management insights into these increasingly complex environments. Tools that deliver unified analytics and consolidated information are essential for helping ensure smooth and efficient infrastructure operations.

The IBM CMC provides a comprehensive view of the Power cloud environment, regardless of the number of systems or data centers that are involved. IBM CMC provides the following functions:

- ▶ Centralized inventory management of systems and virtual components
- ▶ Consolidated performance metrics to help optimize resource usage and system performance across all data centers
- ▶ Aggregated logging and analytics for enhanced operational insights

This unified approach enables IT teams to manage infrastructure more effectively and make informed decisions with confidence. The IBM CMC now supports Power11 processor-based servers, helping ensuring compatibility with the latest advancements in the IBM Power Architecture.

The IBM CMC also monitors and manages PEP2.0 configurations in the enterprise. The IBM CMC Enterprise Pools 2.0 application monitors base and metered capacity across a PEP 2.0. It provides summary and detailed drill-down views of real-time and historical resource consumption by the logical partition (LPAR).

The IBM CMC provides a cloud-based interface to monitor, manage, and optimize an IBM Power infrastructure. It supports AIX, IBM i, and Linux VM workloads, and is accessible through a secure, web-based dashboard. The IBM CMC is a cloud-based portal that enables clients to achieve the following goals:

- ▶ View the health and performance of Power servers.
- ▶ Monitor capacity and performance metrics.
- ▶ Gain insights into firmware levels and configurations.
- ▶ Receive proactive support and alerts.

For organizations with limited IT staff or seeking to streamline infrastructure oversight, the IBM CMC enables essential monitoring and management tasks without requiring complex, locally hosted solutions.

The IBM CMC integrates with the HMC and provides advanced insights into performance, capacity, firmware compliance, fix planning, and system health. With the scalability and built-in virtualization of Power11 processor-based systems, the IBM CMC helps administrators maintain control over increasingly complex hybrid IT environments, whether on-premises or connected to IBM Cloud or other supported clouds.

The IBM CMC is offered through a software as a service (SaaS) model and is licensed on a per-server subscription basis. Power11 processor-based servers are registered to the IBM CMC through their HMCs, which securely transmit telemetry and configuration data to the cloud portal. IBM provides both a no-charge base offering and optional chargeable add-ons that include advanced features such as predictive analytics, fix automation, and extended data retention for historical performance views. Customers can choose monthly or annual billing, and entitlement can be managed through IBM Passport Advantage® or directly through the IBM Cloud portal.

10.2 Red Hat OpenShift

With the growing demand for digital services and the continuous release of new offerings, organizations face the challenge of delivering exceptional customer experiences while maintaining resilient and highly available (HA) services. To meet these demands, they must innovate through developing new applications, modernizing existing ones, and helping ensure seamless scalability.

Organizations can achieve these goals by embracing cloud technologies, whether through full cloud migration, developing cloud-native applications, or adopting a hybrid cloud approach.

Red Hat OpenShift on IBM Power provides developers and IT operations (ITOps) teams with the flexibility and speed that they need to build, deploy, and manage applications across diverse environments, including on-premises, in the cloud, and across multiple infrastructures. This platform accelerates digital transformation by simplifying scalability and enhancing security.

10.2.1 Red Hat OpenShift on IBM Power

Red Hat OpenShift Container Platform is fully supported on Power servers, offering a robust and scalable environment for running cloud-native applications and modernizing traditional workloads on the Power Architecture (ppc64le). This integration enhances the Power software ecosystem by enabling enterprise-grade container orchestration for building, deploying, and managing containerized applications with Red Hat OpenShift.

Red Hat OpenShift on IBM Power provides a flexible and efficient platform for hybrid IT environments, enabling the seamless coexistence of cloud-native applications with traditional VM-based workloads. By using the same infrastructure, it enables colocation of containerized workloads alongside existing applications running on AIX, IBM i, or Linux, helping ensure low-latency communication between applications and data. Power servers provide a resilient and secure platform, making them an excellent choice for cloud implementations that use Red Hat OpenShift.

This section describes the benefits of using a multi-architecture cluster, along with implementation guidelines and best practices for deploying a multi-architecture cluster that uses IBM Power control plane nodes and x86 or AMD-based worker nodes.

Red Hat OpenShift is a leading enterprise Kubernetes platform that provides a robust foundation for developing, deploying, and scaling cloud-native applications. It extends Kubernetes with extra features and tools to enhance productivity and security, making it an ideal choice for businesses seeking to use container technology at scale.

Red Hat OpenShift is a unified platform for building, modernizing, and deploying applications at scale. It offers a complete set of services for bringing applications to market on the infrastructure of choice. Red Hat OpenShift delivers a consistent experience across public cloud, on-premises, hybrid cloud, and edge architectures.

Red Hat OpenShift provides a unified, flexible platform to address various business needs, ranging from an enterprise-ready Kubernetes orchestrator to a comprehensive cloud-native application development platform that can be self-managed or delivered as a fully managed cloud service.

Red Hat OpenShift clusters

A Red Hat OpenShift Container Platform cluster consists of multiple nodes that can run on either physical machines or VMs. At a minimum, the following node requirements apply:

- ▶ Three control plane nodes manage and maintain the cluster.
- ▶ Two worker nodes run containerized workloads.
- ▶ A temporary bootstrap node is used during installation to host configuration and installation files.

In Red Hat OpenShift Container Platform 4.18, the bootstrap and control plane nodes must run Red Hat Enterprise Linux CoreOS (RHCOS), which is a minimal, immutable container host OS that is derived from Red Hat Enterprise Linux (RHEL). Compute nodes can run either RHCOS or standard RHEL, depending on the deployment architecture. RHEL machines are deprecated in Red Hat OpenShift Container Platform 4.16 and will be removed in a future release.

Multiple-Architecture Compute clusters

Enterprises expanding their operations often deploy applications in heterogeneous environments that use different types of hardware. A prime example includes data centers that use x86_64 servers and Power servers, but edge locations might feature ARM-based devices because of their power efficiency. The Power Architecture is an ideal choice for running Red Hat OpenShift clusters. However, some applications within the environment might not be compatible with Power Architecture nodes. Organizations can retain the benefits of the Power Architecture while maintaining simpler cluster management and efficient resource usage by incorporating both Power Architecture and x86_64 architecture nodes in the same cluster.

The release of Red Hat OpenShift 4.14 introduced the Red Hat OpenShift Container Platform Multiple-Architecture Compute feature to the Power Architecture. Multiple-Architecture Compute provides a single heterogeneous cluster that enables fit-for-purpose computing so that clients can align tasks and applications to CPU strengths and software availability rather than a single architecture. This support was expanded in Red Hat OpenShift 4.15, which enabled a Red Hat OpenShift cluster to support a Power Architecture control plane and add x86 architecture worker nodes.

With Multiple-Architecture Compute for Red Hat OpenShift Container Platform, organizations can use a pair of compute architectures, such as ppc64le and amd64, within a single cluster. This feature opens new possibilities for versatility and optimization in composite solutions that span multiple architectures.

Multiple-Architecture Compute offers several benefits:

- ▶ Platform independence

Multiple-Architecture Compute enables applications to function across various hardware platforms, including Intel servers in data centers, ARM-based Raspberry Pis in remote locations, and Power servers in corporate settings.

- ▶ Reduced complexity through standardized application deployment across architectures

Streamlines operations and eliminates the need to maintain separate stacks for different hardware.

- ▶ Cost efficiency

Different hardware architectures provide varying cost-to-performance ratios, which organizations can use to minimize overall infrastructure expenses.

- ▶ Optimal resource usage

Organizations can select the most cost-effective architecture for each specific workload. For example, ARM servers might be more affordable for lightweight services, but x86_64 or Power servers might excel at handling heavy computational tasks.

- ▶ Energy efficiency

Specific architectures, such as ARM, are known for reduced power consumption, which can substantially decrease energy costs, particularly in scale-out scenarios like IoT and mobile services.

- ▶ Scalability and flexibility

Implementing applications on multiple architectures enhances scalability and operational agility, which is vital for businesses experiencing fluctuating loads.

Multiple-Architecture Compute support allows companies to achieve the following goals:

- ▶ Scale elastically across platforms.

Dynamically allocate resources among different types of hardware to handle surges in demand without being restricted to a single architecture.

- ▶ Avoid vendor lock-in.

Reduce dependency on a single supplier or type of hardware, increasing bargaining power during procurement decisions.

- ▶ Optimize performance.

Align application requirements with the strengths of each architecture to maximize performance.

- ▶ Create tailored solutions.

Certain applications might benefit from the high I/O throughput of Power servers, and others might perform optimally on the high-throughput, multi-core processors of x86_64 architectures.

- ▶ Meet specialized computing needs.

Certain tasks might require specialized hardware, such as GPUs for machine learning workflows, which might be more accessible or better supported on a particular architecture.

Single-node clusters

For a simpler and more accessible deployment experience, Red Hat OpenShift now supports a single-node cluster in which all management and worker functions are installed on a single node. This approach is especially suitable for development, testing, or edge computing environments.

It is important to recognize the inherent limitations of a single-node configuration: most notably, the absence of built-in high availability (HA). If the node fails, the entire cluster becomes unavailable, potentially resulting in downtime and data loss. The robust HA features of the Power Architecture, when properly designed with redundant power, networking, and storage, can significantly mitigate these risks.

10.2.2 Red Hat OpenShift AI on Power servers

Red Hat OpenShift AI brings enterprise-grade capabilities for artificial intelligence (AI) and machine learning directly into the Red Hat OpenShift platform, enabling organizations to build, train, deploy, and monitor AI models within a unified, containerized environment. Red Hat OpenShift AI is built on Red Hat OpenShift, an enterprise-grade Kubernetes platform that is designed to support AI and machine learning workloads across hybrid and multi-cloud infrastructures. By integrating AI workflows into the same infrastructure that is used for traditional applications, Red Hat OpenShift AI simplifies operations, accelerates development cycles, and helps ensure consistent governance and security across the entire AI lifecycle. Red Hat OpenShift AI enables organizations to build, deploy, and manage AI-enabled applications at scale across hybrid cloud environments.

Red Hat OpenShift AI Self-Managed is available by installing the Red Hat OpenShift AI Operator and configuring it to manage the stand-alone components of the platform. Red Hat OpenShift AI Self-Managed is supported on Red Hat OpenShift Container Platform across multiple architectures, including ppc64le (Power Architecture), s390x (IBM Z), and x86_64. There is support for the following infrastructure providers:

- ▶ IBM Power (technology preview)
- ▶ IBM Z (technology preview)
- ▶ IBM Cloud
- ▶ Red Hat OpenStack
- ▶ Bare metal
- ▶ Hosted control planes on bare metal
- ▶ Amazon Web Services (AWS)
- ▶ Google Cloud Platform
- ▶ Microsoft Azure
- ▶ VMware vSphere
- ▶ Oracle Cloud

Power servers, renowned for their high performance, scalability, and reliability, are suitable for compute-intensive AI workloads. The integration of Red Hat OpenShift AI with Power servers enables enterprises to use the capabilities of both platforms for robust, production-grade AI deployments.

As of Red Hat OpenShift AI Self-Managed 2.20, the Power Architecture (ppc64le) is available as a technology preview. Currently, model deployments on the Power Architecture are supported only in standard mode. Red Hat OpenShift AI provides an integrated platform for developing, training, serving, and monitoring AI and machine learning models. Support for vLLM is also available on the Power Architecture as a technology preview, with vLLM runtime templates accessible for experimentation and development.

When deployed on Power servers, Red Hat OpenShift AI benefits from the platform's high-performance architecture, which is optimized for data-intensive and compute-heavy workloads. The latest POWER processors feature built-in AI inferencing capabilities, delivering faster insights and reduced latency for real-time applications. This capability will be further enhanced with the planned availability and support of the IBM Spyre Accelerator in Power servers. The IBM Spyre Accelerator is a Peripheral Component Interconnect Express (PCIe)-attached AI card that is designed for low-precision AI arithmetic and enterprise-grade AI algorithms, making it ideal for deploying large-scale models and generative AI frameworks.

Together, Red Hat OpenShift AI and the Power Architecture provide a scalable, production-ready foundation for modern AI initiatives. Whether running models built with PyTorch, TensorFlow, or vLLM, this integrated stack supports hybrid cloud deployments and enables enterprises to implement AI with confidence, helping ensure performance, efficiency, and flexibility across diverse workloads and environments.

10.3 IBM Power Private Cloud Rack Solution

The IBM Power Private Cloud Rack Solution is a preintegrated, turnkey infrastructure offering that is designed to simplify the deployment and management of private cloud environments on Power servers. This solution is built on Power11 processor-based technology and combines compute, storage, networking, and management components into a single rack-based system that is optimized for enterprise workloads, including Db2 Warehouse, SAP HANA, and other data-intensive applications.

10.3.1 Key features and components

- ▶ Preconfigured architecture
 - Includes Power11 processor-based compute nodes, IBM FlashSystem storage, management and fabric switches, and power distribution units (PDUs), all preintegrated and validated for performance and reliability.
- ▶ Cloud-ready software stack
 - Bundled with IBM PowerVC, Red Hat OpenShift, and optional IBM Cloud Pak solutions that support containerized and virtualized workloads.
- ▶ HA and security
 - Designed with built-in redundancy, Secure Boot, and compliance-ready configurations that meet enterprise-grade uptime and data protection requirements.
- ▶ Simplified management
 - Centralized tools enable administrators to monitor and manage multiple Power environments from a single dashboard, with support for automation and self-service provisioning.

This solution is ideal for organizations seeking to modernize their infrastructure with a cloud-like experience on-premises while maintaining control over data locality, security, and performance.

10.3.2 Available configurations

The two available configurations are described in this section.

IBM Power Private Cloud Rack Solution

The Power Private Cloud Rack Solution is offered in a preconfigured setup, providing an optimized full stack for a production-level environment. It uses the unique virtualization technologies of PowerVM to host the entire software stack on just three servers. More nodes can be added to the initial configuration as needed.

A minimum configuration includes the following items:

- ▶ Hardware stack
 - Three Power S1122 servers
 - One FlashSystem 5200 storage enclosure with a minimum of 9.6 TB
 - Two SAN24B-6 switches with 24 Fibre Channel (FC) ports and industry-leading Gen 6 FC technology
 - An optional IBM Ethernet switch with high-performance Gigabit Ethernet Layer 2 and Layer 3 switching, featuring 52 ports
 - One IBM Enterprise slim rack with 42 EIA units of vertical mounting space and a 19-inch rack enclosure
- ▶ Software stack
 - RHEL 8 for Power11 processor-based servers
 - IBM PowerVM Enterprise Edition
 - IBM PowerVC for Private Cloud 2.0
 - Red Hat OpenShift Container Platform
 - IBM Spectrum® Scale Data Access Edition or IBM Spectrum Scale Data Management Edition

If an organization has Red Hat entitlements for Red Hat OpenShift Container Platform or RHEL 8, these components can be deselected from the solution edition in e-config. Proof of entitlement for each software license is required to authorize the manufacturing and shipping of the solution.

IBM Power Private Cloud Starter Solution

The Power Private Cloud Starter Solution includes the following items:

- ▶ Hardware stack
 - At least one Power S1122 server node
 - An optional IBM FlashSystem 5200 storage enclosure
- ▶ Software stack
 - RHEL 8 for Power9 processor-based servers
 - PowerVM Enterprise Edition
 - Optional PowerVC for Private Cloud 2.0
 - Red Hat OpenShift Container Platform

If an organization has Red Hat entitlements for Red Hat OpenShift Container Platform or RHEL 8, these components can be deselected from the solution edition in e-config. Proof of entitlement for each software license is required to authorize the manufacturing and shipping of the solution.

10.3.3 IBM Cloud Pak

IBM Cloud Pak solutions are a suite of modular, containerized software offerings that are designed to accelerate digital transformation. Built on Red Hat OpenShift, they enable organizations to modernize applications, automate business operations, manage data, and integrate AI across hybrid cloud environments with consistency and scalability.

Each IBM Cloud Pak solution includes a combination of IBM middleware, open-source technologies, Kubernetes operators, and enterprise-grade security, providing a robust foundation for innovation and agility in the cloud. IBM Cloud Pak solutions take a bundled approach that enables organizations to accelerate their modernization journey by packaging all the required components to get started.

There are three main benefits of IBM Cloud Pak solutions:

- ▶ They provide comprehensive capabilities.
- ▶ They are supported by IBM.
- ▶ They run in any environment where Red Hat OpenShift is supported.

IBM Cloud Pak on IBM Power

Optimized for deployment on Power servers, IBM Cloud Pak solutions deliver high performance and efficiency for containerized workloads. They support a wide range of use cases, from data and automation to AI and security, while helping ensure seamless integration across on-premises and cloud environments.

The Power platform enhances the value of IBM Cloud Pak solutions by offering superior performance, scalability, and cost efficiency for modern workloads. Combined with Red Hat's open-source ecosystem, the Power platform empowers businesses to modernize applications and build new cloud-native solutions on a unified, enterprise-ready platform.

Several IBM Cloud Pak solutions are supported on IBM Power, each focused on a specific domain:

- ▶ IBM Cloud Pak for Applications
- ▶ IBM Cloud Pak for Data
- ▶ IBM Cloud Pak for Integration
- ▶ IBM Cloud Pak for AIOps

IBM Cloud Pak for Applications

IBM Cloud Pak for Applications is an enterprise-ready, containerized software solution that is designed to modernize applications and develop new cloud-native applications. Built on IBM WebSphere® offerings and Red Hat OpenShift Container Platform, it provides a comprehensive set of tools to help organizations move among public, private, and hybrid clouds.

IBM Cloud Pak for Applications includes IBM Cloud Transformation Advisor, which is an AI-powered tool that helps re-factor and redesign applications. The solution provides automated vulnerability assessment and identification, helping ensure continuous security compliance across all deployment environments. It also automates audit reporting, which simplifies compliance management.

Developers can use their preferred integrated development environments (IDEs) to build and deploy applications, with support for modern run times and DevOps workflows. This integration streamlines the development process and enhances productivity.

IBM Cloud Pak for Data

IBM Cloud Pak for Data enables organizations to unify and simplify data collection, organization, and analysis. It is designed for AI and analytics workloads. IBM Cloud Pak for Data is a unified, preintegrated data and AI platform that helps organizations collect, organize, analyze, and infuse AI into their data. Running natively on the Red Hat OpenShift Container Platform, it supports deployment across various cloud environments, including IBM Cloud, AWS, and Microsoft Azure.

The platform enables secure access to data at its source, eliminates the need for data migration, and reduces data silos, helping ensure seamless data integration. It creates a trusted, business-ready analytics foundation, simplifying data preparation, policy enforcement, security, and compliance while automating data governance and the AI lifecycle. IBM Cloud Pak for Data provides tools for building, deploying, and managing AI and machine learning models, scaling these capabilities consistently across the organization to enable comprehensive data analysis and insights.

By using AI throughout the business with trust and transparency, the platform supports the end-to-end AI workflow, helping ensure effective integration of AI into business processes. Offering a single interface for end-to-end analytics with built-in governance, IBM Cloud Pak for Data simplifies the management of data and AI capabilities. Its scalable Kubernetes environment enables organizations to grow their data and AI capabilities as needed. Supporting multi-cloud deployments, it provides agility and avoids vendor lock-in, making it a powerful tool for accelerating the journey to AI and unlocking the value of data in AI-driven digital transformation.

IBM Cloud Pak for Integration

IBM Cloud Pak for Integration is a comprehensive, AI-powered hybrid integration platform that connects applications, data, systems, and services across any environment. It provides a unified experience with a suite of integration tools that streamline the creation, management, and deployment of integration solutions. Running on Red Hat OpenShift, IBM Cloud Pak for Integration supports both cloud and on-premises deployments, helping ensure scalability and security.

The platform includes components such as IBM API Connect for managing APIs, IBM App Connect for no-code integration, and IBM Event Streams for real-time data processing. By using AI and automation, IBM Cloud Pak for Integration accelerates integration processes, reduces manual workflows, and enhances responsiveness to real-time events.

These capabilities make IBM Cloud Pak for Integration a suitable solution for organizations seeking to modernize integration capabilities and drive digital transformation.

IBM Cloud Pak for Business Automation

IBM Cloud Pak for Business Automation is a modular set of integrated software components that automates work and accelerates business growth. Built for any hybrid cloud, it simplifies complex workflows, facilitates records management, and enhances overall productivity. The platform uses AI to identify gaps and build low-code and no-code automations, streamlining operations. Running on Red Hat OpenShift, IBM Cloud Pak for Business Automation supports containerized deployments across various cloud environments, providing flexibility and scalability.

Key features include automation of case and process workflows, conversion of unstructured content into valuable data, and use of software robots to complete tasks based on AI insights. This comprehensive automation solution enables organizations to improve efficiency, reduce operational costs, and drive continuous process improvements.

IBM Cloud Pak for AIOps

IBM Cloud Pak for AIOps is an advanced, AI-powered platform that enhances ITOps by using AI and machine learning. It integrates seamlessly with existing ITOps toolchains to provide comprehensive visibility, proactive incident management, and automated remediation. By analyzing data from various sources, such as logs, metrics, and events, IBM Cloud Pak for AIOps enables IT teams to predict and resolve issues before they impact business operations.

The platform supports hybrid cloud environments, enabling organizations to manage their IT infrastructure across on-premises, cloud, and containerized environments. Key features include event correlation and compression, anomaly detection, root cause analysis, and automated runbooks, all designed to reduce mean time to resolution (MTTR) and improve overall operational efficiency.

With collaborative tools and real-time insights, IBM Cloud Pak for AIOps enables IT teams to innovate faster, reduce operational costs, and ensure the reliability of mission-critical workloads.

For more information about IBM Cloud Pak solutions, see [IBM Documentation](#).

10.3.4 Other cloud enablement solutions

From a Red Hat software perspective, a comprehensive set of software solutions is available to accelerate modernization efforts, including Red Hat Runtimes, Red Hat 3scale API Management, Red Hat Fuse, and Red Hat AMQ.

10.4 IBM Power Virtual Server

PowerVS is a cloud-based infrastructure as a service (IaaS) offering that enables businesses to run Power server workloads in a flexible, scalable, and secure virtual environment. Built on the same architecture as Power servers that are used on-premises, it enables seamless hybrid cloud integration, making it suitable for enterprises seeking to modernize their IT infrastructure without abandoning existing investments. With support for AIX, IBM i, and Linux OSs, PowerVS provides a versatile platform for mission-critical applications. It provides a platform for moving AIX, IBM i, and Linux on Power workloads to the cloud without requiring a time-consuming and risky migration to a different platform.

One of the key advantages of PowerVS is its ability to deliver high performance and reliability for enterprise workloads. It uses POWER processors, which are optimized for data-intensive tasks and high-throughput computing. The platform also offers features such as dynamic resource allocation, automated scaling, and integrated backup and disaster recovery (DR) options. These capabilities help organizations maintain business continuity and meet demanding service-level agreements (SLAs).

PowerVS integrates with a wide range of IBM Cloud services and third-party tools, enabling organizations to build and manage hybrid cloud environments efficiently. It supports DevOps practices, containerization with Red Hat OpenShift, and AI and machine learning workloads, making it suitable for both traditional enterprise applications and modern cloud-native development. With a pay-as-you-go pricing model and global availability, PowerVS offers a cost-effective and agile solution for businesses seeking to innovate and grow in a digital-first world.

PowerVS enables organizations to run IBM i, AIX, and Linux workloads in a cloud environment, providing fast, self-service provisioning and flexible compute, memory, and storage resources. PowerVS is designed to help organizations modernize their infrastructure, migrate to hybrid cloud models, and optimize their data center resources.

Here are the key features and benefits of PowerVS:

- ▶ Flexible and scalable
Organizations can adjust compute, memory, and storage resources on demand, scaling up or down as needed.
- ▶ Pay-as-you-use
Billing is based on consumption, enabling organizations to manage costs effectively.
- ▶ Hybrid cloud integration
PowerVS facilitates seamless integration of AIX and IBM i workloads with cloud-native applications and services.
- ▶ Performance and security
PowerVS uses the performance and security features of Power servers while offering the agility of cloud computing.
- ▶ Managed infrastructure
PowerVS offers a managed infrastructure in which IBM handles hardware maintenance, enabling organizations to focus on workloads from a software and application standpoint.
- ▶ DR
PowerVS can be used to implement DR solutions for critical applications.
- ▶ Simplified cloud adoption
PowerVS streamlines the cloud adoption process for IBM i, AIX, and other Power server workloads.

10.4.1 IBM Power Virtual Server options

PowerVS is a single offering available in two variations: off-premises, in which the infrastructure components are in IBM data centers, and on-premises, in which the infrastructure components are in the client's data center. The on-premises variation is referred to as IBM Power Virtual Server Private Cloud. Both variations provide a cloud-based consumption model in which organizations pay for resources as they are consumed, and both are managed with the same management interfaces.

Figure 10-1 shows the two different implementations of PowerVS.

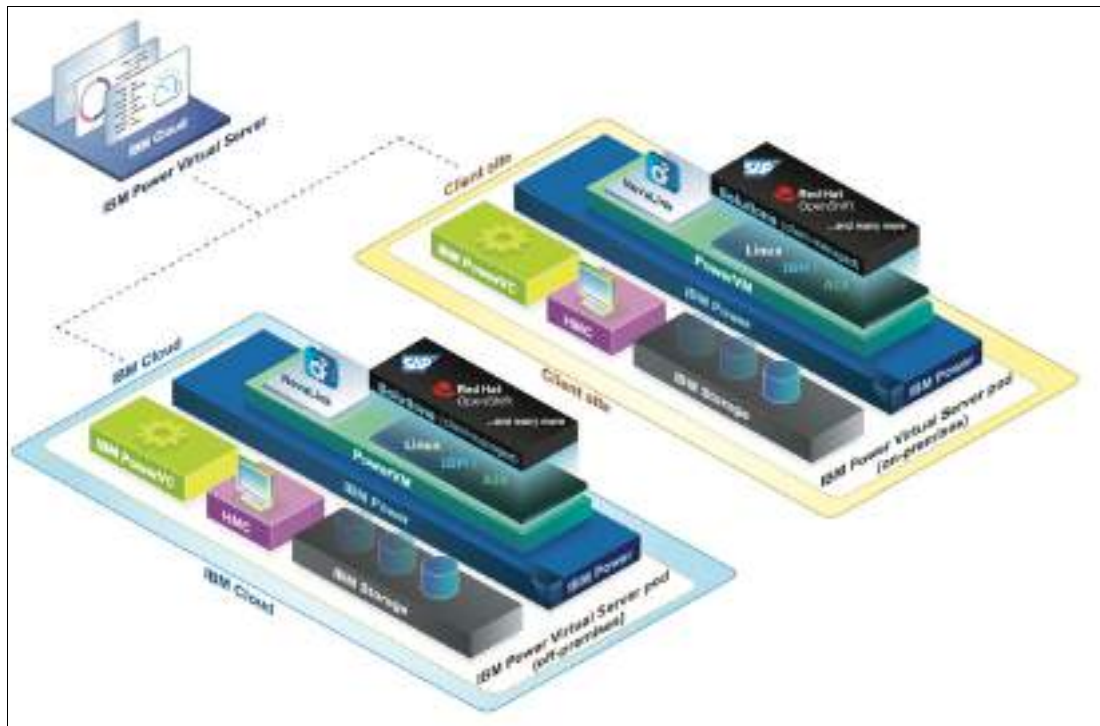


Figure 10-1 IBM Power Virtual Server options

For more information about PowerVS Private Cloud, see 10.4.3, “IBM Power Virtual Server Private Cloud” on page 240.

Use cases for IBM Power Virtual Server

PowerVS supports a wide range of use cases, particularly for enterprises that rely on Power servers to run mission-critical workloads. Here are some of the most common and impactful scenarios:

- ▶ Hybrid cloud modernization

Organizations with on-premises Power servers can extend their infrastructure to the cloud by using PowerVS. This approach enables a hybrid cloud model in which workloads can be moved or replicated to the cloud for scalability, DR, or testing, without requiring replatforming or application rewrites. It is especially useful for businesses running AIX or IBM i environments that plan to modernize gradually.

- ▶ High availability and disaster recovery (HADR)

PowerVS provides a reliable platform for DR by enabling businesses to replicate on-premises workloads to the cloud. If there is a failure, workloads can be restored quickly in the cloud, minimizing downtime. It also supports HA configurations, help ensure business continuity for critical applications.

- ▶ Development and testing environments

Developers can use PowerVS to provision isolated environments quickly for application development, testing, and quality assurance. This capability is valuable for teams working on AIX, IBM i, or Linux-based applications because it eliminates the need for dedicated on-premises hardware and enables faster iteration and deployment cycles.

- ▶ SAP HANA and enterprise resource planning (ERP) workloads

PowerVS is certified to run SAP HANA and other SAP ERP applications. Enterprises can migrate or extend SAP environments to the cloud for improved scalability, performance, and cost efficiency, especially when paired with IBM high-performance Power11 processors.

- ▶ AI and machine learning

PowerVS supports AI and machine learning workloads, particularly ones that require high compute power and memory bandwidth. It integrates with IBM Watson® and Red Hat OpenShift, enabling data scientists to build, train, and deploy models by using familiar tools while using the performance of Power servers.

- ▶ Database hosting

Many organizations use PowerVS to host enterprise-grade databases like IBM Db2, Oracle, or PostgreSQL. The platform offers high I/O throughput and reliability, making it ideal for transaction-heavy applications such as banking, retail, and logistics.

Many organizations use PowerVS to host enterprise-grade databases such as IBM Db2, Oracle, or PostgreSQL. The platform provides high I/O throughput and reliability, making it suitable for transaction-heavy applications in industries such as banking, retail, and logistics.

- ▶ Application hosting

For businesses with applications running on AIX or IBM i, PowerVS offers a cloud-based alternative to on-premises hardware that enables companies to extend the life of critical applications while gradually modernizing infrastructure.

- ▶ Compliance and data sovereignty

Organizations in regulated industries, such as healthcare, finance, and government, use PowerVS to meet strict compliance and data residency requirements. IBM provides regional data centers and compliance certifications to meet these requirements.

10.4.2 IBM Power Virtual Server in the cloud

In this variation, PowerVS resources are in IBM data centers with dedicated networking and storage area network (SAN)–attached storage. Organizations can choose from 21 available data centers, selecting the location that is closest to their users. IBM Power clients that rely on private cloud infrastructure can quickly and economically extend Power IT resources to the cloud.

In these IBM data centers, PowerVS resources are separated from other IBM Cloud servers by using dedicated networks and direct-attached storage. This offering provides flexibility to choose from different Power server models, an IBM-supplied OS image, and various storage tiers based on specific workload requirements. Organizations can also choose their own image under a Bring Your Own License model, enabling them to import preconfigured images from enterprise infrastructure into PowerVS.

This offering also supports distinctive features such as shared processor pools (SPPs), public connectivity for VMs, placement groups, and a global replication service, which can be consumed based on organizational requirements.

For more information about IBM PowerVS, see [Getting Started with IBM Power Virtual Server](#).

10.4.3 IBM Power Virtual Server Private Cloud

Some organizations require a cloud experience that runs on Power servers while maintaining control of data and applications, whether for regulatory, security, or performance reasons. For these organizations, IBM designed an offering that brings the PowerVS service to a data center of the organization's choice.

PowerVS Private Cloud extends all the benefits of PowerVS into an organization's data center or a partner's facility. The enhanced capabilities of PowerVS Private Cloud provide managed IaaS at client locations, with metered consumption and no upfront costs, supporting Hybrid by Design delivery of services.

The PowerVS Private Cloud offering offers the following features:

- ▶ Maintain customer data and workloads on your own site.
Enterprises might have workloads or data that are regulated and cannot be hosted off-premises. Sometimes, enterprises have workloads that are sensitive or have ultra-short latency requirements, which are better served onsite and near other onsite workloads.
- ▶ Maintain customer data in region and specific geographies in the location of their choice.
Country sovereignty regulations require some data and workloads to remain within national borders. According to a recent IBM Institute of Business Value study, 61% of cloud leaders cite security or compliance as reasons for moving certain workloads from public clouds to private clouds or on-premises data centers.¹
- ▶ Provide a seamless hybrid cloud experience.
Enterprises can foster a unified hybrid cloud landscape by integrating PowerVS running at both an IBM site and a client site, with the ability to manage all VMs and infrastructure through a unified user interface. Organizations gain flexibility by using an as-a-service model with intentional workload placement on-premises and off-premises.
- ▶ Deliver a predictable charging model with committed monthly expenditures that is combined with flexible consumption with metered usage-based pricing.
Both the PowerVS off-premises at an IBM location and on-premises at a client site offerings include compute, memory, storage, and OS licenses that are fully metered by the hour, enabling organizations to pay for the resources that they use each month with no upfront payment.
- ▶ Streamline ITOps.
Whether in the cloud or at an enterprise site, IBM manages the infrastructure, enabling enterprises to focus on business outcomes rather than on infrastructure management. IBM owns, delivers, and sets up PowerVS in the data center of the organization's choice, providing a fully managed solution that includes monitoring, security, firmware updates, and infrastructure management.
- ▶ Provide enhanced security and control of data.
PowerVS provides comprehensive security for a Power platform by integrating with IBM Cloud tools to manage security. This integration alleviates the need to manage a Power platform security manually, with the added benefit of maintaining sensitive data and workloads on-premises.

¹ Source:
<https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/data-story-hybrid-cloud-a>
i

The physical infrastructure is delivered as a point of delivery (pod) that is deployed in the client's data center. A pod is the physical component that is within the client data center and contains the compute, storage, and network components. A pod contains one or more racks where each of the components is installed. The racks are interconnected to provide a self-contained infrastructure, including customer-usable components, spare components, and management components.

IBM site reliability engineers maintain the pod, which is managed through the IBM Cloud platform. Each pod is associated with an IBM Cloud Satellite® location that is owned by the client's IBM Cloud account. This architecture enables organizations to scale private cloud infrastructure horizontally by adding more pods to meet workload requirements.

For more information about PowerVS Private Cloud, see *Introduction to IBM Power Virtual Server Private Cloud*, REDP-5745.

10.4.4 Introducing Power11 processor-based servers in the cloud

Power11 processor-based servers will be available on the day of announcement in PowerVS and PowerVS Private Cloud. With the launch of Power11 processor-based servers, organizations and independent software vendors (ISVs) gain immediate access to the latest Power hardware in the cloud, enabling faster innovation and greater flexibility. Whether modernizing applications or enhancing business continuity, Power11 processor-based servers in PowerVS deliver a seamless hybrid cloud experience.

Here are the key benefits of Power11 processor-based servers in the cloud:

- ▶ **Rapid deployment**
Start Power11 processor-based virtual servers in the cloud in under 10 minutes.
- ▶ **Accelerated modernization**
Enable agile development and testing on the latest Power platform.
- ▶ **Enhanced business resiliency**
Strengthen continuity strategies with cloud-based Power11 processor-based resources.
- ▶ **Data sovereignty compliance**
Meet regulatory and sensitive data requirements with on-premises options.
- ▶ **Consistent hybrid cloud experience**
Provide unified operations across on-premises and cloud environments.

Figure 10-2 shows the initial implementation of Power11 processor-based servers in PowerVS on announcement day.



Figure 10-2 Availability of Power11 processor-based servers in IBM Power Virtual Server

Power11 processor-based technology in PowerVS enables organizations to use new functions immediately without requiring initial capital investment. Organizations can implement workloads that use the trusted security capabilities of Power11 processor-based servers:

- ▶ Quantum-safe infrastructure compliance, with optional workload cryptographic inventory discovery
- ▶ A unified cyber resiliency solution that responds to evolving cyberthreats and regulatory standards, with optional third-party services

In addition, organizations can use the improved performance of Power11 processor-based servers, as shown in Figure 10-3.



Figure 10-3 Potential performance benefits from Power11 processor-based servers in PowerVS

Abbreviations and acronyms

ACF	Access Control File	eSCM	entry Single-Chip Module
AER	Autonomous Error Resolution	ESM	Enclosure Services Manager
AES	Advanced Encryption Standard	ESS	Entitled Systems Support
AI	artificial intelligence	EUS	Extended Update Support
AME	Active Memory Expansion	FC	Fibre Channel
AML	anti-money laundering	FHE	Fully Homomorphic Encryption
AMM	Active Memory Mirroring	FICON	Fibre Channel connection
AOC	active optical cable	FLRT	Fix Level Recommendation Tool
API	application programming interface	FOM	fan-out module
ASMI	Advanced System Management Interface	FSP	Flexible Service Processor
AWS	Amazon Web Services	GLVM	Geographic Logical Volume Manager
BMC	Baseboard Management Controller	HA	high availability
BSC	blind swap cassette	HADR	high availability and disaster recovery
CBU	Capacity BackUp	HCL	HashiCorp Configuration Language
CCA	Common Cryptographic Architecture	HDDW	Huge Dynamic DMA Window
CoD	Capacity on Demand	HIS	healthcare information system
CUoD	Capacity Upgrade on Demand	HMC	Hardware Management Console
CVE	Common Vulnerabilities and Exposures	HPC	high-performance computing
DBM	Dedicated Backup Machine	HPT	hardware page table
DCM	Dual-Chip Module	HSM	hardware security module
DDIMM	differential DIMM	IaaS	infrastructure as a service
DDW	Dynamic DMA Window	IaC	infrastructure as code
DLPAR	dynamic logical partition	IASP	independent auxiliary storage pool
DMA	direct memory access	IBM	International Business Machines Corporation
DME	dense math engine	IBM CMC	IBM Cloud Management Console
DR	disaster recovery	IPMI	Intelligent Platform Management Interface
DRO	Dynamic Resource Optimization	IPsec	IP Security
DWPD	Drive Writes Per Day	ISA	Instruction Set Architecture
eBMC	enterprise Baseboard Management Controller	ISC	Integrated Stack Capacitor
ECC	Elliptic Curve Cryptography	iSCSI	internet Small Computer Systems Interface
EEH	enhanced error handling	ISV	independent software vendor
EHR	electronic health record	ITIC	Information Technology Intelligence Consulting
EMR	electronic medical record	ITOps	IT operations
EOL	end-of-life	JBOF	Just a Bunch of Flash
EOS	end of service	KVM	Kernel-based Virtual Machine
ERP	enterprise resource planning	KYC	know-your-customer
ESA	Electronic Service Agent		

LAN	local area network	RHSA	Red Hat Security Advisory
LDAP	lightweight directory access protocol	RMC	Resource Monitoring and Control
LIC	Licensed Internal Code	SaaS	software as a service
LLM	large language model	SAN	storage area network
LPAR	logical partition	SCM	Single-Chip Module
LPM	Live Partition Mobility	SE2	Standard Edition 2
LPP	licensed program product	SEA	Shared Ethernet Adapter
LWE	Learning With Errors	SFP	small form-factor pluggable
MES	Miscellaneous Equipment Specification	SIMD	single instruction, multiple data
MMA	Matrix Math Accelerator	SLA	service-level agreement
MPIO	multipath I/O	SMT	simultaneous multithreading
MSS	Merkle Signature Scheme	SP	Service Pack
MTTR	mean time to resolution	SPP	shared processor pool
NPIV	N_Port ID Virtualization	SPT	System Planning Tool
NTP	Network Time Protocol	SR-IOV	Single Root I/O Virtualization
NVMe	NVMe	SSH	Secure Shell
ODB	Operational Database	SUMA	Service Update Management Assistant
OMI	Open Memory Interface	SWMA	Software Maintenance Agreement
OS	operating system	TAM	Technical Account Manager
PCIe	Peripheral Component Interconnect Express	TCE	translation control entity
PDU	power distribution unit	TCO	total cost of ownership
PEC	PCI Express controller	TL	Technology Level
PEP	Power Enterprise Pools	TME	Transparent Memory Encryption
PFA	Predictive Failure Analysis	TR	Technology Refresh
PFW	partition firmware	UAK	Update Access Key
PHYP	PowerVM Hypervisor	UOV	Unbalanced Oil and Vinegar
PoC	proofs of concept	USB	Universal Serial Bus
PowerVC	Power Virtualization Center	vHMC	Virtual Hardware Management Console
PowerVS	Power Virtual Server	VIOS	Virtual I/O Server
PQC	Post-Quantum Cryptography	VLAN	virtual local area network
PTF	Program Temporary Fix	VM	virtual machine
QSE	Quantum-Safe Encryption	VMRM	VM Recovery Manager
RAG	retrieval augmented generation	XMSS	eXtended Merkle Signature Scheme
RAS	reliability, availability, and serviceability	ZPD	zero planned downtime
RBAC	role-based access control		
REST	Representational State Transfer		
RGA	Resource Groups Advisor		
RHBA	Red Hat Bug Advisory		
RHCOS	Red Hat Enterprise Linux CoreOS		
RHEA	Red Hat Enhancement Advisory		
RHEL	Red Hat Enterprise Linux		

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 book.

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.

- ▶ *Creating Red Hat OpenShift Multiple Architecture Clusters with IBM Power*, SG24-8565
- ▶ *IBM i 7.6 features and function*, SG24-8588
- ▶ *IBM PowerHA SystemMirror and IBM VM Recovery Manager Solutions Updates*, REDP-5694
- ▶ *IBM Power Security Catalog*, SG24-8568
- ▶ *Introduction to IBM Power Virtual Server Private Cloud*, REDP-5745
- ▶ *Modernization Techniques for IBM Power*, SG24-8582
- ▶ *Using Ansible for Automation in IBM Power Environments*, SG24-8551

You can search for, view, download, or order these documents and other Redbooks, 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:

- ▶ Red Hat OpenShift AI: Supported Configurations
<https://access.redhat.com/articles/rhoai-supported-configs>
- ▶ Red Hat OpenShift AI Self-Managed 2.20
https://docs.redhat.com/en/documentation/red_hat_openshift_ai_self-managed/2.20
- ▶ Red Hat OpenShift AI Self-Managed 2.20 - Documentation
https://docs.redhat.com/en/documentation/red_hat_openshift_ai_self-managed/2.20/pdf/release_notes/Red_Hat_OpenShift_AI_Self-Managed-2.20-Release_notes-en-US.pdf
- ▶ Red Hat OpenShift AI Self-Managed Life Cycle
<https://access.redhat.com/support/policy/updates/rhoai-sm/lifecycle>
- ▶ Red Hat OpenShift AI Service Definition
<https://access.redhat.com/articles/7102731>

- ▶ Red Hat OpenShift for IBM Power
<https://www.redhat.com/en/resources/openshift-ibm-power-systems-datasheet#:~:text=Red%20Hat%20OpenShift%20and%20IBM,on%20IBM%20Power%20are%20supported>
- ▶ Red Hat OpenShift Operator Life Cycles
https://access.redhat.com/support/policy/updates/openshift_operators
- ▶ RHEL Versions Utilized by RHEL CoreOS and Red Hat OpenShift Container Platform
<https://access.redhat.com/articles/6907891>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM Power11 Scale-Out Servers: Introduction and Overview

SG24-8590-00

ISBN 0738462322



(0.5" spine)

0.475" x 0.873"

250 <-> 459 pages



SG24-8590-00

ISBN 0738462322

Printed in U.S.A.

Get connected

