Redbooks
ibm.com/redbooks

IBM Power11 E1150 Introduction and Technical Overview

Tim Simon
Henry Vo
Jose Martin Abeleira
Giuliano Anselmi
Kevin Gee
Rob McNelly
Peter Piskutz
Borislav Ivanov Stoymirski
Atanas Yordanov

IBM Power

IBM

Redbooks

IBM Redbooks

Power11 E1150 Introduction

October 2025

SG24-8589-00

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 Power E1150 (9043-MRU)

Contents

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 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®
Concert®
HashiCorp®
IBM®
IBM API Connect®
IBM Cloud®
IBM Cloud Pak®
IBM Cloud Satellite®
IBM Concert®
IBM FlashSystem®
IBM Garage®
IBM Instana®
IBM Research®
IBM Security®
IBM Spyre™
IBM Turbonomic®
IBM Watson®
IBM Z®
Instana®
Micro-Partitioning®
Passport Advantage®
Power8®
Power9®
PowerHA®
PowerVM®
Redbooks®
Redbooks (logo)  ®
Satellite®
Spyre™
System z®
SystemMirror®
Terraform®
Tivoli®
Turbonomic®
watsonx®
WebSphere®

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.

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

SUSE LLC. All rights reserved. SUSE and the SUSE logo are registered trademarks of SUSE LLC 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

The goal of this publication is to provide a hardware architecture analysis and highlight the changes, new technologies, and major features that are introduced in the IBM® Power E1150.

This publication is intended for the following professionals who want to acquire a better understanding of IBM Power server products:

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

This paper expands the set of IBM Power documentation by providing a desktop reference that offers a detailed technical description of the Power E1150.

The IBM Power E1150 is a next-generation enterprise server that is designed to meet the demands of modern, data-intensive workloads with unmatched performance, security, and operational agility. Built on IBM's full-stack innovation, from the Power11 processor to firmware, operating systems, and cloud integration, the E1150 delivers a unified platform for mission-critical computing.

The IBM Power E1150 provides these benefits for clients running critical workloads:

- Resilient operations: The E1150 is engineered for continuous availability with zero planned downtime, automated system maintenance, and spare core failover. It helps ensure business continuity even in the face of hardware faults or cyberthreats.
- Quantum-safe security: With built-in quantum-safe encryption, secure boot, and support for FIPS 140-3 certified cryptographic modules, the E1150 protects sensitive data against both current and emerging threats.
- AI-optimized performance: Integrated Matrix Multiply Assist (MMA) accelerators enable high-throughput AI and machine learning workloads, making the E1150 ideal for organizations scaling AI across their operations.
- Cloud-ready flexibility: The E1150 supports hybrid cloud strategies through seamless integration with IBM Cloud® and Power Virtual Servers, enabling consistent performance and workload portability.
- Efficient resource utilization: With support for Capacity on Demand and Power Enterprise Pools 2.0, organizations can dynamically scale compute and memory resources across systems, optimizing cost and performance.
- Simplified management: Advanced automation, cryptographic inventory through the IBM PowerSC, and energy-efficient modes reduce operational complexity and free IT teams to focus on innovation.

Authors

This book was produced by a working remotely.

Tim Simon is an IBM Redbooks® Project Leader in Tulsa, Oklahoma, US. He has over 40 years of experience with IBM primarily in a technical sales role working with customers to help them create IBM solutions to solve their business problems. He holds a BS degree in Math from Towson University in Maryland. He has worked with many IBM products and has extensive experience creating customer solutions using IBM Power, IBM Storage, and IBM System z® throughout his career.

Henry Vo is an IBM Redbooks Project Leader with 11 years working for IBM. He has technical expertise in business problem solving, risk/root-cause analyze, and writing technical plans for business. During IBM life, he has held multiple roles such as Project management, ST/FT/ETE Test, Back-end Developer, DOL agent for NY and certified in IBM Z®, IBM Power. Henry holds a Master of MIS from the University of Texas at Dallas since 2012.

Jose Martin Abeleira is the Team Leader of the Infrastructure/Network Section at DGI (Uruguay Taxes Collection Agency) working in Montevideo, Uruguay. Martin has more than 30 years of experience in the IT field and is a Platinum Redbooks Author, a Certified Consulting IT Specialist, and an IBM Certified Systems Expert Enterprise Technical Support for IBM AIX® and Linux. Martin worked with IBM for 8 years and has 25 years of AIX experience. His areas of expertise include IBM Power Servers, IBM AIX, UNIX, and LINUX, Live Partition Mobility (LPM), IBM PowerHA® SystemMirror®, SAN and Storage on the IBM DS line, V5000, V7000, HITACHI HUSVM, and G200/G400/G370/E590. He holds an Information Systems degree from Universidad Ort Uruguay. He teaches Systems Administration in the Systems Engineer career at the Universidad Catolica del Uruguay, he also teaches Infrastructure Administration, in the Computer Technologist career created by the joint venture between Universidad del la Republica Uruguay, Universidad del Trabajo del Uruguay and Universidad Tecnologica.

Giuliano Anselmi is a Presales and IT Specialist at Ricca IT, an Italian company that designs, optimizes, and manages secure IT and high-performance computing (HPC) infrastructures. Ricca IT takes a personalized, people-centric approach to help organizations enhance efficiency, sustainability, and security, while also supporting the implementation of AI-driven initiatives. Giuliano began his career at IBM, where he spent over two decades in various technical roles focused on IBM Power Systems. His work included close collaboration with IBM Austin Labs to introduce emerging technologies. He has contributed to IBM Redbooks and remains active in global technical communities. Throughout his career, Giuliano has held roles such as Product Engineer, IT Architect, Team Leader, and Project Manager. During his tenure at IBM Ireland, he played a key role in supporting the European launch of Power10 and contributed to innovative projects as a Digital Technical Specialist.

Kevin Gee is an IBM Champion with over 30 years of IT industry experience in roles spanning from technical sales and enablement, consulting, solution design, systems engineering and architecture, and project management. He is currently a Technical Solutions Architect in the Modern Data Center practice at Pellera Technologies (formerly Mainline Information Systems). Kevin has both broad and deep experience in IBM compute, cloud, and storage infrastructure and performance, high availability and disaster recovery, enterprise backup and recovery and data protection, and product development and go-to-market strategy. He has authored hundreds of white papers, training guides, and presentations for clients, co-authored IBM Redbooks, and helped develop IBM product certification exams. He frequently presents at IBM technical conferences and contributes content to others' presentations. He has also conducted and published original research in AI, machine learning, graph processing, and natural language processing. Kevin holds a M.S. in Computer

Science from The University of Texas at Arlington and a B.S. in Computer Science and a B.A. in Spanish from Brigham Young University.

Rob McNelly is one of Meridian IT's senior Power Systems engineers. Beginning in the late 1980s he has had roles as an IBM customer, IBM employee, and a consultant specializing in IBM Power systems. He has expertise that is anchored by six years with the IBM Global Services team where he supported a variety of clients. At Meridian IT, Rob focuses on Power Systems architecture and engineering, supporting AIX and Linux for a range of clients across multiple industries. Rob was first named as an IBM® Power Champion in 2011 and in 2019 he was named as a Lifetime IBM Champion. Rob's contributions include creating and reviewing multiple IBM certification tests, generating technology videos, writing articles and weekly blog entries on IBM technology, and serving as a technical editor for IBM Systems magazine.

Peter Piskutz is a pre-sales Solutions Architect in the Modern Data Center practice at Pellera Technologies (formerly Mainline Information Systems), an IBM Platinum Business Partner. He has 26 years of experience designing IBM Power solutions for customers across many different industries, from very large enterprises to customers with just one or two systems. Peter also manages their Business Partner Innovation Center, where he creates demos and proof of concepts showcasing IBM Power technology to customers. Over the years he has obtained many certifications related to the IBM Power platform as well as Red Hat and holds a Bachelor of Science degree in Mathematics from Rensselaer Polytechnic Institute with minors in Psychology and Accounting.

Borislav Ivanov Stoymirski is an IBM Power Servers Hardware Product Engineer at IBM Bulgaria, specializing in solving complex hardware and software issues on IBM Power Servers, including IBM AIX, VIOS, HMC, IBM i, PowerVC, PowerVM® NovaLink, PowerVM, Linux on IBM Power Servers, and Red Hat OpenShift. Since joining IBM in 2015, he has provided reactive break-fix, proactive, preventative, and cognitive support to several clients worldwide. Borislav earned a Master's degree in Computer and Software Engineering from the Technical University of Sofia, Bulgaria, a Master's degree in Transport Machinery and Technologies from the Technical University of Sofia, Bulgaria, and another Master's degree in Ecology and Environmental Protection from the University of Chemical Technology and Metallurgy, Bulgaria. He has authored several publications and has led technical training programs both inside and outside IBM.

Atanas Yordanov is an IBM Power Systems Product Engineer based in Bulgaria. He has over 19 years of professional experience in the information technology field, with deep expertise in IBM Power Systems, AIX, PowerVM, Linux and system administration. His career spans both the public and private sectors, including more than a decade of service in government IT infrastructure administration, where he led and supported mission-critical enterprise environments. His areas of expertise include enterprise systems optimization, virtualization, hardware lifecycle support, and cross-platform integration. He is highly skilled in designing resilient infrastructure solutions, implementing virtualization strategies, and maintaining high-availability systems for large organizations. Throughout his career, he has contributed to major projects aimed at improving operational efficiency, infrastructure reliability, and digital transformation within both government bodies and corporate enterprises.

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

Nicole Nett
IBM Chief Architect Power Servers, Austin, TX

Lyudmila Simeonova
IBM AIX IO Device Drivers/PowerVM VIOS, Sofia, Bulgaria

Milen Paiakov
IBM Power Courseware Developer, Sofia, Bulgaria

Ruby Zgabay
IBM Power Systems Chief Engineer, Austin, TX

Kaveh Naderi
IBM Power System Chief Architect, Austin, TX

Paul Bastide
IBM Senior Software Engineer, Cambridge, MA

Tanya Koleva
IBM Technology Lifecycle Services, Sofia, Bulgaria

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

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 E1180 Introduction and Technical Overview,* SG24-8587 and *IBM Power11 Scale-Out Servers: Introduction and Overview,* SG24-8590 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

1

Introduction to Power11

Power11 represents a significant advancement in enterprise computing. Building on the strengths of Power10, Power11 introduces 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 help future-proof critical workloads. The processor is designed to support demanding enterprise applications, particularly those involving AI, analytics, and hybrid cloud environments.

Power11 uses Integrated Stacked Capacitor (ISC) technology and advanced 2.5D packaging to increase performance and efficiency. These innovations, combined with enhanced thermal management, such as more efficient fans and heat sinks, significantly improve system density, cooling effectiveness, and overall compute capability.

A key feature of Power11 is its continued support for AI workloads through the Matrix-Math Assist (MMA) architecture and the integration of the IBM Spyre™ accelerator card, which is optimized for generative AI and complex model inference. This design positions Power11 as an alternative to traditional GPU-heavy AI infrastructures. Additionally, Power11 strengthens its virtualization capabilities with deeper KVM integration, improving compatibility with Linux-native tools and hybrid cloud platforms. This combination makes it a versatile choice for enterprises seeking scalable, AI-ready infrastructure.

The following topics are covered in this chapter:

- IBM Power11: Advancing Enterprise Computing
- Power roadmaps
- Power11 improvements from Power10
- Power11: Trusted, autonomous, and modern
- IBM Power11 processor architecture
- Operating system support
- Firmware and Hardware Management Console
- Rack support

1.1 IBM Power11: Advancing Enterprise Computing

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

Power architecture follows a steady innovation cadence and introduces a new processor generation approximately every three years. Each release delivers substantial advancements, including increased core and thread counts, enhanced energy efficiency, greater memory bandwidth, and expanded I/O capabilities: These advancements help ensure that the platform remains a leader in enterprise computing.

The Power11 platform is built on three foundational principles:

1. Trusted

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

2. Autonomous

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

3. Modern

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

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

1. Entry-level: Power S1122, L1122, S1124, and L1124, optimized for departmental and edge workloads
2. Midrange: Power E1150, ideal for enterprise consolidation and scalable virtualization
3. High-end: Power E1180, designed for large-scale, mission-critical, and AI-intensive workloads

Power11 introduces several key advancements across the portfolio:

- Up to six nines (99.9999%) availability on high-end systems
- Enhanced security with quantum-safe cryptography and secure boot enhancements
- 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, using DDR5 and advanced memory controllers

Power11 introduces serial number-preserving upgrades from Power10 systems, with scheduled availability starting in Q4 2025. These upgrades enable seamless transitions without disrupting asset tracking or licensing.

This IBM Redbooks publication focuses on the Power E1150, the midrange system in the Power11 family. It offers a well-balanced blend of performance, scalability, and cost efficiency,

making it a suitable platform for enterprise workloads, virtualization environments, and AI integration.

1.2 Power roadmaps

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

IBM Power Virtual Server extends these capabilities into the cloud, allowing clients to run Power workloads either on-premises or in IBM Cloud while benefiting from flexible consumption models and meeting data residency and compliance requirements. This hybrid approach ensures seamless workload portability and supports modernization initiatives.

The Power platform also supports a variety of operating systems, including IBM AIX, IBM i, and Linux on Power, which supports distributions such as Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server. This combination makes the platform suitable for open-source and cloud-native environments.

Complementing its technology stack, IBM has built a strategic network of alliances that significantly expand the reach and capabilities of the Power platform. Notably, IBM Power Virtual Server 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 systems, supporting automation and DevOps practices. Additionally, IBM collaborates with major cloud providers, including Microsoft, AWS, and Google Cloud, to deliver hybrid and multicloud solutions that allow Power workloads to extend into public cloud infrastructures.

This section discusses IBM's roadmaps for the evolution of IBM Power-based infrastructure.

1.2.1 IBM Power processor roadmap

The IBM Power roadmap evolves with the introduction of Power11. Building on the innovations of Power10, Power11 introduces enhancements across the processor architecture, improves packaging, and provides advances in energy efficiency.

A key advancement is the integration of an additional silicon layer for improved energy management, which allows for better performance with less power consumption. Power 11n also increases core strength and count, enabling it to manage more demanding workloads. IBM is refining its matrix math accelerators (MMAs), first introduced in Power10, to accelerate AI inferencing directly on the chip, reducing the need for external GPUs in many scenarios.

IBM's roadmap emphasizes full-stack innovation and adaptability to emerging technologies. The Power11 platform is designed to support 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, with a focus on hybrid cloud readiness, AI integration, and support for open-source environments. 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 highlights the ongoing plans for improvement. With the release of Power11, the next generation of Power processors is already in development, and additional processor designs are in the pipeline after that generation.

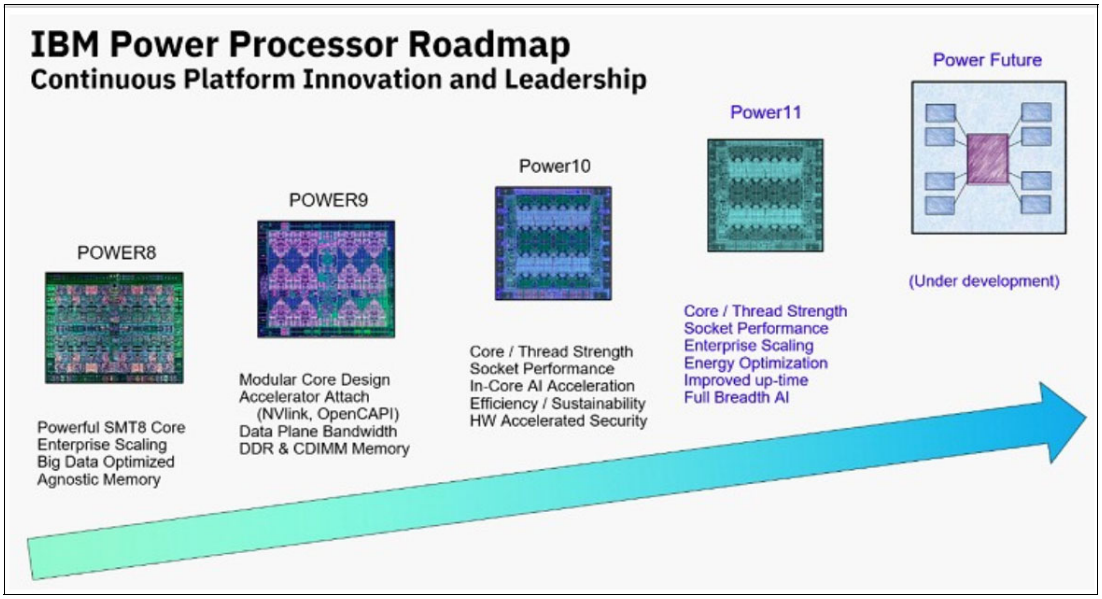

Figure 1-1 Power processor roadmap from IBM Power8® to the future

1.2.2 IBM AIX Roadmap

IBM AIX is a proprietary UNIX-based operating system that is built on open standards and designed specifically for IBM Power servers.

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

AIX's proven binary compatibility allows 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 that includes further innovations and purposeful support extending beyond 2035, ensuring a stable and future-ready platform for enterprise workloads.

The current roadmap is shown in Figure 1-2 on page 5.

For more information about 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 section 6.1, "AIX" on page 120.

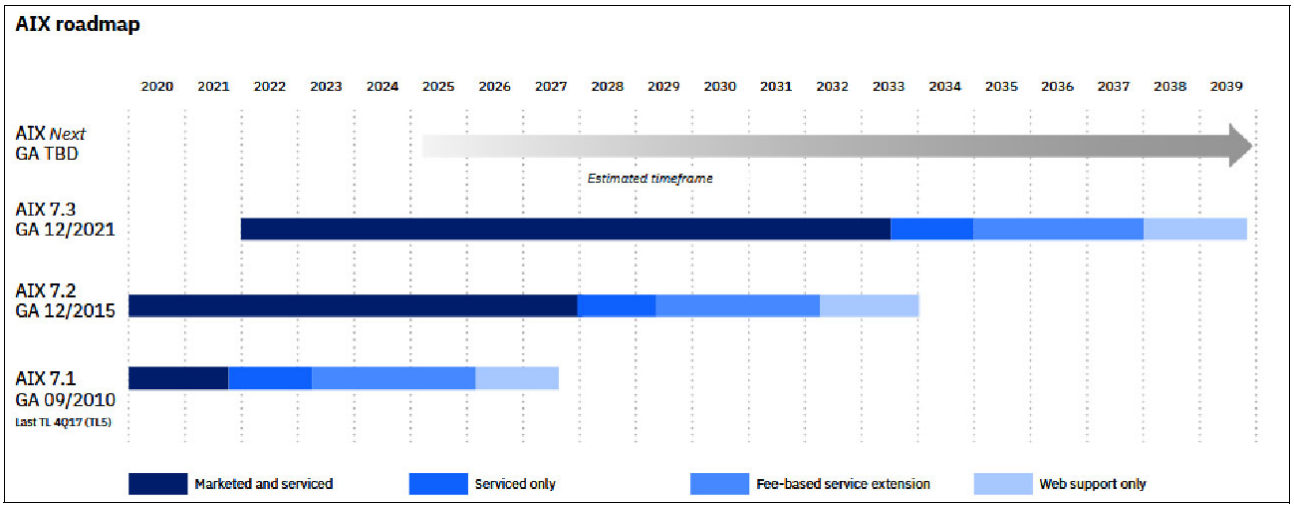

Figure 1-2 Support timeline for future generations of AIX

1.2.3 PowerVM VIOS roadmap

PowerVM Virtual I/O Server (VIOS) is a critical component of the IBM Power virtualization stack, enabling shared access to physical I/O resources for client logical partitions (LPARs). Similar to all IBM software, VIOS follows a defined product lifecycle that includes general availability, support phases, and end-of-support (EOS) milestones. Table 1-1 shows the VIOS release schedule as of July 2025 and includes the following headings:

- End of Fix Support (EoFS): The end of the maintenance period for a VIOS Release. Fix Packs and interim fixes are not created for a VIOS Release after EoFS.
- Latest FP: The most recent FP that is available for the VIOS Release and links to that FP on Fix Central.
- Next FP: The target availability date for the next Fix Pack for the given VIOS Release. After the final Fix Pack for a VIOS Release is available, further fixes for the release are provided as interim fixes.

Table 1-1 VIOS release schedule

VIOS Release	Release Date	End of Fix Support	Latest FP	Next FP
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	Sep-21	30-Sep-24	VIOS_FP_3.1.3.40	None
3.1.2	Nov-20	30-Nov-23	VIOS_FP_3.1.2.60	None

Note: The end of service (EOS) date for a VIOS release is not indicated in this table. EOS dates can be found on the IBM Software Lifecycle page.

For additional details regarding the PowerVM Virtual I/O Server (VIOS) roadmap, see PowerVM VIOS Lifecycle Information and System to PowerVM Virtual I/O Server maps.

1.2.4 Red Hat Enterprise Linux and IBM Power

Red Hat reaffirmed its commitment to the IBM Power platform as both companies align around hybrid cloud, AI, and open-source innovation. This partnership is based on a 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 Red Hat Enterprise Linux (RHEL) on Power. 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 use Red Hat's open hybrid cloud technologies across both 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 Matrix Math Assist (MMA) and the planned Spyre AI Accelerator. This synergy enables enterprises to deploy AI models efficiently on Power infrastructure by using Red Hat's open-source toolchains and platforms.

Furthermore, 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 on Red Hat Enterprise Linux support on IBM Power, see "Red Hat Enterprise Linux" on page 125.

1.2.5 SUSE Enterprise Linux and IBM Power

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

SUSE and IBM have collaborated for more than two decades to provide robust, scalable, and secure Linux solutions on IBM Power infrastructure. Their joint efforts are focused on the following aspects:

- SUSE Linux Enterprise Server for IBM Power
- SLES 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 alignment enables organizations to modernize IT environments and avoids vendor lock-in.

1.3 Power11 improvements from Power10

IBM Power E1150 introduces 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, operating systems, and cloud services, IBM Power delivers a unified platform that is built for

autonomous IT. This holistic approach drives measurable business outcomes across three foundational pillars: business continuity, productivity and efficiency, and growth and scalability.

1.3.1 Business continuity

IBM Power E1150 enhances operational resilience with features that are designed to minimize downtime and help ensure continuous service delivery:

- Planned downtime elimination: By using technologies such as live updates, rolling upgrades, and autonomous patching, the E1150 enables maintenance without taking applications offline, targeting zero hours of planned downtime.
- Spare core technology: At the hardware level, the system includes spare processor cores that can be dynamically activated in response to hardware failures. This proactive fault-tolerance mechanism allows seamless failover without requiring system or application restarts, preserving compute capacity and system integrity.
- Power Cyber Vault: Advanced threat detection can identify ransomware attacks in under a minute.[1] 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 reboots and live partition mobility, 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.

1.3.2 Enhanced productivity and efficiency

IBM Power E1150 is engineered to maximize IT efficiency and output:

- Autonomous IT operations: Built-in automation and AI-driven management reduce manual intervention and operational overhead.
- Integrated monitoring and management: Tools such as HMC and PowerVC provide centralized control, enabling streamlined operations across hybrid environments.

1.3.3 Accelerated growth and scalability

IBM Power E1150 supports dynamic scaling and hybrid cloud integration to meet evolving business demands:

- IBM Power Virtual Servers (PowerVS): Immediate availability of Power11 in IBM Cloud enables seamless extension of workloads to the cloud, 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 Matrix Multiply Assist (MMA), a hardware acceleration unit optimized for AI and machine learning workloads. By running matrix operations directly on-chip, MMA reduces latency and boosts throughput, making the E1150 suitable for deep learning inference and training tasks.

[1] https://community.ibm.com/community/user/blogs/brandon-pederson1/2025/06/26/meet-the-ibm-power11-family

1.3.4 Technical enhancements

With the previously described benefits, Power11-based systems provide performance improvements from faster core frequencies, additional cores per server, and improved memory bandwidth with reduced memory latency. Table 1-2 shows some of the benefits that these improvements can bring to your enterprise.

Table 1-2 Power11 performance benefits

Improvement	Benefit
Core performance increase of 15–25% with both HW and SW planned improvements	Exceed SLAs for response times and batch windows.
System performance increase of 30–45% with the new larger 30-Core Power11 DCMs	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.

1.3.5 Power11 server portfolio

The Power11 server portfolio consists of several offerings to provide clients with choices that are designed to meet specific business requirements.

High-end: IBM Power E1180

The top of the portfolio is the IBM Power E1180, a high-end enterprise server built for the most demanding workloads. It delivers performance, scalability, and resiliency, making it suitable for large-scale ERP systems, core banking, and high-throughput analytics. With support for large memory footprints, advanced RAS features, and robust virtualization capabilities, the Power E1180 is engineered for continuous availability and workload consolidation.

Midrange: IBM Power E1150

The Power E1150 serves as the midrange server of the Power11 family. 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 designed for organizations seeking a balance between performance, scalability, and operational efficiency, especially in hybrid cloud environments.

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

For smaller enterprises or distributed environments, IBM offers the Power S1124 and Power S1122 servers. These entry-level, scale-out systems are optimized for cost-effective deployment of AIX, IBM i, and Linux workloads. They are well suited for departmental applications, edge computing, and cloud-native workloads. Although compact, these systems use the same Power11 processor technology, ensuring consistent performance and security across the portfolio.

Figure 1-3 shows the new Power11 server portfolio.

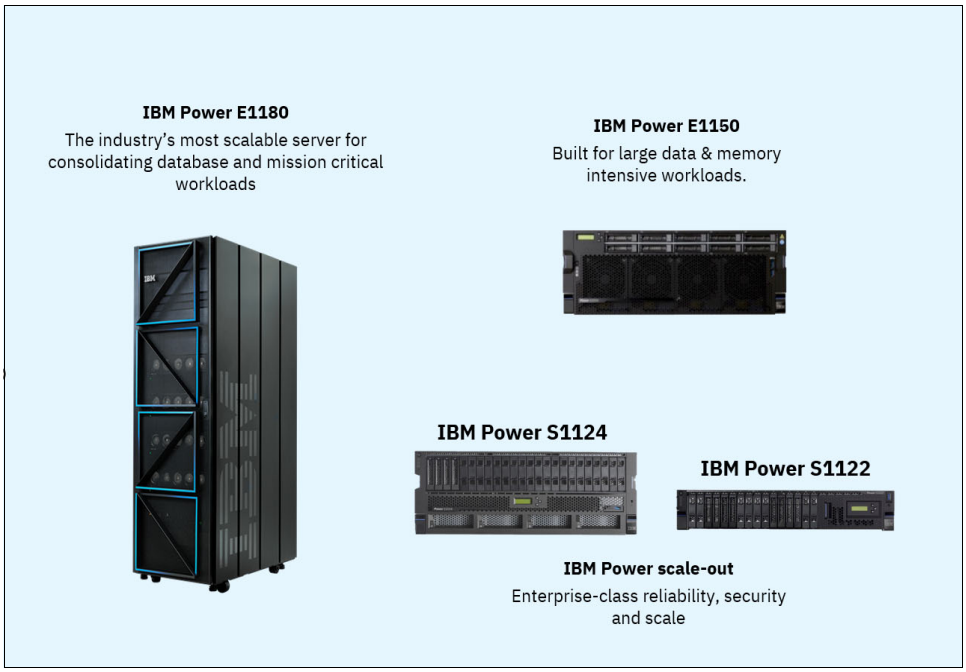

Figure 1-3 Power11 portfolio

Architectural technology Improvements

The Power11 portfolio introduces significant innovations across the product lineup and underlying core technologies that are aimed at delivering enhanced performance and greater value. One of the most notable advancements lies in the processor and memory architecture, where Power11 brings substantial improvements over the previous Power10 generation. 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

Integrated Stacked Capacitor 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 to maximize performance but maintain energy efficiency.

Increased memory frequencies

The IBM 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, Power11 significantly boosts memory throughput and reduces latency. These enhancements provide quicker data access and processing, which is critical for memory-intensive workloads such as real-time analytics, large-scale databases, and AI applications. The combination of increased bandwidth and reduced latency enables smoother performance, better responsiveness, and more efficient use 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 Compare
High-End (4S+):	IBM Power E1180	► 4 nodes (CEC) each 5U with 4 processor modules each ► Up to 64 TB of DDR5 based memory	► 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)	IBM Power E1150	► 2–4 processors in a 4U form factor ► Up to 30 cores per module ► Up to 16 TB of DDR5 based memory	► Up to 25% more core capacity ► Same memory capacity ► Introducing Resource Groups to maximize system utilization ► In-place upgrades with serial number preservation
Scale Out 2S4U	IBM Power S1124 IBM Power L1124	► Up to 2 processors in a 4U form factor ► Up to 30 cores per module ► Up to 8 TB of DDR5 based memory	► Up to 25% more core capacity and same memory capacity ► Introducing Resource Groups to maximize system usage
Scale Out 2S2U	IBM Power S1122 IBM Power L1122	► Integrating the eSCM modules into the S1122 family: 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 of DDR5 based memory	► Up to 50% more core capacity ► Same memory ► Introducing Resource Groups to maximize system usage

Investment protection

IBM Power11 systems are designed with strong investment protection in mind, ensuring that organizations 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 Power E1180 and from the 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, organizations 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 moving to the latest technology but includes the performance, scalability, and security benefits of the Power11 architecture.

Additionally, IBM supports Power Enterprise Pools (PEP) that include both Power10 scale-out servers and Power11 scale-out servers. With this mixed-system support, organizations can gradually migrate workloads from older to newer systems within the same resource pool. It allows for 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

A notable change in the Power11 generation is the streamlining of the scale-out server lineup. IBM reduced the number of scale-out models, simplifying system selection and reducing complexity in deployment and support. This focused approach helps make it easier for organizations to choose the right system for their needs but still benefit from the full-stack integration and innovation that IBM Power is known for.

Figure 1-4 shows how customers can migrate from their existing IBM Power9® or Power10 servers to the new Power11 portfolio.

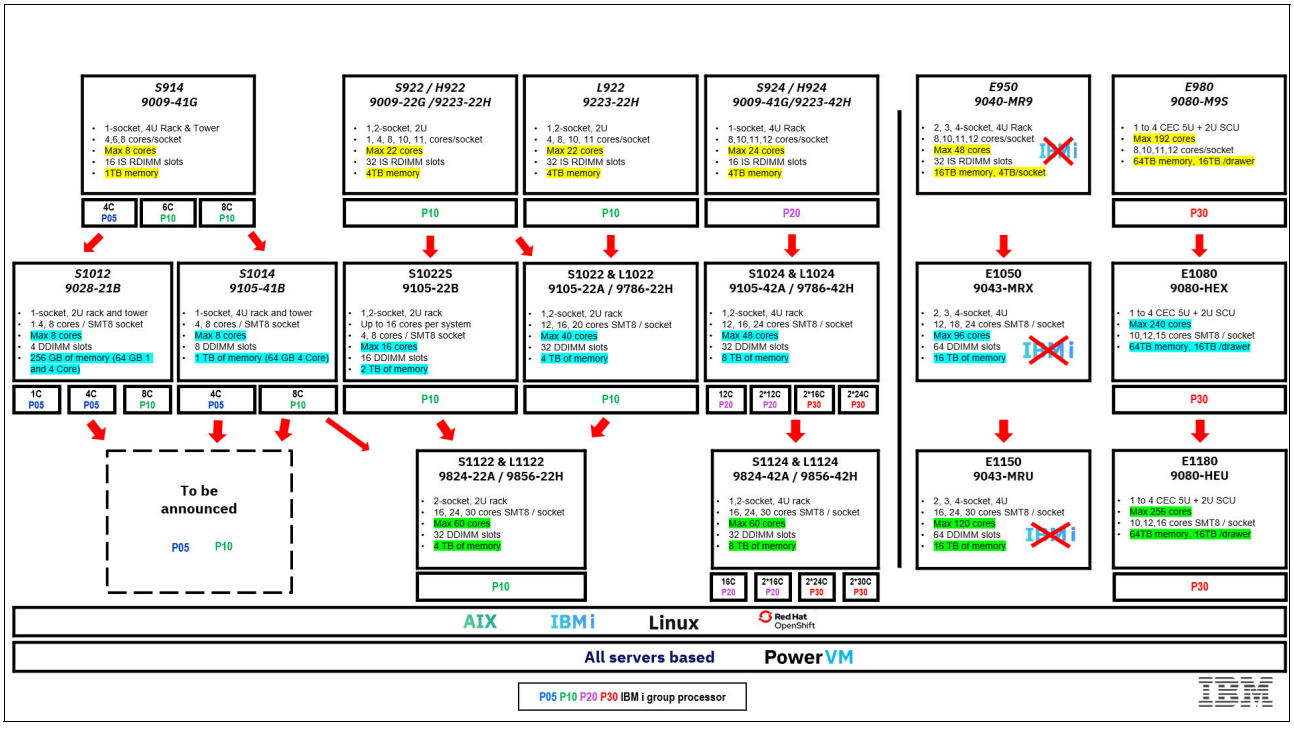

Figure 1-4 Migration path for Power9 or Power10 to Power11

1.4 Power11: Trusted, autonomous, and modern

IBM's vision for Power is autonomous IT, where Power11 processor-based servers can self-manage, self-heal, and scale independently. This capability supports business agility by enabling enterprises to respond more quickly to market changes and customer demands, deploy new services rapidly, and reduce time to market for new products.

IBM is working toward this vision by building automation into the platform and combining it 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 with IBM Terraform®, Application Configuration Management with Red Hat Ansible, and enhancements to existing platform tools such as the Hardware Management Console and Cloud Management Console. For AI, IBM is developing AI-infused workflows with IBM watsonx and new tools such as IBM Concert®.

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

1.4.1 Building a trusted infrastructure with Quantum Safe Encryption

Infrastructure built on IBM Power processor-based servers benefits from robust security technologies that are integrated into both the hardware and software stacks. Power processor-based servers offer advanced security features at every level of the system, helping ensure comprehensive protection for sensitive data and applications. These features include advanced encryption technologies, secure boot capabilities, and integrated firmware updates. Additionally, Power processor-based servers make use of IBM's expertise in securing mission-critical workloads, which can help organizations that are seeking a secure environment for digital assets.

Workloads on Power11 processor-based servers benefit from improved cryptographic accelerator performance compared to previous generations. Specifically, the Power11 chip supports accelerated cryptographic algorithms such as AES, SHA2, and SHA3, resulting in higher per-core performance for these algorithms. This enhancement allows features such as AIX Logical Volume Encryption to operate with minimal impact on system performance.

The processor-core technology of Power11 incorporates integrated security protections for several aspects of security:

- Improved cryptographic performance: Integrated cryptographic support reduces the performance impact of encrypting and decrypting data, enabling pervasive encryption to protect critical information.
- Increased application security: Hardened defenses against return-oriented programming (ROP) attacks.
- Simplified hybrid cloud security: Setup-free hybrid cloud security administration with a single interface.
- Enhanced virtual machine isolation: Provides advanced virtual machine isolation technology to defend against attacks that exploit operating system or application vulnerabilities in one virtual machine to access others or the host system.

Encryption technologies and their applications

Power11 emphasizes comprehensive security throughout its design, offering multiple encryption options. Key among these are transparent memory encryption (TME), fully homomorphic encryption (FHE), and quantum-safe encryption (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's cryptographic engines to perform encryption and decryption efficiently. TME ensures pervasive protection of data in memory with minimal performance impact, as encryption and decryption are performed within the memory chips. Integration into normal operations is seamless and automatic.

Fully homomorphic encryption

With FHE, computations can be run directly on encrypted data without decrypting it first, ensuring that sensitive data remains confidential during processing. FHE operates at the software level and uses advanced mathematical algorithms to enable computations on ciphertexts. Implementing FHE requires specialized libraries and frameworks. However, FHE is computationally intensive and can introduce performance issues compared to hardware-only encryption methods.

Quantum-safe encryption

QSE is designed to resist quantum attacks, securing data against the computational capabilities of future quantum computers that could break current cryptographic algorithms. QSE employs cryptographic algorithms 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 ensure robust security against future quantum advancements. QSE is typically used for securing long-term data, sensitive communications, and critical infrastructure.

The relevant features and differences in these technologies is shown in Table 1-4

Table 1-4 Key Differences

Feature	TME	FHE	QSE
Encryption Scope	Secures data in memory	Allows computations on encrypted data	Prevents against 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 accelerated through the Power11 cryptographic engines and designed to have minimal performance impact	Involves substantial cpu usage	The impact of QSE varies; some quantum-safe algorithms can introduce performance overhead; this 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-safe compliance

Quantum-safe encryption (QSE), also known as Post-Quantum Cryptography (PQC), refers to encryption methods designed to remain secure against both classical and quantum computers. As quantum computing advances, it poses a potential threat to existing cryptographic systems, which could compromise their security. QSE is essential for protecting sensitive data, communication channels, and user identities against attacks from quantum computing.

The urgency of adopting QSE stems from two primary concerns:

1. Advanced quantum computers might enable adversaries to intercept and decrypt protected digital communications through Harvest Now, Decrypt Later (HNDL) strategies, even before reaching Q-Day. Q-Day refers to the anticipated point when quantum supremacy becomes widespread and many current encryption algorithms are no longer effective.
2. Transitioning to QSE might require more than a decade because of the complexities of organizational structures and IT infrastructure.

Organizations should begin evaluating and implementing QSE solutions immediately to ensure continued protection and maintain stakeholder trust. Delaying QSE adoption can have severe consequences. Legacy cryptographic systems that are left unaltered might be compromised during a successful quantum attack, exposing sensitive data and risking confidential business transactions and individual privacy. Financial institutions, critical infrastructure providers, and government agencies might face significant challenges in maintaining operational integrity and confidentiality. Therefore, prioritizing QSE implementation is crucial for long-term cybersecurity resilience.

Power11 is designed to support these quantum-safe algorithms, helping ensure robust security as quantum computing capabilities evolve. Power11's support for quantum-safe features includes the following features:

- Data encryption breakage protection:
  - Risk: Quantum computers might break widely used cryptographic algorithms such as RSA, ECC (Elliptic Curve Cryptography), 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: Power11 supports quantum-safe algorithms such as lattice-based, hash-based, code-based, and multivariate quadratic cryptography, which are believed to be resistant to quantum attacks. The crypto engines in Power11 enhance the performance of these algorithms, ensuring secure encryption and key exchange processes with minimal performance degradation.
- Secure communications:
  - Risk: Quantum computers could intercept and decrypt secure communications, undermining protocols that currently rely on classical encryption methods.
  - Protection: Power11 secures communication channels with quantum-resistant protocols, ensuring data confidentiality throughout the data lifecycle, from storage to transmission.
- Data integrity and authenticity:
  - Risk: Quantum computers could forge digital signatures or tamper with data.
  - Protection: Power11 supports quantum-safe digital signature algorithms such as XMSS (eXtended Merkle Signature Scheme) and UOV (Unbalanced Oil and Vinegar), providing strong security against quantum attacks.

- Long-term data protection:
  - Risk: Sensitive data stored today could be harvested and decrypted in the future as quantum computers become more powerful.
  - Protection: Implementing quantum-safe encryption methods ensures that data remains secure over time. Power11's 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, could expose sensitive data if not adequately protected.
  - Protection: Power11 transparent memory encryption (TME) ensures that data in memory is encrypted, protecting it from physical attacks at runtime.

Quantum-safe algorithms supported by Power11:

The following quantum-safe algorithms are supported by Power11:

- Lattice-based cryptography:
  - Algorithms: ML-DSA (Dilithium), ML-KEM(Kyber) and NTRUEncrypt.
  - Characteristics**:** Secure against quantum attacks. Based on lattice problems – 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), eXtended Merkle Signature Scheme (XMSS), and SPHINCS+.
  - Characteristics: Secure based on hash functions, though generally produces larger signatures and keys.
- Code-based cryptography
  - Algorithms: McEliece Cryptosystem, Bit Flipping Key Encapsulation (BIKE), and Hamming Quasi-Cyclic (HQC).
  - Characteristics: Quantum-resistant based on decoding random linear codes, though public keys can be large.
- Multivariate quadratic equations
  - Algorithms: Unbalanced Oil and Vinegar (UOV), Rainbow.
  - Characteristics: Secure against solving systems of multivariate quadratic equations, generally efficient in signature generation but might involve larger key sizes.

Power11 implementation

Power11 processors support these quantum-safe algorithms with the following methods:

- Crypto engines: Multiple engines per core enable efficient execution of cryptographic operations.
- Software updates: The architecture allows updates to cryptographic libraries, ensuring the integration of new quantum-safe algorithms as they become standardized.

Power11's design and capabilities help ensure robust security against future quantum threats by using hardware acceleration and flexible software updates to maintain high-security standards as the cryptographic landscape evolves.

Encryption enablement in hardware

You have two options for accelerating encryption in an IBM Power11 processor-based server. The first option is to use the built-in encryption acceleration in the Power11 chip. In addition, IBM Power supports a PCIe based encryption accelerator.

On-chip encryption support in Power11

The Power11 processor is designed to support future encryption, including fully homomorphic encryption (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 currently available or will soon be available in the corresponding open source communities.

Workloads on Power11 benefit from cryptographic algorithm acceleration, which enables significantly higher per-core performance than Power10 processor-based servers for algorithms such as Advanced Encryption Standard (AES), SHA2, and SHA3. Features such as AIX Logical Volume Encryption can be activated with minimal performance overhead because of this performance enhancement.

With four times as many AES encryption engines, Power11 processor technology is designed to offer 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 fully homomorphic encryption. The technology also introduces additional improvements to container security. By using hardware features for a seamless user experience, transparent memory encryption simplifies encryption and supports end-to-end security without compromising performance.

1.4.2 Protect, detect, and recover with Power Cyber Vault

Regulatory frameworks such as the U.S. Securities and Exchange Commission (SEC) cybersecurity disclosure rules and the Digital Operational Resilience Act (DORA) are no longer distant challenges; they are active requirements. As regulations increase, each change requires increased 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 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.

Issues with traditional siloed security

Traditional siloed security solutions can no longer keep pace. Most organizations still operate with fragmented security architectures, which 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 can cause the following issues:

- Limited understanding of threats across domains, such as linking identity misuse with network anomalies
- Reactive security postures measured by containment speed, not 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

Additionally, lack of contextual awareness across tools and data streams can allow cyberattacks, such as ransomware attacks that paralyze operations.

Security was not intended to be siloed. It evolved that way because of necessity and legacy architectures. However, that model is no longer sufficient. Forward-looking enterprises are rethinking cybersecurity strategy not only as a defense mechanism but also as 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 frameworks

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 breach, Zero Trust offers a compelling blueprint.

However, the blueprint alone is not sufficient. Applying Zero Trust includes several requirements:

- A deep understanding of the organization's risk surface
- Alignment with regulatory and operational requirements
- Integration with 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: dynamically, digitally, and in a distributed manner.

Impacts of cyberattacks

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

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

According to the Global cyber resilience report, 78% of organizations take more than 100 days to fully recover from a major cyberattack. The recovery cost includes not only downtime but also represents lost trust, revenue, and opportunity.

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

IBM Power Cyber Vault

IBM Power Cyber Vault helps keep businesses running by combining different features:

- Power Cyber Vault protected data

  Automated immutable images stored in the Power Cyber Vault to protect clients' snapshots and backups

- Continuous threat detection

Near real-time filesystem monitoring, FlashSystem inline detection, and recovery data scanning with threat detection time as short as 60 Seconds

- Rapid clean room response

  Automated response to attacks, creating Power Cyber Vault clean rooms for testing and validation with accelerated recovery response

- Cyber Vault Assisted Recovery

  Automated multi-level testing and validation of Power Cyber Vault images to accelerate system recovery with return to full operations in hours versus months

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

- IBM Power11 processor-based servers, which provide the following benefits:
  - Industry-leading threat resistance
  - PowerSC for security and compliance tracking
  - Secure boot and memory encryption built in
  - Scalable and flexible architecture from entry to enterprise servers
- IBM FlashSystems storage, which provides the following benefits:
  - Immutable copies for data protection
  - Ransomware detection for real-time threat detection
  - Flash-speed recovery for rapid restore of snapshot data
  - Storage Insights Pro for comprehensive reporting
- IBM Technology Expert Labs Power Cyber Vault deployment, which includes the following benefits:
  - Design workshop for client-specific requirements
  - Power Cyber Vault deployment to customize and enable the unified Power Cyber Vault to meet client needs
- IBM Expertise Connect, which includes the following benefits:
  - Primary technical point of contact for Power Cyber Vault
  - Dedication to client and client success
  - Availability after installation to help guide client's directions

Power Cyber Vault implementation

The implementation of the solution consists of several phases. Each phase is designed to support risk mitigation and ensure the fastest possible recovery:

1. Identify and assess risks and design mitigation solutions:
   a. IBM Security® and Resilience Assessment:
      - View the organization, processes, and strategies regarding operational and cyber resilience
      - List prioritized recommendations to improve overall resilience

   b. IBM Cyber Vault Solution and Design Workshop for creating detailed solutions based on your cyber resiliency goals, infrastructure, and required services
2. Protect: Safeguarded Cyber Vault copies and backups

   Client specific scheduled creation and protection of immutable snapshots and backups to meet RPO and RTO requirements:
   - Snapshots: FlashSystems and Storage Defender CSM save protected system snapshot state into the Cyber Vault on a regularly scheduled basis.
   - Orchestration for crash consistency

- Application-specific customization available

3. Detect: Real-time threat detection

   Monitoring real-time threats and cyberattack vectors to respond quickly as threats emerge through PowerSC Real Time Compliance orchestration:

   - Zero Trust Execution to detect, prevent, and report unapproved applications from running on Power11 servers
   - FlashSystem's integrated ransomware threat detection, which alerts on active attacks and reports to PowerSC via Storage Insights Pro
   - SEIM Integration to accept alerts and threat reports from SEIM systems

4. Respond and recover: Real time threat response

   Routinely, or when threats are detected, a response phase is triggered to audit and recover the system:

   - Upon receiving an integrity alert, IBM Power Cyber Vault automatically initiates the validation process by creating an immutable backup of the VM within seconds and spinning up a copy of that VM in the clean room within minutes.
   - Various types of Integrity Checks are run in the clean room (including platform-specific tests).
   - Power Cyber Vault reports corrupted copies, and the newest clean copies are found.
   - When approved by the administrator, recovery to production helps meet 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 specifically for virtualized environments running on IBM Power with AIX, Linux, and IBM i. It integrates with the 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 patch management. With tools such as PowerSC Trusted Surveyor, it ensures consistent network configuration and compliance across dynamic virtual environments, helping reduce administrative overhead while maintaining strong security postures.

IBM is actively involved in the development and adoption of quantum-safe cryptographic standards, including collaboration with NIST on the Post-Quantum Cryptography Standardization project. PowerSC supports a seamless transition to NIST's forthcoming standards.

IBM PowerSC offers capabilities to help verify and manage quantum-safe encryption. The Quantum Safety Analysis feature analyzes the quantum-safe status of virtual machines (VMs) to identify areas of concern and ensure compliance with emerging quantum-safe cryptographic standards:

- PowerSC scans endpoints on your server to identify cryptographic artifacts and vulnerabilities related to quantum safety. These scans can also be scheduled.
- It creates an inventory of where different encryption algorithms are implemented on your system.
- The analysis 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.

A Quantum Safety Analysis report in the PowerSC graphical user interface (GUI) lists the discovered cryptographic elements and their assessed strengths.

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

Planned downtime remains one of the biggest challenges in IT operations, especially for critical infrastructure such as Power11. IBM aims to deliver Autonomous IT by embedding advanced automation and AI-infused workflows directly into the platform. The result is a significant reduction in operational effort, minimized disruptions, and improved resilience.

Today's enterprise IT environments face increasing pressure to ensure infrastructure reliability and security without disrupting operations and includes the following challenges:

- Operational complexity: Maintaining firmware, I/O adapters, and Virtual I/O Servers (VIOS) is highly specialized and intricate.
- Service interruptions: Routine maintenance often requires downtime, impacting application availability and business continuity.
- Security risks: Delayed updates expose systems to vulnerabilities (CVEs), weakening compliance and increasing risk.
- Resource drain: IT teams spend 20%–40% of their time on maintenance tasks such as patch evaluation, compatibility testing, and scheduling.
- High downtime costs: With downtime averaging $336,000 per hour, annual maintenance-related outages can exceed $2 million per system.

These issues force organizations into a difficult tradeoff: stay current and risk downtime, or delay updates and increase exposure, which is an unsustainable model in today's always-on digital economy.

Overview of Infrastructure

In this context, infrastructure refers to the Power11 platform, encompassing firmware, I/O components, VIOS, and the operating system. To maintain performance and security, clients activities must include the following actions:

- Apply patches and updates regularly
- Remain on supported versions
- Respond swiftly to security advisories
- Conduct preventive maintenance

However, these actions typically require planned downtime. Businesses must decide to accept regular service interruptions to stay current or defer updates and risk security and performance.

IBM Power11 and autonomous IT

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

At the core of Power11 is the Automated Maintenance Framework, directed and managed by IBM Concert, an AI-driven engine that automates the full maintenance lifecycle. The Automated Maintenance Framework includes the following capabilities:

- Unified update orchestration for system firmware, VIOS, and I/O adapters
- Live Partition Mobility (LPM) to ensure uninterrupted application availability during updates
- Pre-validation checks to confirm system readiness and compatibility
- Flexible update sourcing from IBM repositories, SFTP, NFS, USB, or HMC-local file systems

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

IBM Concert for AI-powered application management

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

- Critical vulnerability exploits (CVEs)
- Certificate management issues
- Application compliance challenges

Built on IBM watsonx, IBM Concert uses advanced generative AI to analyze complex environments and deliver actionable insights. Several features help differentiate IBM Concert:

- Data integration from multiple sources: Concert supports a wide range of data sources across networks, infrastructure, and application architectures, enabling integration without requiring major redesign.
- Hybrid cloud by design: Concert seamlessly supports hybrid cloud environments to meet the realities of modern enterprise IT.
- AI-Infused Orchestration for Proactive Maintenance

  IBM Concert provides the following AI-infused orchestration 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 triggering of automated actions to remediate issues, reducing manual effort and downtime

Automation can be triggered directly from IBM Concert or from the HMC, ensuring flexibility and ease of use.

Key benefits of IBM Power11's autonomous maintenance

IBM Power11's integrated solution for zero planned downtime can deliver measurable benefits across operational, security, and financial dimensions:

- Business continuity: Updates are applied without interrupting workloads, preserving uptime and service-level commitments.
- Operational efficiency: Automation streamlines the maintenance lifecycle, reducing manual effort and administrative overhead.

- Risk reduction: AI-driven orchestration ensures 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 significantly lowers the total cost of infrastructure operations.

Together, these capabilities position IBM Power11 as a foundational platform for secure, scalable, and self-managing enterprise IT infrastructure – aligned 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. This flow includes the following steps:

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

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

Power11 includes a built-in Automated Maintenance Tool that enables platform maintenance without application impact. This automation can be triggered through IBM Concert configuration or manually from the Hardware Management Console (HMC), independent of IBM Concert.

The automated maintenance tool provides the following capabilities:

- Ability to update system firmware, Virtual I/O Servers (VIOS), and I/O adapters from a single update flow
- Support for both concurrent and disruptive updates
- Validation for Live Partition Mobility (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 logical partitions (LPARs) or select a subset of LPARs and define their order
  - Option to choose the order of updates
- Option to either download only or download and update from different sources:
  - IBM website (preferred and recommended 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 1-5 on page 23.

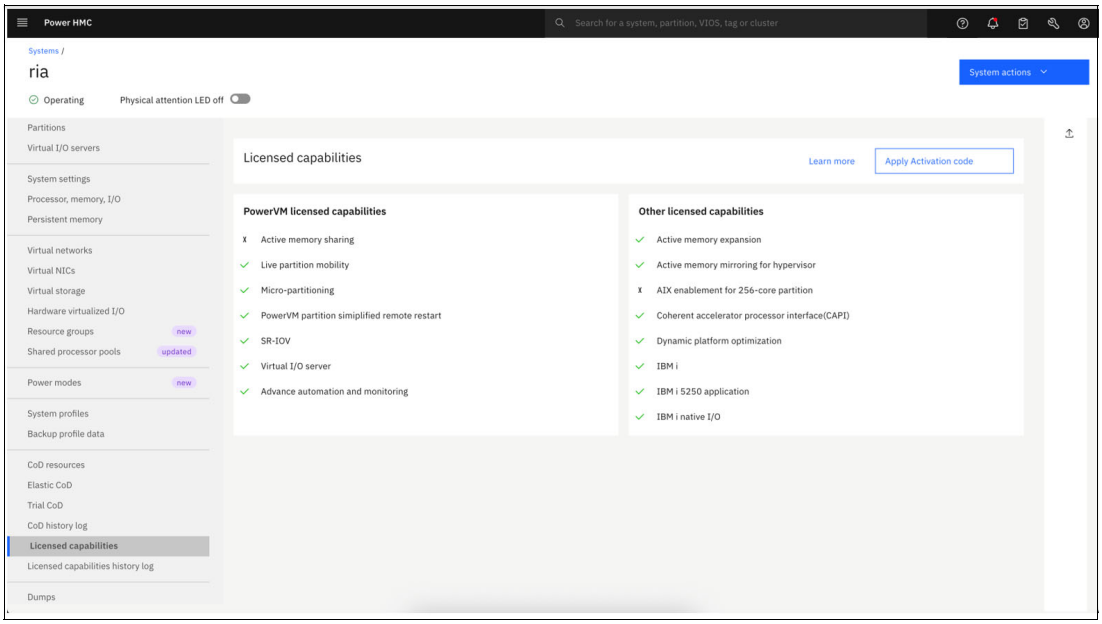

Figure 1-5 Power System Licensed Capabilities

For more detail on using the new automated maintenance tool, see section 5.1.7, “Using the automated maintenance tool” on page 107.

1.4.5 Automated data collection for faster error resolution

Enterprises are increasingly adopting hybrid cloud environments, with 82% reporting usage as of 2024 according to an IBM Cloud Survey.[2] This shift introduces significant 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. More than 90% of mid-size and large enterprises report that a single hour of downtime now exceeds $300,000 in cost.[3]

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. Administrators can initiate support cases directly from the HMC, with relevant first failure data capture (FFDC) logs automatically gathered and transmitted to IBM Support.

Use case: Automated support workflow

Consider a system administrator at a large insurance company running critical backend applications on IBM Power infrastructure. During a routine maintenance task involving Live Partition Mobility (LPM), the administrator encounters a migration validation error. Previously, resolving such issues required manually opening a support case, coordinating with IBM Support, and collecting FFDC data, a process that might take 5–8 hours before investigation can begin.

With Power11, the administrator can open a support case directly from the HMC. The system automatically collects and submits the necessary diagnostic data, eliminating manual steps

2 https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/transformation-index
3 https://uptimeinstitute.com/2019-data-center-industry-survey-results

and reducing turnaround time. IBM Support uses AI-enhanced diagnostic tools trained on Power processor-based 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 Smart energy scheduling options

In recent years, sustainability has become a central strategic priority in the IT industry because of regulatory pressure and business value. Over the past five years, climate commitments among companies have increased by 800%, indicating broader awareness and recognition that sustainable practices can enhance profitability. Currently, 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 products.[4]

Although 95% of companies established operational sustainability goals, only 41% achieved measurable progress. Many companies face challenges in starting the implementation process. This situation indicates a need for more explicit guidance, enhanced governance frameworks, and improved integration of sustainability metrics into core business operations.

Power processor-based servers have consistently improved energy efficiency with each generation, and Power11 continues this trend. Power11 enhances efficiency by reducing energy use, carbon emissions, and the data center energy footprint.

Clients upgrading from Power9 to Power11 can reduce energy consumption by up to 60% for the same performance.[5] Power11 outperforms x86 systems, providing twice the performance per watt. Compared to Power10, Power11 offers up to 33% better performance per watt, with the S1022 showing the largest gain.

Figure 1-6 on page 25 illustrates natural IT efficiency improvements from Power9 to Power11, offering up to a 3:1 consolidation ratio, reducing energy use and carbon emissions by approximately 60%, and cutting data center footprint by two-thirds.

[4] https://www.pwc.com/us/en/services/esg/library/assets/pwc-sustainability-decarbonization-2025.pdf
[5] https://techchannel.com/power11/server-models/

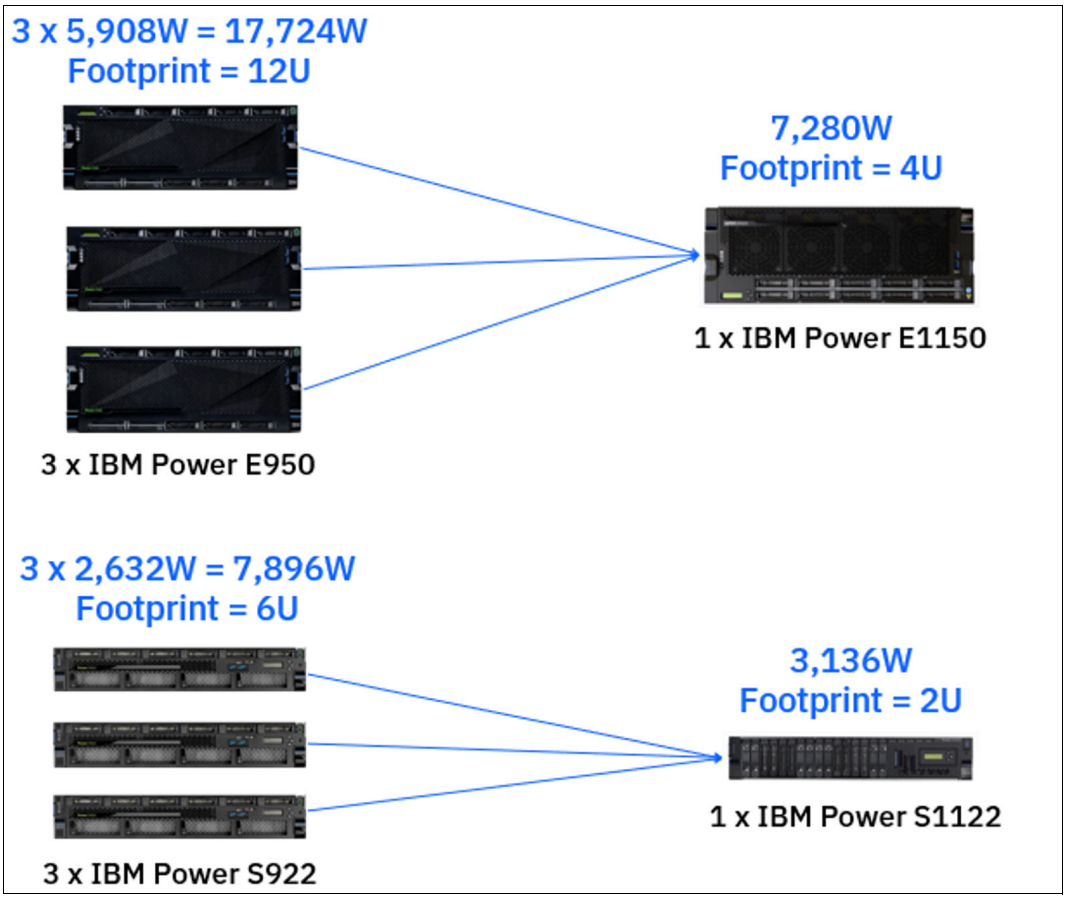

Figure 1-6 Power9 to Power11 energy consumption reduction

Along with the base energy savings delivered with Power11, there are some additional features that are introduced with Power11 to address some of the other feedback received from clients.

Partition level monitoring

In mid-2024, a new monitoring and reporting capability was introduced through HMC to track energy usage, carbon emissions, and other environmental factors. Power11 enhances this capability by providing insights at the partition level as shown in Figure 1-7 on page 26.

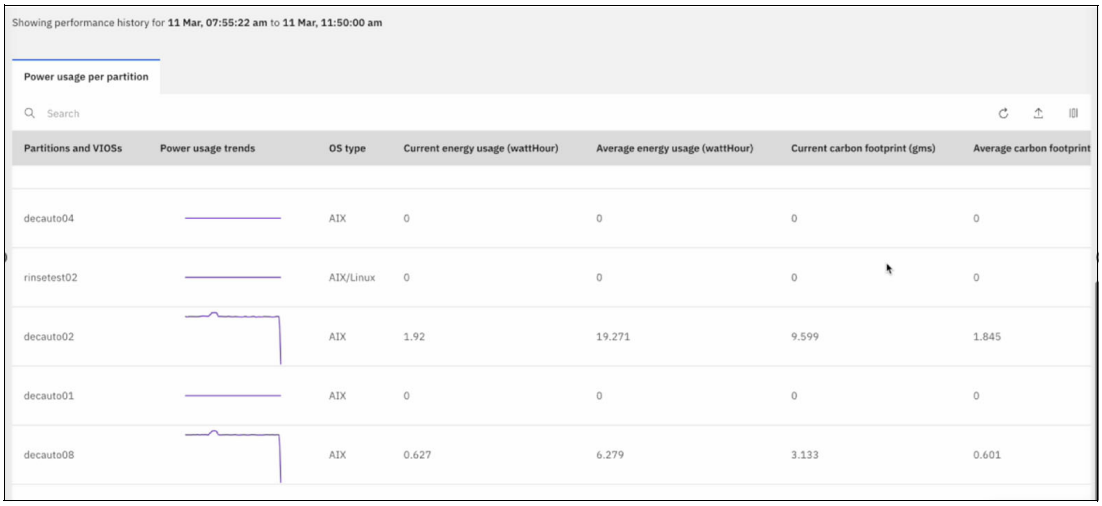

Figure 1-7 Partition level monitoring

Energy Efficient mode

Power11 introduced a new Energy Efficient mode that, when enabled, can reduce energy consumption by up to 30%, with only an approximate 10% reduction in performance. This impact varies depending on the system and configuration. The power mode can be dynamically established through the HMC GUI, as illustrated in Figure 1-8.

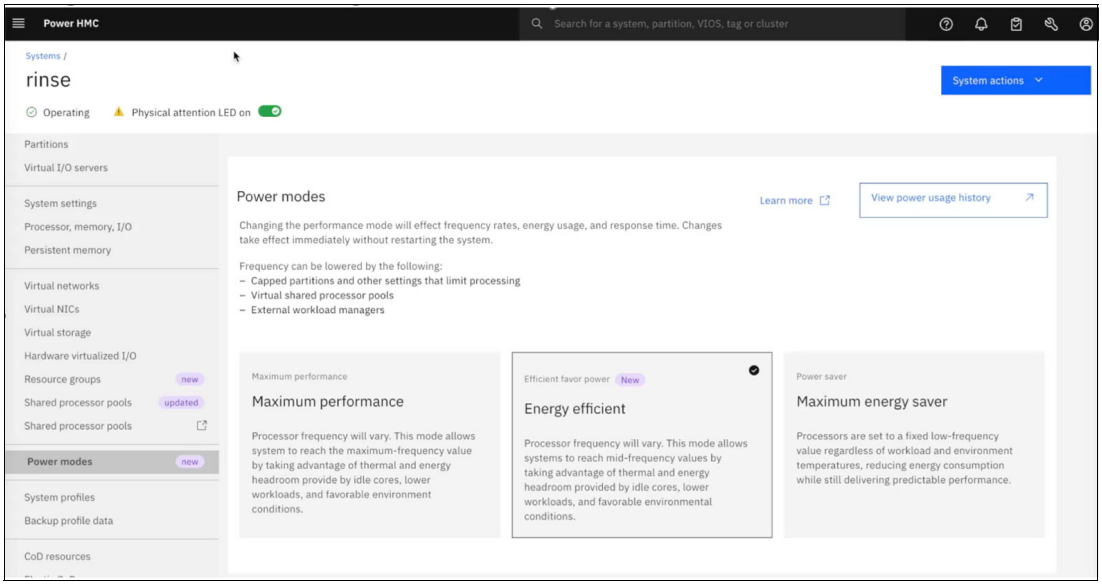

Figure 1-8 Power modes

In the environmental dashboard, it is observed that in maximum performance mode, energy consumption remains almost constant even when core usage decreases.

Figure 1-9 shows the energy usage when the processor is set to maximum performance.

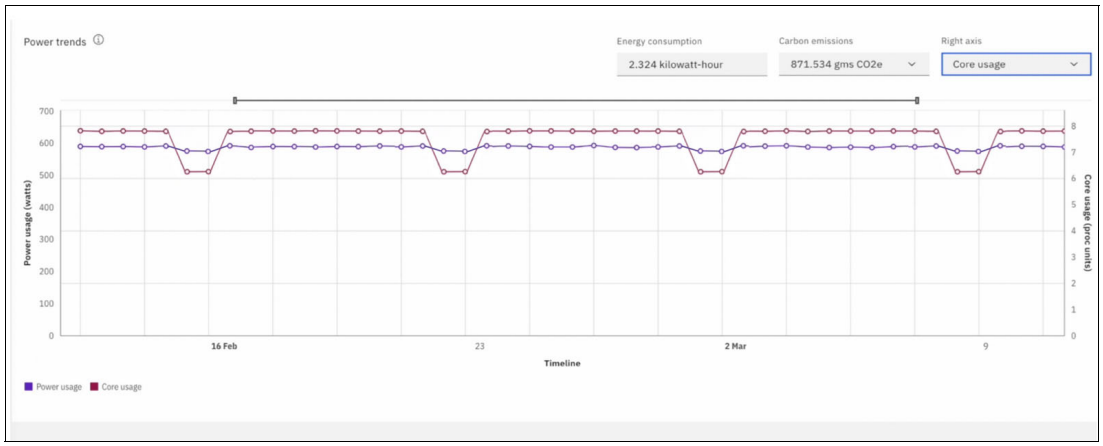

Figure 1-9 Energy consumption in max performance mode

Switching to Energy Efficient mode, which is implemented instantaneously, results in only a 10% reduction in performance, while significantly decreasing energy consumption by up to 30%, as illustrated in Figure 1-10.

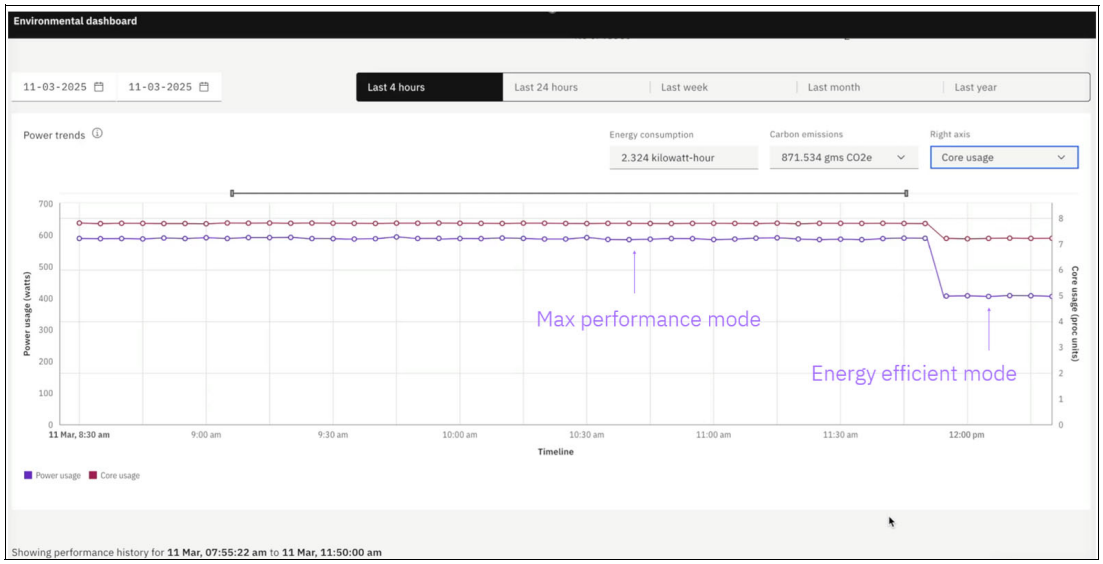

Figure 1-10 Energy Efficient mode

Power mode scheduling

If there is a pattern in system use, you can implement the scheduling power mode feature. This enables the system to activate Energy Efficient mode during designated times.

As illustrated in Figure 1-11, the Energy Efficient mode will be engaged during the weekend, starting from 8:00 PM on Friday until 5:00 AM on Monday.

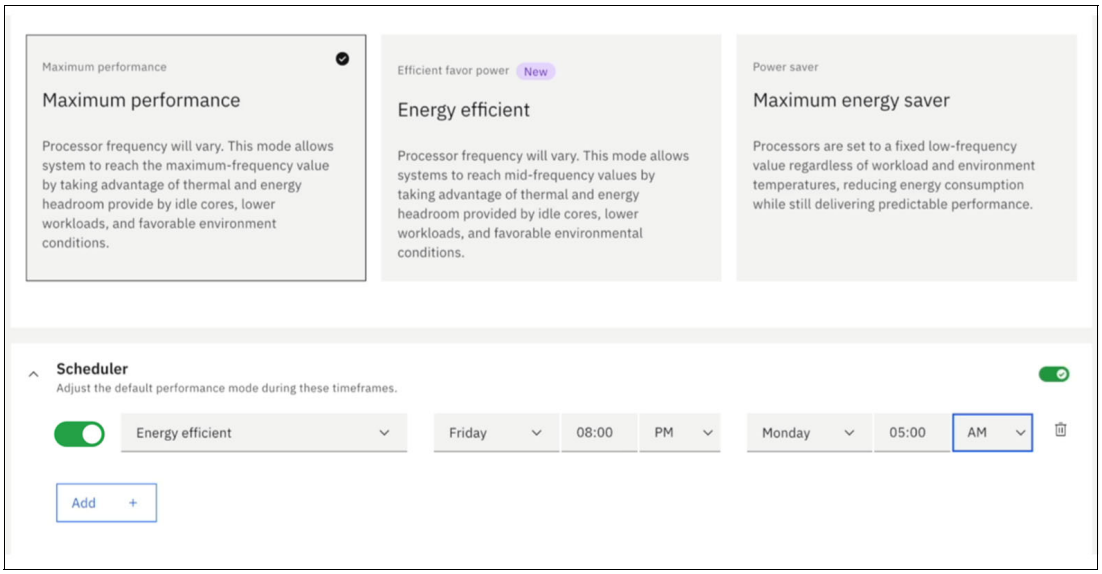

Figure 1-11 Power mode scheduling

The new environmental dashboard enables energy savings whenever potential opportunities arise. The automated scheduling capabilities allow the system to transition into and out of various modes efficiently as needed.

1.4.7 Up to 25% improvement for consolidation of Oracle 19c workloads

At the time of writing, Oracle Database 19c is the latest long-term release and version 19.12 is the latest Release Update for the AIX operating system. Oracle Database 19c includes many features over its previous database versions. Oracle Database versions for IBM AIX are compatible with later versions of AIX Technology Levels (TL) within the same major version. For certification information of Oracle Database with IBM AIX, see Oracle Support in the Certification section.

Oracle Database Instant Client has been supported on Linux on Power (32-bit and 64-bit) since Oracle Database 11g Release 2 (11.2.0.4). The Instant Client version of the Oracle Database is also supported on the 19c version of the database. Oracle Database 12c Release 1 introduced support for little-endian for the Instant Client running on Linux on Power, and 12c Release 2 supports only little-endian.

Oracle Database Instant Client is a lightweight version of Oracle Database Client binaries. It is available on Linux on IBM Power to allow an application deployed on Linux on Power to connect to the Oracle Database regardless of the platform on which the database is deployed.

Although servers based on IBM Power8® processor technology have reached end of support, many customers have not yet completed a technology refresh. IBM has demonstrated and published that migrating Oracle Database workloads from a server with IBM Power8 processors to one with IBM Power11 processors can reduce processor core requirements by half. The reduction in required cores provides a potential reduction in total cost of ownership (TCO) because it might be possible to reduce the number of required licenses.

The IBM Power11 processor introduces significant performance advancements, enabling up to 25% improvement in the consolidation of Oracle Database 19c workloads compared to the previous IBM Power10 processor generation. Because of this performance gain, organizations can further optimize infrastructure efficiency, reduce software licensing costs through better core usage, and improve overall system throughput, all while maintaining enterprise-grade reliability, availability, and serviceability (RAS) features.

Oracle is committed to release their latest products on IBM Power. To view the Oracle statement about support for 23ai on IBM Power, see Power Global community blogs.

1.4.8 Day 1 availability of Power11 in PowerVS with Hybrid Billing

As hybrid cloud adoption reaches 82% among enterprises and the number of applications is projected to surpass one billion by 2028, organizations are under pressure to manage growing complexities, skill shortages, and escalating infrastructure costs. IBM Power11, combined with IBM Power Virtual Server (PowerVS), offers a compelling solution to these challenges by delivering a secure, scalable, and high-performance hybrid cloud platform purpose-built for mission-critical workloads.

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

Using Power11 in PowerVS also unlocks advanced enterprise-grade security features, which include quantum-safe firmware signing, encrypted Live Partition Mobility (LPM), and a significantly fewer number of hypervisor vulnerabilities (CVEs) compared to x86 alternatives.

These capabilities are essential for organizations operating in regulated industries or managing sensitive data. Also, PowerVS integrates natively with IBM watsonx services through IBM Satellite® Connector, enabling enterprises to seamlessly infuse AI into their workloads and accelerate digital transformation initiatives.

PowerVS supports three generations of IBM Power technology:

- Power9:
  - Soft landing for Power9 clients moving to cloud.
  - Substantial support life remaining (through October 2028).
  - Least cost.
  - Special pricing on existing capacity.
  - Clients can continue to provision and use Power9.
- Power10:
  - Workhorse for production applications such as SAP RISE.
  - Broad capacity availability across data centers.
  - Excellent price/performance for AIX and IBM i workloads.
  - Clients can continue to provision and use Power10.
- Power11:
  - Immediate access to Power11 via PowerVS without waiting for on-premises delivery
  - Provision in less than 10 minutes
  - Ideal for development, testing, and application qualification
  - First cloud environment to provide Power11 in both IBM Cloud and on-prem (private)

From a business value perspective, PowerVS delivers up to 25% faster SAP workload migration compared to x86 platforms because of architectural consistency and superior compute performance. Flexible consumption models, shared processor pools, and tiered storage options help reduce total cost of ownership (TCO). Automated disaster recovery, secure backup, and nondisruptive maintenance enhance operational efficiency. The platform also supports regulatory compliance with standards such as SOC 2, HIPAA, and PCI DSS.

The following list provides some common use cases for PowerVS:

- Application modernization through access to more than 250 IBM Cloud services
- Enhancing business resiliency with warm disaster recovery environments
- Optimizing data center operations by consolidating or exiting physical infrastructure
- Supporting SAP transformation initiatives with highly available and scalable solutions for large-scale HANA deployments

As of this writing, PowerVS is available in 22 IBM Cloud data centers globally for public cloud use and can also be deployed in dedicated private cloud environments at client sites to meet stringent regulatory or operational requirements. Power11 is initially available in four global data centers and is enabled for automated disaster recovery operations. This broad availability helps ensure that organizations can adopt Power11 in a way that aligns with their specific business goals and compliance mandates.

1.4.9 Off-chip AI acceleration with IBM Spyre

IBM Spyre™ accelerator card (Figure 1-12) 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 to be available on the Power E1150 in the fourth quarter of 2025.

The card uses a system-on-a-chip architecture optimized for enterprise AI workloads, with 32 low-power AI cores, and supports multi-precision for inference and training: FP16/8 and INT8/4. The card is enabled for foundation models and is planned to support the Red Hat OpenShift Container Platform AI. It will also support popular AI frameworks and libraries such as PyTorch and vLLM. Implemented in 5 nm technology, the IBM Spyre card offers high performance and low-power characteristics, consuming only 75 watts of power.I

Figure 1-12 Spyre card physical view

Because each IBM Spyre card operates at only 75 watts, it significantly reduces the power footprint compared to traditional GPU-based acceleration, which often exceeds 300 watts per device.

The IBM Spyre card is planned to be supported inside the accelerator expansion drawer only. Each accelerator expansion drawer supports a maximum of eight Spyre cards. Accelerator expansion drawer support on the E1150 includes the following details:

- At initial availability, a maximum of one accelerator expansion drawer will be supported on a Power11 processor-based server
- Support for additional expansion drawers is planned in 2026

Note: The Spyre card 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 itself, the Spyre card is designed to access and use the system's main DDR5 memory modules. This design enables the use of terabytes of high-performance memory available within the Power11 processor-based server. This architectural choice provides the following two advantages:

1. Reduced latency for data access because of the shared memory model
2. Elimination of memory duplication, avoiding the need to copy data between host and device memory, as is common with discrete GPUs

The use of system RAM simplifies software development and memory management and contributes to a more energy-efficient and thermally balanced system architecture, an increasingly important factor in modern data centers and edge environments. This design makes the Spyre card particularly well suited 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.

In the fourth quarter of 2025, the Spyre card general availability and support on Power11 processor-based servers includes the minimum supported software stack levels:

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

See Chapter 4, "Artificial intelligence support" on page 93 for more detailed information about use cases and solutions with Spyre card.

1.4.10 Red Hat OpenShift with support for expanded software ecosystem

Red Hat OpenShift on Power provides a robust, enterprise-grade container platform that enables financial institutions to modernize applications, deploy AI workloads, and integrate with a broad variety of software. By running Red Hat OpenShift on Power11 systems, organizations can take advantage of the platform's performance, scalability, and security while maintaining architectural consistency across hybrid cloud environments.

One of the key advantages of Red Hat OpenShift on Power is its support for a greater range of software. This expanded range includes the following components:

- IBM Cloud Paks such as Cloud Pak for Data and Cloud Pak for Integration
- Red Hat Runtimes such as Quarkus and Spring Boot
- A wide range of open-source AI and machine learning (ML) frameworks, including PyTorch, TensorFlow, and ONNX

Additionally, many independent software vendor (ISV) applications are now certified for the Power architecture, so financial institutions can run containerized workloads alongside traditional AIX or IBM i systems. This colocation reduces latency, improves data gravity, and can simplify integration.

Red Hat OpenShift on IBM Power also supports seamless hybrid cloud integration, particularly when deployed on IBM Power Virtual Server (PowerVS). This integration enables 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 in which AI models can 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 optimized for Power. Infrastructure-as-code templates using 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, financial institutions can build and scale AI-driven applications with confidence, agility, and enterprise-grade resilience.

1.5 IBM Power11 processor architecture

The Power11 processor is designed to deliver higher clock speeds and adds up to 25% more cores per processor chip than comparable IBM Power10 systems. The Power11 processor builds upon the key capabilities introduced with Power10, including stronger reliability, availability, and serviceability (RAS) characteristics, better energy efficiency, improved energy management, and enhanced quantum-safe security.

Other than the processor itself, IBM is introducing improvements to packaging, memory architecture, and AI acceleration:

- Packaging innovation

  The Power11 processor uses integrated stack capacitor (ISC) technology and advanced 2.5D packaging, along with innovations in cooling such as improved heat sinks and more efficient fans. These enhancements optimize energy delivery, improve thread and core strength, and increase system capacity.

- Enhanced system architecture

  The memory architecture for Power11 systems is based on the recently released DDR5 DDIMMs and enhanced Open Memory Interface (OMI) interfaces, which enable improved memory reliability, capacity, and bandwidth. Because OMI is not technology-specific, the Power11 portfolio also supports OMI DDR4 memory migrated from Power10 high-end systems, enabling organizations to protect their investments in memory technology.

- AI acceleration

  IBM continues to support a range of emerging enterprise AI use cases with the Matrix Math Assist (MMA) architecture. Improvements to Power11 processor core strength and system capacity enhance MMA performance for inferencing workloads. Furthermore, IBM intends to incorporate the IBM Spyre accelerator into Power11 offerings to provide additional AI inferencing capabilities. IBM Power processors and the IBM Spyre accelerator will work together to enable next-generation infrastructure that can scale demanding AI workloads for businesses.

1.5.1 Power11 processor details

The Power11 includes changes to packaging, computational capabilities, and the data plane:

- Technology and packaging:
  - 654 mm^2 7 nm refined Samsung (30 Billion Devices)
  - Advanced 2.5D packaging technology optimizes space and improve signal integrity.
  - Integrated Stack Capacitor technology provides enhanced power delivery and performance by integrating capacitors directly into the chip package
  - 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 (low latency non-uniform cache architecture management)
  - Enterprise performance focus:
    - 3× the core performance relative to Power9®
    - 2× the thread strength relative to Power9
    - 4× L2 cache, 4× MMU per core relative to Power9
    - 4× crypto engine per core relative to Power9
  - AI computational focus:
    - 2× general SIMD per core relative to Power9
    - 4× matrix SIMD per core relative to Power9
    - New AI Instructions and data types
- Robust data plane:
  - 2 TB/s raw (32 GT/s) PowerAXON + OMI signaling
  - SMP interconnect for up to 16 sockets
  - 2× OMI memory bandwidth relative to Power10
  - 64 TB OMI DDR large system memory capacity

Figure 1-13 on page 34 shows the Power11 processor. Each of the cores has an L1 cache size with 64 KB of data cache and 96 KB of instruction cache.

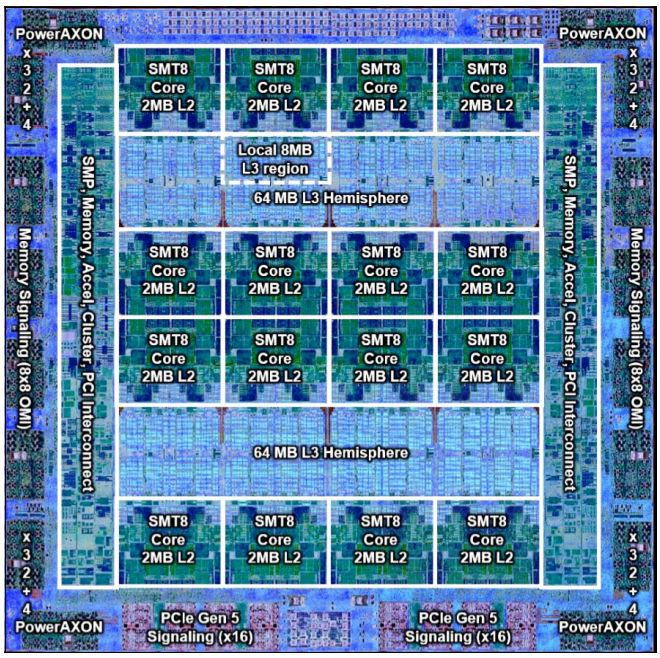

Figure 1-13 Power11 processor diagram

Comparing Power11, Power10, and Power9 processors

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

Table 1-5 Comparison of the Power11 processor technology to prior processor generations

Characteristics	Power11	Power10	Power9
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
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 / 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 SRAMa	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

Characteristics	Power11	Power10	Power9
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 changed its methodology for counting transistors to closer to industry standard.
b. DDR4 only as migration from Power10 during same serial number upgrade
c. DDR5 support added after GA

Chip packaging

To provide flexibility and scalability, the Power11 chip is packaged in multiple configurations.

Packaging options

Configurations include the following packaging options:

- Single chip module (SCM): Used in the high-end E1180, the SCM provides a single chip per socket and supports up to 16 sockets per system.
- Dual chip module (DCM): Used in the midrange E1150, the DCM includes two chips per socket and supports up to four sockets per system. The DCM version is also used in the scale-out server line in the S1122 and S1124, with up to two sockets per system.
- Entry single core module (eSCM): A modified DCM in which the second core in the module does not have processor capability and adds only I/O functionality. The eSCM is used in the entry-level processor features of the S1122, including the 4-core and 10-core options.

The Dual chip module in the E1150

The IBM Power E1150 uses the advanced IBM Power11 processor architecture integrated as a DCM. This DCM configuration allows each socket to house two processor chips, significantly increasing core density and performance. In the Power E1150, the DCM supports configurations with a range of 16– 30 cores per socket, enabling high-throughput computing for enterprise workloads.

This architecture is particularly beneficial for environments that require robust virtualization, large-scale database processing, and AI model training, as it delivers enhanced parallel processing capabilities and energy efficiency. The DCM enables the system to handle demanding workloads while maintaining a compact footprint, making it ideal for data centers with space and power constraints.

The integration of open memory interface (OMI) and differential DIMMs (DDIMMs) further enhances memory performance and reliability, which is critical for mission-critical applications. The E1150 uses the DCM to support IBM's broader goals of flexibility and efficiency in hybrid cloud environments.

The Power11 chips in the DCM not only boost raw performance but also support advanced features such as transparent memory encryption and dynamic resource allocation. These capabilities are essential for modern IT infrastructures that demand both security and adaptability. The DCM-equipped Power E1150 provides a solid foundation for scalable, secure, and high-performance computing in both traditional data centers and hybrid cloud environments.

Figure 1-14 shows the logical layout of the Power11 DCM.

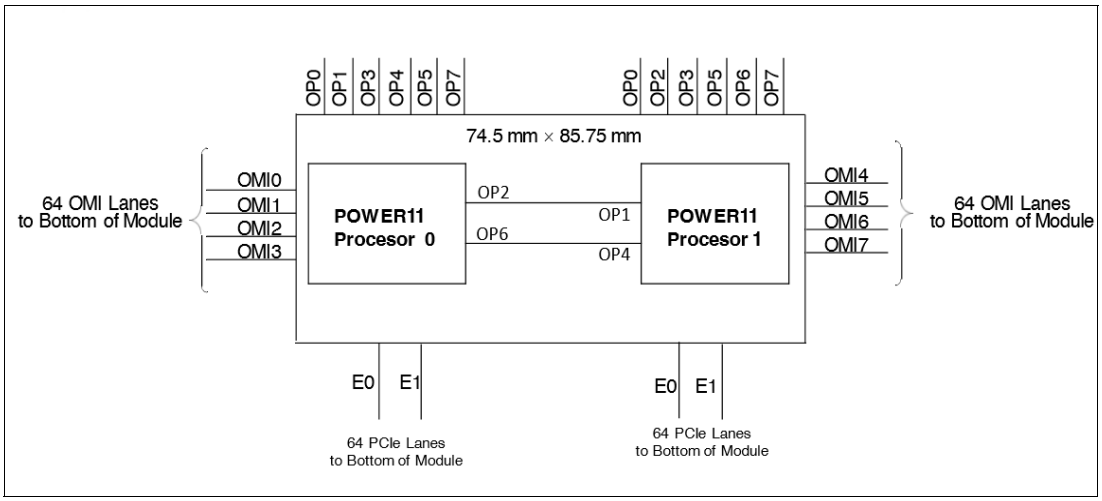

Figure 1-14 Power11 DCM Diagram

1.5.2 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 automatically deconfigured it, and available spare capacity (either unused or unlicensed) was used to maintain system processing capacity. This seamless failover occurred without requiring a system or application restart, ensuring continued operation. This feature remains available in the Power11 processor.

Because the substituted resources were not guaranteed to be on the processor chip with the failing core, the spare might not share the same affinity as the original core. In addition, because the spare core was not a dedicated resource, there was a possibility that the core could eventually be needed or licensed. Therefore, the processor containing the suspect core was flagged for repair. As a result, a planned system outage was typically 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, organizations might decide not to take the system down to repair the processor module because full processor capacity and performance are still maintained for these events.

With Power11, spare cores (one spare in each Single Chip Module [SCM] or four spares per E1180 node and two spares in each Dual Chip Module [DCM]) are allocated on each processor module. The exceptions include the 16-core SCM, which has all cores allocated for client usage, and the eSCM modules, where no spares are allocated. These cores are used only to substitute for a failing core on the same processor module, eliminating affinity issues and ensuring that the spare core would not otherwise be used. As a result, predictive processor core failures in Power11 do not require system downtime to repair the processor module.

1.6 Operating system support

The Power11 processor-based Power E1150 supports the following families of operating systems:

- AIX
- Linux
- Red Hat OpenShift

In addition, the Virtual I/O Server (VIOS) can be installed in special partitions that provide virtualization of I/O capabilities, such as network and storage connectivity. Multiple VIOS partitions can be installed to provide support and services to other partitions running AIX, or Linux, such as virtualized devices and Live Partition Mobility capabilities.

For more information about the Operating System and other software that is available on Power, see IBM Power. The minimum supported levels of IBM AIX and Linux are described in the following sections.

Important: At the time of writing, IBM i is not supported on the Power E1150 server. The operating system support applies to only the specified versions.

IBM periodically announces support for newer versions of its operating systems. To take full advantage of the latest security enhancements, performance improvements, and new feature support, it is strongly recommended that you keep your operating systems up to date.

1.6.1 AIX operating system

Support for the AIX operating system is described in this section.

The Power E1150 server supports the following minimum levels of the AIX operating system when installed with virtual I/O:

- AIX Version 7.3 with Technology Level 7300-02 and Service Pack 7300-02-02-2420
- AIX Version 7.3 with Technology Level 7300-03 and Service Pack 7300-03-00-2446
- AIX Version 7.2 with Technology Level 7200-05 and Service Pack 7200-05-08-2420

The Power E1150 servers support the following minimum levels of the AIX operating system when installed by using direct I/O connectivity:

- AIX Version 7.3 with Technology Level 7300-02 and Service Pack 7300-02-04
- AIX Version 7.3 with Technology Level 7300-03 and Service Pack 7300-03-01
- AIX Version 7.3 with Technology Level 7300-04 and Service Pack 7300-04-00
- AIX Version 7.2 with Technology Level 7200-05 and Service Pack 7200-05-10

Important: In addition, consider the following requirements:

- 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 Linux operating system

Support for the Linux operating system is described in this section for Power E1150. The listed Linux distributions are supported on the Power E1150 server model. Other distributions, including

open source releases, can run on these servers, but do not include any formal Enterprise Grade support.

Red Hat Enterprise Linux

Red Hat Enterprise Linux (RHEL) version 9.6 and any version of RHEL 10 is supported in native Power11 mode, which allows it to access all of the features of the Power11 processor and platform.

At announcement, the Power E1180 servers support the following minimum levels of the Red Hat Enterprise Linux operating system:

- Red Hat Enterprise Linux 8.10 for Power Language Environment, or later (Power10 compatibility)
- Red Hat Enterprise Linux 9.4 for Power LE, or later (Power10 compatibility)
- Red Hat Enterprise Linux 9.6 for Power LE, or later (Power11 native)
- Red Hat Enterprise Linux 10 for Power LE, version 10.0, or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 9 for Power LE version 9.6, or later

Red Hat OpenShift Container Platform

The Red Hat OpenShift Container Platform is supported to run on IBM Power servers, including the Power E1150. To run Red Hat OpenShift Container Platform in native mode on Power11, you must run Red Hat OpenShift Container Platform 4.19, which is based on Red Hat Enterprise Linux 9.6. Previous versions run in Power10 compatibility mode.

The Power S1122 and S1124 servers support the following minimum levels of the operating systems that are supported for Red Hat OpenShift Container Platform:

- Red Hat Enterprise Linux CoreOS 4.18 (Power10 compatibility)
- Red Hat Enterprise Linux CoreOS 4.19 (Power11 native)

SUSE Linux Enterprise Server

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

At announcement, the Power E1180 servers support the following minimum levels of the SUSE Linux Enterprise Server operating system:

- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 15, Service Pack 6, or later
- SUSE Linux Enterprise Server 15, Service Pack 6, or later

1.6.3 VIOS

The minimum required level of VIOS for the Power E1150 is VIOS 4.1.1.10 when running in Power11 mode. The following levels are supported:

- VIOS 4.1.1.10 (Service Pack)
- VIOS 4.1.0.40 (Service Pack) (P10 compatibility mode)
- VIOS 3.1.4.60 (Service Pack) (P9 compatibility mode).

1.7 Firmware and Hardware Management Console

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

Known as the Power Firmware, it includes components like the Flexible Service Processor (FSP), PowerVM Hypervisor (PHYP), and system boot code. This firmware enables advanced features such as logical partitioning (LPARs), dynamic resource allocation, and secure boot processes. It also supports system diagnostics, error reporting, and remote management capabilities, especially when integrated with tools like the Hardware Management Console (HMC). Regular firmware updates from IBM ensure performance improvements, security patches, and compatibility with new hardware or software features.

An HMC is a dedicated appliance to configure and manage system resources on IBM Power. It can be delivered as a physical server or as a virtual appliance running in either PowerVM or an x86 VM. A graphical user interface (GUI), a command-line interface (CLI), or REST API interfaces are available.

1.7.1 HMC Requirements

Supported HMCs are the 7063-CR2 and the Virtual HMC appliances.

The minimum required HMC version for the Power E1150 is V11R1 M1110, which is supported on the 7063-CR2 and the Virtual HMC appliances only.

Restriction: The following restrictions apply:

- The 7063-CR1 and 7042 HMCs are not supported as an HMC for Power11 servers.
- The HMC code at V11R1M1110 cannot manage POWER7 processor-based systems

1.8 Rack support

The Power E1150 server fits a standard 19-inch rack. The server is certified and tested in the IBM Enterprise racks (7965-S42 and 7014-T42). Customers can choose to place the server in other racks if they are confident that those racks have the strength, rigidity, depth, and hole pattern characteristics that are needed. Contact IBM Support to determine whether other racks are suitable.

It is a best practice that the Power E1150 server be ordered with an IBM 42U enterprise rack #ECR0 (7965-S42). This rack provides a high-quality environment for IBM Manufacturing system assembly and testing, and provides a complete package.

If a system is installed in a rack or cabinet that is not from IBM, ensure that the rack meets the requirements.

Important: The customer is responsible for ensuring 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 you order 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 to purchase with a Power E1150 server, but the Power E1150 installed in 7014-T42 is supported.

Vertical PDUs: All PDUs that are installed in a rack that contains a Power E1150 server must be installed horizontally to allow for cable routing in the sides of the rack.

For more information on supported racks, see section 2.3, “Racks” on page 58.

2

Power11 E1150

The IBM Power E1150 is a next-generation enterprise server that is built on the IBM Power11 architecture. It is designed to deliver high performance, scalability, and efficiency for demanding workloads. The system offers a range of processor options, which include configurations with varying core counts and frequencies to match specific performance needs.

The system supports DDR5-based DDIMMs, which provide high memory bandwidth and capacity. It supports configurations that scale to handle large in-memory databases and analytics workloads. The E1150 is engineered for dense compute environments, and it includes advanced cooling and packaging technologies that enhance thermal efficiency and system reliability.

In terms of I/O and storage, the Power E1150 includes multiple PCIe Gen5 slots for high-speed connectivity to storage and networking. It also supports internal NVMe U.2 flash drives, which offer low-latency, high-throughput storage options for performance-critical applications. For environments that require additional expansion, the E1150 is compatible with external expansion drawers, such as the ENZ0 PCIe Gen4 I/O expansion drawer and the NED24 NVMe expansion drawer, which can house up to 24 NVMe devices. These options provide flexible scalability for both compute and storage, which makes the Power E1150 a robust platform for hybrid cloud, AI, and mission-critical enterprise workloads.

The following topics are covered in this chapter:

- Introducing the Power E1150
- Hardware components
- Racks

2.1 Introducing the Power E1150

The Power11 processor represents the culmination of years of architectural innovation and performance enhancements. It powers the IBM Power E1150, a midrange server designed as the successor to the Power E1050. Engineered for both on-premises and cloud deployments, the Power E1150 delivers high performance, flexible capacity, and built-in virtualization capabilities, making it ideal for modern enterprise workloads.

The Power E1150 (model 9043-MRU) features a compact 4U rack-mounted chassis and supports configurations with two, three, or four Dual Chip Modules (DCMs). Each DCM contains up to 30 cores, with additional spare cores available for resiliency and capacity upgrades. The system includes PowerVM virtualization and supports AIX and Linux partitions, with or without the use of the Virtual I/O Server (VIOS), offering robust flexibility for diverse IT environments.

Restriction: IBM i is not supported on the Power E1150.

The base E1150 is configured with two sockets installed. Additional available configurations include three sockets and four sockets. Systems with two and three sockets can later be expanded by adding additional DCMs to the configuration.

Three processor options are available for this server:

- Sixteen cores (+2 spares) running at 3.40–4.20 GHz (maximum) frequency range
- Twenty-four cores (+2 spares) running at a 3.05–4.15 GHz (maximum) frequency range
- Thirty cores (+2 spares) running at a 2.8–3.95 GHz (maximum) frequency range

A Power E1150 server with four 30-core DCMs offers the maximum of 120 cores, and processor cores can run a maximum of eight simultaneous threads to deliver greater throughput. All sockets must be populated with the same processor modules. Table 2-1 shows the number of cores and the number of threads available in the different configurations of the IBM Power E1150.

Table 2-1 Available cores and threads in the different E1150 configurations

Number of sockets	16 Core (EPEX)		24 Core (EPEY)		30 Core (EPEZ)	
	Number of cores	Number of threads	Number of cores	Number of threads	Number of cores	Number of threads
2	32	256	48	384	30	480
3	48	384	72	576	90	720
4	64	512	96	768	120	960

New features in the Power E1150

The Power E1150 includes a number of new features:

- Power11 DCM processor with 16, 24, or 30 cores per socket
- Spare cores for increased availability (2 spare cores per socket)
- Energy Efficient mode with enhanced scheduling capabilities
- 64 DDR5 DDIMM slots with 25% higher bandwidth that provides up to 16 TB max memory
- Quantum safe encryption for secure boot and LPM (Logical Partition Migration)
- Serial Number preserving upgrade support from Power10 systems (4Q 2025 Availability)
- The E1150 is powered by a Titanium class 2300W power supply (PSU).

Power11 introduced a faster memory bus, the Open Memory Interface (OMI), to the DDIMMs. The OMI now runs at max 32 Gbps (with DDR5) compared with the OMI memory channel in the E1050, which runs at 25.6 Gbps. The SMP xBus runs at 32Gbs.

Figure 2-1shows the front of the Power E1150.

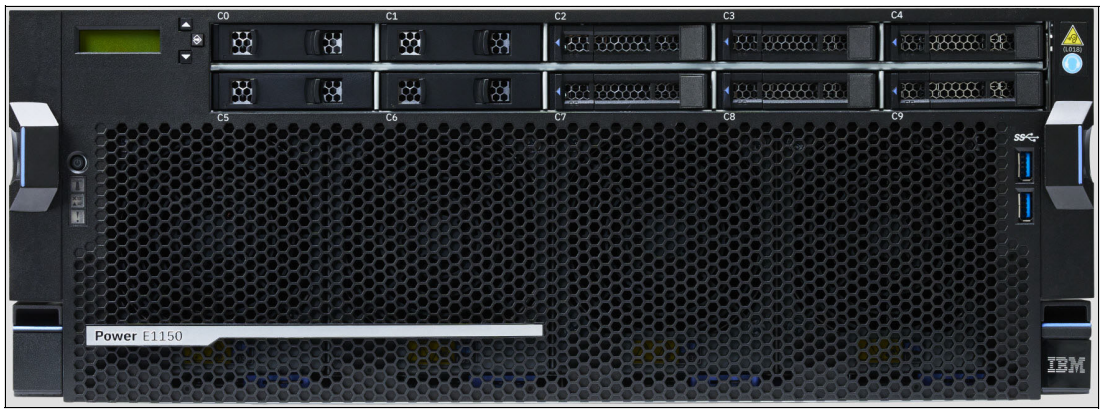

Figure 2-1 Front view of the Power E1150

Power E1150 compared to the E1050

With each Power Server evolution, new features and improvements are added,

- more cores
- more memory
- better speed for memory or buses

Table 2-2 compares E1050 to E1150, which shows the improvements provided by the Power 1150.

Table 2-2 Comparison between the E1050 and the E1150

Features	E1050	E1150
Processor module/max socket/max cores	P10 DCM/4S/96 Cores 74.5 mm x 85.75 mm; 5387 pin	P11 DCM/4S/120 Cores 74.5 mm x 85.75 mm; 5387 pin
Processor power	455W	545W
SMP xBus	xBus 32Gbs	xBus 32Gbs
Memory Slot/System Capacity	64 DDIMM, 4U System Max 16 TB	64 DDIMM, 4USystem Max 16 TB
Memory Bus	2666 Mbps or 3200 Mbps per OMI Channel 409 GB/s per CPU Socket (DCM)	4000 Mbps per OMI Channel 512 GB/s per CPU Socket (DCM)
Service Processor	eBMC	eBMC
System PCIe I/O	176 lanes Gen4 OR 64 lanes Gen5 and 64 lanes Gen4	176 lanes Gen4 OR 64 lanes Gen5 and 64 lanes Gen4
Cooling Fans and Cooling Domains	4 92-mm fans, 1 cooling domain	4 92-mm fans, 1 cooling domain
Processor VRM N+1 Phase Redundancy	Yes	Yes

Features	E1050	E1150
Memory Voltage N+1 Phase Redundancy	Yes	Yes
IO, Standby VRM N+1 Phase Redundancy	Yes	Yes
Number of PCIe slots	11	11
Concurrent Maintenance Serviceability Enhancements for PCIe Adapters	Rear Blind Swap IO Cassette. Cable Bracket	Rear Blind Swap IO Cassette. Cable Bracket
Internal SAS Controllers/Cable	None	None
Internal Storage	10-NVMe PCIe Gen4 only	10-NVMe PCIe Gen4 only
Concurrent Maintenance HDD/SSD/NVMe/PSU/Fan	Yes	Yes
Power Supplies	4× Titanium Class 2300W	4× Titanium Class 2300W
System Depth and Cable Management	35.5” + Cable Bracket	35.5” + Cable Bracket

2.1.1 Server configurations

The Power E1150 is a single machine type-model 9043-MRU delivered in a 4U enclosure. The minimum system has two sockets populated with one of the three available processor feature codes as shown in Table 2-3, Additional configurations are available with three or four sockets populated. The two-socket system can be upgraded in the field to either three sockets or four sockets. The three-socket system can be upgraded in the field to a four-socket system.

Processor options

There are three processor feature codes available at the time of announcement. These features are differentiated by the number of cores and the frequency range of the cores. These characteristics are shown in Table 2-3.

Table 2-3 E1150 socket and core options

Feature Code	Feature Name	Processor Type	Max System Config	Frequency Range
EPEZ	30W P11 DCM	30 Cores +2 Spare	120 Cores	3.0 to 4.1 GHZ
EPEY	24W P11 DCM	24 Cores +2 Spares	96 Cores	3.25 to 4.2 GHZ
EPEX	16W P11 DCM	16 Cores +2 Spares	64 Cores	3.5 to 4.2 GHZ

SAP HANA support

There is no plan to release a separate SAP HANA model for the E1150. Instead, SAP HANA is sold on the E1150 with a special Linux price structure.

SAP HANA will be certified on the 24-core IBM Power11 processor, with support also planned for the 16-core, and 30-core Power11 configurations. IBM intends to support SAP HANA in production environments on the IBM Power System E1150 (9043-MRU), following the official certification of the platform. This support will include specific Linux operating systems that meet SAP’s certification requirements, ensuring optimized performance and reliability for SAP HANA workloads on Power11 infrastructure.

- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 15 Service Pack 6 or later.
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 9.6 for Power LE, or later.

Memory configuration options

There are several memory feature options available in the E1150, as shown on Table 2-4

Table 2-4 E1150 memory options

Feature Code	Feature Name	DDIMM Size	Max System Config	Frequency	Peak Bandwidth /Socket
EM5P	64 GB DDR5	2× 32 GB DDIMM	2 TB	4000 MHZ	512 GBS
EM5Q	128 GB DDR5	2× 64 GB DDIMM	4 TB	4000 MHZ	1024 GBS
EM5R	256 GB DDR5	2× 128 GB DDIMM	8 TB	4000 MHZ	1024 GBS
EM5S	512 GB DDR5	2× 256 GB DDIMM	16 TB	4000 MHZ	1024 GBS

2.1.2 RAS features

RAS stands for Reliability, Availability, and Serviceability. These features are designed to help ensure that a system operates consistently, remains available when needed, is easier to maintain, minimizes downtime, and enhances overall performance.

System reliability begins at the component level. Devices and subsystems are engineered for high dependability and undergo rigorous verification and integration testing throughout the design and development phases. During manufacturing, each system is subjected to comprehensive testing to meet the highest standards of product quality.

The Power11 processor–based Power E1150 system includes an advanced suite of RAS capabilities that improve system uptime and streamline maintenance. One enhancement is the inclusion of spare cores within each Dual Chip Module (DCM). These spare cores are automatically activated when a core failure occurs, providing redundancy. This design enables uninterrupted system operation, reduces the need for immediate maintenance, and improves overall system resilience.

Power11 processor RAS

The Power11 processor supports instruction retry for transient errors and a core-contained checkstop for certain solid faults. The fabric bus design includes cyclic redundancy check (CRC) and retry capabilities. CRC codes verify data integrity on the bus and enable retry of faulty operations.

Cache availability

The L2 and L3 caches in the Power11 processor, located in the memory buffer chip, are protected by double-bit detect and single-bit correct error detection code (ECC). When a threshold of correctable errors is detected on cache lines, the system purges the affected data and removes the cache lines from operation without requiring a reboot in the PowerVM environment.

Modified data is handled through Special Uncorrectable Error (SUE) handling. L1 data and instruction caches support retry for intermittent errors and include a cache set delete mechanism for solid failures.

Special uncorrectable error handling

SUE handling prevents uncorrectable errors in memory or cache from immediately terminating the system. The system tags the affected data and determines whether to use it again. If the error is irrelevant, it does not trigger a checkstop. If the data is used, I/O adapters

controlled by an I/O hub controller freeze when data is transferred to an I/O device. Otherwise, termination can be limited to the program or kernel, or to data not owned by the hypervisor.

PCI extended error handling

PCI extended error handling (EEH)–enabled adapters respond to a special data packet generated by the affected PCI slot hardware. System firmware examines the affected bus, allows the device driver to reset it, and continues operation without requiring a system reboot. For Linux, EEH support extends to most commonly used devices, although some third-party PCI devices might not support EEH natively.

Uncorrectable error recovery

When the auto-restart option is enabled, the system can automatically restart after an unrecoverable software error, hardware failure, or environmentally induced failure such as AC power loss.

Spare core configuration

The system includes two spare cores per DCM, which maintain total processor core capacity after a predictive core failure.

Runtime processor diagnostics

This capability allows the PowerVM hypervisor to test processor cores without affecting production workloads. The service processor supports multiple operation modes, which can be selected through the Advanced System Management Interface (ASMI) menu.

Service processor

The Power E1150 system includes a redesigned service processor based on a Baseboard Management Controller (BMC) architecture. It provides firmware access through open-source, industry-standard APIs such as Redfish. The enhanced ASMI web interface maintains essential RAS functions and offers a more intuitive user experience. The BMC service processor provides the following functions:

- Error monitoring: Recoverable errors from the processor chipset are monitored by the system processor. Fatal errors are handled by the service processor.
- Independent operation: The service processor operates on a separate power boundary, allowing it to function independently of the system processor.
- System surveillance: It monitors connections to the Hardware Management Console (HMC) and system firmware (hypervisor), and supports remote power control, environmental monitoring, resets, restarts, and remote maintenance.

Additional RAS Capabilities

The Power E1150 system includes the following additional RAS capabilities:

- Open Memory Interface with differential DIMMs
- Enterprise-grade BMC for system management and service
- Active Memory Mirroring for the hypervisor
- Concurrent maintenance for NVMe drives and PCIe adapters
- Redundant, hot-plug cooling, and power systems
- Redundant voltage regulators
- Concurrent maintenance for time-of-day battery
- Lightpath diagnostics with enclosure and Field Replaceable Unit (FRU) LEDs
- Clear service and FRU labeling

- Installation by client or IBM
- Proactive support with automated call-home capabilities
- Flexible servicing options (client or IBM)

2.1.3 HMC requirements

The E1150 supports attachment to one or more HMCs or vHMCs for use when operating the system with PowerVM. This is the default configuration for servers supporting logical partitions with dedicated or virtual I/O. In this case, all servers have at least one logical partition.

Functions provided by the HMC

The HMC is either a hardware appliance or virtual appliance that can be used to configure and manage your systems. The HMC connects to one or more managed systems and provides capabilities for following primary functions:

- Systems management functions: Such as Power off, Power on, system settings, Capacity on Demand (CoD), Enterprise Pools, shared processor pools, performance and capacity monitoring, and starting Advanced System Management Interface (ASMI) for managed systems.
- Provide virtualization management capabilities including the creation, management, and deletion of LPARs; support for Live Partition Mobility (LPM) and remote restart; configuration of SR-IOV; management of virtual I/O servers; dynamic resource allocation; and access to operating system consoles.
- 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 Electronic Service Agent (ESA) for faster support.
- Provides appliance management capabilities for configuring network, administrating users on the HMC, and updating and upgrading the HMC.

Alternatives to the HMC on the E1150

An HMC is not required for the E1150, but is the recommended and default option for system management. There are two service strategies available if you choose not to implement an HMC in your environment:

- Full-system partition with PowerVM

  A single partition owns all the server resources and only one operating system can be installed. The primary service interface is through the operating system and the service processor.

- Partitioned system with NovaLink

  In this configuration, the system can have more than one partition and can be running more than one operating system. The primary service interface is through the service processor.

Supported consoles

Supported HMC options for the Power E1150 are the 7063-CR2 and the Virtual HMC appliances.

HMC 7063-CR2

The 7063-CR2 IBM Power HMC is a second-generation Power processor-based HMC.

The CR2 model includes the following features:

- 6-core Power9 130 W processor chip
- 64 GB (4×16 GB) or 128 GB (4×32 GB) of memory RAM
- 1.8 TB with RAID1 protection of internal disk capacity
- 4-port 1 Gb Ethernet (RH-45), 2-port 10 Gb Ethernet (RJ-45), two USB 3.0 ports (front side) and two USB 3.0 ports (rear side), and 1 Gb IPMI Ethernet (RJ-45)
- Two 900 W PSUs
- Remote Management Service: IPMI port (OpenBMC) and Redfish application programming interface (API)
- Base Warranty: Support for 1 year, 8am–5pm, Monday–Friday with available optional upgrades

A USB Smart Drive is not included.

The 7063-CR2 is compatible with flat panel console kits 7316-TF3[1], TF4[1], and TF5.

Virtual HMC

Originally, the IBM HMC was available exclusively as a physical hardware appliance, preinstalled with HMC firmware. However, IBM expanded its offerings to include a virtual HMC appliance. This virtual version can be deployed on PowerPC 64-bit Little Endian (ppc64le) architectures and on x86 platforms, providing greater flexibility for system administrators and data center environments.

Customers with a valid support contract can download the virtual HMC appliance from the Entitled Systems Support (ESS) website. Alternatively, it can be included as part of the initial order for a Power E1150 system.

The virtual HMC supports the following hypervisors:

- On x86 processor-based servers:
  - Xen
  - VMware
  - KVM
- On Power processor-based servers:
  - PowerVM

The following minimum requirements must be met to install the virtual HMC:

- 16 GB of memory
- Four virtual processors
- Two network interfaces (maximum 4 allowed)
- One disk drive (500 GB available disk drive)

For an initial Power E1150 order with the IBM configurator (e-config), the HMC virtual appliance can be found by selecting **Add software** → **Other System Offerings** (as product selections) and then selecting one of the following options:

- 5765-VHP for IBM HMC Virtual Appliance for Power V11
- 5765-VHX for IBM HMC Virtual Appliance x86 V11

For more information, see Installing and configuring the Hardware Management Console.

For more information about how to install the virtual HMC appliance and all requirements, see Installing the HMC virtual appliance.

[1] The 7316-TF3 and 7316-TF4 are withdrawn from marketing.

2.2 Hardware components

This section includes a detailed overview of the hardware components of the Power E1150 system, including its processors, memory, PCI slots, adapters, I/O drawers, and more.

2.2.1 Processors

The E1150 requires a minimum of two DCM processor modules. The system supports a maximum of four modules. You can choose the number of processor modules during the initial configuration. If you start with fewer than four modules, you can add modules later through a MES order. The MES upgrade requires scheduled downtime for installation.

All processor modules in a single server must use the same processor module Feature Code (F/C) number. Processor F/C cannot be mixed in the E1150 server. On the initial order and when adding additional modules with an MES, all processor types must be the same. MES replacements of existing features are not supported. Table 2-5 shows the processor options available.

Table 2-5 E1150 Processor Features

Feature	EPEZ	EPEY	EPEX
Package	DCM	DCM	DCM
#Cores Active per Module	30	24	16
Spare Cores per Module	2	2	2
Maximum Cores Per System	120	96	64
Max DDR speed	4000	4000	4000
Power Save Freq	2.4 GHz	2.4 GHz	2.4 GHz
Fixed Freq	3000	3300	3500
WOF Base Freq	3250	3550	3750
Max Freq with all cores active	4100	4200	4200
Ultra Turbo Freq	4100	4200	4200
Mod Max Power (@27C)	545 W	535 W	490 W
Processor Activation	EPUZ	EPUY	EPUT

Processor activations

Permanent CoD processor core activations are required for the first processor module in the configuration and are optional for the second, third, and fourth modules.

You can optionally use the temporary CoD capabilities of processor cores that are not permanently activated.

Note: An HMC is required for temporary CoD.

The E1150 can also be in a Power Enterprise Pool with other E1150 and E1050 systems when using Power Enterprise Pools with Shared Utility Capacity. Shared Utility Capacity on Power E1150 systems provides enhanced multisystem resource sharing and by-the-minute tracking and consumption of compute resources across a collection of Power E1150 and E1050 systems within a single Power Enterprise Pool that is running version 2.0. Shared Utility Capacity provides a complete range of 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 Power Enterprise Pools 2.0 of Power E1050 systems can add one or more Power E1150 systems into their pool and migrate to them at the rate and pace of their choosing. This is because any Power E1150 and Power E1050 systems can inter-operate and share compute resources within the same pool. Clients with Mobile Capacity on a Power E1150 can upgrade their system to support Power Enterprise Pools 2.0 to use Shared Utility Capacity resources.

Workload Optimized Frequency

All Power11 systems are shipped with Workload Optimized Frequency (WOF) enabled.

WOF is implemented for several reasons:

- Increasing peak performance
- Guaranteeing Turbo frequency
- Delivering Ultra Turbo frequency when possible
- Minimizing run-to-run and part-to-part variations

There are several factors that are used to determine the maximum CPU frequency

- CPU usage: Lighter workloads run at higher frequencies
- Number of active cores: Smaller number of active cores run at higher frequencies
- Environmental conditions: Lower ambient temperatures run at higher frequencies

WOF takes advantage of processor frequency and system capabilities, allowing fewer active workloads with fewer active cores to run at higher frequencies.

The frequency for a given workload is capped only by the maximum attainable frequency for the technology and the thermal design point (TDP) of the processor module supported by the system. Each processor module has its own unique best performance capability, different from other module, but product consistency is the goal.

The IBM Power11 architecture incorporates an advanced On-Chip Controller (OCC) that enables dynamic frequency scaling based on real-time system conditions. This controller continuously monitors key operational parameters, which include voltage levels, current draw, and the number of active processor cores. Using these inputs, the OCC references pre-calibrated lookup tables to determine the optimal voltage-frequency operating point for each processor module.

When system workloads are light, the OCC increases processor frequencies to enhance performance. Conversely, as workloads intensify or thermal and power constraints are approached, frequencies are scaled down to maintain system stability and efficiency. This dynamic adjustment ensures that each system operates within its optimal performance range, contributing to uniform performance across a distributed infrastructure.

By intelligently managing power and performance, the OCC not only improves workload responsiveness but also supports energy-efficient computing, which is an increasingly critical requirement in modern data centers.

2.2.2 Memory subsystem

The IBM Power E1150 server uses IBM's Open Memory Interface (OMI) for memory connectivity. OMI is a high-speed, low-latency memory interface developed by IBM to decouple memory technology from the processor interface. Unlike traditional memory interfaces that are tightly coupled to specific memory types (like DDR4 or DDR5), OMI provides a technology-independent approach. Therefore, IBM Power systems, such as the Power E1150, can support various memory technologies through Differential DIMMs (DDIMMs).

OMI provides several benefits:

- Flexibility: The ability to support multiple memory technologies on the same platform
- Scalability: Enables high memory bandwidth and capacity up to 16 TB in the Power E1150
- Performance: Reduces latency and increases throughput by using a high-speed serial interface
- Future-proofing: Simplifies the adoption of next-generation memory technologies without redesigning the processor-memory interface

OMI plays a crucial role in enhancing the performance and adaptability of IBM Power systems, especially in environments demanding high memory bandwidth and capacity, such as AI, analytics, and large-scale enterprise workloads.

The E1150 is designed to support DDR5 memory and the OMI interfaces are tuned for DDR5. Only DDR5 DDIMMs are offered on new systems. However, IBM intends to support the same serial number upgrades to existing Power10 based E1050 systems. In this case, the E1150 supports the integration of the existing memory on those E1050s, including existing DDR4-based DDIMMs.

Memory Capabilities

Each Power11 DCM has 16 OMI busses or channels, 8 from each of the two chips. Each OMI channel connects to one 4U DDIMM. The Power E1150 supports up to 64 OMI slots when all four processor sockets are populated. This configuration allows for a maximum memory capacity of 16 TB, providing substantial scalability for memory-intensive workloads. Table 2-6 lists the available memory features for the E1150.

Table 2-6 E1150 4U DDIMM offering

Feature Code[a]	DIMM Capacity	Size	DDR Type	DDR Speed	OMI Speed[b]
EM5P/EMGT[c]	32 GB	4U	DDR5	4000 MHz	32 Gb/s
EM5Q	64 GB	4U	DDR5	4000 MHz	32 Gb/s
EMFL	256 GB	4U	DDR5	4000 MHz	32 Gb/s
EM5S	256 GB	4U	DDR5	4000 MHz	32 Gb/s

a. Each feature code provides 2 DDIMMs.
b. DDIMM supports up to 4800 MHZ / 38.4 Gb/s OMI. The E1150 system design is limited to 4000 MHz / 32 Gb/s OMI
c. Feature code for healthcare solution

Minimum/Maximum Orderable System Memory Sizes

The following table lists the minimum and maximum of DDIMM and system memory sizes based on the number of sockets populated in the E1150.

Table 2-7 Memory minimums and maximums

DDIMM and Sys Info	2S	3S	4S
Min # of DDIMM	8	12	16
Max # of DDIMM	32	48	64
DIMM Type	4U DDIMM	4U DDIMM	4U DDIMM
Minimum DDIMM Size	32 GB	32 GB	32 GB
Maximum DDIMM Size	256 GB	256 GB	256 GB
Minimum System Memory	8 x 32 GB 256 GB	12 x 32 GB 384 GB	16 x 32 GB 512 GB
Maximum System Memory	32 x 256 GB 8192 GB	8 x 256 GB 12288 GB	64 x 256 GB 16348 GB

OMI bus speed vs DDIMM Sizes

DDIMM must be of the same size and speed behind each DCM socket. The actual OMI bus speed is set by firmware per DCM socket boundary based on the speed of the DDIMMs installed for that DCM socket.

Memory subsystem RAS

The Power11 processor-based E1150 system introduces a new 4U tall differential DIMM (DDIMM), which has an open CAPI memory interface known as OMI for resilient and fast communication to the processor. This new memory subsystem design delivers solid RAS as described below.

- Memory buffer: The DDIMM contains a memory buffer with key RAS features, including protection of critical data and address flows by using CRC, ECC, parity, a maintenance engine for background memory scrubbing and memory diagnostic. It also includes a Fault Isolation Register (FIR) structure, which enables firmware attention-based fault isolation and diagnostics.
- Open memory interface (OMI): The OMI between the memory buffer and processor memory controller is protected by dynamic lane calibration, as well as a CRC retry and recovery facility to retransmit lost frames to survive intermittent bit flips. A complete lane fail can also be survived by triggering a dynamic lane reduction from 8 to 4, independently for both up and downstream directions. A key advantage of the OMI is that it simplifies the number of critical signals that must cross connectors from processor to memory compared to a typical industry- standard DIMM design.
- Memory ECC: The DDIMM includes a robust 64-byte Memory ECC, with 8-bit symbols, capable of correcting up to five symbol errors (one x4 chip and one additional symbol), as well as retry for data and address uncorrectable errors. An x4 chip is a DRAM with a 4-bit data width.
- Dynamic row repair: To further extend the life of the DDIMM, the dynamic row repair feature can restore full use of a DRAM for a fault contained to a DRAM row, while the system continues to operate.
- Spare temperature sensors: Each DDIMM provides spare temperature sensors, such that the failure of one does not require a DDIMM replacement.

- Spare DRAMs: 4U DDIMMs include two spare x4 memory modules (DRAMs) per rank. These can be substituted for failed DRAMs during runtime operation. Combined with ECC correction, the 2 spares allow the 4U DDIMM to continue to function with 3 bad DRAMs per rank, compared to 1 (with single device data correct) or 2 (with double device data correct) bad DRAMs in a typical industry standard DIMM design. This extends self-healing capabilities beyond what is provided with dynamic row repair capability.
- Spare power management integrated circuits: 4U DDIMMs include power management integrated circuits such that the failure of one power management integrated circuit does not require a DDIMM replacement.

2.2.3 Pervasive memory encryption

The Power11 processor integrates a Memory Controller Unit (MCU) that serves as the interface between the on-chip SMP interconnect fabric and the OMI links. This architecture makes the MCU an ideal location for implementing memory encryption logic.

The Power11 on-chip MCU encrypts and decrypts all traffic to and from system memory using Advanced Encryption Standard (AES) technology. This ensures that data in memory is protected from unauthorized access, not only during operation but also when memory modules are removed from the system.

Supported AES encryption modes

There are two encryption modes built-in to the MCU.

AES XTS mode (for persistent memory)

- XTS (Xor–Encrypt–Xor Tweaked Codebook with Ciphertext Stealing) is a strong block cipher mode ideal for encrypting persistent memory.
- Persistent DIMMs retain data even when power is off, posing a risk if physically removed.
- AES XTS protects against data theft from stolen or serviced DIMMs by encrypting memory contents.
- Although persistent memory is not yet standard in IBM Power, AES XTS support is built in for future readiness.

AES Counter mode (for volatile memory)

- Counter (CTR) mode offers low-latency encryption, making it suitable for volatile memory.
- Although not as strong as XTS, CTR mode is optimized for performance and is effective against physical attacks, especially in cloud environments.
- Each Power10 or Power11 processor holds a 128-bit encryption key that is used by the MCU to encrypt data on DDIMMs connected via OMI links.
- The MCU cryptoengine is integrated into the data path, ensuring no performance degradation during encryption.
- Pervasive memory encryption using AES CTR is enabled by default and cannot be disabled through any administrative interface.

The pervasive memory encryption of the Power11 processor does not affect the encryption status of a system dump content. All data that is coming from the DDIMMs is decrypted by the MCU before it is passed onto the dump devices under the control of the dump program code. This statement applies to the traditional system dump under the OS control and the firmware assist dump utility.

The PowerVM LPM data encryption does not interfere with the pervasive memory encryption. Data transfer during an LPM operation uses the following general flow:

1. On the source server, the Mover Server Partition (MSP) provides the hypervisor with a buffer.
2. The hypervisor of the source system copies the partition memory into the buffer.
3. The MSP transmits the data over the network.
4. The data is received by the MSP on the target server and copied to the related buffer.
5. The hypervisor of the target system copies the data from the buffer into the memory space of the target partition.

To facilitate LPM data compression and encryption, the hypervisor on the source system presents the LPM buffer to the on-chip NX unit as part of the processing in step 2. The reverse decryption and decompress operation is applied on the target server as part of the process in step 4.

The pervasive memory encryption logic of the MCU decrypts the memory data before it is compressed and encrypted by the NX unit on the source server. The logic also encrypts the data before it is written to memory but after it is decrypted and decompressed by the NX unit of the target server.

Active memory mirroring

The Power E1050 server can mirror the Power Hypervisor code across multiple memory DDIMMs. If a DDIMM that contains the hypervisor code develops an uncorrectable error, its mirrored partner enables the system to continue to operate uninterrupted.

Active memory mirroring (AMM) is an optional feature (#EM81).

The hypervisor code logical memory blocks are mirrored on distinct DDIMMs to enable more usable memory. There is no specific DDIMM that hosts the hypervisor memory blocks, so the mirroring is done at the logical memory block level, not at the DDIMM level. To enable the AMM feature, the server must have enough available memory to accommodate the mirrored memory blocks.

In addition to the hypervisor code itself, other components that are vital to the server operation are also mirrored:

- Hardware page tables (HPTs), which are responsible for tracking the state of the memory pages that are assigned to partitions
- Translation control entities (TCEs), which are responsible for providing I/O buffers for the partition's communications
- Memory that is used by the hypervisor to maintain partition configuration, I/O states, virtual I/O information, and partition state

Activating or deactivating AMM is done through the HMC. You can view the current status or modify the status from the Memory Mirroring section of the General Settings window for the E1150.

After a failure occurs on one of the DDIMMs that contains hypervisor data, all the server operations remain active and the enterprise baseboard management controller (eBMC) service processor isolates the failing DDIMMs. The system stays in the partially mirrored state until the failing DDIMM is replaced.

Memory that is used to hold the contents of platform dumps is not mirrored, and AMM does not mirror partition data. AMM mirrors only the hypervisor code and its components to protect

this data against a DDIMM failure. With AMM, uncorrectable errors in data that is owned by a partition or application are handled by the existing special uncorrectable error (SUE) handling methods in the hardware, firmware, and OS.

SUE handling prevents an uncorrectable error in memory or cache from immediately causing the system to stop. Instead, the system tags the data and determines whether to use it again. If the error is irrelevant, it does not force a checkstop. If the data is in use, termination can be limited to the program, kernel or hypervisor that owns the data. If the data must be transferred to an I/O device, the system can instead freeze the I/O adapters managed by the I/O hub controller.

Active memory expansion

Active memory expansion (AME) is an advanced memory optimization feature available on IBM Power systems that are running the AIX operating system. By using AME, the effective memory capacity of a system can exceed its physical memory by compressing memory contents in real time. This can result in memory expansion of 100% or more, enabling a partition to support more users or workloads without requiring additional physical memory. It also allows a server to run more logical partitions (LPARs), increasing overall system efficiency.

AME operates by using CPU resources to compress and decompress memory data. The effectiveness of this tradeoff between memory capacity and processor cycles depends on how compressible the memory content is and whether there is sufficient spare CPU capacity available. To improve this process, Power11 processors include a dedicated hardware accelerator that enhances AME performance while reducing the load on CPU cores. This accelerator benefits from Power11's higher memory bandwidth and lower latency, making AME more efficient.

Administrators have control over AME at the partition level. Each AIX partition can independently enable or disable AME, and configuration parameters allow users to specify the preferred level of memory expansion. Enabling AME on a partition requires an initial program load (IPL). When it is active, you can monitor AME performance by using standard AIX tools such as lparstat, vmstat, topas, and svmon.

To assist with planning, AIX includes a tool that analyzes actual workloads to estimate how much memory can be expanded and how much CPU resource is required. This tool can run on any Power system. Additionally, IBM offers a one-time, 60-day AME trial through the Power Capacity on Demand website, so users can evaluate memory expansion and CPU usage more precisely.

AME is enabled through a chargeable hardware feature code (EMAM), which can be ordered with the system or as a miscellaneous equipment specification (MES) order. A software key is provided for each system node when the feature is purchased. This key is permanent, specific to the system node, and does not require an IPL to activate. However, it cannot be transferred to another server. The CPU resources that are used for AME are part of the partition's assigned capacity and are subject to standard licensing requirements.

OMI to DDIMM wiring

Each of the two processor chips in the DCM has 8 of its 16 OMP PHY ports brought to the DCM pin. The left (front) DCM0 and DCM3 are placed in 180-degrees rotation compared to DCM1 and DCM2.

Note: CPx also means DCMx. They are used interchangeably.

Figure 2-2 shows the OMI to DDIMM wiring.

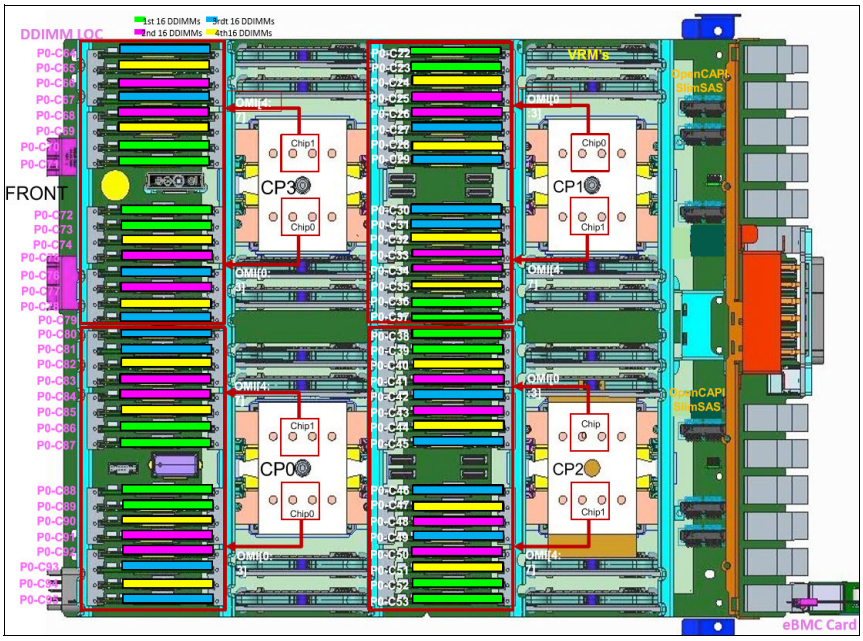

Figure 2-2 OMI to DIMM wiring

OMI PHY ports to DDIMM location codes cross-references

Figure 2-2 on page 47 provides cross-references for DDIMM location codes, offering a visual guide to help identify and map memory module positions within the system.

1st 16 DDIMMs 3rd 16 DDIMMs
2nd 16 DDIMMs 4th16 DDIMMs

OMI Channels to DDIMM Location Codes Cross Reference

(Sorted by DDIMM Location Codes per DCM)

#	CP0			CP1			CP2			CP3		
	DIMM	Port	Location	DIMM	Port	Location	DIMM	Port	Location	DIMM	Port	Location
1	0	7A	P0-C80	0	0A	P0-C22	0	0A	P0-C38	0	7A	P0-C64
2	1	7B	P0-C81	1	0B	P0-C23	1	0B	P0-C39	1	5B	P0-C65
3	2	5B	P0-C82	2	1A	P0-C24	2	1A	P0-C40	2	6B	P0-C66
4	3	6B	P0-C83	3	2A	P0-C25	3	2A	P0-C41	3	7B	P0-C67
5	4	6A	P0-C84	4	2B	P0-C26	4	3B	P0-C42	4	6A	P0-C68
6	5	5A	P0-C85	5	3B	P0-C27	5	2B	P0-C43	5	5A	P0-C69
7	6	4B	P0-C86	6	1B	P0-C28	6	1B	P0-C44	6	4A	P0-C70
8	7	4A	P0-C87	7	3A	P0-C29	7	3A	P0-C45	7	4B	P0-C71
9	0	0A	P0-C88	0	7A	P0-C30	0	7A	P0-C46	0	0B	P0-C72
10	1	0B	P0-C89	1	7B	P0-C31	1	5B	P0-C47	1	0A	P0-C73
11	2	1A	P0-C90	2	5B	P0-C32	2	6B	P0-C48	2	1A	P0-C74
12	3	2A	P0-C91	3	6B	P0-C33	3	7B	P0-C49	3	2A	P0-C75
13	4	2B	P0-C92	4	6A	P0-C34	4	6A	P0-C50	4	3B	P0-C76
14	5	3B	P0-C93	5	5A	P0-C35	5	5A	P0-C51	5	2B	P0-C77
15	6	1B	P0-C94	6	4B	P0-C36	6	4B	P0-C52	6	1B	P0-C78
16	7	3A	P0-C95	7	4A	P0-C37	7	4A	P0-C53	7	3A	P0-C79

Figure 2-3 DDIMM Location Codes

Memory Placement Rules

Each Power11 chip requires a minimum of two installed DDIMMs. Because each processor socket has a dual-chip module, four (2×2 DDIMMs) must be installed per socket. Because each Power E1150 server requires a minimum of two sockets that are populated, a minimum of eight DDIMMs must be installed. Using the smallest 32 GB DDIMMs, there is a minimum of 256 GB per server in a 2-socket configuration.

The OMI sockets that are numbered as P0-C22Q through P0-C95 must be populated in a defined order. Figure 2-4 shows the plugging rules with colors. First, populate the green slots, then the pink slots, then the blue slots, and then the yellow slots.

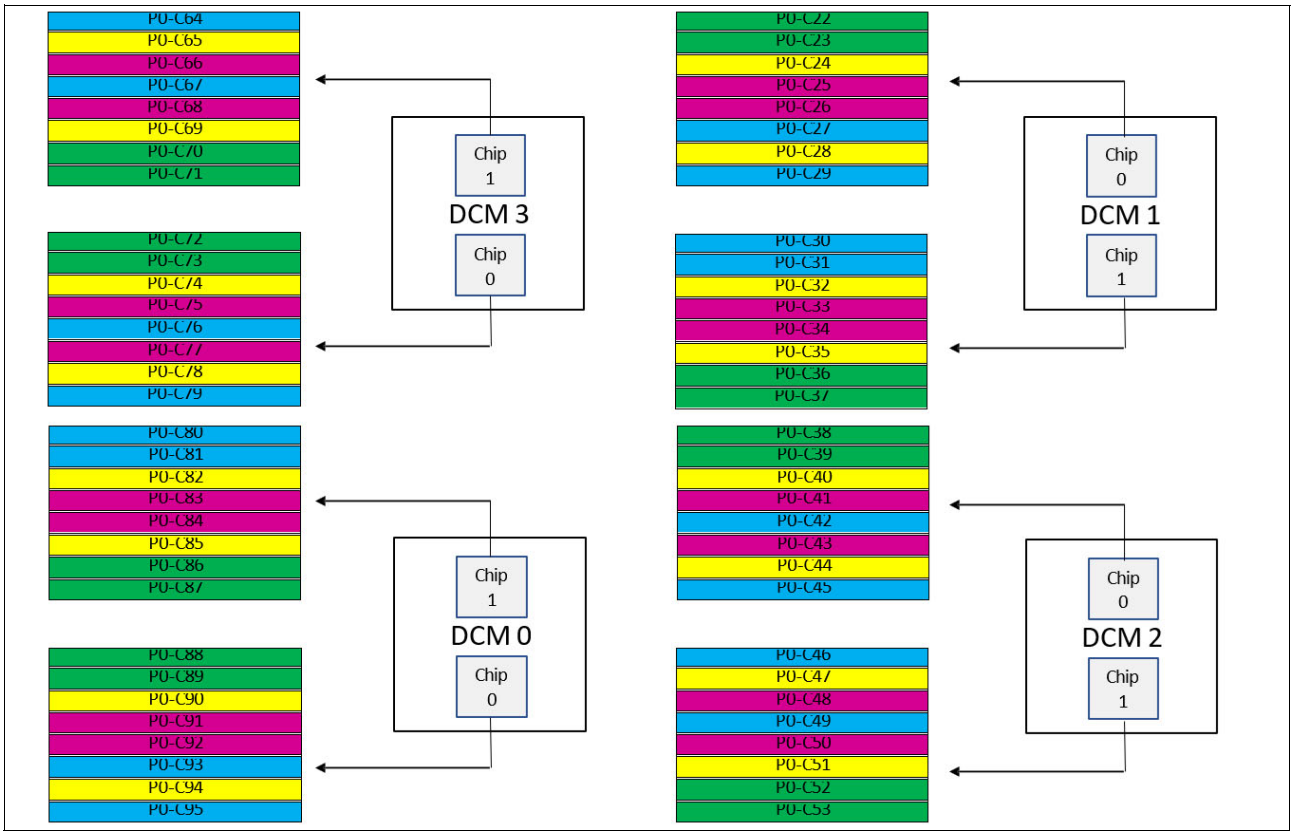

Figure 2-4 Memory placement order

Table 2-8 provides a textual view of the memory placement order.

Table 2-8 Memory placement order

Set	Location codes			
	DCM0	**DCM1**	**DCM2**	**DCM3**
First set of DDIMMs (green in Figure 2-4)	P0-C86 P0-C87 P0-C88 P0-C89	P0-C22 P0-C23 P0-C36 P0-C37	P0-C38 P0-C39 P0-C52 P0-C53	P0-C70 P0-C71 P0-C72 P0-C73
Second set of DDIMMs (pink in Figure 2-4)	P0-C83 P0-C84 P0-C91 P0-C92	P0-C25 P0-C26 P0-C33 P0-C34	P0-C41 P0-C43 P0-C48 P0-C50	P0-C66 P0-C68 P0-C75 P0-C77

Set	Location codes			
	DCM0	DCM1	DCM2	DCM3
Third set of DDIMMs (blue in Figure 2-4)	P0-C80 P0-C81 P0-C93 P0-C95	P0-C27 P0-C29 P0-C30 P0-C31	P0-C42 P0-C45 P0-C46 P0-C49	P0-C64 P0-C67 P0-C76 P0-C79
Fourth set of DDIMMs (yellow in Figure 2-4)	P0-C82 P0-C85 P0-C90 P0-C94	P0-C24 P0-C28 P0-C32 P0-C35	P0-C40 P0-C44 P0-C47 P0-C51	P0-C65 P0-C69 P0-C74 P0-C78

2.3 Racks

The Power E1150 must be mounted in a supported rack for operation. The rack provides the ability to install and maintain the system unit and any attached expansion drawers, providing the power connections and allowing connectivity to other devices providing network connectivity and storage.

IBM supports two racks:

- IBM Enterprise 42U Slim Rack (7965-S42)

  The IBM 7965-S42, also known as the Enterprise Slim Rack, is a compact rack designed for environments where space efficiency is a priority. The S42 offers 42U of rack space and supports a wide range of IBM systems, including Power servers and storage solutions. It is engineered for optimal thermal performance and includes features like front and rear door perforations for improved airflow. The S42 is well suited for modern data centers that require high-density computing in a smaller footprint, and it supports advanced cable routing and power management options to maintain a clean and efficient setup.

  The S42 is supported for IBM integration of new orders and for migration of existing systems.

- IBM 42U Enterprise Rack Enclosure (7014-T42)

  The IBM 7014-T42 rack is a robust, enterprise-grade enclosure that is designed to house a wide range of IBM Power systems and associated IT equipment. The T42 rack is 2.0 meters tall and provides 42 EIA units (42U) of usable vertical space, making it ideal for dense server and storage deployments. Its design accommodates high-power and high-heat systems, with features that support efficient airflow and cable management. The T42 is often used in data centers where scalability and reliability are critical, and it includes options for power distribution units (PDUs), side panels, and enhanced security features. Because of its height and weight, special handling and shipping considerations might apply.

  The T42 rack has been withdrawn from marketing, but it is supported for field integration of the Power E1150 systems.

Other OEM racks are supported assuming that they meet the requirements specified in section 2.3.5, “Original equipment manufacturer racks” on page 63.

2.3.1 IBM Enterprise 42U Slim Rack 7965-S42

The 2.0-meter (79-inch) Model 7965-S42 is compatible with past and present IBM Power servers and provides a 19-inch rack enclosure for your data center. Its 600 mm (23.6 in.) width combined with its 1100 mm (43.3 in.) depth plus its 42 EIA enclosure capacity provides great footprint efficiency for your systems. It can be placed on standard 24-inch floor tiles.

Compared to the 7965-94Y Slim Rack, the Enterprise Slim Rack provides extra strength and provides additional shipping and installation flexibility.

The 7965-S42 rack includes space for up to four PDUs in side pockets. Extra PDUs beyond four are mounted horizontally and each uses 1U of rack space.

The Enterprise Slim Rack comes with options for the installed front door:

- Basic Black/Flat (#ECRM)
- High-End appearance (#ECRF)
- OEM Black (#ECRE)

All options include perforated steel, which provides ventilation, physical security, and visibility of indicator lights in the installed equipment within. All options come with a lock and mechanism included that is identical to the lock on the rear doors. Only one front door must be included for each rack ordered. The basic door (#ECRM) and OEM door (#ECRE) can be hinged on the left or right side.

Orientation: #ECRF must not be flipped because the IBM logo would be upside down.

At the rear of the rack, either a perforated steel rear door (#ECRG) or a Rear Door Heat Exchanger (RDHX) can be installed. The only supported RDHX is IBM Machine Type 1164-95X, which can remove up to 30,000 watts (102,000 BTU) of heat per hour by using chilled water. The no additional charge Feature Code #ECR2 is included with the Enterprise Slim Rack as an indicator when ordering the RDHX.

The basic door (#ECRG) can be hinged on the left or right side, and includes a lock and mechanism identical to the lock on the front door. Either the basic rear door (#ECRG) or the RDHX indicator (#ECR2) must be included with the order of a new Enterprise Slim Rack.

Because of the depth of the Power E1150 server model, the 5-inch rear rack extension (#ECRK) is required for the Enterprise Slim Rack to accommodate this system. This extension expands the space that is available for cable management and allows the rear door to close safely.

Lifting considerations

Three to four people are required to manually remove or insert a system unit into a rack because of its dimensions, weight, and content. To avoid the need for this many people, you can use a lift tool.

The Power E1150 server has a maximum weight of 70.3 kg (155 lb). However, by temporarily removing the power supplies, fans, and RAID assembly, you can reduce the weight to a maximum of 55 kg (121 lb).

When lowering the Power E1150 server onto its rails in the rack, the server must be tilted on one end about 15 degrees so that the pins on the server enclosure fit onto the rails. This equates to lifting one end of the server 4 cm (1.6 in.). This task can be done by using a tip plate on a lift tool, manually adjusting the load on a lift tool, or tilting during the manual lift. Consider the optional feature #EB2Z lift tool.

2.3.2 AC power distribution unit and rack content

The IBM high-function PDUs provide more electrical power per PDU compared to earlier PDUs, and they offer better PDU footprint efficiency. In addition, they are intelligent PDUs (iPDUs) that provide insight to actual power usage by receptacle and also provide remote power on/off capability for support by individual receptacle. The latest PDUs can be ordered as #ECJJ, #ECJL, #ECJN, and #ECJQ. Table 2-9 summarizes the high-function PDU Feature Codes for 7965-S42 followed by a descriptive list.

IBM Manufacturing integrates only the newer PDUs with the Power E1150 server. IBM Manufacturing does not support integrating earlier PDUs, such as #7188, #7109, or #7196. You can choose to use older IBM PDUs in your racks, but IBM does not install them in the 7965-S42.

Table 2-9 High-function PDUs that are available with IBM Enterprise Slim Rack (7965-S42)

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

a. The Power E1050 server has an AC power supply with a C19/C20 connector.

Power sockets: The Power E1150 server takes IEC 60320 C19/C20 mains power and not C13. Ensure that the correct power cords and PDUs are ordered or available in the rack.

- High Function 9 C19 PDU plus (#ECJJ)

  This intelligent, switched 200–240 V AC PDU includes nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, which makes the nine C19 receptacles easily accessible. For comparison, this PDU is most like the earlier generation #EPTJ PDU.

- High Function 9 C19 PDU plus 3-Phase (#ECJL)

  This intelligent, switched 208 V 3-phase AC PDU includes nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, which makes the nine C19 receptacles easily accessible. For comparison, this PDU is most like the earlier generation #EPTL PDU.

- High Function 12 C13 PDU plus (#ECJN)

  This intelligent, switched 200–240 V AC PDU includes 12 C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, which makes the 12 C13 receptacles easily accessible. For comparison, this PDU is most like the earlier generation #EPTN PDU.

- High Function 12 C13 PDU plus 3-Phase (#ECJQ)

  This intelligent, switched 208 V 3-phase AC PDU includes 12 C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, which makes the 12 C13 receptacles easily accessible. For comparison, this PDU is most like the earlier generation #EPTQ PDU.

The PDU receives power through a UTG0247 power-line connector. Each PDU requires one PDU-to-wall power cord. Various power cord features are available for various countries and applications by varying the PDU-to-wall power cord, which must be ordered separately.

Each power cord provides the unique design characteristics for the specific power requirements. To match new power requirements and save previous investments, these power cords can be requested with an initial order of the rack or with a later upgrade of the rack features.

Table 2-10 shows the available wall power cord options for the PDU features, which must be ordered separately.

Table 2-10 PDU-to-wall power cord options for the PDU features

Feature Code	Wall plug	Rated voltage (V AC)	Phase	Rated amperage	Geography
#6489	IEC309, 3P+N+G, 32 A	230	3	32 amps per phase	Europe, Middle East, and Asia (EMEA)
#6491	IEC 309, P+N+G, 63 A	230	1	63 amps	EMEA
#6492	IEC 309, 2P+G, 60 A	200–208, 240	1	48 amps	US, Canada, Latin America (LA), and Japan
#6653	IEC 309, 3P+N+G, 16 A	230	3	16 amps per phase	Internationally available
#6654	NEMA L6-30	200–208, 240	1	24 amps	US, Canada, LA, and Japan
#6655	RS 3750DP (watertight)	200–208, 240	1	24 amps	US, Canada, LA, and Japan
#6656	IEC 309, P+N+G, 32 A	230	1	24 amps	EMEA
#6657	PDL	230–240	1	32 amps	Australia and New Zealand
#6658	Korean plug	220	1	30 amps	North and South Korea
#6667	PDL	380–415	3	32 amps	Australia and New Zealand

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

The Universal PDUs are compatible with previous models. To better enable electrical redundancy, the Power E1150 server has four power supplies that must be connected to separate PDUs, which are not included in the base order. For maximum availability, a best practice is to connect power cords from the same system to two separate PDUs in the rack, and to connect each PDU to independent power sources.

For more information about the power requirements of and the power cord for the 7965-94Y rack, see Planning for the 7953-94X and 7965-94Y rack.

2.3.3 Rack-mounting rules

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

- The system can be placed at any location in the rack. For rack stability, start filling the rack from the bottom.
- As as best practice, use an IBM approved lift tool for the installation of systems into any IBM or third-party rack.
- IBM does not support the installation of server nodes higher than the 29U position.
- Any remaining space in the rack can be used to install other systems or peripheral devices. Ensure that the maximum permissible weight of the rack is not exceeded and the installation rules for these devices are followed.
- Before placing the system into the service position, follow the rack manufacturer's safety instructions regarding rack stability.

Order information: The racking approach for the initial order must be 7965-S42 or #ECR0. If an extra rack is required for I/O expansion drawers, an MES to a system or an #0553 must be ordered.

2.3.4 Useful rack additions

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

IBM System Storage 7226 Model 1U3 Multi-Media Enclosure

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

For more information on the 7726-1U3, see "IBM System Storage 7226 Model 1U3 multi-media enclosure" on page 91.

Flat panel display options

The IBM 7316 Model TF5 is a rack-mountable flat panel console kit that can also be configured with the tray pulled forward and the monitor folded, which provides full viewing and keying capability for the HMC operator.

The Model TF5 is a follow-on product to the Model TF4 and offers the following features:

- A slim and lightweight monitor design that occupies only 1U (1.75 in.) in a 19-inch standard rack.
- An 18.5 inch (409.8 mm x 230.4 mm) flat panel TFT monitor with accurate images and virtually no distortion.
- The ability to mount the IBM Travel Keyboard in the 7316-TF5 rack keyboard tray.
- Support for the IBM 1U 8-port Rack Console Switch (#4283) IBM Keyboard/Video/Mouse (KVM) switches.

  The #4283 is an 8-port Console Switch that fits in the 1U space behind the TF5. It is a CAT5-based switch. It contains eight analog rack interface (ARI) ports for connecting PS/2 or USB console switch cables. It supports chaining of servers that use an IBM Conversion Options switch cable (#4269). This feature provides four cables that connect a KVM switch to a system, or can be used in a daisy-chain scenario to connect up to 128 systems to a single KVM switch. It also supports server-side USB attachments.

2.3.5 Original equipment manufacturer racks

The Power E1150 can be installed in a suitable OEM rack if that the rack conforms to the EIA-310-D standard for 19-inch racks. This standard is published by the Electrical Industries Alliance. For more information, see Rack installation specifications for racks that are not purchased from IBM.

IBM Documentation provides the general rack specifications, including the following information:

- The rack or cabinet must meet the EIA Standard EIA-310-D for 19-inch racks, which was published August 24, 1992. The EIA-310-D standard specifies internal dimensions, for example, the width of the rack opening (width of the chassis), the width of the module mounting flanges, and the mounting hole spacing.
- The front rack opening must be a minimum of 450 mm (17.72 in.) wide, and the rail-mounting holes must be 465 mm plus or minus 1.6 mm (18.3 in. plus or minus 0.06 in.) apart on center (horizontal width between vertical columns of holes on the two front-mounting flanges and on the two rear-mounting flanges).

  Figure 2-5 is a top view showing the rack specification dimensions.

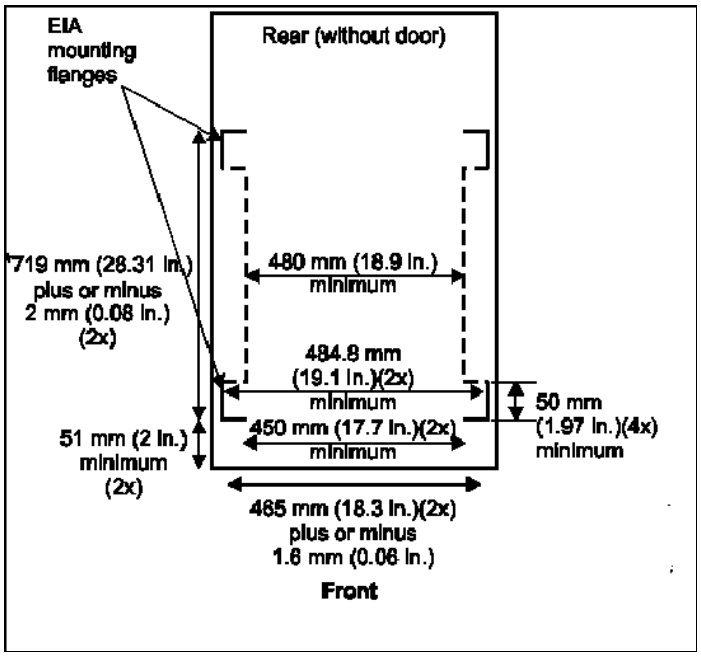

Figure 2-5 Rack specifications (top-down view)

- The vertical distance between mounting holes must consist of sets of three holes that are spaced (from bottom to top) 15.9 mm (0.625 in.), 15.9 mm (0.625 in.), and 12.7 mm (0.5 in.) on center, which makes each three-hole set of vertical hole spacing 44.45 mm (1.75 in.) apart on center.

Figure 2-6 shows the vertical distances between the mounting holes.

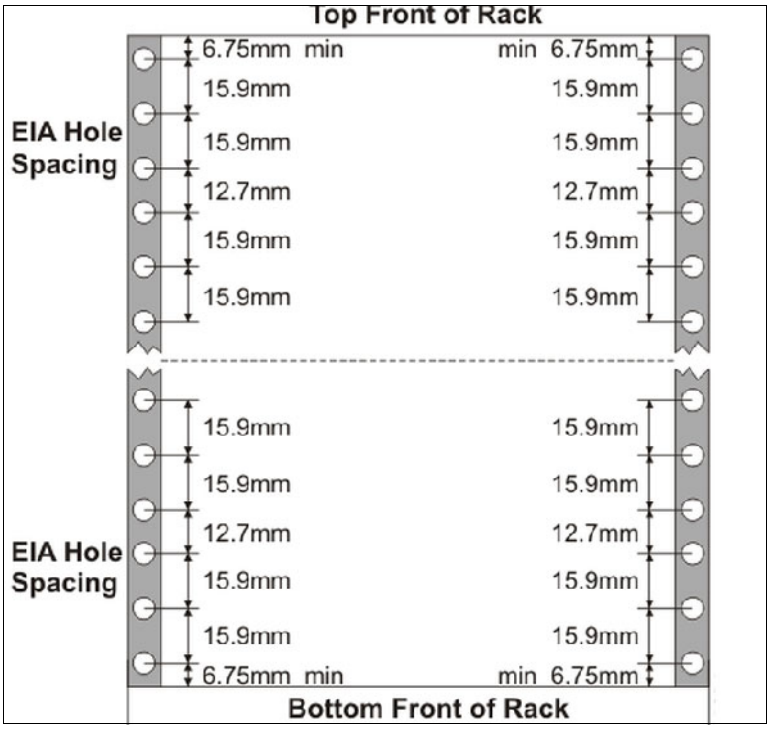

Figure 2-6 Vertical distances between mounting holes

- The following rack hole sizes are supported for racks where IBM hardware is mounted:
  - 7.1 mm (0.28 in.) plus or minus 0.1 mm (round)
  - 9.5 mm (0.37 in.) plus or minus 0.1 mm (square)

The rack or cabinet must be capable of supporting an average load of 20 kg (44 lb.) of product weight per EIA unit. For example, a four EIA drawer has a maximum drawer weight of 80 kg (176 lb.).

3

I/O subsystem

This chapter explores the I/O subsystem within IBM Power, highlighting its design for mission-critical applications. Its modular architecture and PowerVM virtualization facilitate massive, dynamic scalability of I/O resources, including network and storage. The discussion also includes the inherent reliability features, such as NPIV and VIOS and includes advanced error handling, ensuring continuous operation. The chapter examines how direct I/O assignments and optimized shared I/O through VIOS contribute to performance for data-intensive workloads. This chapter includes the following topics:

- Internal I/O
- Enhancing I/O scalability with expansion drawers
- List of supported adapters on IBM Power E1150
- Other device support

3.1 Internal I/O

The I/O subsystem in IBM Power is designed from the ground up to deliver exceptional scalability, reliability, and performance for mission-critical workloads. Scalability is achieved through a highly modular and flexible architecture. This allows for massive expansion of I/O resources, including a large number of PCIe slots, high-speed network adapters (Ethernet, InfiniBand), and a large capacity for internal and external storage, including NVMe drives. PowerVM virtualization, a cornerstone of IBM Power, enables granular allocation and sharing of physical I/O resources across multiple logical partitions (LPARs), allowing administrators to dynamically scale I/O to meet the changing demands of virtualized workloads without physical reconfiguration.

Reliability is paramount in the IBM Power I/O subsystem, built upon decades of experience in mission-critical computing. Features like N_Port ID Virtualization (NPIV) provide highly available Fibre Channel connectivity by allowing multiple LPARs to share a single physical adapter while maintaining independent storage access. Virtual I/O Server (VIOS) plays a crucial role, providing shared I/O capabilities and enabling failover mechanisms for network and storage adapters. Advanced error detection and correction (RAS features), such as predictive failure analysis and dynamic component sparing for I/O resources, are deeply integrated into the hardware and firmware.

The ability to directly assign I/O adapters to LPARs (dedicated I/O) or use hardware-assisted virtualization like SR-IOV for near-native performance further optimizes throughput and reduces CPU usage. For shared I/O, the Virtual I/O Server (VIOS) is continuously optimized to handle high volumes of virtualized traffic, with best practices and tuning options available to maximize performance for various workloads, including those with intensive network and storage demands. The overall architecture is designed to handle immense I/O concurrency, making IBM Power ideal for data-intensive applications such as databases, analytics, and AI workloads that require rapid access to large datasets

3.1.1 Internal PCIe Gen 5 subsystem

The E1150 system provides a total of eleven internal PCIe slots, labeled C1 through C11. All eleven slots are located in the rear. Each of the eleven slots is a general-purpose PCIe slot and is physically x16 in size. In addition to the PCIe subsystem, the E1150 reserves additional PCIe Gen4 slots for connectivity for up to ten NVMe drives that are installed in the system. The internal NVMe configuration is described in section 3.1.4, “Internal NVMe storage subsystem” on page 69.

The E1150 supports the following slot configurations:

- Six slots (C2, C3, C4, C5, C8, and C11) support x8 Gen5 or x16 Gen4 operation
- Two slots (C7 and C10) support x8 Gen5
- Three slots (C1, C6, and C9) support x8 Gen4.

Note: Slots C1, C2, C3, C6, C7, and C8 are designated as OpenCAPI slots. However, OpenCAPI adapters are not concurrently maintainable, and currently, no OpenCAPI adapters are planned for the E1150.

Figure 3-1 shows the PCIe subsystem in the E1150.

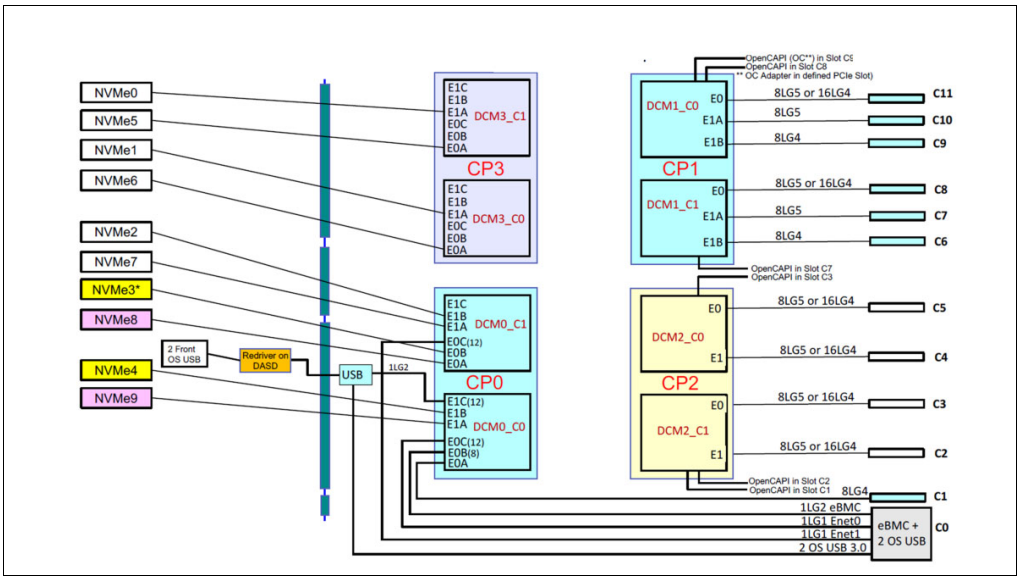

Figure 3-1 Overview of the PCIe subsystem of the E1150

E1150 has 10 out of 11 I/O slots from the right 2 dual chip modules (DCM1 and DCM2), and all 10 NVMe are attached to the two 2 left dual chip modules (DCM0 and DCM3). This means that systems with only two populated sockets are limited to seven usable PCIe slots (P0-C1 and P0-C6 to P0-C11) and six NVMe drives. To make all the slots available, at least three processor sockets must be populated and to enable all of the NVMe drives, all of the sockets must be populated.

Table 3-1 lists the number of PCIe slots and NVMe drives available based on the number of sockets populated.

Table 3-1 Available PCIe slots and NVMe drives

Number of sockets populated	PCIe slots[a]	NVMe drives
2	7	6
3	11	6
4	11	10

a. One slot is dedicated for network card installation.

The Power11 chip design achieves full I/O bandwidth from any I/O slot in a flat-8 SMP (1-hop) design.

Slot C0 is a special-purpose slot that is reserved for the eBMC Service Processor Card. This slot is not a standard PCIe slot and is not concurrently maintainable. The eBMC card is installed in a dedicated I/O cassette. The BMC card is described in section 8.4, “eBMC card” on page 152.

There is no Ethernet connectivity on the E1150 system board, so Slot C1, although a general-purpose x8 PCIe slot, is typically used for the base Ethernet card. The default Ethernet adapter for the E1150 is the 2-Port 25/10 Gb NIC & ROCE SR/Cu PCIe 3.0 Adapter (Feature Code EC2U), but other adapters can be used.

3.1.2 PCIe slot properties

The internal I/O subsystem of the Power E1150 server is connected to the PCIe Express controllers on a Power11 chip in the system. A Power11 chip has two PCI Express controllers (PECs) of 16 lanes each for a total of 32 Gen5/Gen4 lanes per chip and 64 Gen5/Gen4 lanes per DCM.

Each PEC supports up to three PCI host bridges (PHBs) that directly connect to PCIe slots or devices. Both PEC0 and PEC1 can be configured as follows:

- One x16 Gen4 PHB or one x8 Gen5 PHB
- One x8 Gen5 and one x8 Gen4 PHB
- One x8 Gen5 PHB and two x4 Gen4 PHBs

Table 3-2 shows the PCIe port capabilities and is listed by physical position from right to left of the rear view.

Table 3-2 PCIe slot characteristics

Name	Width	P11 CP	PCIe bus
C11	x8 G5/x16 G4	DCM1/C0	E0
C10	x8 G5/G4	DCM1/C0	E1A
C9	x8 G4	DCM1/C0	E1B
C8	x8 G5/x16 G4	DCM1/C1	E0
C7	x8 G5/G4	DCM1/C1	E1A
C6	x8 G4	DCM1/C1	E1B
C5	x8 G5/x16 G4	DCM2/C0	E0
C4	x8 G5/x16 G4	DCM2/C0	E1
C3	x8 G5/x16 G4	DCM2/C1	E0
C2	x8 G5/x16 G4	DCM2/C1	E1
C1	x8 G4	DCM0/C0	E0A

Note: The eBMC service processor card occupies slot C0. Slot C1 is generally reserved for a network adapter.

The x16 slots can 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 a PCIe Gen4 slot, and PCIe Gen4 slots can support up to twice the bandwidth of a PCIe Gen3 slot, assuming an equivalent number of PCIe lanes.

Note: Although some slots provide a x8 connection only, all slots have an x16 connector.

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

All PCIe adapter slots support hardware-backed network virtualization through single-root I/O virtualization (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 instructed to free extra memory and retry the operation.

For energy efficiency and cooling, the Power E1150 server senses which IBM PCIe adapters are installed in their PCIe slots. If an adapter requires higher levels of cooling, the speed of the fans automatically increases to increase airflow across the PCIe adapters. Faster fans increase the sound level of the server.

3.1.3 PCIe cassette

The E1150 PCIe subsystem uses a cassette to facilitate the installation and concurrent maintenance of the PCIe adapters. PCIe adapters are installed into the I/O cassette before the cassette is inserted into the system from the rear. Each I/O cassette includes a single x16 PCIe connector. The PCIe cassettes are shipped with the system and every slot from C1 to C11 contains an I/O cassette, either populated with an adapter or with a blank. All standard PCIe slots (C1 through C11) support concurrent maintenance by using the I/O cassette mechanism, but Slot C0 does not. The PCIe cassette is shown in Figure 3-2.

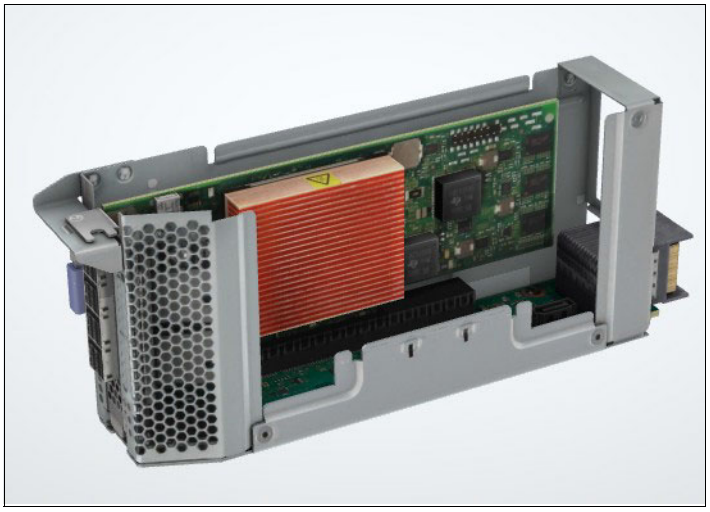

Figure 3-2 PCIe cassette

3.1.4 Internal NVMe storage subsystem

The E1150 system is equipped with an NVMe U.2 backplane. Drives that support both U.3 and U.2 interfaces operate in NVMe U.2 mode only because the E1150 does not support the SAS protocol. Additionally, the system is limited to single-port NVMe mode and does not support dual-port NVMe configurations.

The E1150 includes ten internal NVMe drive bays that are integrated into a backplane assembly. This backplane is a standard component in all E1150 systems. In addition to the ten NVMe bays, the NVMe backplane also houses the Operator Panel Base, Operator Panel LCD, and two USB 3.0 ports. This is shown in Figure 3-3 on page 70.

Note: Although ten NVMe drive bays are available, the number of NVMe drives supported depends on the number of sockets populated in the E1150:

- For 2 and 3 socket systems, 6 NVMe devices are supported.
- For 4 socket systems, 10 NVMe devices are supported,

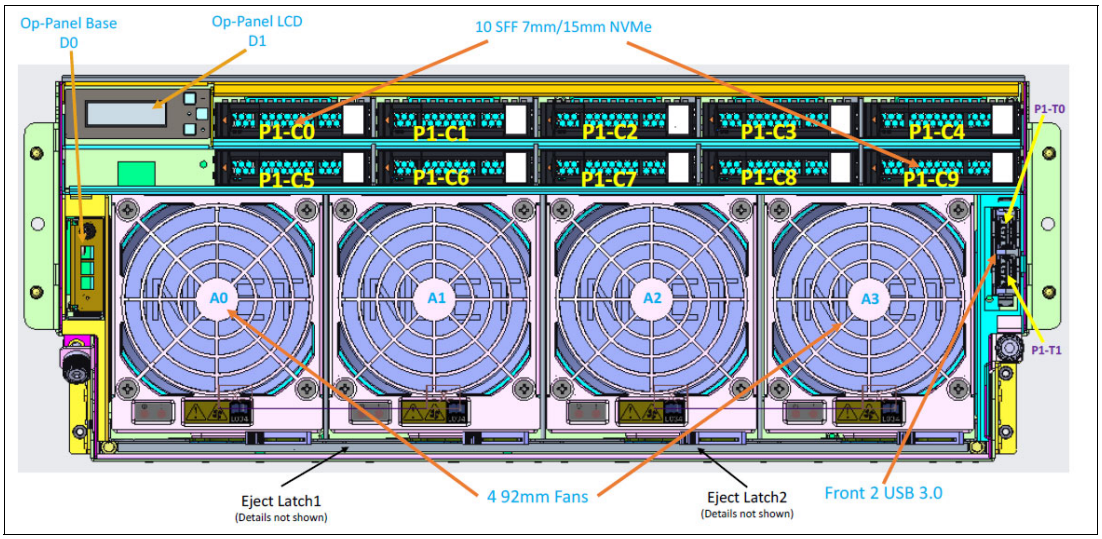

Figure 3-3 Front view of the E1150 showing NVMe drive slots

The NVMe backplane connects directly to the system board by using three connectors and a dedicated power connector. There are no cables between the backplane and the system board, ensuring a clean and reliable connection.

The wiring strategy and materials used in the backplane are specifically selected to support PCIe Gen4 signaling across all NVMe drives. Signal Integrity analysis has confirmed that neither re-timers nor re-drivers are required to maintain Gen4 performance. The NVMe connectors that are used are compliant with Gen4 PCIe specifications.

Table 3-3 provides a cross-reference between the NVMe drive number and its location code.

Table 3-3 Location codes for NVMe drives

NVMe drive	Location Code
NVMe0	P1-C0
NVMe1	P1-C1
NVMe2	P1-C2
NVMe3	P1-C3
NVMe4	P1-C4
NVMe5	P1-C5
NVMe6	P1-C6
NVMe7	P1-C7
NVMe8	P1-C8
NVMe9	P1-C9

Plug rules in the E1150 system unit

The following NVMe plug rules are recommended to provide the most redundancy in hardware for operating system mirror support.

Table 3-4

Number of NVMe drives[a]	Location
2	NVMe3 NVMe4 or NVMe8 NVMe9
4	NVMe3 NVMe4 NVMe8 NVMe9
6	NVMe3 NVMe4 NVMe8 NVMe9 NVMe2 NVMe7
More than 6[b]	NVMe3 NVMe4 NVMe8 NVMe9 NVMe2 NVMe7 NVMe5 NVMe6 NVMe0 NVMe1

a. If odd number of NVMe drives, install by order shown.
b. Only available in four socket configuration

NVMe availability considerations

When you load either VIOS or an operating system on the internal NVMe drives, and you intend to mirror the drives, use the recommended plug rules to attempt to ensure separate hardware paths for the mirrored pairs of disk. The E1150 NVMe thermal design supports 18 W for 15 mm NVMe.

The NVMe drives can be in an OS-controlled RAID0, RAID1 array. Hardware RAID is not supported on the NVMe drives.

Note: It is recommended, though not required, that the two NVMe drives that are used in a mirrored pair be the same capacity. When drives of different capacities are used, the effective usable capacity of the mirror is limited to the smaller of the two drives. In other words, the mirror (secondary) drive must have a capacity equal to or greater than that of the primary drive to ensure proper mirroring without data loss.

For mirrored OS support use the following slots:

- NVMe3 and NVMe4
- NVMe8 and NVMe9

For dual VIOS configurations, use the following slots:

- VIOS1: NVMe 3 and NVMe9
- VIOS2: NVMe4 and NVMe8

Note: When there are multiple running partitions in the system, the mirror pairs selections depend on drives available and allocated to each partition.

3.2 Enhancing I/O scalability with expansion drawers

Adding I/O drawers to an IBM Power E1150 server enhances the system scalability, flexibility, and overall performance in enterprise environments. The Power E1150 supports PCIe Gen4 I/O expansion drawers. The expansion drawers provide increased throughput and bandwidth for connected devices, enabling the system to handle more demanding workloads and a broader range of peripherals.

These drawers expand the available PCIe slots beyond what the system board natively supports. You can add network adapters, storage controllers, and accelerators without compromising existing configurations. This approach is particularly beneficial for data-intensive applications such as AI inferencing, high-speed networking, and large-scale database operations. Also, the modular nature of I/O expansion 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 that can be installed, these drawers enable more flexible system configurations. This capability is especially beneficial in virtualized environments where multiple virtual machines or containers require dedicated I/O resources. The additional capacity supports higher VM density and more granular resource allocation, improving overall system usage. IBM expansion drawers include enterprise-grade reliability features such as hot-plug capabilities, redundant paths, and integration with PowerVM and the Hardware Management Console (HMC). These features 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 Power E1150 design goals of resilience, security, and performance. The system architecture, including its I/O subsystem, supports mission-critical workloads with features such as redundant paths, error recovery, and secure boot. The drawers integrate with IBM HMC and PowerVM virtualization, ensuring centralized control and simplified management. This combination makes the Power E1150 with I/O drawers an ideal platform for enterprises seeking to modernize their infrastructure while maintaining high availability and robust performance.

3.2.1 Supported I/O drawers

The E1150 supports the following I/O expansion drawers:

- PCIe Gen4 I/O expansion drawer (#ENZ0)

  The PCIe Gen4 I/O expansion drawer (#ENZ0) is a 4U high, 19-inch wide, PCIe Gen4 based rack mountable I/O drawer that is available as a feature of Power11 Servers. The PCIe Gen4 I/O expansion drawer (#ENZ0) replaces the PCIe Gen3 I/O expansion drawer (#EMX0). There is no upgrade path from PCIe Gen3 I/O expansion drawer (#EMX0) to PCIe Gen4 I/O expansion drawer (#ENZ0).

- NED24 NVMe expansion drawer (#ESR0)

The NED24 NVMe expansion drawer (#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.

Within the Power E1150 server, the PCIe slots are enabled to support the PCIe x16 to CXP Converter Card (FC EJ2A) that is used to attach expansion drawers. Two PCIe x16 to CXP Converter Cards are used to attach each expansion drawer. Depending on the number of sockets populated in the E1150 you purchase, the maximum number of supported drawers are as detailed in Table 3-5.

Table 3-5 Maximum I/O Drawer configuration

Configuration	Max I/O Drawers	Max NED24	Max ENZ0	Comments
E1150 (3 or 4 Sockets)	4	2[a]	4[b]	Slots C7 and C10 are Gen 5 x8[c]
E1150 (2 Sockets)	2	1[d]	2[e]	Slots C7 and C10 are Gen 5 x8[c]

a. E1150 Slots for NED24: C8/C11, C2/C4, or C3/C5
b. E1150 Slots for ENZ0: C2/C4, C3/C5, C8/C11
c. Using the x8 slots yields performance that is similar to the performance of the EMX0 PCIe3 expansion drawer. Using x16 slots yields improved performance.
d. E1150 (2 Socket) for NED24: C8/C11
e. E1150 (2 Socket) for Nimitz: C8/C11, C7/C10

3.2.2 Nonsupported drawers

The following drawers are not supported by Power11 E1150:

- EMX0:

  PCI Gen 3 I/O expansion drawer (#EMX0) is not supported by Power11. None of the Fanout Modules EMXF/EMXG/EMXH are supported within the new PCI Gen 4 I/O expansion drawer (#ENZ0).

- IBM EXP24SX:

  The EXP24SX SAS storage enclosure (#ESLS/ ESLL) is not supported by Power11. For additional internal storage a new NED24 NVMe expansion drawer (#ESR0) is available, populated with up to 24 NVMe drives.

3.2.3 PCIe Gen 4 I/O expansion drawer

The 19-inch 4 EIA (4U) PCIe Gen 4I/O expansion drawer (#ENZ0) and two PCIe Fanout Modules (#ENZF) provide 12 PCIe I/O full-length, full-height Gen4 PCI slots. One Fanout Module provides 6 PCIe slots that are labeled C0–C5. Slots C0–C3 are x16 slots. Slots C4 and C5 are x8 slots. PCIe Gen1, Gen2, and Gen3 full-high adapters are also supported.

Important: PCI Gen 3 I/O expansion drawer (#EMX0) is not supported by Power11. None of the Fanout Modules EMXF/G nor EMXH are supported within the new PCI Gen 3 I/O expansion drawer (#ENZ0).

A blind swap cassette (BSC) houses the full-high adapters that go into these slots. The BSC is the same as the one used with the previous generation EMX0 drawers. The drawer is shipped with a full set of BSCs.

A PCIe CXP converter adapter (#EJ24) that occupies one of the PCIe Gen5 slots in the system node and a pair of active optical cables (AOCs) or copper cable are used for system node to fanout module connection. Both fanout modules are independent PCIe domains. They can be serviced independently of one another. A minimum of one fanout module is required in the ENZ0 drawer in location P0 being placed at the left side of the drawer when viewed from behind. Each PCIe Gen 4 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 operating system support utilities.

Balancing of I/O, assigning adapters through redundant ENZ0 expansion drawers, and connectivity to different system nodes can help ensure high availability for I/O resources that are assigned to LPARs.

Figure 3-4 shows a PCIe Gen 4 I/O expansion drawer front view.

Figure 3-4 PCIe Gen 4I I/O expansion drawer front view

Figure 3-5 shows the rear view of the PCIe Gen 4 I/O expansion drawer.

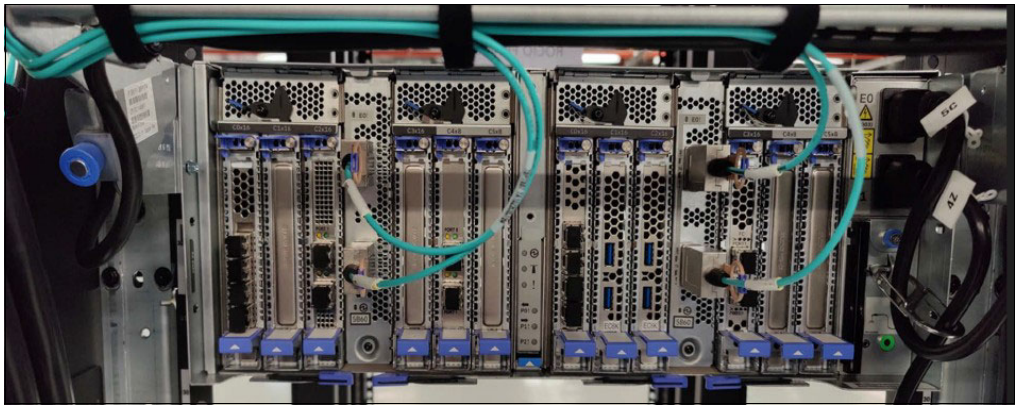

Figure 3-5 Rear view of a PCIe Gen 4 I/O expansion drawer

Supported PCIe adapters for ENZ0 expansion drawer

Table 3-6 lists the PCIe adapters that are supported in the ENZ0 expansion drawer.

Table 3-6 Adapters supported on the ENZ0 expansion drawer.

Feature code	CCIN	Description	FRU
EC2U	58FB	PCIe3 2-port 25/10 Gb NIC & RoCE SFP28 adapter	01FT753

Feature code	CCIN	Description	FRU
EC6K	590F	PCIe2 2-Port USB 3.0 Adapter	02JD518
EC72	2CF9	PCIe4 2-port 25/10/1 Gb RoCE SFP28 adapter	03HD074
EC74	2CF8	PCIe4 2-port 25/10/1 Gb RoCE SFP28 with Crypto adapter	03HD078
EJ2B	57F2	PCIe3 x8 SAS quad-port 12 Gb tape adapter	03MT000
EJ37	C0AF	4769-001 Cryptographic Coprocessor	02JD572
EN1A	578F	PCIe3 x8 2-port Fibre Channel (32 Gb/s)	01FT704
EN1J	579C	PCIe4 x8 2-port Fibre Channel (32 Gb/s)	02CM909
EN1L	2CFC	PCIe4 x8 4-port Fibre Channel (32 Gb/s)	03HD014
EN1N	2CFD	PCIe4 x8 2-port Fibre Channel (64 Gb/s)	03HD020
EN26	EC2A	PCIe4 x16 4-port 25/10/1 GbE RoCE SFP28 adapter	03HD066
EN2L	2F06	PCIe4 x16 4-port Fibre Channel (32 Gb/s)	03JP004
EN2N	2F05	PCIe4 x8 2-port Fibre Channel (64 Gb/s)	03JP010
EN2W	2F04	PCIe3 4-port 10 GbE adapter	03JP016
EPG6	C138	4770 Cryptographic Coprocessor	03JP117

Note: For additional requirements regarding the ENZ0 expansion drawer, supported adapters, and operating system levels, see the Sales Manual.

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 (#ESR0) is a storage expansion enclosure with 24 U.2 NVMe bays.

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

Figure 3-6 shows the front view of a NED24 drawer with 4 NVMe Drives at slots C8–C11.

Figure 3-6 NED24 NVMe drawer front view

Figure 3-7 on page 76 shows the rear view of the NED24 drawer with two power supplies and two Expansion Service Manager (ESM)

Figure 3-7 NED24 NVMe drawer rear view

Each of the 24 NVMe bays in the NED24 drawer is separately addressable and can be assigned to a specific LPAR or Virtual I/O Server (VIOS) to provide native boot support for up to 24 partitions. At the time of writing, each drawer can support up to 153 TB.

Up to 24 U.2 NVMe devices can be installed in the NED24 drawer by using 15 mm Gen3 carriers. The 15-mm carriers can accommodate either 7 mm or 15 mm NVMe devices.

The NED24 drawer is supported in the Power E1150 by using the same interconnect card that is used for the PCIe Gen 4 expansion drawer. A maximum of two NED24 NVMe expansion drawers is supported per system when the E1150 has three or four sockets populated. Because of the reduced number of PCIe slots available in the two socket E1150, the two socket configuration supports only one NED24 drawer.

When mixing the different expansion drawers, the maximum number of drawers that are supported is based on the number of EJ24 fanout cards that are supported.

Drive types supported

The IBM Power E1150 supports two main drive types for NVMe U.2 SSDs, each optimized for different workload profiles:

- Enterprise NVMe SSDs:
  - Designed for high-performance, high-endurance workloads such as databases, analytics, and virtualization
  - PCIe Gen4 NVMe U.2
  - 15 mm form factor
  - Higher write endurance (measured in Drive Writes Per Day - DWPD)

  See Table 3-7 on page 77 for the feature codes and capacities of enterprise SSD PCIe4 NVMe U.2 modules for AIX and Linux.
- Mainstream NVMe SSDs
  - Optimized for read-intensive or mixed workloads, such as boot drives, application servers, or general-purpose storage.
  - PCIe Gen3 or Gen4 NVMe U.2
  - 15 mm or 7 mm form factor
  - Lower write endurance, suitable for workloads with fewer write cycles.

  See Table 3-7 on page 77 for the feature codes and capacities of the mainstream NVMe drives for AIX and Linux.

Additional Notes:

- With the carrier conversion kit (EC7X), you can use 7 mm drives in 15 mm bays, ensuring compatibility across the system.
- All NVMe drives are hot-pluggable.

- Concurrent maintenance is supported for NVMe drives, enabling nondisruptive replacement or upgrades.

Table 3-7 shows the NVMe drives supported in the E1150 and the NED24 (ESR0) NVMe expansion drawer

Table 3-7 Supported NVMe drives

Feature Code	Capacity	Drive Type	Drive Size	Operating System Support
EC5V	6.4 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
EC5X	800 GB	Mainstream PCIe3 NVMe	7 mm	AIX/Linux
EC7T	800 GB	Mainstream PCIe3 NVMe	7 mm	AIX/Linux
ECT9	15.3 TB	Mainstream PCIe4 NVMe	15 mm	AIX/Linux
EKF3	1.6 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
EKF5	3.2 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
EKF7	6.4 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
EKF9	800 GB	Enterprise PCIe4 NVMe	7 mm	AIX/Linux
ES1E	1.6 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES1G	3.2 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES3B	1.6 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES3D	3.2 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES3F	6.4 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES3H	800 GB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES4B	1.6 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES4D	3.2 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES4F	6.4 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES5A	800 GB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES5C	1.6 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES5E	3.2 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux
ES5G	6.4 TB	Enterprise PCIe4 NVMe	15 mm	AIX/Linux

3.2.4 PCIe expansion card

The PCIe4 Expansion Cable adapter (FC EJ2A, CCIN: 6B99) is a full-height PCIe Gen4 cable adapter that is supported on the IBM Power E1150 server. It provides two ports for connecting expansion drawer cables.

A single adapter supports either one Enclosure Services Manager (ESM) in an NED24 NVMe expansion drawer or one PCIe4 6-slot fanout module in an ENZ0 PCIe4 expansion drawer.

The supported cable adapter for IBM Power E1150 server is listed in Table 3-8 on page 78. Figure 3-8 on page 78 shows the PCIe Gen4 cable adapter (#EJ24).

Table 3-8 Supported cable adapters on the IBM Power E1150 server.

Feature code	CCIN	Description	FRU
EJ2A	6B92	PCIe Gen 4 PCIe Expansion Cable Card	02WF534

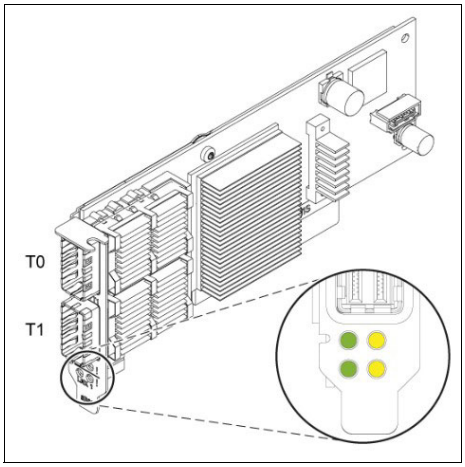

Figure 3-8 PCIe Gen4 cable adapter with FC EJ24

The EJ2A adapter card (Feature Code EJ2A, CCIN 6B99) is a PCIe Gen4 x16 cable adapter that is designed to provide high-speed connectivity between the system and supported I/O expansion drawers. It is a full-height, half-length card featuring two ports for connecting expansion drawer cables. The adapter can be installed in either PCIe Gen4 x16 or x8 slots. However, although using x8 slots allows for greater flexibility in connecting more drawers, it delivers lower performance compared to installation in full x16 slots. See Table 3-5 on page 73 for the recommended adapter locations depending on the expansion drawer being attached.

The EJ2A adapter is a key component for extending the PCIe bus from the IBM Power E1080 system to external I/O resources, enabling scalable configurations for high-performance workloads. It is equipped with 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 be used with either optical or copper cables to connect to PCIe Gen4 expansion drawers. A single adapter connects to one fan-out module in the Gen4 expansion drawer or to one of the ESMs within the NED24 drawer. To use both fan-out modules in a Gen4 expansion drawer, two EJ2A adapters are required. Two adapters are also required when connecting a NED24 expansion drawer,

Internally, the EJ2A includes a built-in PCIe switch, which enables it to extend the E1080's internal PCIe Gen4 bus to external I/O drawers while maintaining bandwidth and performance integrity. The adapter’s physical layout and connectivity are illustrated in Figure 3-9 on page 79.

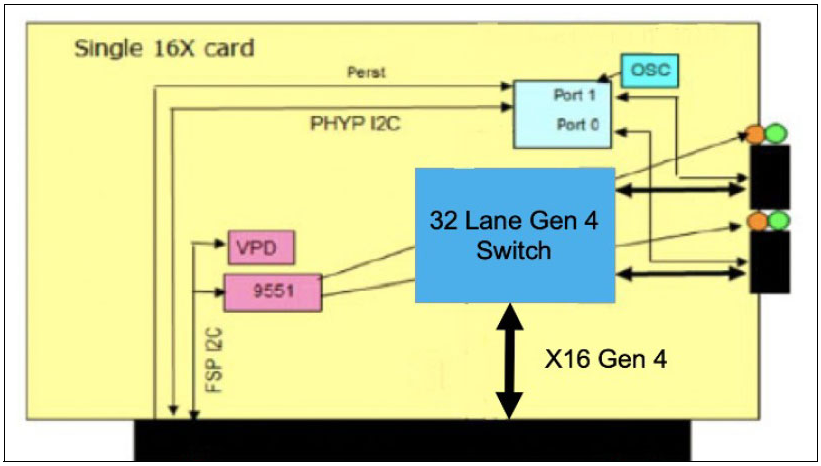

Figure 3-9 PCIe Expansion card

The number of supported cable card adapters for Power E1150 systems are listed in Table 3-5 on page 73.

Cables

Two cables are required to attach each PCIe expansion card to the expansion drawer. Cables are available as either copper or active optical (AOC) with different lengths available. The two cables that are used to connect a card to an expansion drawer must be the same length.

Cables are ordered with a single feature code, which includes two physical cables, either copper or optical.

Restriction: You cannot mix copper and optical cables on the same expansion drawer. Either both fan-out modules use copper cables or both use optical cables.

Table 3-9 lists the supported cables for the ENZ0 PCIe Gen4 I/O expansion drawer and IBM Power11 systems.

Table 3-9 Supported cables for ENZ0 PCIe Gen4 I/O expansion drawer.

FC	CCIN	FRU	Description
ECLS	C1B0	03NG620	3-meter expansion drawer cable (copper)
ECLX	C1B3	78P7688	3-meter active optical cable (AOC)
ECLY	C1B2	78P7689	10-meter active optical cable (AOC)

The supported cables for the NED24 NVMe expansion drawer and IBM Power11 systems are listed in Table 3-10

Table 3-10 Supported cables for NED24 NVMe expansion drawer.

FC	CCIN	FRU	Description
ECLS	C1B0	03NG620	3-meter expansion drawer cable (copper)
ECLX	C1B3	78P7688	3-meter active optical cable (AOC)
ECLY	C1B2	78P7689	10-meter active optical cable (AOC)

The cable feature codes are:

- (#ECLS) - 3.0 M CXP x16 Copper Cable Pair for PCIe4 expansion drawer

  This 3.0-meter cable pair connects a PCIe4 fan-out module in the PCIe Gen4 I/O expansion drawer to a PCIe4 Optical Converter Adapter in the system unit. There are two identical copper cables in the cable pair, each with two CXP connectors. One of the cables attaches to the top CXP port of the PCIe4 fan-out module and to the top CXP port of the PCIe4 Optical Converter Adapter. The other cable attaches to the bottom CXP port.

- (#ECLX) - 3.0 M Active Optical Cable x16 Pair for PCIe4 expansion drawer

  The 3.0-meter active optical cable (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. There are two identical cables in the cable pair, each with two CXP connectors. One of the cables attaches to the top CXP port of the PCIe4 module and to the top CXP port of the PCIe4 Optical Converter Adapter. The other cable attaches to the bottom CXP port.

- (#ECLY) - 10 M Active Optical Cable x16 Pair for PCIe4 expansion drawer

  The 10-meter active optical cable (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. There are two identical cables in the cable pair, each with two CXP connectors. One of the cables attaches to the top CXP port of the PCIe4 module and to the top CXP port of the PCIe4 Optical Converter Adapter. The other cable attaches to the bottom CXP port.

The choice of copper or optical cables depends on your specific configuration and length of cables required. Optical AOC cables are much thinner and can be longer such as the features ECLX (3 M optical), ECLY (10 M optical).

Active Optical Cables

The AOC is a fiber cable and two active electrical-to-optical converter modules that are combined into one assembly, also known as CXP converters. The AOC has a minimum-allowed bend radius of 1 in (25 mm).

Figure 3-10 shows the AOC connector.

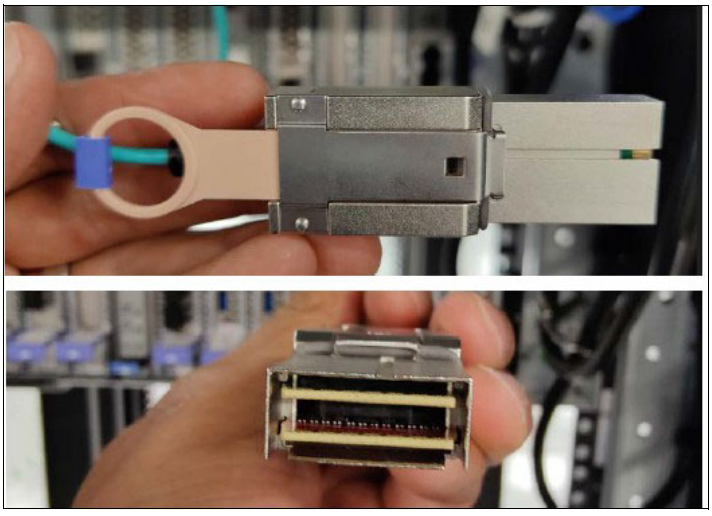

Figure 3-10 AOC connector

Note: Use the 3 m cables for intra-rack installations. Use the 10 m cables for inter-rack installations.

Connectivity to PCIe expansion drawer

A PCIe Gen4 I/O expansion drawer with two I/O fanout modules is connected to one host system node via two PCIe4 cable adapters with four expansion drawer cables (two expansion drawer cable pairs). One pair is used for each of the PCIe4 6-slot fanout modules.

Figure 3-1illustrates the connection of two expansion drawer cable pairs for two PCIe4 6-slot fanout modules.

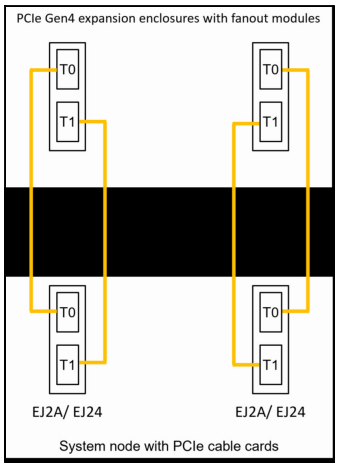

Figure 3-11 Cabling setup for PCIe Gen4 expansion drawer

Connectivity to NED24 expansion drawer

Figure 3-12 shows the connectivity to the NED24 expansion drawer. Both connections to the drawer must be populated and be from the same system.

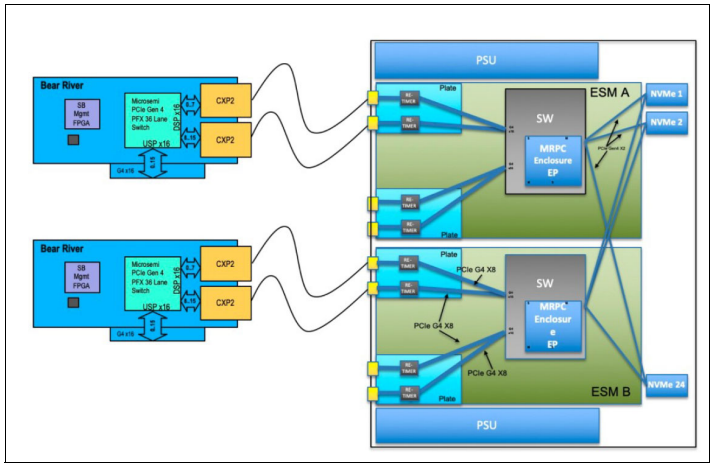

Figure 3-12 Connectivity to NED24 drawers

3.3 List of supported adapters on IBM Power E1150

This section describes the various types and functions of the PCIe adapters that are supported by the IBM Power E1150 server.

This list is subject to change as more PCIe adapters are tested and certified or when listed adapters are no longer available.

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

The Order type table column in the following subsections is defined as:

Supported Denotes that the feature is supported, but no longer orderable with a new system.

Both Denotes the orderability of a feature as part of new and MES upgrade purchases.

The adapters that are supported by the E1150 server are listed in the Table 3-11.

Table 3-11 includes each adapter's Feature Code (FC), description, Customer Card Identification Number (CCIN) and FRU number.

Table 3-11 Adapters supported on the IBM Power E1150 server.

Feature code	CCIN	Description	FRU	Order type
EC2U	58FB	PCIe3 2-port 25/10 Gb NIC & RoCE SFP28 adapter	01FT753	Supported
EC76	2CFB	PCIe4 2-port 100 GbE RoCE x16 adapter	02CM921	Both
EC72	2CF9	PCIe4 2-port 25/10/1 Gb RoCE SFP28 adapter	03HD074	Both
EN26	EC2A	PCIe4 x16 4-port 25/10/1 GbE RoCE SFP28 adapter	03HD066	Both
EC86	EC2C	PCIe5 x16 2-port 200 GbE RoCE adapter	03HD082	Both
EN2W	2F04	PCIe3 4-port 10 GbE adapter	03JP016	Both
EN1J	579C	PCIe4 x8 2-port Fibre Channel (32 Gb/s)	02CM909	Both
EN1A	578F	PCIe3 x8 2-port Fibre Channel (32 Gb/s)	01FT704	Both
EN1L	2CFC	PCIe4 x8 4-port Fibre Channel (32 Gb/s)	03HD014	Both
EN1N	2CFD	PCIe4 x8 2-port Fibre Channel (64 Gb/s)	03HD020	Both
EN2L	2F06	PCIe4 x16 4-port Fibre Channel (32 Gb/s)	03JP004	Both
EN2N	2F05	PCIe4 x8 2-port Fibre Channel (64 Gb/s)	03JP010	Both
EC6K	590F	PCIe2 2-Port USB 3.0 Adapter	02JD518	Both
EJ35	C0AF	4769-001 Cryptographic Coprocessor	02JD572	Supported
EJ37	C0AF	PCIe3 Crypto Coprocessor BSC-Gen3 4769	02JD572	Supported
EPG5	C138	4770 Cryptographic Coprocessor	03JP117	Both
EJ2A	6B99	PCIe4 cable adapter	02WF534	Both

Transceiver replacement for fibre channel and Ethernet adapter features

IBM supports the replacement of transceivers (SFPs) for certain Fibre Channel 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.

Additionally, some Fibre Channel 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 the latest information on adapters that support replacement of the transceivers and those that do not see the following IBM support articles:

- 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.

3.3.1 Fibre Channel adapters

IBM Power11 processor-based Enterprise Midrange servers support connectivity to Fibre Channeldevices, either directly or through a storage area network (SAN). A variety of PCIe-attached FC adapters are available, offered in both low-profile and full-height form factor

All supported Fibre Channel adapters use LC-type connectors. If you are connecting to a switch or device that uses an SC-type fiber connector, you need either an LC-to-SC 50-micron fiber converter cable (feature code #2456) or an LC-to-SC 62.5-micron fiber converter cable (feature code #2459), depending on the fiber specification.

Table 3-12 lists the Fibre Channel adapters that are supported on the IBM Power E1150 server.

Table 3-12 Fibre Channel adapters supported on the IBM Power E1150 server.

Feature code	Description	OS support[a]
EN1J	PCIe4 x8 2-port Fibre Channel (32 Gb/s); (FC EN1J and FC EN1K; CCIN 579C); Adapter part number: 02CM909	AIX and Linux
EN1A	PCIe3 x8 2-port Fibre Channel (32 Gb/s); (FC EN1A and EN1B); CCIN 578F); Adapter part number: 01FT704	AIX and Linux
EN1L	PCIe4 x8 4-port Fibre Channel (32 Gb/s); (FC EN1L and FC EN1M; CCIN 2CFC); Adapter part number: 03HD014	AIX and Linux
EN1N	PCIe4 x8 2-port Fibre Channel (64 Gb/s); (FC EN1N and FC EN1P; CCIN 2CFD); Adapter part number: 03HD020	AIX and Linux
EN2L	PCIe4 x16 4-port Fibre Channel (32 Gb/s); (FC EN2L and EN2M; CCIN 2F06); Adapter part number: 03JP004	AIX and Linux
EN2N	PCIe4 x8 2-port Fibre Channel (64 Gb/s); (FC EN2N and FC EN2P; CCIN 2F05); Adapter part number: 03JP010	AIX and Linux

a. Check for specific AIX or Linux OS requirements in the E1150 Sales Manual

3.3.2 Network adapters

To connect IBM Power E1150 server models to a local area network (LAN), supported LAN adapters can be installed in the system's PCIe slots. A variety of connection speeds and physical interfaces are available, depending on the selected adapter.

Table 3-13 on page 84 lists the LAN adapters that are supported on IBM Power E1150 server.

Table 3-13 List of supported LAN adapters compatible with the IBM Power E1150 server.

Feature code	Description	OS support[a]
EC2U	PCIe3 2-port 25/10 Gb NIC & RoCE SFP28 adapter (FC EC2T and EC2U; CCIN 58FB); Adapter part number: 01FT753	AIX and Linux
EC76	PCIe4 2-port 100 GbE RoCE x16 adapter (FC EC75 and FC EC76; CCIN 2CFB); Adapter part number: 02CM921	AIX and Linux
EC72	PCIe4 2-port 25/10/1 Gb RoCE SFP28 adapter (FC EC71 and FC EC72; CCIN 2CF9); Adapter part number: 03HD074	AIX and Linux
EN26	PCIe4 x16 4-port 25/10/1 GbE RoCE SFP28 adapter (FC EN24 and FC EN26; CCIN EC2A); Adapter part number: 03HD066	AIX and Linux
EC86	PCIe5 x16 2-port 200 GbE RoCE adapter (FC EC85 and FC EC86; CCIN EC2C); Adapter part number: 03HD082	AIX and Linux
EN2W	PCIe3 4-port 10 GbE adapter (FC EN2W and EN2X; CCIN 2F04); Adapter part number: 03JP016	AIX and Linux

a. Check for specific AIX or Linux OS requirements in the E1150 Sales Manual

3.3.3 NVMe options

This section lists the supported NVMe drives supported on the IBM Power E1150. Table 3-14 lists the Feature Codes and provides the size and OS support.

Table 3-14 List of supported NVMe drives compatible with the IBM Power E1150 server.

Feature code	Description	FRUs	OS support	Order types
EC5X	Mainstream 800 GB SSD PCIe3 NVMe U.2 module;	02YC615	AIX and Linux	Supported
EC7T	800 GB Mainstream NVMe U.2 SSD 4k;	02YC661	AIX and Linux	Supported
ES5A	Enterprise 800 GB SSD PCIe4 NVMe U.2 module;	TBD*	AIX and Linux	Both
ES5C	Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module;	TBD*	AIX and Linux	Both
ES4B	Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module;	02YC896 or 02YC912*	AIX and Linux	Both
ES1E	Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module;	01LU967	AIX and Linux	Supported
ES3B	Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module;	02YC725 or 02YC739* (FC EKF3)	AIX and Linux	Supported
ES5E	Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module;	TBD*	AIX and Linux	Both

Feature code	Description	FRUs	OS support	Order types
ES4D	Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module;	02YC897 or 02YC913*	AIX and Linux	Both
ES1G	Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module;	01LU968	AIX and Linux	Supported
ES3D	Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module;	02YC726 or 02YC740* (FC EKF5)	AIX and Linux	Supported
ES5G	Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module;	TBD	AIX and Linux	Both
ES4F	Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module;	02YC898 or 02YC914*	AIX and Linux	Both
EC5V	Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module;	01LU969	AIX and Linux	Supported
ES3F	Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module;	02YC727 or 02YC741*	AIX and Linux	Supported
ECT9	15.3 TB Mainstream NVMe U.2 SSD 4k;	02YC749 or 02YC940* (FC EKF7)	AIX and Linux	Both

Differences between mainstream and enterprise SSD

It is important to note that some of the NVMe options that are listed in Table 3-14 on page 84 are listed as mainstream drives and some are listed as enterprise drives. This section explores the key differences between enterprise NVMe SSDs and mainstream NVMe SSDs. It is important to match the NVMe drive to your workload requirements, specifically the number of writes anticipated to be made to the device.

- Endurance and use case

  Enterprise NVMe SSDs are designed for high-endurance workloads, supporting up to 3 Drive Writes Per Day (DWPD). Mainstream NVMe SSDs, by contrast, usually support ~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 impacts performance and lifespan during write-heavy operations.

- Cost

  Mainstream NVMe SSDs are more cost-effective because of lower over-provisioning and endurance specs, offering lower cost per GB but reduced performance in write-intensive tasks.

- Monitoring and lifecycle management

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

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

For more information, Checking the amount of remaining life in VMe devices.

3.3.4 Internal USB ports and USB adapters

This section discusses the available USB ports and USB adapters and details the available USB devices.

USB ports

In the IBM Power E1150 server, the first DCM (DCM0) also hosts the USB controller, which is connected via four PCI host bridges (PHBs), although only a single lane is used. DCM0 provides a total of four USB 3.0 ports: two on the front and two on the rear of the system.

The two front USB ports are routed from the USB controller on the TPM card, through the DASD backplane, and then to the system backplane, where the front USB cable is connected. These front ports support up to 1.5 A of current, primarily to accommodate devices such as the external USB DVD drive (Feature Code EUA5).

The two rear USB ports are routed from the same USB controller on the TPM card, passing through the system system board and the eBMC card to the eBMC tailstocks. Each of the rear USB ports supports up to 0.9 A of current.

Note: The USB controller is placed on the trusted platform module (TPM) card due to space reasons.

If needed, you can deactivate the USB ports by using the ASMI menu. For more information, see section 8.5, “Managing the system by using the ASMI GUI” on page 152.

Devices supported

The stand-alone USB DVD drive (FC EUA5) is an optional, stand-alone external USB-DVD device. It requires high current at 5 V and must use the front USB 3.0 port on the system or a USB port on a USB 3.0 adapter.

The following media are compatible with the Power E1150 server:

- Read support: CD-ROM, CD-R, CD-RW, DVD-R, DVD+R, DVD-RW, DVD+RW, DVD-ROM, and DVD-RAM discs
- Write support: 4.7 GB DVD-RAM discs
- Read speeds:
  - CD: maximum of 24×
  - DVD-ROM: maximum of 8×
  - DVD-RAM: maximum of 5×
- Write speed (DVD-RAM): maximum of 5×
- Buffer size: 0.75 MB (non-configurable)

The stand-alone USB DVD drive includes the following specifications:

- Interface: USB
- Connector: USB 2.0
- Loading tray: supports 12 cm and 8 cm discs
- Operating positions: horizontal only (on a flat stable surface or floor)
- Form factor: Stand-alone USB DVD drive
- DVD video playback: not supported

Note: A USB extension cable (P/N 32N1311) is included. It is intended for use when there are no safe, flat surfaces available within the rack. The extension allows the drive to reach the floor. Alternative or additional extension cables are not supported. The total USB cable length must not exceed 3 meters.

PCIe2 2-Port USB 3.0 adapter

The PCIe2 2-Port USB 3.0 adapter is a high-performance, PCI Express (PCIe) Generation 2 expansion adapter. FC EC6K is a full-height adapter and is supported on the E1150 server and S1124. FC EC6J is a low-profile adapter and is supported in S1122. FC EC6K is listed in 3.3.5, “IBM PCIe Cryptographic Coprocessor” on page 87. The adapters provide the following features and capabilities:

- Compliant with the PCIe Base Specification Revision 2.0
- Single-lane (x1) PCI Express interface with a throughput of up to 5 Gbps
- Single-slot, half-height, half-length PCIe2 form factor (EC6J)
- FCC Class A compliant
- Provides two external, downstream SuperSpeed USB 3.0 ports with Type A connectors
- Compatible with USB 2.0 and USB 1.1 devices
- Supports simultaneous operation of multiple USB 3.0, USB 2.0, and USB 1.1 devices

3.3.5 IBM PCIe Cryptographic Coprocessor

The IBM 4769 and 4770 Cryptographic Coprocessors, based on PCI Express (PCIe) Gen3 x4 adapters, are supported on Power E1150 servers. These adapters use the IBM Common Cryptographic Architecture (CCA) to accelerate cryptographic workloads. For more information, see IBM PCIe Cryptographic Coprocessor.

PCIe Gen3 4770 Cryptographic Coprocessor

At the time of writing, the 4770 is the fastest generation of PCIe-based hardware security modules (HSMs). It delivers top-tier security processing and high-speed cryptographic performance, providing high throughput to minimize latency and eliminate bottlenecks. Additionally, the 4770 supports FPGA updates and offers acceleration for Dilithium-based post-quantum cryptography.

The 4770 Cryptographic Coprocessor provides the following features:

- Enhanced hardware to support over 300 asymmetric, symmetric and hashing algorithms including hardware support for Quantum Safe Algorithms
- Quantum Safe protected firmware through the use of parallel signatures (ECDSA and CRYSTALS-Dilithium)

The Secure Key Adapter combines both cryptographic coprocessor and accelerator functions in a single PCIe card. The 4770 PCIe Cryptographic Coprocessor is designed for applications that require high-speed, security-sensitive cryptographic operations, such as RSA acceleration, data encryption, and digital signing.

This adapter is also well suited for secure key management and custom cryptographic applications. It provides tamper-resistant, secure hardware storage for cryptographic keys and is designed to comply with FIPS 140-2 Level 4 security requirements.

Note: The next Gen HSM, the PCIe Gen3 4770 adapter, operates in dedicated mode only.

This adapter is available in both low-profile and full-height form factors. Feature Code (FC) EPG4 corresponds to a low-profile adapter and is supported in the S1122. Feature codes EPG5 and EPG6 are full-height adapters are supported in the E1150 and S1124. FC EPG5 is a full-height adapter without a blind-swap cassette. FC EPG6 is a full-height adapter that includes a Generation 3 blind-swap cassette.

PCIe Gen3 4769 Cryptographic Coprocessor

The 4769 Cryptographic Coprocessor is a PCI Express (PCIe) Gen3 x4 adapter that provides both cryptographic coprocessor and cryptographic accelerator functions in a single card. It is designed for applications requiring high-speed, security-sensitive cryptographic operations, such as RSA acceleration, data encryption, and digital signing.

This adapter is also well suited for secure management and usage of cryptographic keys and for custom cryptographic applications. It offers secure key storage within a tamper-resistant hardware security module, which is engineered to meet FIPS 140-2 Level 4 security requirements.

Note: The 4769 operates exclusively in dedicated mode.

Feature codes EJ35 and EJ37 refer to identical adapters with the same CCIN (C0AF). The difference between them lies in the cassette configuration:

- FC EJ35 is provided without a blind-swap cassette.
- FC EJ37 includes a Generation 3 blind-swap cassette.

Table 3-15 summarizes the supported cryptographic coprocessor and accelerator adapters for IBM Power E1150 servers.

Table 3-15 Cryptographic Adapter Features for IBM Power E1150 server.

Feature code	Description	OS support
EJ35	4769-001 Cryptographic Coprocessor	AIX and Linux
EPG5	4770 Cryptographic Coprocessor	AIX and Linux

3.3.6 PCIe adapters that require increased cooling

Some PCIe adapters generate higher levels of heat and are often designated as *hot adapters*. These adapters often have strict temperature limits. The Power E1150 server supports a variety of such adapters, and to ensure adequate cooling, the minimum fan speed is increased when these adapters are installed. The minimum fan speed varies depending on the specific adapter and the ambient temperature.

Note: Higher fan speeds can result in a noticeable increase in overall system noise.

Fan activity is often triggered by temperature readings from sensors located on the operator panel. The IBM Power E1150 server is equipped with three sensors located on the operator panel to measure ambient temperature. You can view the current readings from these sensors through the ASMI (eBMC) interface by navigating to Hardware Status and selecting Sensors, as shown in Figure 3-13.

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 3-13):

- Ambient 0 Temp = 25.062°C
- Ambient 1 Temp = 25.422°C
- Ambient 2 Temp = 25.022°C

The firmware selects Ambient 0 Temp (25.062°C) as the controlling value because it represents the median reading among the three sensors.

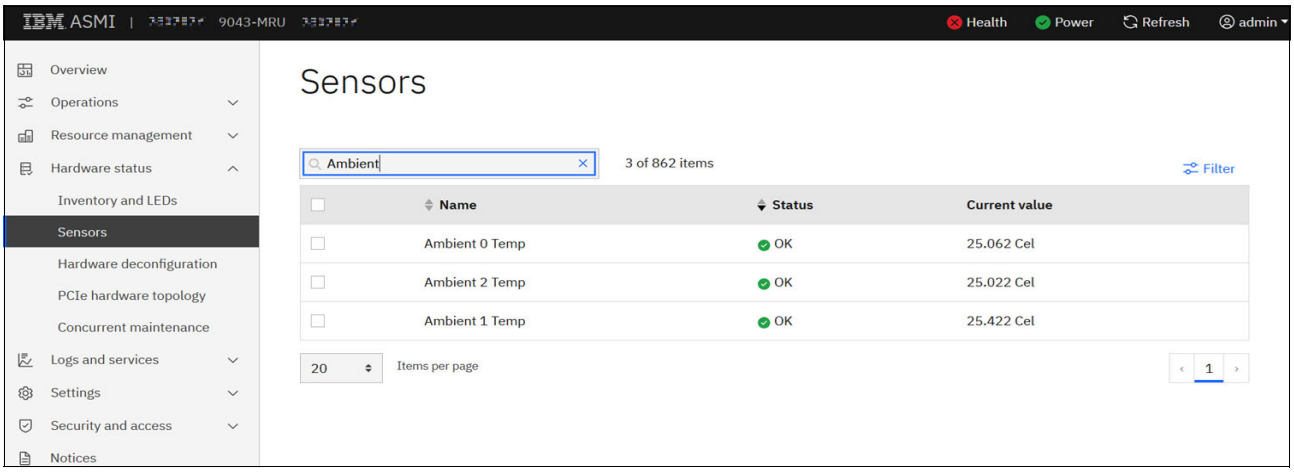

Figure 3-13 Ambient temperature readings viewed through the ASMI GUI (eBMC).

On the E1150 server, if the intake air temperature measured by the front-mounted sensors exceeds 25°C but remains below 30°C, then the fan speed target is set to 5000 rpm as illustrated in Figure 3-14 on page 90. This assumes that no hot adapters are installed.

Note: The fan speed target and the actual speed readings (in rpm) that is displayed in the ASMI GUI (eBMC) under the Sensors tab can vary slightly from the expected values.

If hot adapters are present, the fan speed behavior is adjusted based on the installed adapters and the corresponding ambient temperature range. For more information, see https://github.com/openbmc/phosphor-fan-presence/blob/master/control/config_files/p10bmc/com.ibm.Hardware.Chassis.Model.Fuji/events.json

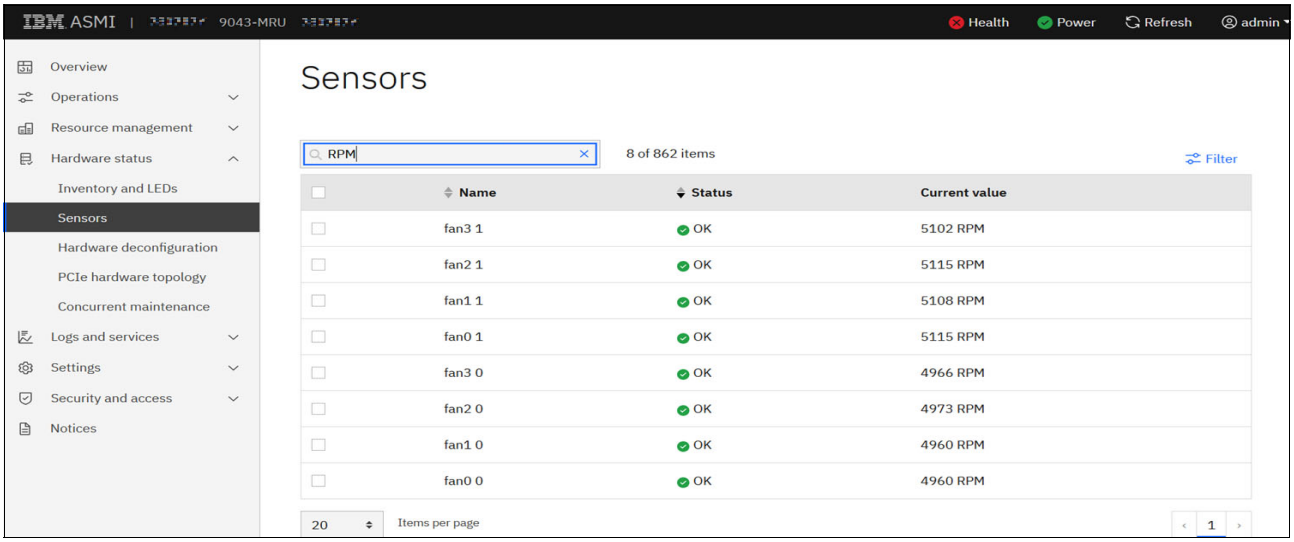

Figure 3-14 Fan speed readings in rpm as viewed through the ASMI GUI (eBMC) under the Sensors tab.

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

Table 3-16 provides a list of adapters that run on a Power11 server and that require increased air flow. Not all of the adapters listed in the table are supported in the E1150, but run in other Power11 servers. When any of these adapters are installed in an E1150, the system will automatically adjust the minimum fan speed in the system to increase the air flow to protect the system.

Table 3-16 PCIe adapters that require increased cooling.

Feature code	CCIN	Description	FRU
EJ2A	6B99	PCIe4 cable adapter	02WF534
EC2S	58FA	PCIe3 2-port 10 Gb NIC & RoCE SR/Cu adapter	01FT759
EC2U	58FB	PCIe3 2-port 25/10 Gb NIC & RoCE SFP28 adapter	01FT753
EC66	2CF3	PCIe4 2-port 100 GbE RoCE x16 adapter	01FT742
EC76	2CFB	PCIe4 2-port 100 GbE RoCE x16 adapter	02CM921
EJ35	C0AF	4769-001 Cryptographic Coprocessor	02JD572
EJ37	C0AF	4769-001 Cryptographic Coprocessor	02JD572
EC86	EC2C	PCIe5 x16 2-port 200 GbE RoCE adapter	03HD082
EN2N	2F05	PCIe4 x8 2-port Fibre Channel (64 Gb/s)	03JP010
EN2L	2F06	PCIe4 x16 4-port Fibre Channel (32 Gb/s)	03JP004
EN26	EC2A	PCIe4 x16 4-port 25/10/1 GbE RoCE SFP28 adapter	03HD066
EN1L	2CFC	PCIe4 x8 4-port Fibre Channel (32 Gb/s)	03HD014

Feature code	CCIN	Description	FRU
EN1N	2CFD	PCIe4 x8 2-port Fibre Channel (64 Gb/s)	03HD020
EC63	2CF1	PCIe4 x16 1-Port EDR 100 GB IB ConnectX-5 CAPI Capable Adapter	00WT179
EC65	2CF2	PCIe4 x16 2-Port EDR 100 GB IB ConnectX-5 CAPI Capable Adapter	00WT176
EC3M	2CEC	PCIe3 2-port 100 GbE NIC & RoCE QSFP28 Adapter	00WT078

3.4 Other device support

This section provides an overview of the available media features with Power E1150

IBM System Storage 7226 Model 1U3 multi-media enclosure

The IBM System Storage 7226 Model 1U3 Multi-Media Enclosure can accommodate up to two tape drives, or up to four DVD-RAM drives.

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

The IBM System Storage 7226 Multi-Media Enclosure is compatible with most Power8, Power9, Power10 and Power11 processor-based systems that offer current level AIX and Linux operating systems.

Figure 3-15 shows the 7226 configured with two DVDs and one Tape Drive

Figure 3-15 7226-1U3 media drawer

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

- IBM Power9 processor-based systems
- IBM Power10 processor-based systems
- IBM Power11 processor-based systems

The IBM System Storage 7226 Multi-Media Enclosure offers several drive feature options, as listed in Table 3-17.

Table 3-17 Supported drive features for the 7226-1U3

Feature Code	Description	Status
#8541	LTO Ultrium 8 Half High SAS Tape Drive	Available
#8546	LTO Ultrium 8 Half High Fibre Channel Tape Drive	Available
#8646	LTO Ultrium 9 Half High Fibre Channel Tape Drive	Available

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

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

For more information, see the sales manual pages.

RDX removable hard disk cartridge

The RDX removable hard disk Cartridge is *not* supported and cannot be purchased with Power E1150.

Alternative solutions to an RDX based Backup strategy are offered:

- Cloud-based Backup Services or
- On-Premise Entry Tape Drives and Libraries

4

Artificial intelligence support

Artificial intelligence (AI) is becoming a cornerstone of digital transformation across industries. It enables organizations to automate processes, gain deeper insights, and deliver more personalized experiences. Predictive analytics, natural language processing, computer vision, and generative AI are driving rapid growth in demand for AI-driven solutions. To support these workloads effectively, enterprises require infrastructure that handles the computational intensity and data throughput that AI demands. IBM Power11 processor-based servers fulfill this requirement.

IBM Power supports AI workloads with its high-performance architecture. The latest Power11 processors feature integrated AI acceleration. These processors handle low-precision arithmetic operations commonly used in AI models, which significantly boost inferencing speed without compromising accuracy. IBM Power also supports advanced technologies such as Power11 on-chip AI inferencing and IBM Spyre Accelerator, a PCIe-attached AI card optimized for enterprise AI workloads. These technologies enable real-time decision-making at scale. These capabilities make Power processor-based servers suitable for deploying AI models in production environments where performance, reliability, and scalability are essential.

IBM Power integrates with Red Hat OpenShift AI so that you can build, train, and deploy AI models within a containerized, hybrid cloud environment. This integration supports consistent DevOps practices, efficient resource usage, and simplified management of AI workflows. IBM Power supports popular AI frameworks such as PyTorch, TensorFlow, and vLLM. This chapter contains the following topics:

- On-chip support
- AI Acceleration with IBM Spyre adapter
- AI solutions on IBM Power11

4.1 On-chip support

With the upcoming availability of the Spyre card on IBM Power-based servers, starting with Power11 systems, organizations can take advantage of a new class of hardware acceleration that is designed to support a wide spectrum of enterprise workloads. The Spyre card delivers significant performance benefits, particularly for compute-intensive and AI-driven use cases, while aligning with the architectural strengths of the IBM Power platform. The Spyre card is not a GPU, but a dedicated hardware device 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 that involves artificial intelligence. This perception often leads to infrastructure decisions that prioritize GPU integration by default, regardless of the nature or scope of the AI tasks to be run. However, GPUs are not a prerequisite for all AI use cases. Their inclusion in every server configuration might result in disproportionate acquisition costs, elevated power consumption, and increased complexity in thermal and workload management, without guaranteed benefits for the workload at hand.

Beginning with IBM's Power10 chip and improved with Power11, each core integrates four Matrix Math Accelerators (MMAs) that can 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 well suited to 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, and audio processing.

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

Also, when the IBM Spyre card is available on IBM Power11 processor-based systems, you can use its additional acceleration capabilities, complementing the on-chip features. IBM Power11 processor-based systems with their integrated AI accelerators and future support for IBM Spyre card present a highly optimized, energy-aware, and secure platform for deploying AI across a broad range of use cases, without the default dependency on discrete GPUs.

4.2 AI Acceleration with IBM Spyre adapter

Improvements to Power11 processor core strength and system capacity increase the performance of the Matrix-Math Accelerator (MMA) for inferencing workloads. Adding IBM Spyre accelerator adapters to Power11 provides additional AI inferencing capabilities. Together, IBM Power processors and the IBM Spyre accelerator enable next-generation infrastructure to scale demanding AI workloads for businesses.

The IBM Spyre card extends the capabilities of Power11 processor-based systems by providing a low-power, high-efficiency acceleration solution for workloads that require frequent memory access and streamlined data movement.

Because of the advanced virtualization and workload consolidation capabilities of IBM Power processor-based servers, Spyre-based applications can run while other mission-critical services run within the same physical server. This configuration allows AI inference engines, real-time analytics, or data preprocessing workloads that use the IBM Spyre card to operate 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 delays typically associated with cross-node or cross-platform communication.

The IBM Spyre card is positioned as one of the only AI solutions that combines the following features:

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

Important: Support for the Spyre adapter on IBM Power is expected to be announced and available late in 2025. Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.

4.2.1 Deploying AI in the enterprise

Artificial intelligence (AI) transforms enterprise operations across industries. It enables organizations to optimize decision-making, streamline processes, and deliver personalized experiences at scale. However, integrating AI into enterprise environments is complex. Successful adoption requires not only appropriate algorithms and models but also a robust, secure, and scalable infrastructure designed for diverse AI workloads.

IBM Power, an AI-optimized platform, addresses this challenge by offering both on-chip and off-chip acceleration capabilities that support varying 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 processed using the on-chip AI acceleration embedded in IBM Power11 processors. This includes Matrix Math Assist (MMA) engines and SIMD vector instructions, combined with high memory bandwidth, which is a critical factor for feeding data into AI models quickly. On-chip acceleration enables real-time inference near the data source, which minimizes latency and reduces the need for additional hardware. This approach is effective for traditional machine learning models, data warehouse analytics, and vector database operations, such as retrieval-augmented generation 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 significantly. At this stage, enterprises require a more flexible and powerful acceleration option, especially as workloads transition from experimentation to production.

4.2.2 Fit for purpose AI acceleration

To address the need for a more flexible acceleration option, IBM introduced IBM Spyre, a dedicated off-chip accelerator designed specifically for AI-intensive workloads. Compared to the on-chip accelerators in Power11, Spyre offers significantly higher throughput, parallelism, and model capacity because of its dedicated memory architecture, optimized data paths, and ability to offload processing from the CPU. Spyre is ideal for large-scale transformer models, such as those

used in generative AI applications, including 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 configuration. Initial AI workloads can begin on existing Power processor-based servers with no added hardware by using the built-in accelerators for cost efficiency. As demands grow, the Spyre accelerators can be integrated to boost performance without requiring a platform change or software rewrite. This flexibility protects existing investments and enables future scaling, as shown in Figure 4-1.

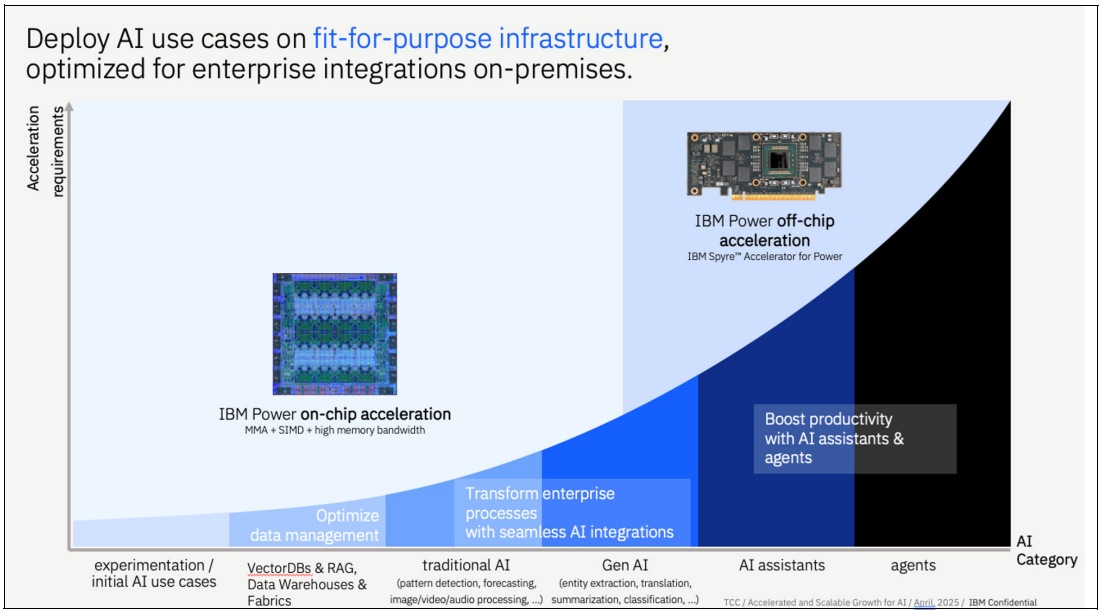

Figure 4-1 IBM's fit for purpose AI architecture for Power

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, which exceed the practical limits of CPU-based inference alone. IBM Spyre is built for such scenarios, and when paired with Power11, it provides a hybrid architecture that is secure, performant, and manageable within enterprise IT constraints.

Additionally, this fit-for-purpose infrastructure aligns well with existing on-premises environments, especially those where data sovereignty, latency, and compliance are nonnegotiable. Unlike cloud-only solutions, the Power11 and Spyre 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 you are starting with lightweight models or scaling to multimodal AI agents, you can use IBM Power11 processor-based servers to deploy AI on your terms. With on-chip acceleration for efficient starting points and Spyre off-chip accelerators for high-performance growth, businesses gain the ability to move from AI exploration to enterprise-wide transformation, all on a platform built for the future of AI.

4.2.3 IBM Spyre accelerator adapter

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 onboard, and each Spyre accelerator is mounted on a PCIe card. Jointly designed by IBM Research and IBM Infrastructure, the IBM Spyre™ architecture is designed for more efficient AI computation. Notably, the chip sends data directly from one compute engine to the next, leading to an efficient use of energy. This family of processors also uses a range of lower precision numeric formats, such as int4 and int8, to make running an AI model more energy efficient and far less memory intensive.

IBM Spyre Is designed with a system-on-a-chip architecture that is optimized for enterprise AI workloads. Spyre is implemented in 5 nm technology with a high-performance and low-power design with 75W consumption and provides the following capabilities:

- 32 low-power AI cores.
- Supports multi-precision for inference training: FP16/8, INT8/4.
- Supports Foundation Models enabled in the Red Hat OpenShift Container Platform AI
- Supports popular AI Framework and libraries (PyTorch, vLLM)

Figure 4-2 shows the IBM Spyre adapter.

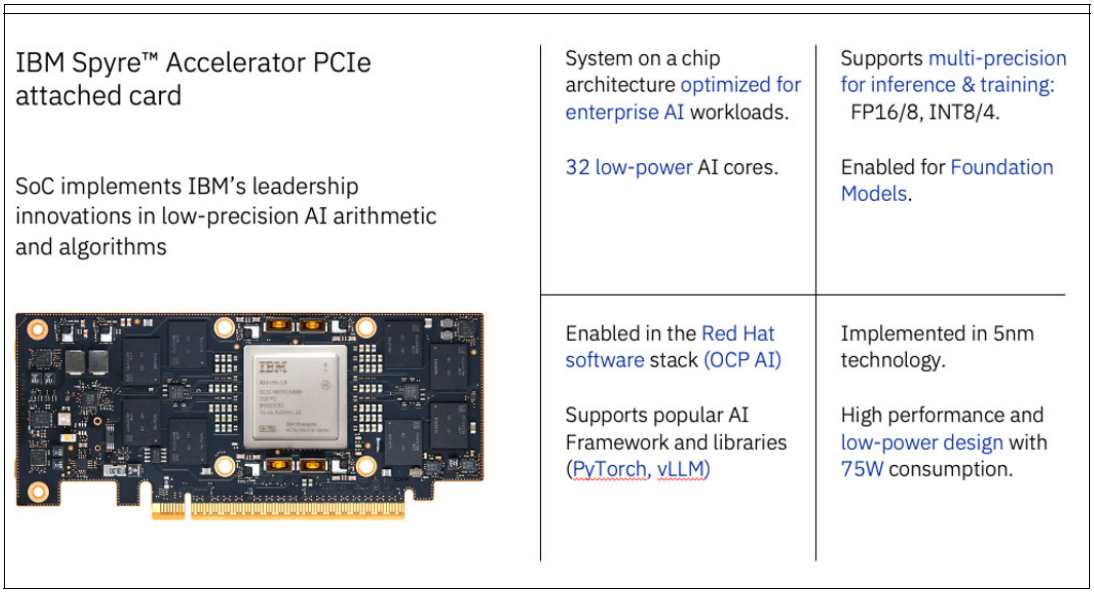

Figure 4-2 The IBM Spyre adapter

This dedicated, enterprise-grade AI acceleration chip sits on a 75W PCIe adapter, with 128 GB of LPDDR5 memory to hold a wide variety of LLMs in support of the heterogeneous workloads that are typically seen on IBM Power. A single Spyre adapter cannot provide enough compute capacity for most use cases. Therefore, the IBM solution uses current I/O expansion technology to attach a cluster of eight Spyre adapters in a single I/O expansion drawer, creating a logical cluster. The firmware on those eight cards coordinates the distribution of compute, and transfer of data among the cards. Because of the architecture, the software sees the clusters as one high-performance compute engine, with 1 TB of memory and 1.6 TB/s of memory bandwidth (Figure 4-3 on page 98).

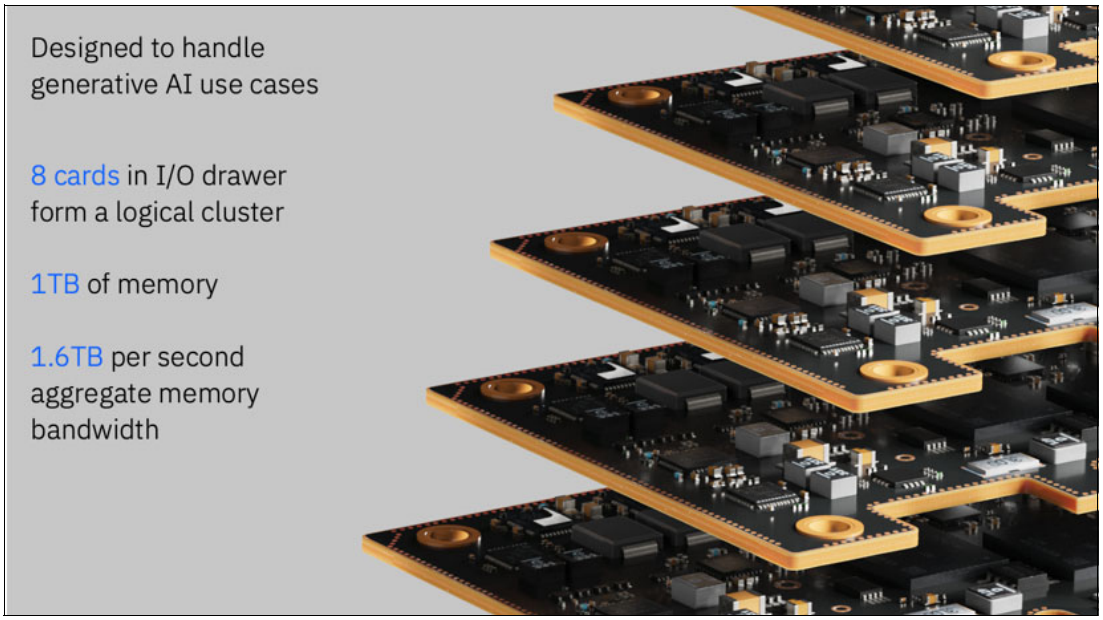

Figure 4-3 Spyre cluster in an expansion drawer

This expansion drawer with Spyre adapters can be attached to any of the announced IBM Power11 servers. The number of supported drawers depends on the Power11 server that it is attached to.

At announcement, Spyre is planned to be available in a fixed configuration which includes eight Spyre adapters installed in the PCIe Gen4 I/O expansion drawer. The Spyre expansion drawer can be attached to any Power11 server: S1122, S1124, E1150, E1180. All models support one expansion drawer except the E1180, which supports up to two. New I/O drawer components such as power supplies and fan-out modules are used to support the additional power requirements of the Spyre adapter configuration.

Initial support includes the following supported software and hardware:

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

4.3 AI solutions on IBM Power11

Artificial intelligence is no longer an emerging trend; it is a foundational pillar of competitive advantage. According to the IBM Institute for Business Value, at the time of writing, approximately 72% of top-performing CEOs identify advanced generative AI as essential to their future success. However, only a fraction of enterprises have successfully moved beyond pilot projects and proofs of concept (PoCs) into full production deployments.

The IBM Power11 processor-based server provides the AI-optimized infrastructure and integrated software stack required to scale AI workloads, secure enterprise data, and modernize mission-critical processes.

4.3.1 Benefits and challenges of implementing AI

Organizations across industries are realizing tangible benefits from deploying AI on IBM Power:

- Logistics service providers can reduce order processing time by embedding generative AI into their ERP systems.
- Financial institutions can accelerate anomaly detection by four times while reducing total cost of ownership (TCO).
- Utility providers can reduce equipment downtime and lower energy consumption by using AI-powered visual inspection.

Benefits of implementing generative AI

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 spanning key sectors:

- Finance and ERP: Fraud detection, anti-money laundering (AML), order processing, and invoice compliance
- Healthcare: Medical transcription, claims and EHR matching, diagnostic imaging, and assistant bots
- IT and development: Predictive IT operations (ITOps), code assistants (RPG, Ansible), transcription, and 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.

Challenges to implementing AI

Although the promise of AI is great, common enterprise challenges persist. Table 4-1 provides some of the challenges to adoption of AI.

Table 4-1 challenges to scale AI

Barrier	Description
Data complexity	Fragmented formats, distributed sources, lack of standardized connectors
AI integration difficulty	Legacy workflows, lack of AI skills, 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

AI on IBM Power11

IBM Power11 addresses these challenges by offering a complete AI platform, integrated into existing systems, with hardware acceleration, data fabric, and AI assistants.

To scale AI in production, IBM Power provides an AI-ready portfolio and includes the following features:

- AI-Ready infrastructure: Built-in accelerators, support for hybrid and cloud-native deployment (for example, PowerVS)
- Optimized software stack: Red Hat, IBM, open source, and certified 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 and maintain governance, performance, and security.

Table 4-2 describes outcome examples that are enabled by the implementation of AI-solutions on IBM Power processor-based servers.

Table 4-2 outcomes examples enabled by IBM Power

Objective	Outcome example
Transform processes	Accelerate ERP or financial workflows integrated with AI services
Boost productivity	Roll out new application features faster with GenAI 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 an enterprise scale.

4.3.2 Uniqueness of the platform

With its AI-optimized infrastructure and enterprise-proven architecture, IBM Power is uniquely positioned to help organizations move from experimentation to real AI-driven transformation. Its scalable ecosystem (spanning hardware, software, data, and automation) makes it a reliable platform for deploying production-grade AI at scale.

A key differentiator of the IBM Power11 processor-based platform lies in its ability to deliver AI capabilities without requiring external GPUs. Unlike conventional assumptions that all AI workloads demand dedicated GPU acceleration, IBM Power11 uses advanced on-core, which is combined with a system-wide infrastructure design optimized for AI and data-intensive tasks.

This architectural approach enables the deployment of a wide range of AI use cases, such as inferencing, predictive analytics, and cognitive automation, while maintaining enterprise-grade attributes including resilience, reliability, performance, security, and sustainability. By eliminating or reducing dependency on external GPUs, organizations benefit from simplified infrastructure, reduced energy consumption, and lower total cost of ownership (TCO), all without compromising AI readiness. The IBM Power11 platform is purpose-built to support scalable and secure AI solutions across diverse industry environments.

5

Automation and management

Managing IBM Power-based processor servers effectively requires a combination of tools that provide automation, orchestration, and visibility across the infrastructure. The Hardware Management Console (HMC) is at the core of this management stack and serves 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 command-line interface, a graphical interface, and a REST API, which allows integration with automation tools for more 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 patching, user management, and software deployment. It supports IBM Power through modules and collections tailored for AIX, IBM i, and Linux on Power. Terraform complements Ansible by enabling declarative provisioning of infrastructure components, including Power Virtual Servers and PowerVC-managed resources. Together, they allow teams to automate the full lifecycle of infrastructure, from provisioning to configuration, ensuring consistency and reducing manual effort.

Power Virtualization Center (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, which enables seamless orchestration of virtual machines 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

Configure and manage your systems by using the Hardware Management Console (HMC), which can be a hardware or virtual appliance. 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
  - Enterprise pools
  - Shared Processor Pools
  - Performance and capacity monitoring
- Launch Advanced System Management Interface (ASMI) for managed systems
- Deliver virtualization management through the following functions:
  - Creation, management, and deletion of logical partitions
  - Live partition mobility
  - Remote restart
  - Configuring SRIOV
  - Managing virtual i/o servers
  - Dynamic resource allocation
  - Operating system terminals
- Act 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 Electronic Service Agent for faster support
- Provide appliance management capabilities for configuring network and users on the HMC and for updating and upgrading the HMC

5.1.1 HMC options

Power11 servers can be connected to either the 7063-CR2 HMC or to a Virtual HMC.

Restriction: The 7063-CR1 is not supported by the Power11 servers.

HMC 7063-CR2

The 7063-CR2 IBM Power HMC (see Figure 2-1) is a second-generation Power processor-based HMC. It includes the following features:

- 6-core IBM Power9 130W processor chip
- 64 GB (4×16 GB) or 128 GB (4×32 GB) of memory
- 1.8 TB of internal disk capacity with RAID1 protection
- Four 1 Gbps Ethernet (RJ-45) ports, two 10 Gbps Ethernet (RJ-45) ports, two USB 3.0 ports (front side) and two USB 3.0 ports (rear side), and 1 Gbps IPMI Ethernet (RJ-45)
- Two 900 W power supply units
- Remote Management Service: IPMI port (OpenBMC) and Redfish application programming interface (API)

The base warranty is 1-year 9×5 (8am–5pm, Monday–Friday) with available optional upgrades. A USB Smart Drive is *not* included.

Figure 5-1 is a picture of a 7063-CR2 HMC.

Figure 5-1 HMC 7063-CR2

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

Virtual HMC

Initially, the HMC was sold only as a hardware appliance, including the HMC firmware installed. However, IBM extended this offering to allow the purchase of the hardware appliance or a virtual appliance that can be deployed on ppc64le architectures or x86 platforms.

Any customer with a valid contract can download the HMC from the Entitled System Support, or it can be included within an initial Power S1122 or S1124 order.

The virtual HMC supports the following hypervisors:

- On x86 processor-based servers:
  - KVM
  - Xen
  - VMware
- On Power processor-based servers: IBM PowerVM

The following minimum requirements must be met to install the virtual HMC:

- 16 GB of Memory
- 4 virtual processors
- 2 network interfaces (maximum 4 allowed)
- 1 disk drive (500 GB available disk drive)

For an initial Power S1122 or S1124 order with the IBM configurator (e-config), locate the HMC virtual appliance by selecting **Add software** → **Other System Offerings** (as product selections) and then choosing:

- 5765-HMD for IBM HMC Virtual Appliance for Power V11
- 5765-HMU for IBM HMC Virtual Appliance x86 V11

For more information and an overview of the Virtual HMC, see Virtual HMC appliance (vHMC) overview. For more information about how to install the virtual HMC appliance and all requirements, see Installing the HMC virtual appliance.

5.1.2 BMC network connectivity rules for 7063-CR2 HMC

The 7063-CR2 HMC features a baseboard management controller (BMC), which is a specialized service processor that monitors the physical state of the system by using sensors. OpenBMC that is used on 7063-CR2 provides a graphical user interface (GUI) that can be accessed from a workstation that includes network connectivity to the BMC. This connection requires an Ethernet port to be configured for use by the BMC.

The 7063-CR2 provides two network interfaces (eth0 and eth1) for configuring network connectivity for BMC on the appliance.

Each interface maps to a different physical port on the system. Different management tools name the interfaces differently. The HMC task **Console Management** → **Console Settings** → **Change BMC/IPMI Network Settings** modifies only the Dedicated interface.

The BMC ports are listed in Table 5-1.

Table 5-1 BMC ports

Management tool	Logical port	Shared/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 BMC/IPMI Network settings	lan2	Dedicated	Management port only

Figure 1-15 shows the BMC interfaces of the HMC.

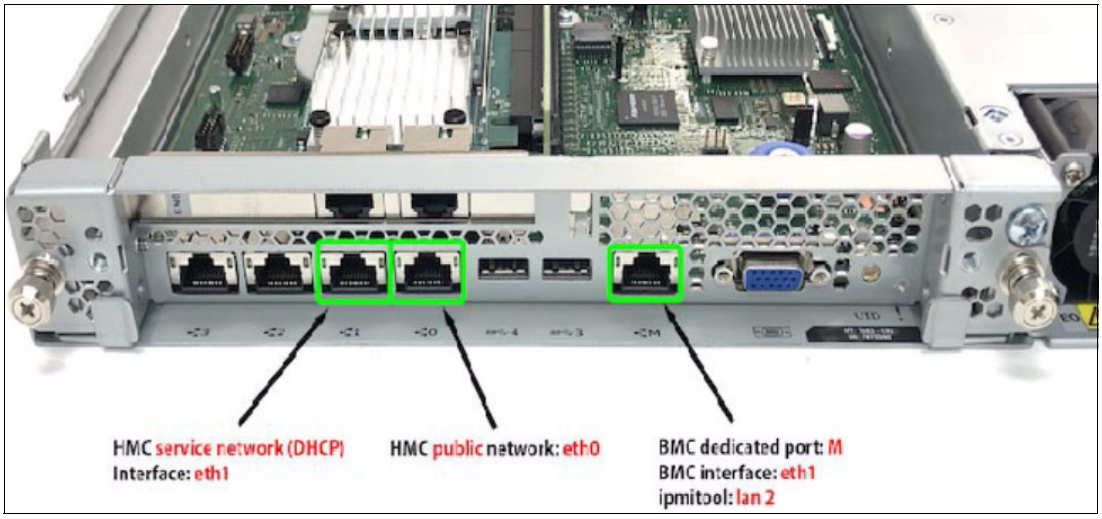

Figure 5-2 BMC interfaces

The primary difference is that the shared and dedicated interface to the BMC can coexist. Each has its own LAN number and physical port. Ideally, the customer configures one port, but both can be configured. The rules for connecting IBM Power to the HMC remain the same as for previous versions.

5.1.3 High availability HMC configuration

For the best manageability and redundancy, a dual HMC configuration is suggested. This configuration can be two hardware appliances, or one hardware appliance and one virtual appliance or two virtual appliances. A dual HMC configuration has the following requirements:

- Two HMCs are at the same version.
- The HMCs use different subnets to connect to the BMCs.
- The HMCs can communicate with the servers' partitions over a public network to allow for full synchronization and function.

5.1.4 HMC code level requirements

The minimum required HMC version for the Power E1150 is V11R1 M1110. Version V11R1 M1110 is supported on 7063-CR2, and Virtual HMC appliances only. It is *not* supported on the 7063-CR1 or the 7042 machine types.

Note: HMC with V11R1M1110 cannot manage POWER7 processor-based systems.

An HMC that is running V11R1M1110 includes the following features:

- Support for managing Power11 systems
- Support for new I/O adapters
- VIOS management enhancements:
  - Resource groups
  - Increase in platform keystore size
  - Remove support for vTPM 1.2
  - Quantum safe LPM
  - Minimum affinity score and actions
- Console management and user experience improvements:
  - User experience improvements:
    - Network Topology
    - Trusted keystore
    - Import Certificate
    - Multi-factor authentication allow list
  - Ability to advertise device information via LLDP
- Power infrastructure maintenance and automated Power platform updates:
  - Run Power platform updates with minimal touchpoints to enhance simplicity and reduce the risk of human error
  - Experience seamless platform updates with one-touch solutions
  - Automatic operational recovery and resiliency
- Autonomous error resolution
  - Ability to collect platform logs (FW, hypervisor, HMC, VIOS) from a single interface
  - Ability to create a case and upload logs to the case from the HMC

- Sustainability
  - New Energy Efficient mode with increased performance per watt
  - Partition level energy monitoring, which includes real-time monitoring and reporting of energy and carbon emissions at the VM or partition level
  - Scheduling of energy modes

5.1.5 *HMC currency*

In recent years, cybersecurity emerged as a national security issue and an increasingly critical concern for CIOs and enterprise IT managers.

The IBM Power processor-based architecture has always ranked highly in terms of 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 ensuring that the HMC is current and fully supported including hardware, software, and Power firmware updates.

Outdated or unsupported HMCs represent a technology risk that can quickly and easily be mitigated by upgrading to a current release.

5.1.6 New features

The minimum level of the HMC required to support Power11 is V11 R1 M1110. V11 of the HMC is not supported on the 7063-CR1. HMC V11 runs on the 7063-CR2 or the virtual HMC. Additionally, HMC V11 does not support Power8 or earlier servers.

HMC V11.1.1110 was generally available in July 2025. Specific new features include support for managing Power11 systems and support for new I/O adapters.

Virtualization management includes the following new features:

- Resource Groups
- Increase in platform keystore size
- Removal of support for vTPM 1.2
- Quantum safe LPM
- Minimum affinity score and actions

The release includes the following console management and user experience improvements:

- User experience improvements for the following configurations:
  - Network topology
  - Trusted keystore
  - Import certificate
  - Multi-factor authentication list
- Ability to advertise device information via LLDP

The release includes the following Power infrastructure maintenance and automated Power platform updates:

- Run Power platform updates with minimal touchpoints to enhance simplicity and reduce the risk of human error, which includes the ability to update system firmware, VIOS, and I/O adapters from a single update flow.

- Experience seamless platform updates with advanced one-touch solutions. Designed for both evacuation and return or in-place updates, each has automatic operational recovery and resiliency.
- Validate LPM and VIOS redundancy (VIOS maintenance readiness check).
- Automatically migrate partitions and return as part of the update process.

The release also includes autonomous error resolution:

- Reduce problem resolution time with ability to collect platform logs such as FW, Hypervisor, HMC, and VIOS, from a single Interface
- Ability to create a case and upload logs to the case from the HMC

HMC V11 includes the following sustainability features:

- Energy Efficient mode with increased performance per watt
- Partition level energy monitoring: Real-time monitoring and reporting of energy and carbon emissions at the VM and partition level.
- Scheduling of energy modes

5.1.7 Using the automated maintenance tool

This section describes using the automated maintenance tool that is provided in the IBM Power11 HMC.

Launch point

Before you launch the automated maintenance, ensure the system has the most recent 11.10 firmware installed.

1. To start the process, select **Update system** → **VIOS** → **adapter levels**. The menu was introduced with Firmware 11.10. This is shown in Figure 5-3.

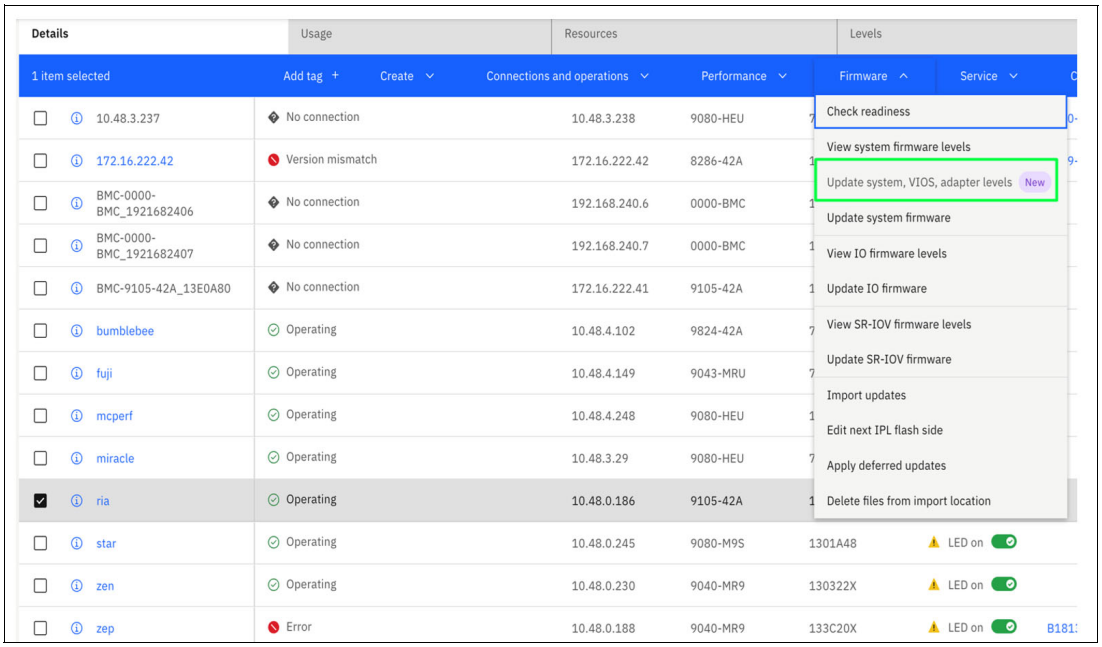

Figure 5-3 Start update

This starts a wizard for the following process:

- Import update files option: “Import files to HMC filesystem and perform update” or “Import files only”
- Check System readiness: System must be in Ready state to proceed with update process
- Select the source file location

2. From the drop-down Update menu, choose **Update** or **Upgrade** as shown in Figure 5-4.

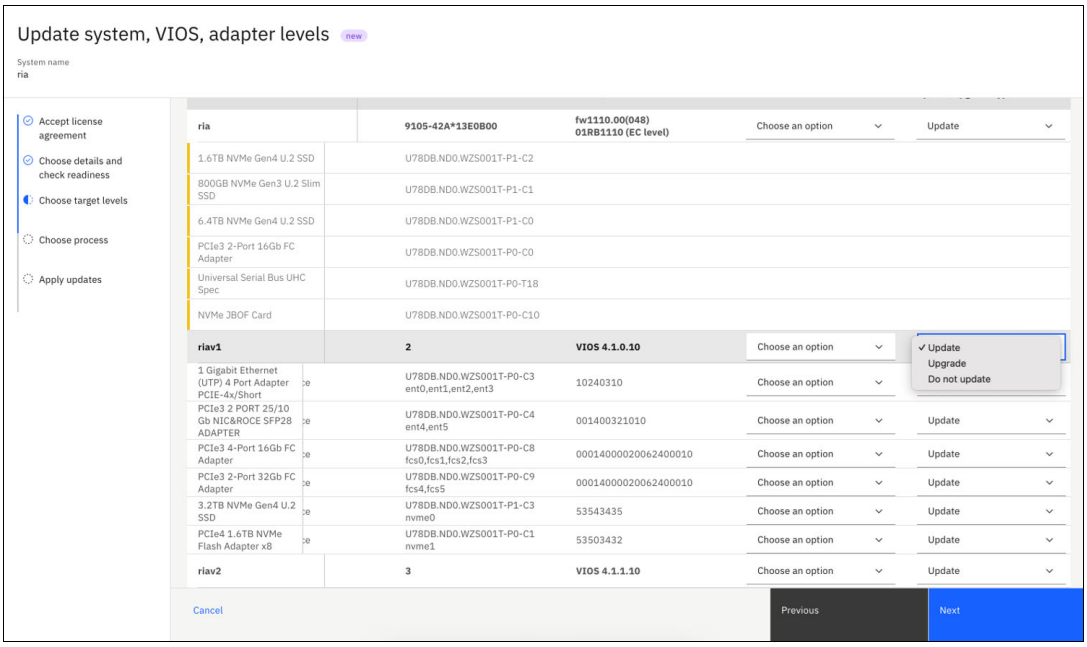

Figure 5-4 Upgrade type

3. To preserve the system availability, for disruptive upgrades and updates, the partitions migration panel is shown as in Figure 5-5.

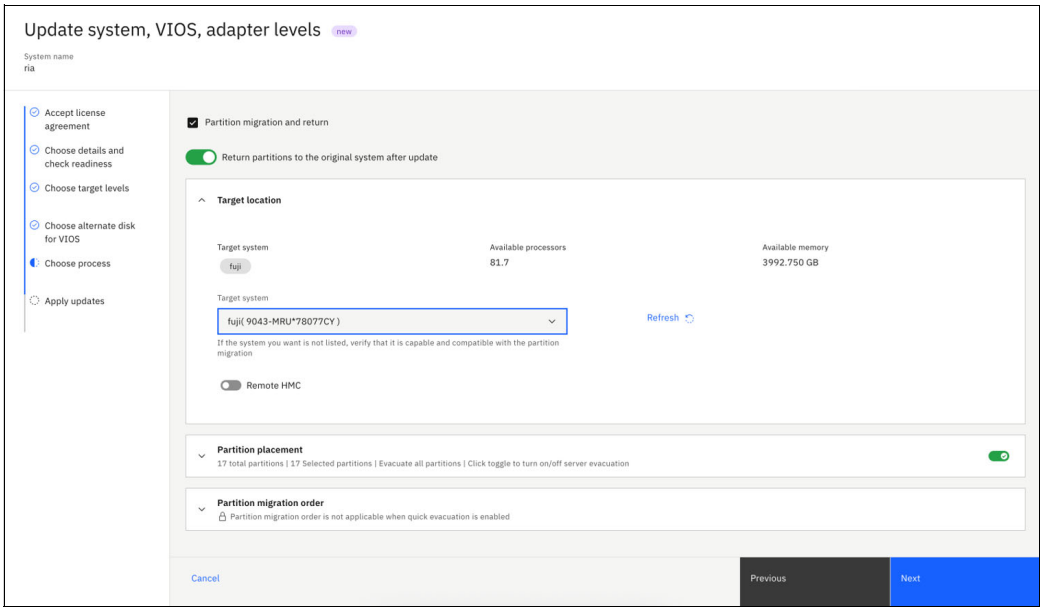

Figure 5-5 Choose Process

You can choose the following options:

- Remote HMC switch: Enables migration of the partitions to a system managed by remote HMC. When selected, partitions are not automatically migrated back to the system after the update.
- Partition placement section: All partitions are selected for migration by default. You can select individual partitions by individually selected to migrate by toggling the server evacuation switch.
- If you select partitions for migration, you can specify the partition migration sequence for a set of partitions.

4. After making your selections, click **Next**.
5. Verify the results (Figure 5-6).

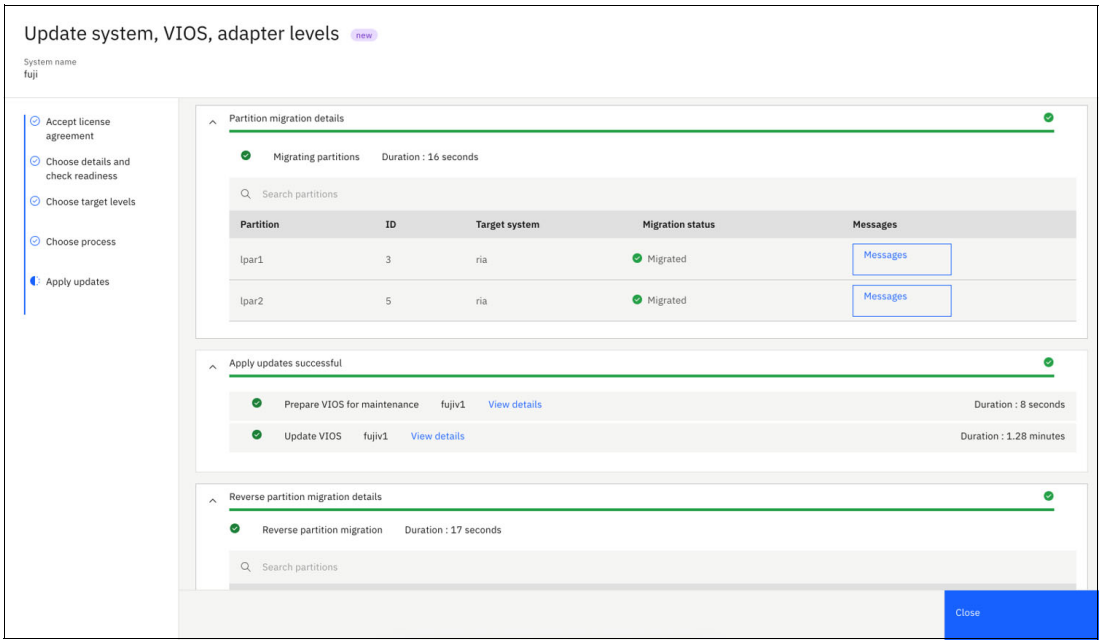

Figure 5-6 Figure 5 – Update status

This solution provides business resiliency with zero planned downtime that is delivered through platform automation capabilities that are integrated with IBM Concert's AI-infused workloads. It enables faster and more frequent maintenance updates, keeping systems secure, stable, and compatible with evolving software and hardware requirements, lowering risks of performance degradation, security breaches, or unplanned downtimes.

5.2 Ansible

Ansible is an open-source, cross-platform tool for resource provisioning automation that DevOps professionals use for continuous deployment (CD) of software code by using an Infrastructure as Code (IaC) approach. The Ansible automation platform changed and delivers automation solutions for operators, administrators, and IT decision-makers across various technical disciplines. It is an enterprise automation solution with open-source software. It operates on several UNIX like platforms, and can manage systems such as UNIX and Microsoft architectures. It comes with descriptive language for describing system settings.

Ansible is gaining broad acceptance of the Ansible platform because of its open-source design and its wide support for many devices and platforms. However, it is also common to use other automation tools with Ansible to do more complex automation. For example, many companies use Ansible with Terraform to provide automatic provisioning of their infrastructure.

As shown in Figure 5-7, the Ansible architecture consists of an Ansible Controller and one or more Ansible client hosts. The controller runs automation tasks and houses Ansible collections, which contain modules, plug-ins, and roles defining the actions Ansible can perform on client nodes.

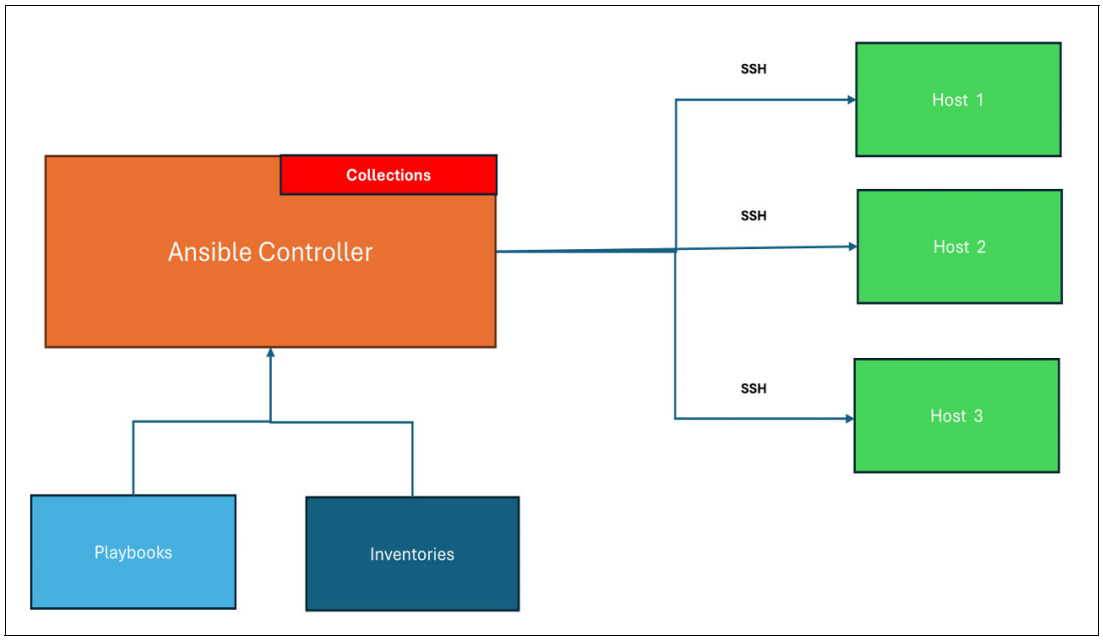

Figure 5-7 Simplified Ansible architecture

The heart of Ansible Automation Ansible playbooks are YAML files that define sequences of tasks to run on remote hosts. These tasks can range from installing packages to configuring services or copying files. Playbooks enable IT teams to automate infrastructure provisioning, configuration management, application deployment, and more.

Benefits of Ansible

Ansible offers numerous benefits for IT professionals seeking to improve efficiency, scalability, and consistency in their infrastructure and includes some key advantages:

- Versatility: Ansible supports a wide range of devices and can scale to accommodate growing environments and automation needs.
- Agentless architecture: Ansible manages devices by using Secure Shell (SSH), which eliminates the need for agents on target systems.
- Flexibility: Ansible can be used for simple CLI tasks and complex workflows that are defined in playbooks.
- Extensive module library: Ansible provides a rich collection of modules for managing various systems, cloud infrastructures, and OpenStack.
- Declarative approach: With the Ansible declarative syntax, you can define the state of a system, and Ansible takes the necessary steps to achieve it.

- Ease of learning: The Ansible YAML syntax and minimal learning curve make it accessible to IT professionals at all levels.

Ansible is a powerful automation tool that can help organizations improve efficiency, scalability, and reliability in their IT infrastructure. By implementing Ansible playbooks, IT teams can streamline routine tasks, automate complex workflows, and help ensure consistent configurations across their environments.

Options for implementing Ansible

As you decide to implement Ansible for IT management, select the correct product and support level to meet your organization's needs. This section describes some of the options that are available to you.

Ansible Community

The community versions of Ansible primarily include the following:

- Ansible Core

  Ansible Core is a fundamental part of Ansible. It provides the core automation engine. It is an open-source tool that includes the 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

  AWX is the upstream, open-source project that serves as the community version of Red Hat Ansible Tower. AWX provides a web-based UI, Representational State Transfer (REST) API, and task engine for managing Ansible automation at scale. AWX offers role-based access control (RBAC), job scheduling, graphical inventory management, and more. It helps users manage and scale automation efforts.

- Ansible Collections

  Ansible Collections are prepackaged modules, roles, and plug-ins that are created and shared by the community. With Collections, users can extend Ansible functions with more content that is often maintained by the community or specific organizations. Collections can be downloaded from Ansible Galaxy, a community hub for sharing and discovering Ansible content.

- Ansible Galaxy

  Ansible Galaxy is a repository for sharing and discovering Ansible roles and collections. It is a community-driven platform where users can find reusable Ansible content to simplify automation tasks. It provides a searchable repository of roles and collections that are created by the Ansible community, which can be integrated into your automation workflows.

These community versions are suitable for individual users, small teams, and development environments but 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 over 20 community projects into a fully supported automation platform. Ansible Automation Platform provides curated, certified, and validated Ansible Collections and roles from partners like IBM, Juniper, Cisco, and public cloud providers.

Consider the following key factors when choosing Ansible Automation Platform:

- Support level: Ansible Automation Platform offers enterprise-grade support, which includes SLAs for security, compatibility, and upgrades. Community options might have limited support.
- Features: Ansible Automation Platform includes features beyond Ansible Core, such as a web interface and integration with other tools.
- Cost: Ansible Automation Platform is a subscription-based product, but community options are available at no charge.
- Scale and complexity: For large organizations with complex automation needs, Ansible Automation Platform might be the better choice due to its enterprise-grade features and support.

By carefully evaluating these factors, you can select the Ansible offering that best 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 brings significant value to managing IBM Power. Several key use cases demonstrate how Ansible enhances efficiency, consistency, and scalability in Power environments:

1. System configuration and provisioning

   Ansible automates the setup and configuration of AIX, IBM i, and Linux on Power. This includes tasks like user and group management, network configuration, software installation, and system tuning. By using Ansible playbooks, administrators can ensure consistent configurations across multiple systems, reducing manual errors and speeding up provisioning.

2. Patch management and compliance

   Keeping systems up to date is critical for security and stability. Ansible can automate the patching process for AIX and IBM i, including downloading patches, applying them, and verifying system health post-update. It also supports compliance checks by validating system configurations against predefined baselines, which helps organizations meet regulatory and security standards.

3. Integration with HMC and PowerVC

   Ansible collections for IBM Power include modules that interact with the Hardware Management Console (HMC) and PowerVC. This enables automation of tasks such as LPAR creation, VIOS configuration, and virtual machine lifecycle management. Combined with dynamic inventory capabilities, Ansible can discover and manage Power infrastructure components in real time.

For more information on 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 originally developed by HashiCorp® and now owned by IBM. It is written in the Go programming language and compiles into an executable file named Terraform. Terraform is an infrastructure as code tool that lets you build, change, and version both 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 that is involved in calling the API. The definition of which APIs to call is defined in configuration files. These configuration files are the code that is referenced in Infrastructure as code.

Figure 5-8 Terraform functionality

You can use Terraform to define and provision infrastructure by using a high-level configuration language called HashiCorp Configuration Language (HCL). With Terraform, infrastructure components such as servers, databases, networking, and storage can be described in code, to enable version control, collaboration, and repeatability. This approach eliminates the need for manual setup and reduces the risk of different configurations across environments.

Terraform operates through a workflow that includes writing configuration files, initializing the working directory, planning changes, and applying them. The **`terraform init`** command sets up the environment by downloading the necessary provider plug-ins. The **`terraform plan`** command creates an execution plan, showing what actions Terraform will take to reach the wanted state. To implement the plan, the **`terraform apply`** command makes the actual changes to the infrastructure. Terraform maintains a state file that tracks the current state of the infrastructure, so you can determine what to change for future runs.

One of Terraform's key strengths is its provider ecosystem, which allows it to manage resources across a wide range of platforms, including PowerVS, AWS, Azure, Google Cloud, Kubernetes, and many others. This makes it a powerful tool for managing hybrid and multi-cloud environments. Additionally, Terraform 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 on page 114 shows how Terraform works. Define your intended end point through configuration files. Place the configuration files into a plan. Apply the plan through infrastructure providers.

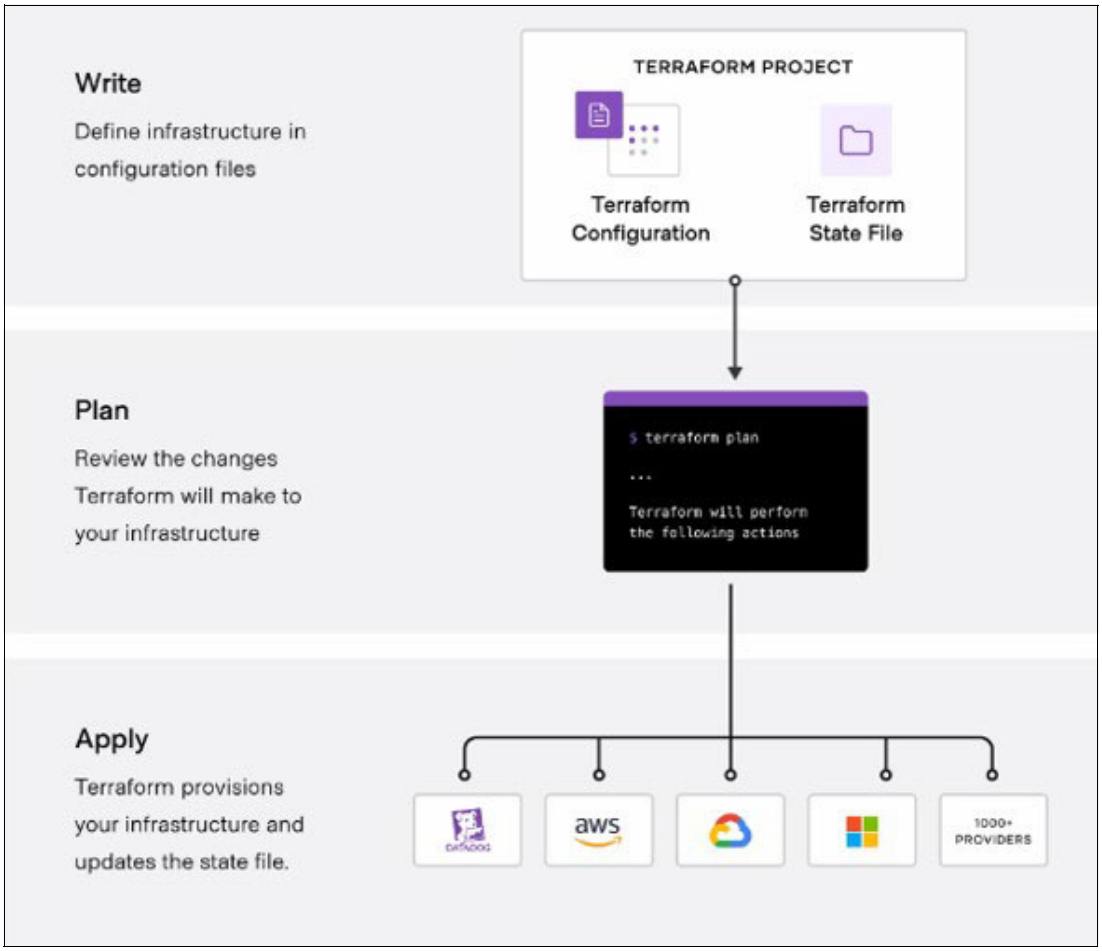

Figure 5-9 Terraform process

You can use Terraform to perform the following tasks:

- Track your infrastructure

  Terraform generates a plan and prompts you for your approval before modifying your infrastructure. The state of your infrastructure is kept in a file named `terraform.tfstate`, which can be held in Git, Gitlab or HCP Terraform to version, encrypt, and securely share it with your team. This acts as the single source for your environment.

- Automate Changes

  Terraform configuration files are declarative, describing the end state and can be automated by using tools such as Ansible.

- Standardize configurations

  Terraform provides standardization in modules. A module consists of a collection of `.tf` and `.tf.json` files kept together in a directory. Modules are the main way to package and reuse resource configurations with Terraform.

- Collaborate

  Terraform can be distributed as configuration files and version-controlled in applications like git, GitHub, and HCP Terraform. These are ideal locations to share and collaborate.

Combining Ansible and Terraform

Terraform connects to any provider, such as Ansible, to manage your infrastructure. Browse the Terraform registry for providers. The Terraform Provider for Ansible provides a more straightforward and robust means of running Ansible automation from Terraform rather than a local-exec[1]. Figure 5-10 shows the Ansible provider entry in the registry.

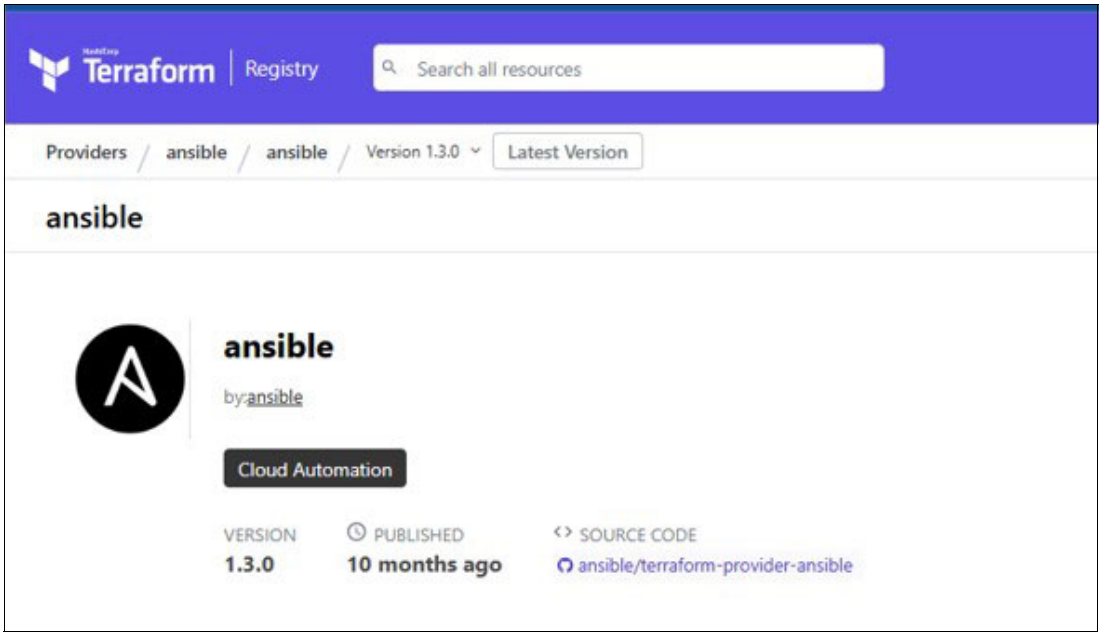

Figure 5-10 Terraform Ansible provider

Install Ansible and its prerequisites:

1. Install Go

   For installation instructions, see Download and install.

2. Install Terraform

   Install the ppc64le version for operation on IBM Power.For installation instructions, see GitHub registry.

3. To install Ansible, see Installing Ansible.

For more information on Terraform in the IBM Power environment, see the IBM Redbooks publication *Modernization Techniques for IBM Power*, SG24-8582

5.4 PowerVC

IBM PowerVC is an advanced virtualization and cloud management platform built on OpenStack, specifically designed for managing virtualized workloads on IBM Power. It supports AIX, IBM i, and Linux operating systems and provides a unified interface for deploying, managing, and automating virtual machines (VMs) across Power environments. PowerVC helps simplify the creation of private clouds and integrates seamlessly with higher-level cloud orchestrators, such as Ansible, Terraform, and Red Hat OpenShift, enabling hybrid cloud strategies.

[1] https://developer.hashicorp.com/terraform/language/resources/provisioners/local-exec

PowerVC offers two editions: Standard Edition and Private Cloud Edition. The Standard Edition includes the following core virtualization management features:

- VM image capture, import/export, and deployment
- Policy-based VM placement for optimized resource utilization
- Snapshots and cloning for backup and testing
- Live VM mobility and remote restart for high availability
- Role-based access control and automated I/O configuration

The Private Cloud Edition builds on this by adding the following features:

- A self-service portal for end users to provision VMs
- Approval workflows for provisioning requests
- Pre-built deployment templates
- Cloud management policies and metering for chargeback

PowerVC supports Dynamic Resource Optimization (DRO), which automatically balances workloads based on CPU and memory usage. PowerVC supports Simplified Remote Restart, which ensures that virtual machines (VMs) restart on alternate hosts after when necessary. These features reduce administrative overhead and improve system resilience.

NovaLink is a lightweight, Linux-based virtualization management interface that acts as a bridge between PowerVC (IBM's cloud and virtualization manager) and PowerVM (the hypervisor on IBM Power). It enables direct, scalable, and efficient control of VMs and system resources without relying solely on the traditional HMC.

Installed on the Power System, NovaLink runs as a service on a dedicated Linux logical partition (LPAR) on the same Power server that it manages. This LPAR communicates directly with the PowerVM hypervisor. NovaLink has direct access to the hypervisor APIs, which provides faster and more granular control of virtualization tasks such as VM creation, deletion, migration, and resource allocation. This design makes PowerVC more responsive and scalable, especially in environments with many VMs or frequent provisioning changes. Table 5-2 compares the features provided by PowerVC management to the features provided by HMC.

Table 5-2 Comparisons of PowerVC and HMC management:

Feature/ Capability	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
Dynamic Resource Optimization (DRO)	Automatically balances workloads across hosts	Not available
Simplified Networking & Storage setup	Automated configuration of virtual I/O, storage, and networking	Manual configuration required

Feature/ Capability	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	Best suited for smaller or static environments
User Interface	Modern, web-based UI with dashboards and templates	Traditional interface, more administrative

6

Operating systems

The IBM E1150 supports two operating systems: AIX and Linux. Each operating system addresses distinct enterprise computing needs.

AIX, IBM's UNIX operating system, is designed for high-performance, scalable enterprise workloads. AIX was built on a UNIX System V base and incorporates IBM innovations. It offers advanced features such as dynamic system tuning, workload partitioning, and robust security. Organizations widely use AIX in environments that demand stability and performance, such as ERP systems, databases, and analytics platforms. AIX is optimized for IBM Power hardware, ensuring seamless integration and efficient resource utilization.

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 AI or machine- learning development. Running Linux on Power enables organizations to use the performance advantages of the Power architecture and maintain compatibility with Linux software. Together, these operating systems make IBM Power a versatile platform that supports a wide range of modern and legacy workloads.

Red Hat OpenShift complements these operating systems. It is a Kubernetes-based container platform that runs natively on IBM Power. Red Hat OpenShift enables organizations to build, deploy, and manage containerized applications across hybrid cloud environments with consistency and scalability. Red Hat OpenShift provides a unified DevOps experience and supports modern application architectures, including microservices and AI workloads. Because of this integration, enterprises can modernize their IT infrastructure and use the performance, security, and reliability of IBM Power.

This chapter includes the following topics:

- AIX
- Linux on IBM Power
- Red Hat OpenShift
- PowerVM Virtual I/O Server
- Setting your LPAR compatibility mode

6.1 AIX

The AIX operating system is a secure, scalable, and robust open standards-based UNIX system. For over thirty years, AIX has been the cornerstone of mission-critical computing for enterprise organizations in highly complex industries, evolving to introduce a wealth of new hybrid cloud and open-source capabilities.

At the time of writing, AIX 7.3 is the latest AIX release that is available in the market. It builds on a solid foundation by offering new functions and capabilities that further enhance performance, scalability, availability, and security, and it preserves application-binary compatibility to safeguard IT investments in AIX.

Coupled with the IBM Power11 processor-based systems, AIX 7.3 provides an optimized and more resilient computing platform that adapts to changing business demands, including new cloud use cases and improved economics.

6.1.1 AIX 7.3 key features

AIX 7.3 includes the following key features:

- Workload scalability and automation

  AIX 7.3 provides enhanced workload scalability and improved cloud automation through Ansible. It includes more than 300 open-source packages in the AIX Toolbox for Open Source Software, enabling modern application development.

- Live kernel update

  Introduced in AIX 7.2 and enhanced in AIX 7.3, this feature applies interim fixes, service packs, and technology level updates without requiring system reboots. It supports PowerVC-managed environments and Power Enterprise Pool systems for optimized resource usage.

- High availability and disaster recovery

  IBM PowerHA and VM Recovery Manager deliver automated recovery and multi-site replication. These capabilities help to minimize downtime and ensure business continuity in hybrid or public cloud environments.

- AI integration

  AIX workloads provide a valuable source for AI development. These systems host large volumes of high-quality data on customer behavior and transactional information. Organizations can use this data for machine learning and deep learning, enabling actionable insights on a unified platform.

- Cloud flexibility

  AIX supports private, on-premises cloud transformation with PowerVC and offers hybrid cloud functionality for seamless AIX virtual machine import and export. It provides software-defined infrastructure for SAN-less DevOps environments. Available through IBM Power Virtual Server, 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, ensuring reliability and efficiency.

  AIX 7.3 includes the following key features:

- Workload scalability and automation: AIX 7.3 offers enhanced workload scalability, improved cloud automation by using Ansible, and more than 300 open-source packages through the AIX Toolbox for Open Source Software, enabling modern application development.

6.1.2 Supported levels

When AIX is installed by using direct I/O connectivity, the IBM Power11 processor-based server supports the following minimum levels of the AIX operating system:

- AIX 7.3 with the 7300-03 Technology Level and Service Pack 1 or later
- AIX 7.3 with the 7300-02 Technology Level and Service Pack 4 or later
- AIX 7.2 with the 7200-05 Technology Level and Service Pack 10 or later

When AIX is installed using virtual I/O, the IBM Power11 processor-based server supports the following minimum levels of AIX operating system:

- AIX 7.3 with the Technology Level 7300-03 and Service Pack 0 or later
- AIX 7.3 with the Technology Level 7300-02 and Service Pack 2 or later
- AIX 7.2 with the Technology Level 7200-05 and Service Pack 8 or later

Important:

For AIX Release 7.3, Technology Level 3 Service Pack 1 and later, AIX logical partitions operate in native Power11 processor compatibility mode, using the advanced capabilities of Power11.

Logical partitions running versions earlier than AIX Release 7.3 Technology Level 3 Service Pack 1 operate in Power9 or Power10 processor compatibility mode on Power11 systems.

Logical partitions running AIX Release 7.2 can operate in Power9 or Power10 processor compatibility mode.

6.1.3 AIX maintenance levels

IBM periodically releases maintenance packages (service packs (SPs) or technology levels (TLs) for the AIX operating system. For more information about these packages, downloading, and obtaining the installation packages, see IBM Fix Central. For more information about hardware features compatibility and the corresponding AIX Technology Levels, see IBM Support.

The Service Update Management Assistant (SUMA) can help you automate the task of checking and downloading operating system downloads and is part of the base operating system. 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 to verify the current health of a system.

The IBM AIX Operating System Service Strategy and Best Practices is a free resource available to AIX clients and gives insight into the AIX service strateg. It also provides helpful lifecycle information to best maintain your version of AIX. For more information on AIX, see AIX on IBM Power.

For additional information, see the following websites:

- AIX support Lifecycle information
- System Software Maps
- System to AIX maps

6.1.4 Licensing

The AIX operating system is available in the following editions:

- AIX Standard Edition
- AIX Enterprise Edition
- AIX Cloud Edition

The Enterprise and Cloud Editions include Power-related software that is typically required to manage larger IBM Power environments, such as those in hybrid clouds.

There are two licensing models for AIX:

1. CPU-based licensing
2. Subscription licensing

CPU-based licensing

CPU-based AIX licenses can either be ordered together with the server, or purchased later as an MES upgrade for the server. The license grants you the right to use AIX on a specific server. If you require support for AIX, you must have a valid Software Maintenance Agreement (SWMA) for AIX.

Subscription licensing

Subscription licensing is a new model that offers greater flexibility in AIX acquisition. The subscription license includes access to IBM software maintenance for a specified subscription term (1 or 3 years). After this term, you can renew the subscription if you continue using AIX.

This approach aligns with modern IT consumption trends, so businesses can scale usage up or down based on workload demands and simplify budgeting and procurement processes. Subscription licensing also includes access to updates, patches, and IBM support, which helps ensure that systems remain secure and up to date.

AIX Standard Edition, AIX Enterprise Edition, IBM Private Cloud Edition, and IBM Private Cloud Edition with AIX are available under a subscription licensing model that provides access to an IBM program and IBM software maintenance for a specified subscription term (1 or 3 years). The subscription term begins on the start date and ends on the expiration date, which is reflected at the IBM ESS website.

This model provides flexible and predictable pricing over a specific term with lower upfront acquisition costs. Another benefit is that the licenses are customer-number-entitled, meaning you can use the licenses on any IBM Power server you have in your environment. You can even move them between on-premises, your private cloud, and public cloud if needed.

These subscriptions are available for deployment on IBM Power, including Power11, and are well suited for hybrid cloud environments where agility and operational efficiency are key. By adopting subscription licensing, organizations can modernize their licensing strategy and continue to use the reliability and performance of AIX.

The product IDs for the subscription licenses are listed in Table 6-1.

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)

You can order the subscription licenses through IBM configuration tools. The AIX perpetual and monthly term licenses for standard edition are still available.

6.1.5 New AIX editions

The IBM AIX operating system is an open standards-based UNIX operating system that has been the foundation of mission-critical workloads and databases for tens of thousands of customers for over 35 years. AIX provides an enterprise-class IT infrastructure that delivers the reliability, availability, performance, and security that is required for organizations to be successful in a global economy. AIX Editions are packaged offerings that include IBM AIX and relevant AIX software products such as PowerSC, VM Recovery Manager, and Cloud Management Console.

IBM offers the following updates and enhancements to the AIX operating system and products that include AIX for IBM Power8, Power9, Power10, and Power11 servers:

- IBM AIX 7 Enterprise Edition 1.13
- IBM Private Cloud Edition 1.13
- IBM Private Cloud Edition with AIX 1.13
- VM Recovery Manager 1.9

IBM AIX 7 Enterprise Edition 1.13

IBM has updated the AIX 7 Enterprise Edition and its corresponding subscription offering to version 1.13. The bundled software components offered with AIX 7 Enterprise Edition 1.13 (5765-CD3 and 5765-2E1) now include:

- IBM AIX 7.3 TL3 or IBM AIX 7.2 TL5
- IBM PowerSC 2.3
- IBM PowerVC for Private Cloud 2.3.1
- IBM VM Recovery Manager HA 1.9
- IBM Tivoli® Monitoring 6.3

The following bundle components are updated:

- AIX 7.3 TL3 is updated with service pack 1.
- AIX 7.2 TL5 is updated with service pack 10.
- IBM PowerSC 2.3 is updated from 2.2 to 2.3.
- IBM PowerVC for Private Cloud is updated from 2.3.0 to 2.3.1.
- IBM VM Recovery Manager HA is updated from 1.8 to 1.9.

Additional information might apply to your configurations:

- Clients with active Software Maintenance (SWMA) or subscriptions 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 Entitled Systems Support website.
- Clients can choose either AIX 7.3 TL3 or AIX 7.2 TL5.
- Clients selecting AIX 7.2 TL5 can later upgrade to AIX 7.3 TL3 at any time, provided 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 has updated the Private Cloud Edition and its corresponding subscription offering to version 1.13. The bundled software components offered with Private Cloud Edition 1.13 (5765-ECB and 5765-6C1) now include:

- IBM PowerSC 2.3
- IBM PowerVC for Private Cloud 2.3.1
- IBM VM Recovery Manager DR 1.9
- IBM Tivoli Monitoring 6.3

The bundle components are updated as follows:

- IBM PowerSC 2.3 is updated from 2.2 to 2.3
- IBM PowerVC for Private Cloud is updated from 2.3.0 to 2.3.1
- IBM VM Recovery Manager DR is updated from 1.8 to 1.9
- Cloud Management Console (CMC) is updated from 1.22 to 1.23.

IBM Private Cloud Edition with AIX 1.13

IBM has updated Private Cloud Edition with AIX and its corresponding subscription offering to version 1.13. The bundled software components offered with Private Cloud Edition with AIX 1.13 (5765-CBA and 5765- 2C1) now include:

- IBM AIX 7.3 TL3 or IBM AIX 7.2 TL5
- IBM PowerSC 2.3
- IBM PowerVC for Private Cloud 2.3.1
- IBM VM Recovery Manager DR 1.9
- IBM Tivoli Monitoring 6.3

The bundle components are updated as follows:

- AIX 7.3 TL3 is updated with service pack 1
- AIX 7.2 TL5 is updated with service pack 10
- IBM PowerSC 2.3 is updated from 2.2 to 2.3
- IBM PowerVC for Private Cloud is updated from 2.3.0 to 2.3.1
- IBM VM Recovery Manager DR is updated from 1.8 to 1.9
- Cloud Management Console (CMC) is updated from 1.22 to 1.23.

IBM Private Cloud Edition includes the following features and options:

- Private Cloud Edition 1.13 and Private Cloud Edition 1.13 with AIX 7 include an entitlement for subscription to IBM Cloud Management Console (5765-CMT) for the same term as their SWMA.
- Clients with active SWMA or subscriptions for earlier versions of Private Cloud Edition or Private Cloud with AIX are entitled to upgrade at no charge. To update, download, or install, see the IBM Entitled Systems Support website.
- Clients can choose either AIX 7.3 TL3 or AIX 7.2 TL5.
- Clients selecting AIX 7.2 TL5 can later upgrade to AIX 7.3 TL3 at any time, provided their SWMA or subscription is current.
- Clients with AIX Enterprise Edition can move to Private Cloud Edition with AIX.

VM Recovery Manager 1.9

VM Recovery Manager provides automated VM management in the data center and disaster recovery management between sites. In IBM Power Virtual Server public the product is called DR Automation. DR Automation provides automated disaster recovery management between geographically dispersed Power Virtual Server data centers.

6.2 Linux on IBM Power

Linux is a powerful, open-source, cross-platform operating system that runs on a broad spectrum of hardware, from embedded devices to mainframes. It delivers a consistent, UNIX-like environment across diverse architectures. Its compatibility with IBM Power provides a flexible and cost-effective option for running a wide range of applications while taking full advantage of the platform's performance, availability, and reliability.

As a free and open-source solution, Linux reduces the total cost of ownership 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 and agile IT environments. Its high portability allows it to run on architectures such as x86, ARM, and IBM Power. Linux is supported by a large, active global community that provides extensive documentation, tools, and peer support.

IBM has been a leading advocate for Linux for more than two decades, integrating it deeply into its enterprise portfolio, including IBM Power, IBM Z mainframes, and IBM Cloud. IBM collaborates closely with major Linux distributions such as Red Hat, SUSE, and Canonical to ensure optimized performance and seamless integration. Through initiatives such as LinuxONE and ongoing contributions to the Linux kernel and open-source projects, IBM continues to position Linux as a cornerstone of secure, modern infrastructure.

Security is another key strength of Linux. Its robust user privilege model, combined with community-driven updates and tools such as SELinux and AppArmor, helps protect systems from vulnerabilities and malware. Linux is also efficient and lightweight, capable of running on everything from low-powered IoT devices to high-performance enterprise servers. Its modularity allows organizations to tailor the system to specific needs, whether that means a minimal installation or a fully featured desktop environment.

Linux also provides access to a vast array of tools and software, some of which are exclusive to Linux or better supported on Linux than on other operating systems. Its open-source nature allows for greater flexibility and adaptability, making it easier to align with specific business or technical requirements. Linux serves as a solid foundation for hybrid cloud infrastructure, supporting application modernization and deployment at scale. Its scalability, reliability, and cost-effectiveness make it well suited for large enterprise environments and mission-critical workloads.

6.2.1 Supported distributions

The Linux distributions that are described in this section are supported on the Power E1150. Other distributions, including open source releases, can run on these servers, but do not include any formal Enterprise Grade support.

Red Hat Enterprise Linux

RHEL on IBM Power servers delivers scalability and performance for high-demand workloads such as big data, cloud, and SAP HANA. The efficiency of IBM Power allows organizations to optimize resource usage and reduce costs by using fewer processor cores.

Red Hat can be deployed in both on-premises and cloud environments, giving organizations full control over their data and infrastructure with the flexibility to choose their preferred platform.

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

At launch, IBM Power11 servers support the following versions of RHEL:

- RHEL 10 is the latest Red Hat Enterprise Linux release with native Power11 support, making full use of the architecture's advanced features.
- RHEL 9.6 for Power LE or later (Power11 native support).
- RHEL 9.4 or 9.5 for Power LE (Power10 processor compatibility mode).
- RHEL 8.10 for Power LE or later (Power10 processor compatibility mode).

Red Hat Enterprise Linux roadmap

IBM Power servers support Linux workloads while using Power hardware's performance, reliability, and availability. This section outlines supported enterprise Linux distributions. Although other distributions might run on Power, only select enterprise versions offer formal support.

Red Hat Enterprise Linux (RHEL) is a leading open-source platform for Linux, hybrid cloud, containers, and Kubernetes. It provides secure, transparent, and proactive lifecycle management, supporting autonomous operations. Lifecycle planning is essential for customers, partners, ISVs, and the broader market. Starting with RHEL 8, the lifecycle includes three phases: Full Support, Maintenance Support, and Extended Life. Red Hat also shares projected release timelines and extended support details for minor versions. RHEL versions 8, 9, and 10 each offer a 10-year lifecycle across the first two phases, followed by Extended Life.

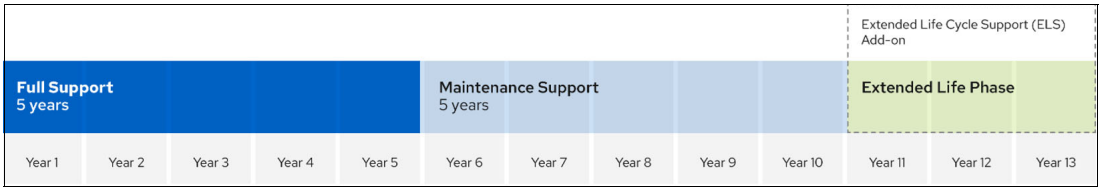

Figure 6-1 Red Hat Enterprise Linux support lifecycle

Note: The Red Hat Enterprise Linux lifecycle phases are designed to minimize changes within each major release over time and ensure predictable availability and content.

Red Hat includes the following types of support:

- Full support phase

  During this phase, Red Hat provides updates for the following issues:

  - Security issues (CVEs with CVSS greater or equal to 7) via RHSAs
  - Urgent and select high-priority issues via RHBAs
  - Additional errata as needed

  New or improved hardware support and select software enhancements might be included, typically in minor releases. These 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 includes the following releases:

  - Security updates (CVEs with CVSS)

– Urgent and selected high-priority bug fixes

New features and hardware support are not included during this phase.

- Extended life phase

  Customers retain access to existing content via 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 to existing installations only, and Red Hat can end support at its discretion.

Release frequency

To provide predictability, minor RHEL releases are scheduled every six months during the full support phase. Details on extended update support (EUS) and SAP-specific services are included per release. Figure 6-2 shows the planned release schedule for RHEL8.

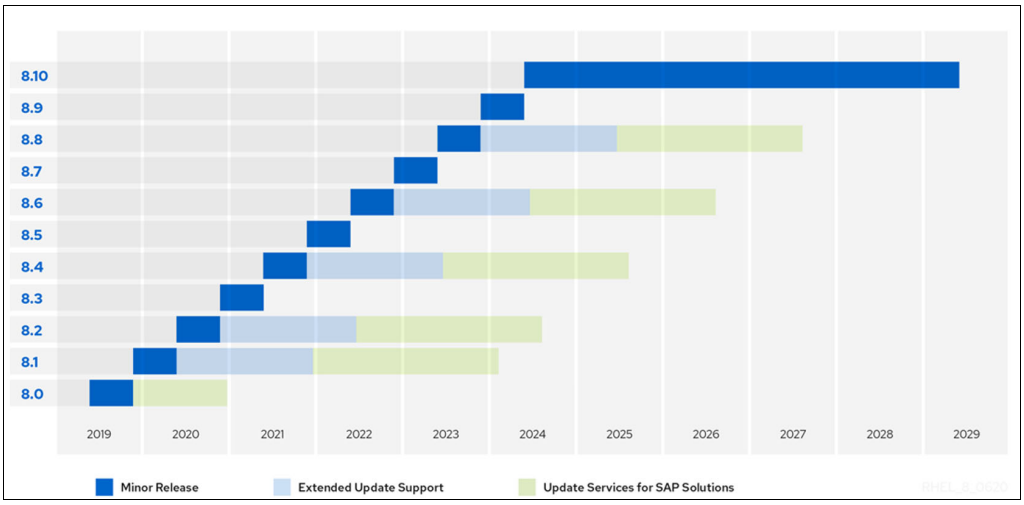

Figure 6-2 Release schedule for RHEL8

Figure 6-3 shows the planned release schedule for RHEL9.

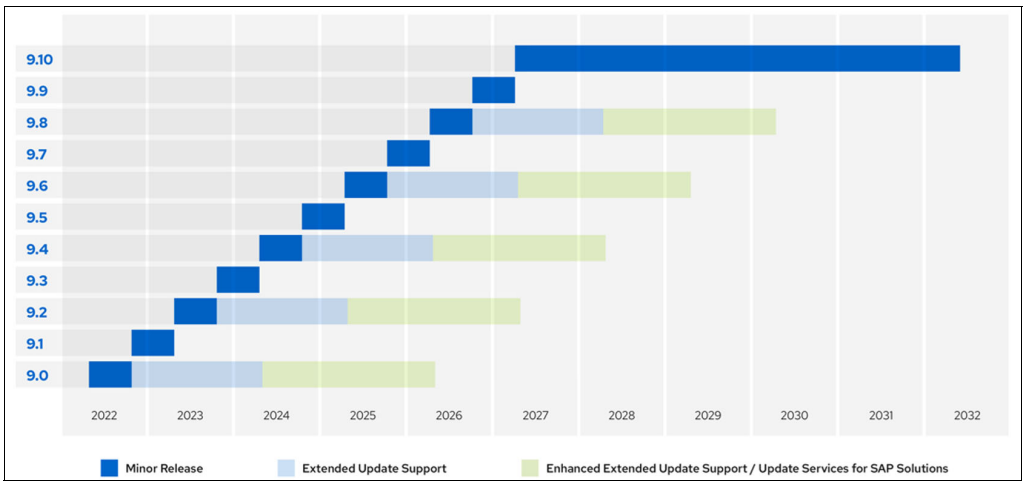

Figure 6-3 Release schedule for RHEL9

Figure 6-4 shows the release schedule for RHEL10.

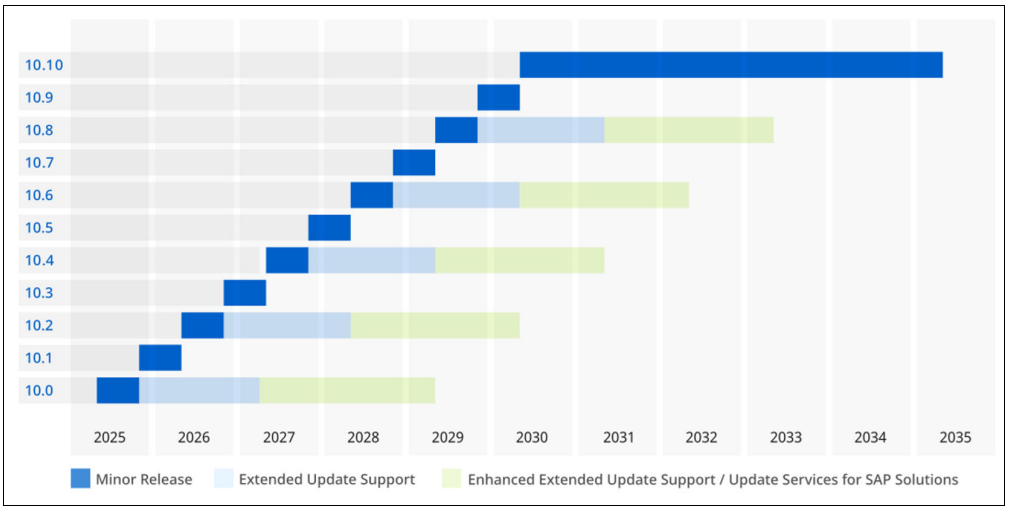

Figure 6-4 Release schedule for RHEL10

For additional details regarding the RHEL roadmap, see Red Hat Enterprise Linux Life Cycle. To find the latest Red Hat certified IBM Power servers, see 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, which allows it to access all of the features of the Power11 processor and platform.

At announcement, the Power E1150 servers support SUSE Linux Enterprise Server 15 Service Pack 6 or later running native Power11.

SUSE roadmap

SUSE Linux Enterprise Server is designed with long-term enterprise stability and support in mind. Its product lifecycle provides predictable and reliable maintenance phases, enabling organizations to plan deployments, updates, and migrations effectively.

SUSE Linux Enterprise Server offers a total lifecycle of 13 years per major release, which is divided into two main phases:

- General support (10 years): This phase includes full maintenance, security patches, bug fixes, hardware enablement, and new certified third-party software. It is ideal for production environments requiring continuous updates and active support.
- Extended support (3 years): After General Support ends, Extended Support offers critical security updates and selected bug fixes, allowing customers more time to transition to newer versions while maintaining operational security.

Release frequency

Major Releases are scheduled for every 4 years and they introduce new features, platform support, and architectural improvements. Service Packs (SPs) are released every 12 to 14 months and include incremental updates, feature enhancements, and hardware enablement.

Each Service Pack is supported for six months after the release of the subsequent SP, providing customers time to validate and upgrade within a consistent and predictable window. Figure 6-5 on page 129 shows the service pack schedule.

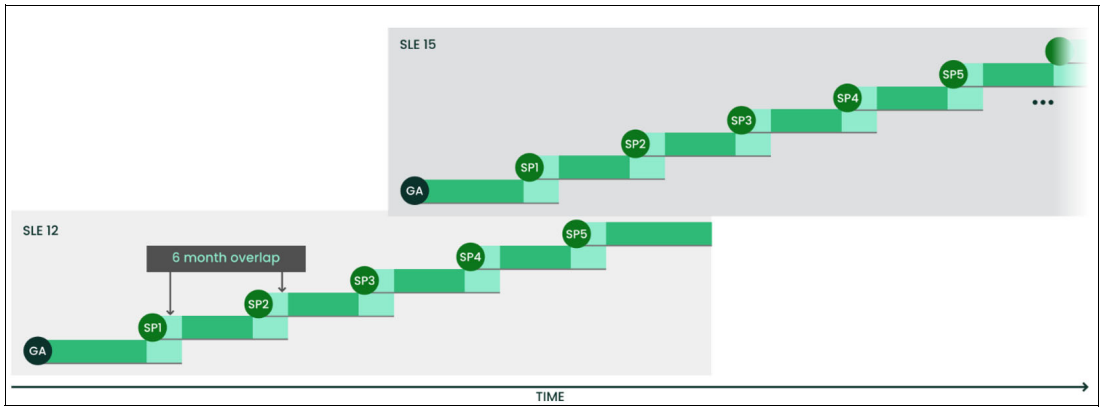

Figure 6-5 SLES major releases and service packs

Figure 6-6 shows the release lifecycle, including long-term service pack support.

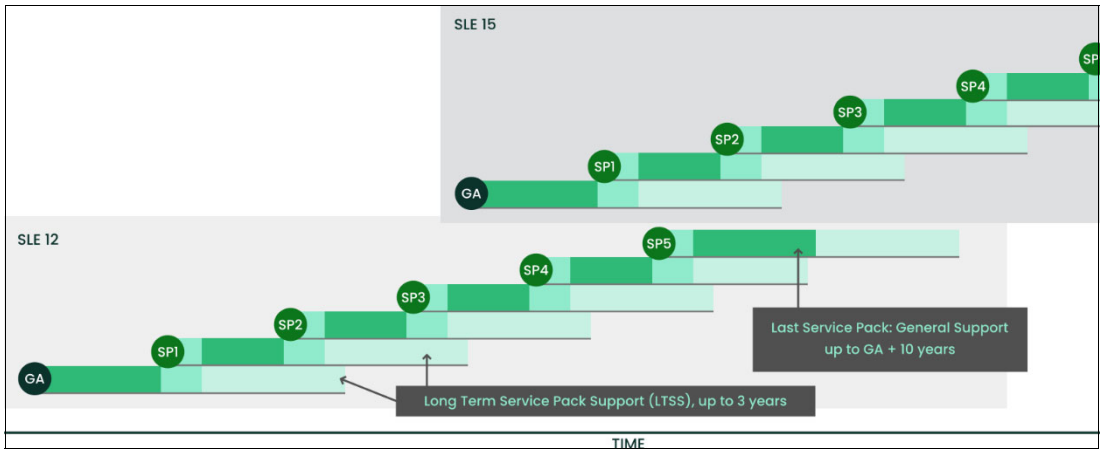

Figure 6-6 SLES long-term service pack support.

For additional details regarding the SUSE roadmap, see SUSE Linux Enterprise Server Documentation and Product Support Lifecycle. To find the latest SUSE certified IBM Power servers, see 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 the IBM Advance Toolchain for Linux version 15.0, which allows customers and developers to use the Power11 processor-based technology instructions when programming. Cross-module function call uses a PC-relative addressing mode.

One specific benefit of Power11 technology is a 10–20 times advantage over Power9 processor-based technology for AI inferencing workloads because of increased memory bandwidth and new instructions. One example is the special purpose-built matrix math accelerator (MMA) that is tailored for the demands of machine learning and deep learning inference. It also supports many AI data types.

6.2.2 Licensing

Linux licensing on IBM Power follows the same foundational principles as on other platforms but with considerations that are tailored to enterprise environments. The Linux kernel itself is licensed under the GNU General Public License version 2 (GPLv2), which ensures that the software remains free and open-source, allowing users to run, study, modify, and distribute it. On IBM Power, several major Linux distributions are supported, including Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server, and Ubuntu. Each of these distributions adheres to open-source licensing models but can differ in how they package and deliver their software.

For example, RHEL and SUSE Linux Enterprise Server provide access to source code under open-source licenses but require a paid subscription for access to precompiled binaries, updates, and enterprise support. This model allows organizations to benefit from open-source flexibility while receiving the stability and support needed for mission-critical workloads.

On IBM Power, these distributions are optimized to take advantage of the architecture's performance, scalability, and reliability features. Licensing terms typically include support for virtualization technologies like PowerVM.

Red Hat licensing

Red Hat Enterprise Linux is sold on a subscription basis, with initial subscriptions and support available for one, three, or five years. Support is available directly from Red Hat or IBM Technical Support Services.

Red Hat Enterprise Linux 8 for Power LE subscriptions covers a maximum of four cores and a maximum of four LPARs, and can be stacked to cover a larger number of 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 Entitled Systems Support (ESS), you use it to establish proof of entitlement and download the software from Red Hat.

For more information, see 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 one, three, or five years. Support is available directly from SUSE or from IBM Technical Support Services.

SUSE Linux Enterprise Server 15 subscriptions include a maximum of one socket or one LPAR, and can be stacked to cover a larger number of sockets or LPARs.

When you order SUSE Linux Enterprise Serve from IBM, a subscription activation code is automatically published in Entitled Systems Support (ESS), you use it to establish proof of entitlement and download the software from SUSE.

Clients are required to register the Linux offering purchased at the Distributor's website with the activation code. After they are registered, clients can download the software packages and obtain the latest upgrades available for the product purchased.

For more information, see SUSE Licensing.

6.2.3 Introduction to KVM support

KVM stands for Kernel Virtual Machine, and it's a widely spread virtualization technology in X86_64. KVM can also be used in Power Servers.

With KVM on Power, IBM offers both KVM and PowerVM to meet the varying virtualization needs of Linux on Power customers. However, KVM virtualization technology represents an opportunity to simplify a customer's virtualization infrastructure with a single hypervisor and management software across multiple platforms.

The KVM guests can run within a PowerVM LPAR. KVM can be used to run KVM guests. These guests are essentially VMs that run on the LPAR and use the existing resources of the LPAR. This approach combines the advantages of PowerVM virtualization capabilities with the power and flexibility of KVM.

KVM guests running in a PowerVM LPAR have a unique runtime architecture, different from other virtualization mechanisms on IBM Power systems. For more information on KVM support on IBM Power see section 9.2, "KVM support" on page 192.

6.2.4 Other Linux distributions available for IBM Power

IBM officially supports AIX, IBM i, Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server for the IBM Power architecture. A variety of other Linux distributions that are released by the open-source community and third parties can also run on IBM Power servers.

These alternative distributions, although not officially supported by IBM, can be installed and run on IBM Power systems, particularly within LPARs (Logical Partitions) configured to run Linux workloads. In addition to supporting enterprise-grade operating systems, the IBM Power platform also provides an open environment for developers and enthusiasts to experiment with other Linux distributions. This flexibility allows users to explore their preferred Linux variants on IBM Power and take full advantage of the platform's performance, reliability, and security features.

By enabling access to alternative Linux distributions, IBM Power fosters innovation and broadens the community of contributors who can test, optimize, and run workloads on Power architecture, while continuing to benefit from the system's advanced infrastructure capabilities.

A useful tool for identifying Linux distributions available for IBM Power is the search engine at DistroWatch.com, which allows filtering by architecture. To find distributions compatible with IBM Power, it is recommended to select one of the following architectures during the search:

- Power
- PowerPC
- ppc64le
- ppc64el

At the time of writing, the following Linux distributions are released to run on IBM Power LPARs (listed in order of popularity):

- Debian
- Ubuntu
- Fedora
- OpenSUSE
- AlmaLinux
- CentOS

- Alpine Linux
- Rocky Linux
- ALT Linux
- Chimera Linux
- Gentoo Linux
- T2 SDE
- Trisquel GNU/Linux
- OpenEuler
- Bedrock Linux
- Vine Linux
- Adélie Linux
- UOS (Unity Operating System)

6.3 Red Hat OpenShift

Red Hat OpenShift on IBM Power provides a powerful platform for building, deploying, and managing containerized applications across hybrid cloud environments. Red Hat OpenShift uses the performance and reliability of IBM Power architecture to enable enterprises to modernize their IT infrastructure with cloud-native technologies such as Kubernetes, containers, and microservices. Because Red Hat OpenShift supports AIX, IBM i, and Linux, you can colocate containerized workloads alongside traditional applications to reduce latency and improve data access. Red Hat OpenShift integrates with IBM Cloud services and automation tools to simplify cluster provisioning, scaling, and lifecycle management, which can make it easier for developers to focus on innovation rather than infrastructure.

Red Hat OpenShift 4.18 introduces enhanced capabilities to IBM Power, including improved networking with user-defined networks (UDNs), advanced Operator Lifecycle Management (OLM), and deeper GitOps integration. These features support more secure, scalable, and flexible deployments, whether on-premises or in the IBM Power Virtual Server cloud. Red Hat OpenShift on Power also supports IBM Cloud Paks and AI workloads, which enables enterprises to accelerate digital transformation while maintaining enterprise-grade security and performance. This combination of technologies empowers you to innovate faster, optimize resource usage, and build a resilient hybrid cloud foundation.

For more details on how Red Hat OpenShift can improve your hybrid cloud experience see section 10.2, “Red Hat OpenShift” on page 198.

Licensing

Red Hat OpenShift Container Platform is offered through subscription, with terms available for 1, 3, or 5 years. Subscriptions are sold in increments covering two processor cores and can be stacked to match workload requirements. Support is available either directly from Red Hat or through IBM Technical Support Services.

When ordering OpenShift Container Platform for Power from IBM, the client receives a subscription activation code through the IBM Entitled Systems Support (ESS) portal. This code serves as proof of entitlement and is used to download the software directly from the Red Hat Customer Portal.

For further information and documentation, see Red Hat OpenShift subscription editions.

Supported levels

The IBM Power E1150 server supports Red Hat OpenShift Container Platform 4.18 and later. Red Hat OpenShift 4.18 supports IBM Power11 when running with Power10 Compatibility mode at GA.

Full support for IBM Power11 native mode is planned for Red Hat OpenShift version 4.19.

Red Hat OpenShift Container Platform roadmap

Red Hat provides a defined product life cycle for RedHat OpenShift Container Platform to help customers and partners plan, deploy, and support their infrastructure. This lifecycle is published for transparency, though exceptions can occur.

Red Hat OpenShift v4 follows a phased, time-based lifecycle, with at least four minor versions supported concurrently. Each minor version has a fixed support period, offering varying levels of maintenance. Red Hat targets a release cadence of every four months to support customer planning. All errata remain available to active subscribers throughout the lifecycle.

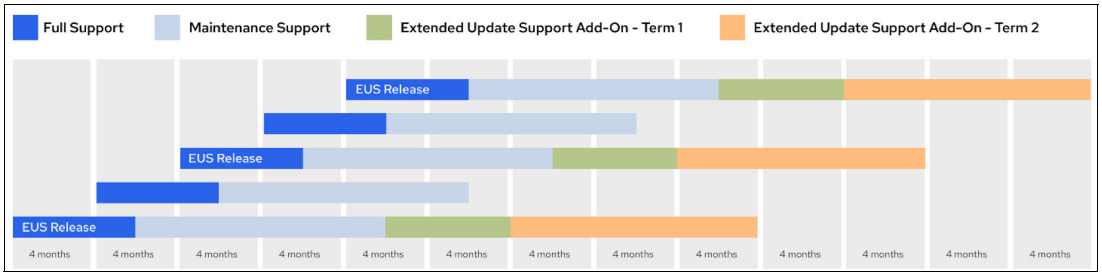

Figure 6-7 Figure 4. RedHat OpenShift Container Platform Lifecycle phases

Full support

Full support begins at GA of a minor version and ends 6 months later or 90 days after the next minor version GA-whichever is later and includes the following support:

- Full and development support are provided per the Scope of Coverage and SLA.
- Critical and important security fixes (RHSAs) are released as needed.
- Urgent and high-priority bug fixes (RHBAs) are released promptly; others might be included in periodic updates.
- Customers must stay on a supported micro version (for example, 4.x.z) to receive updates.

Maintenance support

Starts after full support and ends 18 months after the minor version GA and includes the following support:

- Critical and Important RHSAs continue.
- Select urgent RHBAs might be released.
- Other fixes and enhancements (RHEAs) can be issued at Red Hat's discretion.
- No technical support is provided after this phase, except for assistance with upgrading.
- Access to hosted services for unsupported versions is not guaranteed.

Extended update support (EUS)

Even-numbered minor versions, such as 4.8, 4.10, are designated as EUS releases, offering extended support phases to simplify upgrades and reduce node reboots. Extended support terms include the following terms:

- EUS Add-On: Term 1 (6 Months)

Optional support following Maintenance Support:

- Includes Critical/Important security updates and urgent bug fixes
- Allows customers to stay on a minor release for up to 24 months
- Included with Premium subscriptions (x86_64) and available as an add-on to Standard. 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 OpenShift Container Platform components.

With both Term 1 and Term 2 support, you can extend support for a maximum of 36 months for stable, mission-critical environments.

Table 6-1shows the dates for current versions of Red Hat OpenShift Container Platform.

Table 6-1 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 OpenShift Container Platform 4 full support						
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 OpenShift Container Platform 4 maintenance support						
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 support OpenShift Container Platform 4 extended Support						
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 (RHOCP) 4 uses Red Hat Enterprise Linux CoreOS (RHCOS) as its managed node operating system. RHCOS is updated during cluster upgrades, which can include changes between RHEL minor versions.

For additional details about the Red Hat OpenShift Container Platform roadmap, see Red Hat OpenShift Container Platform Life Cycle Policy and Product Life Cycles.

6.4 PowerVM Virtual I/O Server

IBM PowerVM software is a virtualization environment that can run AIX, IBM i, and Linux virtual machines on IBM 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 server virtualization environment for your applications, built upon 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, ensuring the reliability of mission-critical applications. With robust automation capabilities, PowerVM accelerates service delivery by streamlining the provisioning and management of virtual machines and storage resources, making it well suited for cloud-based infrastructures. It also enhances operational efficiency and maximizes return on investment because of features like live partition mobility, which enables zero-downtime workload migration and because of resource optimization tools that improve the usage of compute and storage infrastructure.

The Virtual I/O Server (VIOS) is part of the PowerVM Editions hardware feature. The VIOS is software that is located in a logical partition and facilitates the sharing of physical I/O resources between client logical partitions within the server.

Currently, the IBM Power11 processor-based server supports the following minimum levels of the PowerVM VIOS operating system:

- VIOS 4.1.1.10
- VIOS 4.1.0.40
- VIOS 3.1.4.60

6.5 Setting your LPAR compatibility mode

IBM Power Processor Compatibility Mode is a feature that enables newer IBM Power systems to run applications and operating systems originally compiled for earlier generations of Power processors. This is especially important when moving to the latest hardware, such as Power11, where some operating systems and applications might not yet support native execution. In such cases, compatibility mode allows these workloads to run in a Power10 environment, ensuring continuity and minimizing disruption during upgrades.

The compatibility mode is defined at the logical partition (LPAR) level during the creation of the LPAR. Administrators can specify which processor generation the partition should emulate so that multiple workloads can operate in different compatibility modes on the same physical server. This flexibility is useful in mixed environments where not all software has been updated to support the latest architecture.

Although running in a previous-generation compatibility mode helps maintain application portability and eases the migration process, it might restrict access to newer processor features and performance enhancements. For this reason, IBM recommends using compatibility mode as a transitional solution. The recommended long-term goal is to recompile applications and upgrade operating systems to fully take advantage of the advanced capabilities and performance improvements of the latest Power processor architecture.

Figure 6-8 on page 136 shows how to set compatibility mode in your LPAR definition.

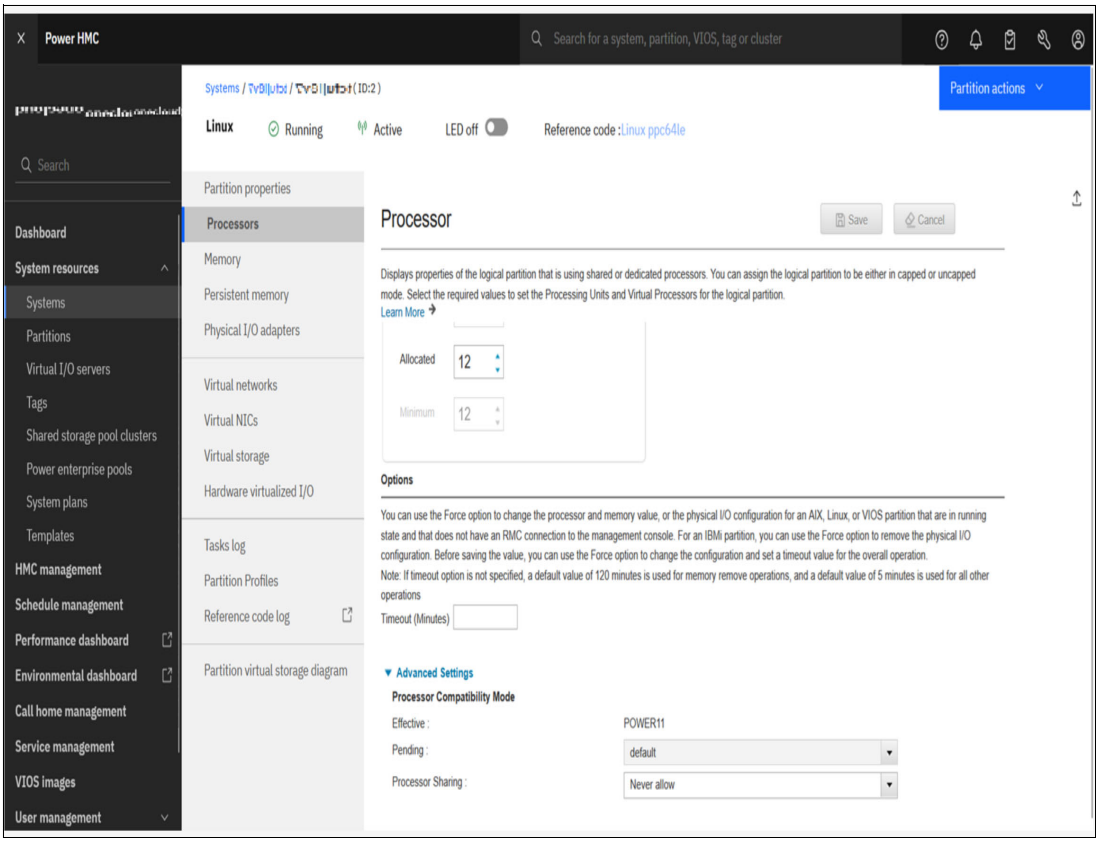

Figure 6-8 Set LPAR processor compatibility mode.

7

Enterprise solutions

IBM Power servers are engineered to support a wide range of enterprise workloads that demand high availability, 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, Power servers help ensure minimal downtime and seamless workload management. They are also well suited for enterprise resource planning (ERP) and customer relationship management (CRM) systems, offering the processing power and memory capacity needed to handle large-scale, real-time data operations.

In addition to traditional workloads, IBM Power servers are optimized for modern enterprise needs, including artificial intelligence (AI), machine learning, and hybrid cloud deployments. Organizations can use the built-in accelerators and support for containerized environments like Red Hat OpenShift to run AI models and cloud-native applications efficiently. These servers also excel in high-performance computing (HPC) scenarios, such as financial modeling, scientific research, and health care analytics because of their high core counts and memory bandwidth. With robust security features like transparent memory encryption and secure boot, IBM Power is a trusted platform for industries with strict compliance requirements. Overall, Power servers provide 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
- Health care

7.1 High availability and disaster recovery solutions

IBM offers a suite of high availability and disaster recovery (HA/DR) solutions that are tailored to its Power platform, each designed to meet different operational needs and workloads. PowerHA SystemMirror for AIX is a clustering solution that helps ensure application uptime through an active-standby model, enabling fast failover and minimal disruption during outages or maintenance. For IBM i environments, PowerHA SystemMirror for i provides tightly integrated HA/DR capabilities by using features like independent auxiliary storage pools (IASPs) and real-time data replication to protect mission-critical workloads. Complementing these is VM Recovery Manager (VMRM), a flexible solution that supports AIX, IBM i, and Linux workloads by managing full LPAR recovery across systems. VMRM is ideal for environments where full system recovery is acceptable, offering automated failover and disaster recovery without requiring application-level configuration. Together, these tools provide a comprehensive HA/DR strategy across the IBM Power ecosystem.

7.1.1 PowerHA SystemMirror for AIX

PowerHA SystemMirror for AIX, generally known as PowerHA, is a high availability clustering solution to ensure application uptime and fast recovery. It is considered as an active-standby or active-passive system of application availability. It is also known as an application restart model, where an application and its resources are logically combined and operate as a single entity on a VM or LPAR (known as a node) in a cluster. The resources include file systems with data, an IP address, application scripts, or a volume group. That group of resources is moved as a unit to a standby VM or LPAR, and the application is restarted. During a planned outage, PowerHA can gracefully quiesce the application on the primary node. For either scenario, the outage lasts as long as it takes to acquire the resources and restart the application. This might be measured in minutes. It is also considered to be a warm standby method of high availability because the target node already has a running version of AIX.

PowerHA for AIX can also be used across data centers, where data is replicated through the storage subsystem or over IP to the target node by using the AIX Geographic Logical Volume Manager (GLVM) subsystem.

Within or across data centers, you can automate the failover process to reduce application outages. You can move applications to another running node to allow for maintenance, with near-zero downtime although the PowerHA application can be maintained with a live kernel update.

PowerHA for AIX is tightly integrated with the AIX operating system's inherent reliability and with IBM Power hardware features such as virtual I/O, concurrent hardware maintenance, and dynamic resizing of LPARs. PowerHA supports critical enterprise apps, like Oracle, SAP, Db2, and Epic systems, which are all common workloads. Licensing is on an n+1 model per cluster, where n is the number of cores in the production online copy of the application, with 1 additional core added for the standby. For example, on a 24-core Power system, if two clusters are created, one supporting an LPAR with 4 cores and another with 5, then you must license a total of 11 cores: 4+1=5 and 5+1=6.

7.1.2 IBM PowerHA SystemMirror for i

PowerHA for IBM i is IBM’s high-availability and disaster recovery (HA/DR) solution that is designed specifically for the IBM i operating system. It provides automated failover capabilities and data replication to help ensure business continuity during planned or unplanned outages. PowerHA integrates tightly with the native features of IBM i, such as the

independent auxiliary storage pool (IASP), to enable fast and efficient switching between systems with minimal downtime.

At the core of PowerHA is its ability to replicate data in real time between primary and backup systems using technologies like Geographic Mirroring or IBM's Storage-based replication. This ensures that a critical application and its data remain available even if the primary system becomes unavailable. PowerHA supports both local high availability within a single data center and remote disaster recovery across geographically dispersed locations, offering flexibility for various business continuity strategies.

PowerHA also includes tools for monitoring, automation, and management, making it easier for IT administrators to configure and maintain HA/DR environments. With features like role-based access, cluster resource groups, and automated failover policies, PowerHA helps reduce the complexity of managing high-availability systems. It is a key component for organizations that rely on IBM i to run mission-critical workloads and require continuous access to their data and applications.

7.1.3 VM Recovery Manager (VMRM)

IBM VM Recovery Manager (VMRM) uses an active-inactive model in which the entire definition of a VM or LPAR is re-created on a target Power system and then restarted. VMRM supports any workload running on a Power System, including AIX, IBM i, and Linux (as noted elsewhere, the E1150 supports only AIX and Linux). Recovery takes longer than an application restart model because the entire LPAR must restart after being redefined, instead of just restarting the application.

In HA mode, VMRM uses Live Partition Mobility to dynamically move an LPAR to a system that shares an 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 any given time, except in limited circumstances when you test a disaster exercise on the target system.

As with PowerHA, VMRM helps automate disaster recovery. It can detect when a VM or an entire Power System is down and initiate a policy-based failover to a single target or to a group of targets. VMRM can coexist with PowerHA, where a cluster of LPARs can fail over as a group to Power systems at an alternate site. You can often use VMRM to provide simple DR support for Linux VMs because no software installation or OS configuration is required. The Linux VM restarts automatically.

The ideal use cases for VM Recovery Manager include development, test, and sandbox environments that might not require the same RTO requirements as production systems, and Linux on Power. VMRM integrates well with solutions such as SAP HANA and Oracle. No configuration work is required on the individual operating systems being managed, although you can install some agents to check for application failures and trigger an LPAR failover.

Similarly to PowerHA, VMRM is licensed based on the number of production cores in use, not on an entire Power system. However, unlike PowerHA, no standby licenses are required. On a 24-core Power system, if the VMs that need VMRM support use 9 cores, then you need to procure only 9 cores.

7.2 IBM Db2

In today's data-driven world, organizations across industries, such as manufacturing, health care, 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 high availability while keeping operational costs under control.

By combining IBM Db2 with IBM Power built on the advanced Power11 architecture, organizations can meet these demands. This integrated, workload-optimized stack delivers exceptional performance for both transactional and analytical workloads. Db2 uses cutting-edge database innovations to maximize efficiency and throughput, earning top rankings in industry benchmarks such as TPC-C, TPC-H, and SAP SD 3-Tier. Power systems provide the robust hardware foundation, offering superior price/performance and the ability to process more than a million transactions per minute at a cost of less than $1 per transaction. Db2 automatically exploits Power11's parallelism and large page sizes, simplifying deployment and enabling cost-effective scaling for web applications, messaging backbones, and workload consolidation.

Also, the combination of Db2 and IBM Power ensures high availability and operational efficiency. Power systems are engineered 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 needs, and IBM PowerVM virtualization, which allows for efficient resource usage and workload consolidation. These capabilities reduce hardware, energy, and management costs, enabling a more agile and cost-effective IT infrastructure. Together, Db2 and IBM Power offer a powerful, scalable, and resilient platform for modern data-driven enterprises.

For more information on Db2 high availability, see the IBM Redbooks publication *IBM PowerHA SystemMirror for AIX Cookbook*, SG24-7739

7.3 Oracle

For over 35 years, clients have relied on IBM Power to deploy their Oracle database and application workloads. Organizations, both large and small, can take advantage of the reliability and security of Power with its advanced recovery, self-healing, and diagnostic capabilities that are designed to reduce application downtime.

7.3.1 Benefits of running Oracle on Power11

IBM Power11 servers are fully certified by Oracle software. Organizations can use Power11 servers to consolidate multiple workloads on fewer servers, which increases overall system usage and lowers overall costs. This higher usage provides the customer with the opportunity to maximize their Oracle license use while reducing the IT costs by managing fewer physical machines.

The cost of security breaches continues to grow, averaging at USD 4 million per breach. IBM Power technology is built to protect businesses against cyberthreats with end-to-end security protocols, including new transparent memory encryption that doesn't impact performance. Power servers can be considered 60 times more secure than unbranded white box servers. By combining IBM Power servers with IBM AIX Trusted Execution (TE), you can verify the integrity of your system and implement advance security policies to enhance the trust level of the complete system.

As data continues to rapidly expand, organizations might struggle when modernizing their IT infrastructure. IBM Power servers and the IBM AIX operating system (OS) create a solid

foundation for the modernization of traditional Oracle database workloads, new application developments, and workload consolidation. With the Red Hat Ansible Automation Platform, clients can manage Oracle workloads on IBM Power servers as part of their wider enterprise automation strategy.

Mission-critical workloads need a server and OS with reliability, high availability, and the ability to scale without impacting mission-critical performance. IBM Power servers can offer 2.5 times better performance per core when compared to x86 servers, enabling Oracle workloads to grow linearly without introducing bottlenecks.

Oracle certifies its products on Power, delivering a host of benefits, including comprehensive end-to-end support, portability, and efficiency. IBM Power servers provide 99.999% of reliability to maintain maximum availability. The combined design of the AIX OS on Power servers with IBM PowerHA technology can bring clients stunning uptime, and the ability to manage and monitor availability to prevent both planned and unplanned outages.

Running your Oracle database workloads on IBM Power11 includes the following benefits:

- Industry-leading security

  IBM Power and AIX keep your critical Oracle workloads protected and available while reducing costs.

- Simplified management

  Automatically deploy a cloud-ready OS that can meet any organization's private cloud requirements with IBM PowerVC.

- Unmatched uptime

  IBM Power supports the most demanding workloads and provides 99.999% of reliability to maintain maximum availability.

- Improved workload cost-effectiveness

  By using Power LPAR and DLPAR eligibility for Oracle hard partitioning, clients can license only the Power cores available to Oracle software.

- Standards-based automation

  Improve manageability and scalability to ensure consistent and repeatable outcomes by using the enhanced and expanded automation portfolio for Oracle workloads, built on Ansible.

- Data Protection

  Benefit from top-to-bottom security in Power with trusted boot, main memory encryption, run-time verification of OS files, and role-based access and run control with PowerSC

For additional information on running Oracle on Power, see *Oracle on IBM Power Systems*, SG24-8485.

7.3.2 Running Oracle Standard Edition 2 on IBM Power

Oracle Standard Edition 2 (SE2) is a database option suitable for many business needs and can be used on IBM Power servers. SE2 offers a full-featured database with ease of use, power, and performance. SE2 includes features like relational, JSON, and XML data handling, and Real Application Clusters for clustering services. SE2 is licensed on servers with up to two sockets, with licensing costs remaining the same regardless of the number of cores within each socket.

Oracle Standard Edition 2 running on IBM Power includes the following benefits:

- Licensing

  SE2 can be licensed with 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 system compatibility

  Oracle Database Standard Edition 2 is compatible with IBM Power. IBM has designed specific Power10 servers (Power S1012 and Power S1014) to compete with x86 engines in the two-socket space, catering to Oracle's need for SE2 to be cost-competitive with SQL Server.

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, and advanced recovery, self-healing, and diagnostic capabilities:

- High availability

  Starting with Oracle Database 19c Release Update 19.7, Standard Edition High Availability is supported on IBM AIX on Power. This feature provides cluster-based database failover for Standard Edition Oracle Databases 19c.

- Other features

  SE2 offers a range of features for building business applications, including support for relational, JSON, XML, spatial, graph, and unstructured data, and Oracle Multitenant Architecture.

Licensing and implementation

Consider the following issues when you plan for your Oracle Standard Edition 2 implementation:

- Although the number of cores in each socket does not affect the license price, it is important to adhere to the maximum number of CPUs allowed for the cluster, which is not per node.
- Standard Edition high availability provides cluster-based failover for single-instance Standard Edition Oracle Databases using Oracle Clusterware:
  - Oracle Standard Edition High Availability benefits from the cluster capabilities and storage solutions that are already part of Oracle Grid Infrastructure, such as Oracle Clusterware, Oracle Automatic Storage Management (Oracle ASM), and Oracle ASM Cluster File System (Oracle ACFS).
  - The use of integrated, shared, and concurrently mounted storage, such as Oracle ASM and Oracle ACFS for database files and for unstructured data, enables Oracle Grid Infrastructure to restart an Oracle Database on a failover node faster than any cluster solution that relies 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 (64-bit).

7.4 SAP HANA

IBM Power has options to support clients on SAP S/4HANA. For both SAP HANA and S/4HANA, IBM Power offers on-premises, off-premises, or fully-managed (RISE with SAP) solutions. The solutions are built to run mission-critical applications like SAP to help accelerate your ERP and application deployments and help maximize the impact they can have on your data management, data integration, automation, and business processes.

IBM and SAP have partnered for over 50 years:

- Over 4,800 clients are running SAP HANA on IBM Power servers
- Over 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 current infrastructure while simultaneously supporting digital transformation. IBM Power servers are built for data-intensive applications such as SAP HANA and S/4HANA that require large amounts of in-memory computing. By using IBM Power servers, you can also maintain the high availability and flexibility required for your hybrid cloud.

With global data volumes set to grow to more than 180 zeta bytes in 2025, organizations across every sector are facing tremendous pressure to manage, process, store and extract valuable insights from their critical data. By running SAP HANA on IBM Power servers, businesses can gain the following benefits:

- Provision faster

  Simplify system management and boost business agility. Create new environments by allocating resources in small increments, starting with 0.01 cores and 1 GB memory.

- Maximize uptime

  Minimize disruption to business-as-usual activities. #1 Best-in-class reliability for 15 years.[1]

- Cut energy usage

  Reduce data center costs and enhance environmental sustainability. IBM Power E1150 provides comparable performance and requires half the amount of energy that is used by a comparable x86-based server. IBM Power E1180 uses 15% less energy and provides 54% more performance at maximum input power than a comparable x86-based server.

- Scale affordably

  Reduce the risk of over-provisioning with flexible hybrid cloud solutions, instant scaling, and pay-per-use consumption options with 40 TB scale-up capacity, which is the largest supported for SAP S/4HANA and SAP BW.

- Strengthen security

  Protect critical data and applications from cyberthreats with end-to-end security, including new transparent memory encryption with no performance impact. Power servers are considered more secure than other servers based on the number of security incidents reported.

- Gain faster insights

  Make rapid decisions to maximize business efficiency with 2.5 times better per core performance when compared to x86 servers.[2]

[1] https://www.itjungle.com/2024/10/21/ibm-nears-the-end-of-the-road-for-server-reliability-improvements/

[2] https://newsroom.ibm.com/2021-09-08-IBM-unveils-new-generation-of-IBM-Power-servers-for-frictionless,-scalable-hybrid-cloud

7.5 Banking

As financial institutions accelerate digital transformation, IT architects must design platforms that are secure, scalable, and AI-ready. IBM Power systems, particularly Power10 and Power11, 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, you can use Power systems enable to modernize legacy infrastructure while embedding intelligence into core banking workflows.

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. IBM Power systems support hybrid cloud deployments by using several features:

- IBM Power Virtual Server (PowerVS) for public cloud scalability
- Red Hat OpenShift for container orchestration across environments
- IBM Cloud Pak® for Data and watsonx for AI lifecycle management

Banking use cases that use AI on IBM Power

The following use cases are examples of the use of AI in banks and other financial institutions.

On-premises AI training with cloud-based inference

Banks must often train AI models on-premises to comply with strict data residency and regulatory requirements. Using IBM Power10 or Power11 systems, you can securely train models on sensitive customer or transactional data. After training, you can deploy these models to IBM Power Virtual Server (PowerVS) for inference, which enables 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 but rapid response times are also 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 IBM PowerVS while maintaining secure access to on-premises data sources. You can use this setup to use the scalability and flexibility of the cloud without moving sensitive data offsite. For example, you might deploy a generative AI model on PowerVS to automate the summarization of complex financial reports, pulling data securely from internal systems. This approach ensures compliance with data governance policies while accelerating insights and reporting.

Distributed AI microservices across hybrid cloud

Some banks adopt a microservices architecture, distributing AI components across both on-premises IBM Power systems and PowerVS in the cloud. This model supports modular, scalable AI applications that operate across environments. For instance, 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 across geographies or 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 IBM Power systems, you can continuously monitor transactions, detect 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

IBM Power systems also support the integration of AI directly into core banking platforms. This enables real-time intelligence within mission-critical applications such as loan underwriting, treasury operations, and fraud analytics. For example, you can use AI models running on Power10 to detect unusual transaction patterns or predict liquidity needs, which enhances operational efficiency and decision-making. This model ensures that AI capabilities are embedded where they deliver the most immediate business value.

Integrating IBM watsonx with IBM Power

IBM watsonx is a modular AI and data platform designed to accelerate the deployment of enterprise-grade AI. When integrated with IBM Power10 and Power11 systems, watsonx enables financial institutions to build, deploy, and govern AI models across hybrid cloud environments – including IBM PowerVS.

IBM watsonx includes the following components:

- watsonx.ai: Model training and inference
- watsonx.data: Data lakehouse integration
- watsonx.governance: Model risk and compliance
- watsonx.orchestrate: AI agents and workflow automation

IBM Power systems are high-performance servers optimized for data-intensive workloads. The integration between watsonx and IBM Power provides several benefits:

- Accelerated AI workloads: Power10 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, so businesses can keep sensitive data on-premises while using cloud-native AI capabilities.
- Enterprise-grade security and reliability: Power systems provide robust security and uptime, aligning with watsonx's focus on trustworthy AI.
- 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 Health care

IBM Power systems provide a secure, high-performance, and resilient infrastructure that enables health care organizations to manage and protect digital health data, drive data-informed decisions, and meet regulatory requirements with confidence.

Note: IBM offers a cost-effective Power Solution Edition for Healthcare for the E1150, which is intended to reduce initial investments and enhance the economics of IT infrastructure for the healthcare industry. This healthcare solution is intended to be available for US and Canada only.

Also, IBM intends to offer Power Solution Edition for Healthcare for the E1180, S1124 and S1122.

7.6.1 Healthcare and IBM Power

Designed for mission-critical workloads, IBM Power empowers health care providers to improve patient outcomes, enhance operational efficiency, and reduce risk.

Security and compliance

Health care organizations handle sensitive patient data that must be protected under regulations such as HIPAA. IBM Power delivers advanced security features including secure boot, role-based access control, and encryption at rest and in transit. With quantum-safe firmware signing and encrypted Live Partition Mobility (LPM), Power ensures data integrity and privacy. According to ITIC, organizations using Power experience an average of just 3.3 minutes of unplanned downtime per year because of security issues, demonstrating its industry-leading resilience.

Unmatched reliability for clinical systems

IBM Power is engineered for continuous availability of mission-critical applications such as EHRs, PACS, and health care information systems. Features such as predictive failure analysis, 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[3]. Power has held the title of the most reliable non-mainframe server platform for more than 15 years.

High-performance computing for health care analytics

Health care workloads, from medical imaging to genomics and real-time analytics, demand exceptional compute power. Power10-based servers can deliver 2.5 times better performance per core than previous generations, enabling faster insights for research, drug discovery, and personalized medicine, and Power11 adds an additional performance benefit to that.

Scalable and Flexible Infrastructure

Health care environments often face unpredictable demand. IBM Power supports dynamic scaling to accommodate workload fluctuations, whether driven by seasonal surges or new clinical initiatives. Its support for multiple operating systems and virtualization technologies allows health care IT teams to run diverse workloads efficiently on a single platform.

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. Power S1022 servers can process 42% more batch queries per second than comparable x86 systems under peak load, with subsecond inferencing latency, which is ideal for applications like diagnostic imaging, predictive analytics, and treatment optimization.

Seamless Integration with Health care applications

IBM Power integrates with a wide range of health care applications, including EMRs, imaging platforms, analytics tools, and telemedicine systems. It supports interoperability and data exchange across systems, and works closely with leading ISVs such as Epic to certify hardware for industry-standard health care software.

[3] https://www.itjungle.com/2024/10/21/ibm-nears-the-end-of-the-road-for-server-reliability-improvements/

7.6.2 Epic

Epic is a health care information system (HIS) provider that develops and delivers a comprehensive electronic medical record (EMR) system covering all aspects of the medical health care profession. The Epic solution includes a variety of applications that cover areas, such as medical billing, emergency room, radiology, outpatient, inpatient, and ambulatory care.

IBM and Epic Systems maintain a strategic partnership that enables health care organizations to run Epic's electronic health record (EHR) platform on IBM infrastructure, particularly IBM Power and IBM Storage solutions. This collaboration is designed to support the performance, security, and compliance needs of health care providers managing mission-critical clinical workloads.

IBM Power systems are certified to run Epic operational database (ODB) workloads, with best practices published for tuning AIX environments to meet Epic's stringent performance requirements. These configurations are optimized for high throughput, low latency, and high availability – critical for real-time clinical operations. At the time of writing, IBM Power10 servers are certified by Epic Systems, supporting key infrastructure components on both AIX and Linux for large health care deployments, primarily in the database and backend services tier. These platforms are used to host InterSystems Caché or IRIS for Health, which serve as the core databases behind Epic's EHR system.

It is expected that IBM and Epic will certify Power11 configurations after the Power11 systems are available.

Supported Products

Here is a list of currently supported Epic products:

- InterSystems Caché / IRIS for Health:
  - Officially supported by Epic on both AIX and Linux (RHEL) on Power10
  - Deployed in production for the Epic operational database (Chronicles)
- Epic production and reporting database tiers: Most often run on AIX or RHEL on Power10 for performance and reliability.
- Other backend services, such as database copies used for Clarity extracts, can also run on AIX or Linux, depending on the customer's architecture.

Note: The Epic application layer, for example, Hyperspace and Interconnect, typically runs on Windows or Linux on x86. Power-based AIX or Linux deployments are focused on the infrastructure and database tiers.

Required OS and Firmware Levels for Power10:

- AIX: Certified by both Epic and InterSystems for Power10 compatibility
  - AIX 7.2 TL5 SP4 or later
  - AIX 7.3 TL1 or later.
- 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

  High-performance SAN such as IBM FlashSystem is commonly used, with tuning for Caché/IRIS.
- HA/DR:
  - PowerHA for AIX,
  - Linux HA tools (including IBM VM Recovery Manager),
  - InterSystems mirroring for resilience.

8

Servicing Power11

The goal of serviceability is to enable efficient system repair while minimizing or avoiding any disruption to operations. It includes system installation, upgrades, downgrades, and ongoing maintenance or repair activities. Depending on the system configuration and warranty agreement, you, an IBM technician or an authorized service provider can perform service tasks.

IBM Power systems are designed with advanced serviceability features to support a highly efficient maintenance environment. These features help 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
- eBMC card
- Managing the system by using the ASMI GUI
- Entitled system support
- System firmware

8.1 IBM maintenance

IBM provides a comprehensive global maintenance and support framework for Power to ensure high availability, rapid issue resolution, and minimal disruption to business operations. Maintenance services are available under various service level agreements (SLAs) that are tailored to meet the needs of different environments, from standard business hours to mission-critical 24×7 operations. IBM's support infrastructure includes remote diagnostics, onsite service, and proactive monitoring, all backed by a global network of skilled service professionals and parts depots.

You can choose between 8×5 and 24×7 service coverage. The 8×5 option provides support during standard business hours (typically 8 AM to 5 PM, Monday through Friday, excluding holidays), which is ideal for noncritical systems or environments with internal IT support. For systems that require continuous uptime, the 24×7 service 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 have significant operational or financial consequences.

IBM also defines response time targets based on the selected service level. For example, for a 24×7 agreement, IBM might 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 IBM's extensive global logistics and service infrastructure, which includes strategically located parts centers and field engineers in more than 170 countries.

This worldwide coverage ensures that IBM Power customers receive consistent, high-quality support regardless of their geographic location. IBM's maintenance services are further enhanced by features such as call-home diagnostics, automated error reporting, and remote problem determination, which help accelerate issue resolution and reduce the need for manual intervention. Combined, these capabilities provide a robust and reliable maintenance structure that supports both traditional on-premises deployments and modern hybrid cloud environments.

8.2 IBM Expert Care

Managing your IBM Power systems should be seamless and efficient. Power Expert Care delivers immediate access to a curated bundle of high-value services that are trusted by the Power community, eliminating the delays and complexity of traditional procurement processes. With this preselected service package, clients benefit from streamlined support, which is enhanced by AI-driven tools that accelerate response times, improve case resolution, and elevate overall satisfaction.

8.2.1 Expert Care Premium

IBM Power Expert Care has three tiers: Basic, Advanced, and Premium. Premium focuses on mission-critical environments where downtime is not an option. Backed by the same experts who helped design and build the Power11 platform, this service tier ensures your infrastructure is supported by experts in technical depth and operational excellence.

Expert Care Premium includes the following key features:

- 30-minute response time for both hardware and software support cases
- 4-hour onsite response target for urgent issues
- System and microcode compatibility guidance to help ensure optimal performance

- Dedicated Technical Account Manager and mission-critical support resources
- Automated case creation and log analysis for faster issue resolution
- Health checks for hardware, OS, and applications (available as add-ons)
- Zero planned downtime support with proactive planning and TAM assistance

For non-mission-critical systems, IBM offers the Advanced Expert Care bundle providing 24×7 standard response times on repair and maintenance cases.

8.2.2 Technical Account Manager

The Technical Account Manager (TAM) plays a vital role in delivering proactive, product-based support for IBM Power systems. Acting as the primary point of contact, the TAM provides strategic guidance and direct engagement for both hardware and software within the scope of your support agreement. TAM services are delivered in English during the client's business hours, with support in other languages available upon request and mutual agreement, subject to availability.

The TAMs help streamline operations and reduce downtime through a wide range of responsibilities:

- Enabling Call Home for proactive error reporting, autonomous error resolution (AER), and zero planned downtime (ZPD), available on Power11
- Activating Support Insights for predictive analytics and delivering monthly reports
- Providing firmware and microcode compatibility analysis, which is 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 systems (onsite code load available separately for Power11)
- Supporting change management by communicating planned events to relevant teams

For clients running SAP HANA on Power10 or Power11, 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 IBM Power servers relies on a robust combination of proactive maintenance, intelligent diagnostics, and modern remote management tools, all designed to help ensure high availability and system reliability. The Hardware Management Console (HMC) serves 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 further 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 additional troubleshooting capabilities. For modern, automated environments, the service processor also supports DMTF Redfish APIs, which enable 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 IBM Power systems remain secure, efficient, and continuously operational.

8.4 eBMC card

The IBM Power E1150 server uses the enterprise baseboard management controller (eBMC) as a central component for system service management, monitoring, maintenance, and control. The eBMC is a specialized service processor that operates independently from the main system processors and continuously monitors the physical state of the system through onboard sensors. It provides access to critical diagnostic data, including system event log (SEL) files, which are essential for identifying and resolving hardware issues. This independent operation ensures that you can monitor system health even when the main system is powered down or unresponsive.

You can interact with the eBMC through multiple interfaces. It integrates seamlessly with the Hardware Management Console (HMC) for centralized system management and supports direct access through the Advanced System Management Interface (ASMI) for low-level configuration and troubleshooting. Additionally, the eBMC supports DMTF Redfish APIs, which enable secure, modern, and automated management through RESTful interfaces. This capability makes it easier to integrate IBM Power servers into broader data center orchestration and monitoring frameworks. Designed with serviceability in mind, the eBMC card is field-replaceable, which helps minimize downtime and maintain high system availability.

8.5 Managing the system by using the ASMI GUI

The ASMI is the GUI to the eBMC. It is similar in terms of its function to the ASMI of FSP-managed servers (for example, the Power E1180 server). However, it is a complete redesign of the UI, which is a result of customer feedback from a Design Thinking workshop.

8.5.1 Accessing the ASMI

The following steps describe how to access and use the ASMI:

1. To open the ASMI GUI from the HMC, select the server, and then select **Operations** → **Launch Advanced System Management**, as shown on Figure 8-1 on page 153.

   A window opens and shows the name of the system, MTM and serial number, and the IP address of the service processor (eBMC).

2. Click **OK** to open the ASMI window.

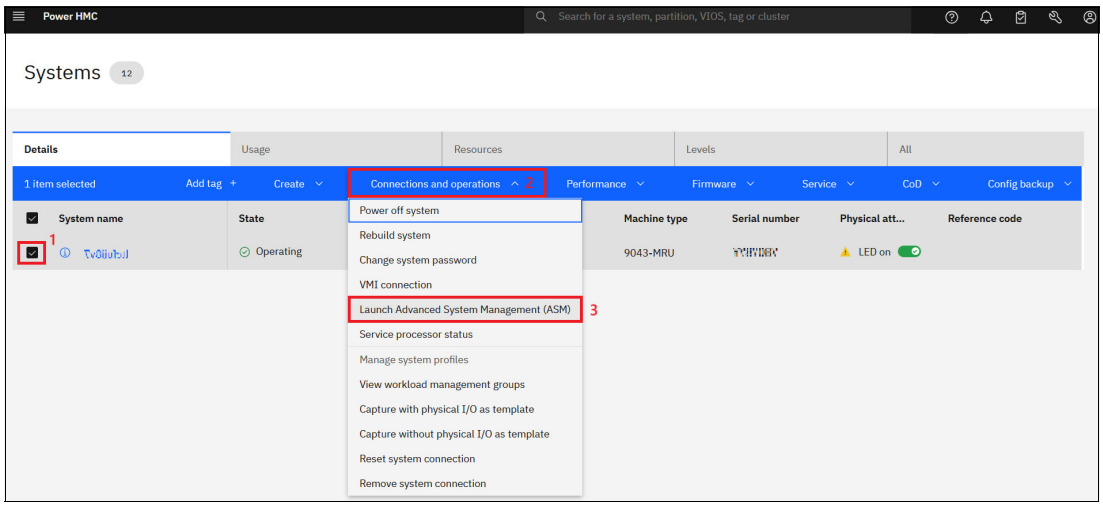

Figure 8-1 Launch Advanced System Management (ASMI).

3. If the eBMC is not in a private, reachable network, you can connect directly by entering https://<eBMC IP> into your web browser. Figure 8-2 show the login window for the ASMI.

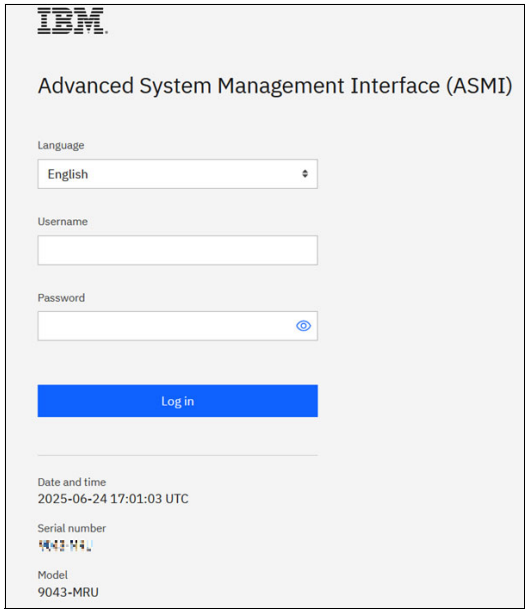

Figure 8-2 ASMI login window.

The eBMC has the following default login credentials:

- username: admin
- password: admin

When you first log in or after performing a factory reset, the default password is immediately invalidated, and you are required to set a new, secure password. This policy is in place to prevent the system from being left in a vulnerable state with a well-known default password. The new password must meet strong security standards and should not be something easily guessable, such as abcd1234. For detailed information on password complexity requirements and best practices, see Setting the password.

After you log in, the Overview page opens with server, firmware, network, power, and status information.

Figure 8-3 shows the system overview page on the ASMI.

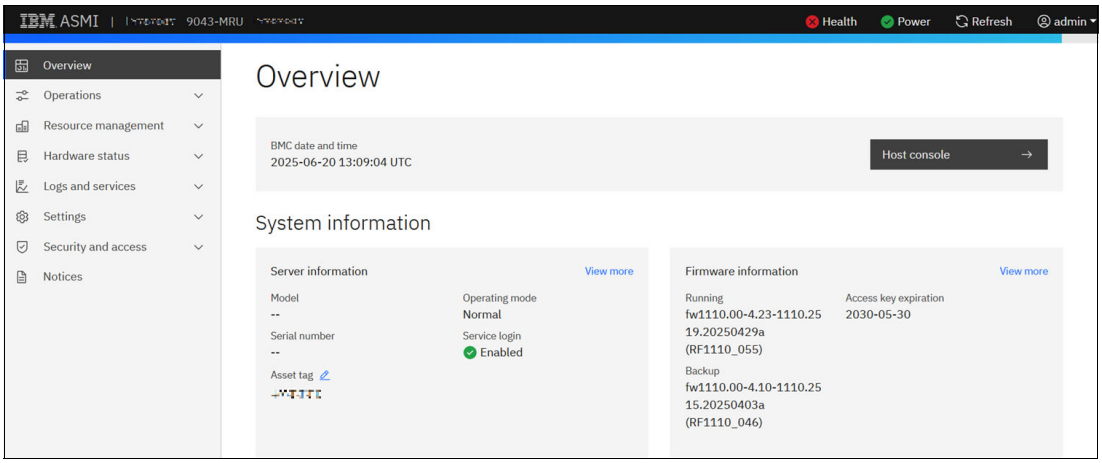

Figure 8-3 ASMI Overview window.

8.5.2 Features of the ASMI

The ASMI for eBMC managed servers includes the following features:

- System firmware updates

  It is possible to install a firmware update for the server by using the ASMI GUI, even if the system is managed by an HMC. However, the firmware update is always disruptive when you use the ASMI GUI. To install a concurrent firmware update, you must use the HMC.

- Download of dumps

  Dumps can be downloaded by using the HMC, but if necessary, you can also download them from the ASMI menu.

  It is also possible to initiate a dump from the ASMI by selecting **Logs** → **Dumps**, selecting the dump type, and clicking **Initiate dump**. You can view the following dump types:

  - Baseboard management controller (BMC) dump (nondisruptive)
  - Resource dump
  - System dump (disruptive)

- Network Time Protocol (NTP) server support
- Lightweight directory access protocol (LDAP) for user management
- Host console

  By using the host console, you can watch the boot process of the server. The host console can also be used to access the OS if there is only a single LPAR that uses all resources.

  Note: The host console can also be accessed by using an SSH client over port 2200 and logging in with the admin user.

- User management

  In the eBMC, you can create your own users. You can also create an individual user that can be used for the HMC to access the server.

  There are two types of privileges for a user: Administrator or ReadOnly. As the name indicates, with the ReadOnly privileges, you cannot modify anything except the password

of that user, and a user with ReadOnly privilege cannot be used for HMC access to the server.

- Access Control for eBMC-Managed Systems

  For systems that use the eBMC, such as the IBM Power E1150, access to the service processor is managed through a secure and auditable process. In scenarios where administrative credentials are unavailable, such as a lost or forgotten password, IBM Support can generate a time-limited, digitally signed Access Control File (ACF).

  The ACF must be uploaded to the server through the eBMC interface to temporarily enable administrative access. This mechanism ensures that access is tightly controlled, traceable, and compliant with security best practices, without requiring permanent credential resets or exposure.

- Jumper reset

  It is possible to reset everything on the server by using a jumper. This reset is a factory reset that resets everything, including LPARs, eBMC settings, and the NVRAM.

Real-time progress indicator

The ASMI of an eBMC server also provides a real-time progress indicator to see the operator panel codes. To open the windows that show the codes, select **Logs** → **Progress logs**, and then click **View code in real time**.

Inventory and LEDs

Under the **Hardware status** → **Inventory and LEDs** menu in the eBMC interface, you can view the status of most hardware components, including their associated identification LEDs. This section also provides access to the System Identify LED and the System Attention LED.

All identification LEDs can be toggled on or off to assist with physical component identification. However, only the System Attention LED can be turned off manually.

Additionally, you can view detailed information for each hardware component. This includes specifications such as the size of a DIMM or part numbers, which is particularly useful when planning hardware replacements or upgrades.

Sensors

The ASMI has many sensors for the server, which you can access by selecting **Hardware status** → **Sensors**. A progress bar indicates the status as the sensors are loading.

Note: Although the progress bar might be finished, it might take additional time until the sensors appear.

Figure 8-4 on page 156 shows an example of the Sensors window.

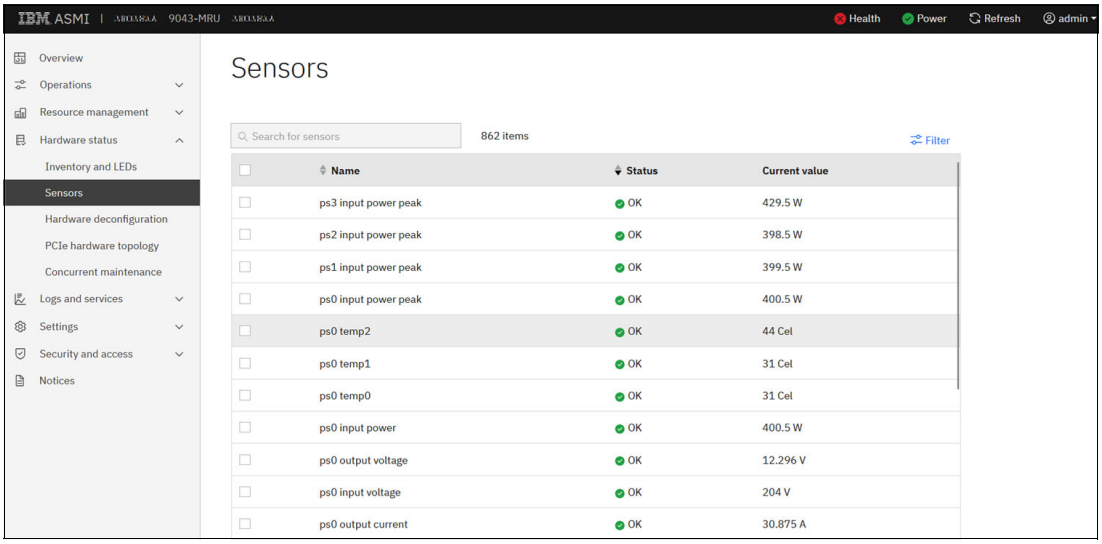

Figure 8-4 ASMI Sensors window.

8.5.3 Network settings

The default network setting for the two eBMC ports is DHCP. Therefore, when you connect a port to a private HMC network where the HMC is connected to a DHCP server, the HMC provides an IP address to the newly connected system during the startup of the firmware. Then, the new system automatically appears in the HMC and can be configured. As a best practice, use DHCP to attach a server to the HMC.

If you do not use DHCP and want to use a static IP address, you can set the IP address in the ASMI GUI. However, because there are no default IP addresses that are the same for every server, you first must discover the configured IP address.

To discover the configured IP address, use the operator panel and complete the following steps:

1. Use the Increment or Decrement buttons to scroll to function 02.
2. Press Enter until the value changes from N to M, which activates access to function 30.
3. Scroll to function 30 and press Enter. Function 30** appears.
4. Scroll to 3000 and press Enter, which shows you the IP address of the eth0 port.
5. If you scroll to 3001 and press Enter, you see the IP address of eth1.
6. After you discover 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 panel, see Function 30: Service processor IP address and port location.

After you discover the IP address, you can connect any computer with a web browser to an IP address in the same subnet (class C) and connect the computer with the correct Ethernet port of the Power E1150 server.

Hint: Most connections work by using a standard Ethernet cable. If you do not see a link when using a standard Ethernet cable, try a crossover cable, where the send and receive wires are crossed.

After connecting the cable, you can use a web browser to access the ASMI by using the URL `https://<IP address>`, and then you can configure the network ports. To configure the network ports, select **Settings** → **Network** and select the correct adapter to configure. Figure 8-5 shows an example of changing eth1. Before you can configure a static IP address, turn off DHCP. It is possible to configure several static IP addresses on one physical Ethernet port.

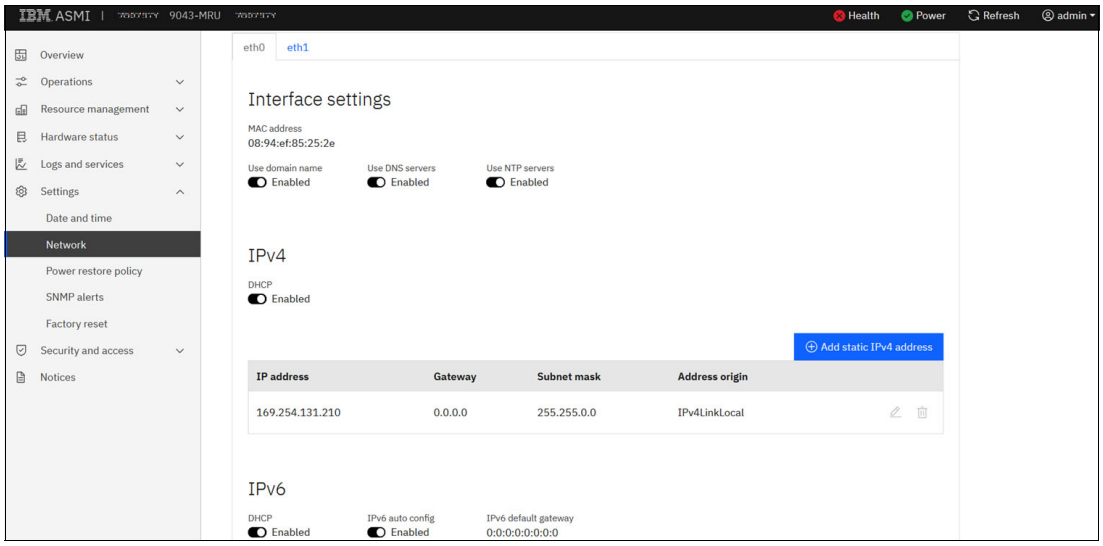

Figure 8-5 ASMI network settings

You cannot configure the address of the Virtualization Management Interface (VMI) in the ASMI network settings. The VMI address is another IP address that is configured on the physical eBMC Ethernet port of the server to manage the virtualization of the server. The VMI address can be configured only 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. An ACF is a digital certificate that is provided by IBM Support 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 the Increment or Decrement buttons to scroll to function 74.
2. Press Enter, and select **00** to accept the function (**FF** rejects it).
3. The ACF function now is active for 30 minutes. To use it, access the ASMI login window.
4. To upload the ACF into the system and access the ASMI, click **Upload service login certificate**.

For more information, see Function 74: Authentication Override for ACF upload.

Policies

In the **Security and access** → **Policies** menu, you can turn on and off security-related functions. For example, you can specify whether you can manage your server by using Intelligent Platform Management Interface (IPMI).

Some customers require that the USB ports of the server be disabled. You can accomplish this task by clearing **Host USB enablement**.

8.5.4 Managing the system by using DMTF Redfish API

eBMC-based systems can be managed by using the DMTF Redfish application programming interfaces (APIs). Redfish is a Representational State Transfer (REST) API that is used for platform management and is standardized by the Distributed Management Task Force, Inc. For more information, see Redfish.

There are multiple ways of working with Redfish. One possibility is to use the OS command **`curl`**. The following examples show how to work with **`curl`** and Redfish.

Before you can get data from the server or run systems management tasks by using Redfish, you must authenticate against the server. After you authenticate by using a user ID and password, you receive a token from the serve that you use later.

Example 8-1 shows how to obtain a token.

Example 8-1 Obtaining a token from Redfish

```
# export eBMC=<IP>
# export USER=admin
# export PASSWORD=<PW>

# export TOKEN=`curl -k -H "Content-Type: application/json" -X POST https://${eBMC}/login
-d "{\"username\" : \"${USER}\", \"password\" : \"${PASSWORD}\"}" | grep token | awk
'{print $2;}' | tr -d '"'`
```

With the token, you can get data from the server. First, request the data of the Redfish root by using `/Redfish/v1`. You get data with other branches in the Redfish tree, for example, Chassis (uppercase C). To dig deeper, use the newly discovered `odata.id /Redfish/v1/Chassis`, as shown in Example 8-2.

Example 8-2 Getting chassis data from Redfish

```
#curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/redfish/v1
{
  "@odata.id": "/redfish/v1",
  "@odata.type": "#ServiceRoot.v1_12_0.ServiceRoot",
  "AccountService": {
    "@odata.id": "/redfish/v1/AccountService"
  },
  "Cables": {
    "@odata.id": "/redfish/v1/Cables"
  },
  "CertificateService": {
    "@odata.id": "/redfish/v1/CertificateService"
  },
  "Chassis": {
    "@odata.id": "/redfish/v1/Chassis"
  },

# curl -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/Redfish/v1/Chassis
{
  "@odata.id": "/Redfish/v1/Chassis",
  "@odata.type": "#ChassisCollection.ChassisCollection",
  "Members": [
    {
```

```
      "@odata.id": "/Redfish/v1/Chassis/chassis"
    }
  ],
  "Members@odata.count": 1,
  "Name": "Chassis Collection"

# curl -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/Redfish/v1/Chassis/chassis
{
  "@odata.id": "/Redfish/v1/Chassis/chassis",
  "@odata.type": "#Chassis.v1_16_0.Chassis",
  "Actions": {
...
"PCIeSlots": {
    "@odata.id": "/Redfish/v1/Chassis/chassis/PCIeSlots"
  },
...
"Sensors": {
    "@odata.id": "/Redfish/v1/Chassis/chassis/Sensors"
  },
...
```

Under Chassis (uppercase C), there is another chassis with a lowercase c. Use the tree with both (`/Redfish/c1/Chassis/chassis`). After running the chassis, you can see in Example 8-2 on page 158 for example, that there are `PCIeSlots` and `Sensors`, among other resources of the server.

In Example 8-3, you can see what is under `Sensors`, and you can find the same sensors as in the ASMI GUI (see Figure 8-4 on page 156). In the output, for example, you can view the sensor, `total_power`. When you request details about that sensor, as shown in Example 8-3, you can see that the server needed 1.426 watts at the time of running the command.

Example 8-3 Getting sensor data from Redfish

```
# curl -k -H "X-Auth-Token: $TOKEN" -X GET
https://${eBMC}/Redfish/v1/Chassis/chassis/Sensors
{
  "@odata.id": "/Redfish/v1/Chassis/chassis/Sensors",
  "@odata.type": "#SensorCollection.SensorCollection",
  "Description": "Collection of Sensors for this Chassis",
  "Members": [
    {
      "@odata.id": "/Redfish/v1/Chassis/chassis/Sensors/Altitude"
    },
    {
      "@odata.id": "/Redfish/v1/Chassis/chassis/Sensors/1V1CS_0167_p1_rail_iout"
    },
...
    {
      "@odata.id": "/Redfish/v1/Chassis/chassis/Sensors/total_power"
    },

# curl -k -H "X-Auth-Token: $TOKEN" -X GET \
https://${eBMC}/Redfish/v1/Chassis/chassis/Sensors/total_power
{
  "@odata.id": "/Redfish/v1/Chassis/chassis/Sensors/total_power",
  "@odata.type": "#Sensor.v1_0_0.Sensor",
  "Id": "total_power",
  "Name": "total power",
  "Reading": 1426.0,
  "ReadingRangeMax": null,
```

```
  "ReadingRangeMin": null,
  "ReadingType": "Power",
  "ReadingUnits": "W",
  "Status": {
    "Health": "OK",
    "State": "Enabled"
  }
```

Some other useful strings to request are shown in Example 8-4.

Example 8-4 Other useful Redfish data

```
#Type and model of the server
   # curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/Redfish/v1/Systems/system
   | grep Model | grep -v SubModel | grep \ -v \"\"
   "Model": "9043-MRU",
#Serial number
   # curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/redfish/v1/Systems/system
   | grep SerialNumber \
   "SerialNumber": "<SN>",
#Type, model, and serial number
   #curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/redfish/v1/Systems/system |
   grep AssetTag
   "AssetTag": "Server-9043-MRU-<SN>",
#System indicator LED
   #curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/redfish/v1/Systems/system |
   grep IndicatorLED
   "IndicatorLED": "Off",
#Total memory
   # curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/redfish/v1/Systems/system
   | grep TotalSystemMemoryGiB
   "TotalSystemMemoryGiB": 8192
#Power state
   # curl -s -k -H "X-Auth-Token: $TOKEN" -X GET https://${eBMC}/redfish/v1/Systems/system
   | grep PowerState
     "PowerState": "On",
```

It is also possible to run operations on the server by using the **POST** method with the Redfish API interface. In Example 8-5, you can see the **curl** commands that can start or stop the server.

Example 8-5 POST command examples for Redfish

```
#Power on server
   # curl -k -H "X-Auth-Token: $TOKEN" -X POST https://${eBMC}/redfish/v1/Systems/
   system/Actions/Reset -d '{"ResetType":"On"}'
#Power off server
   # curl -k -H "X-Auth-Token: $TOKEN" -X POST https://${eBMC}/redfish/v1/Systems/
   system/Actions/Reset -d '{"ResetType":"ForceOff"}'
```

For more information about Redfish, see Managing the system by using DMTF Redfish APIs. For more information about how to work with Redfish in IBM Power servers, see Managing Power systems servers by using DMTF Redfish APIs.

8.5.5 Managing the system by using the IPMI

The Power E1150 server can also be managed by using the IPMI, but the IPMI is disabled by default on your server. Inherent security vulnerabilities are associated with the IPMI, so consider using Redfish APIs or the GUI to manage your system.

If you want to use IPMI, you must first enable it. To do so, select **Security and access** → **Policies**. Then, enable the policy **Network IPMI (out-of-band IPMI)**.

A list of common IPMI commands can be found at Common IPMI commands.

8.6 Entitled system support

IBM Enterprise Storage Server is available to view and manage Power and Storage software and hardware. In general, most products that are offered by IBM Systems that are purchased through our IBM Digital Sales representatives or Business Partners are accessed on this site when the IBM Configurator is used.

The site features the following three main sections:

1. My entitled software

   Activities are listed that are related to Power and Storage software, including the ability to download licensed, free, and trial software media, place software update orders, and manage software keys.

2. My entitled hardware

   Activities are listed related to Power and Storage hardware including the ability to renew Update Access Keys, buy and use Elastic Capacity on Demand, assign or buy credits for new and existing pools in a Power Private Cloud environment (Enterprise Pools 2.0), download Storage Capacity on-Demand codes, and manage Hybrid Capacity credits.

3. My inventory

   Activities related to Power and Storage inventory including the ability to browse software license, software maintenance, and hardware inventory, manage inventory retrievals by way of Base Composer or generate several types of reports.

8.7 System firmware

System firmware provides low-level control for the system hardware. New features and fixes are introduced with new firmware release levels. Fixes are often bundled into service packs. A service pack is referred to as an update level. A new release is referred to as an upgrade level. All system firmware is available for download in IBM FixCentral.

8.7.1 Terminology and strategy for firmware updates

IBM uses the following terms when describing firmware changes:

Release Level A major new function (introduction of new hardware models and significant function and features enabled via firmware). This firmware upgrade is disruptive.

Service Pack (SP)	Primarily firmware fixes and minor function changes applicable to a specific Release Level. These firmware updates are usually concurrent.
Concurrent	A code update that allows the operating system(s) running on the Power system to continue running while the update is installed and activated.
Deferred	A code fix that is installed concurrently, but does not activate until the system is restarted.
Partition Deferred	A code fix that is installed concurrently but does not activate until the partition is restarted.
Disruptive	A code fix which requires a system restart during the code update process.

Service pack severity

The severity classification is specific for each service pack that becomes available in FixCentral. All types are listed below:

NEW	New features and functions. This is considered as a new release level for a product.
PE	PTF in error: This service pack addresses minor issues. It can be installed when convenient.
ATT	Attention: This service pack addresses low impact and low potential issues. It should be installed at the customer's earliest convenience.
SPE	Special attention: This Service Pack addresses high impact but low potential issues. It should be installed at the earliest convenience.
HIPER	High impact / pervasive: This service pack addresses high impact and/or pervasive issues with significant customer impacts, and therefore should be installed as soon as possible.

For more details, see IBM Support Glossary.

System firmware update strategy for IBM Power11 systems

The system management model determines the appropriate firmware update strategy for IBM Power11 systems.

- For managed systems, the recommended method for updating system firmware is by using the Hardware Management Console (HMC).
- For unmanaged systems, the preferred method is to update the system firmware directly through the operating system.

Although these are the recommended approaches, it is also possible to update the system firmware y using the eBMC.

The update method impacts the type of update:

- Firmware updates from the HMC can be either concurrent or disruptive, depending on the firmware version.
- Firmware updates from the operating system or eBMC are always disruptive.

A concurrent firmware update does not require a system restart. This is indicated in the firmware description file under the Service Pack Summary, which categorizes the update as either Disruptive Service Pack, Deferred Service Pack, or Concurrent Service Pack.

However, any fixes marked as DEFERRED in the Service Pack Summary do not take effect until the next system IPL. Fixes labeled as DEFERRED: PARTITION_DEFERRED require a partition reboot to take effect.

Depending on the system management model, the following firmware update options are available:

- HMC-managed systems: Firmware updates can be performed using the HMC or the eBMC ASMI interface.
- Co-managing a system with both PowerVM NovaLink and the HMC: Firmware updates can be performed using the HMC, the NovaLink partition, or the eBMC ASMI interface. However, to perform the update through the HMC, the system must first be transitioned from PowerVM NovaLink management to HMC management.
- Unmanaged systems running IBM i: Updates can be applied using PTFs or by using the eBMC ASMI interface.
- Unmanaged systems running AIX: Firmware updates are supported through AIX system diagnostics or the eBMC ASMI interface.
- All IBM Power11 systems also support system firmware updates through the eBMC USB port.

Note: The responsibility for performing system firmware updates resides with the customer. If the customer requests that an IBM Support Service Representative (SSR) perform the firmware update, the update is considered a billable service unless the customer has a valid support agreement that explicitly includes on-site firmware update coverage.

The firmware description file 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 service pack, it is strongly recommended to review the firmware description package.

For cross-version compatibility details, see the IBM Power systems Fix Level Recommendation Tool, Power code matrix.

8.7.2 Updating system firmware

There are several strategies and methods for updating the system firmware in IBM Power.

Updating system firmware from the HMC

Use the following steps to update or upgrade the system firmware using the HMC:

1. In the System View, select the system you want to update.
2. Click the **Firmware** menu.
3. Choose **Update System Firmware**.

   The Update System Firmware Wizard guides you through the necessary steps to complete the firmware update process, as shown in Figure 8-6 on page 164.

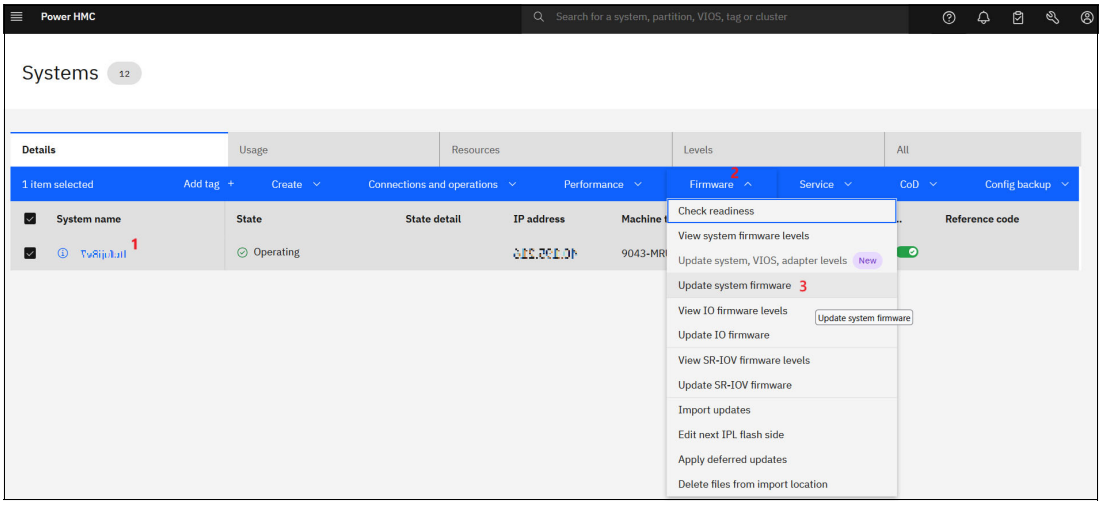

Figure 8-6 Launch the Update System Firmware wizard.

4. Accept the license agreement and click **Next**.
5. After the Readiness Check is finished, provide the required firmware files and select the firmware level to be applied.
6. The final step involves monitoring the update progress and confirming its successful completion. Figure 8-7 through illustrate these key steps in the update process.

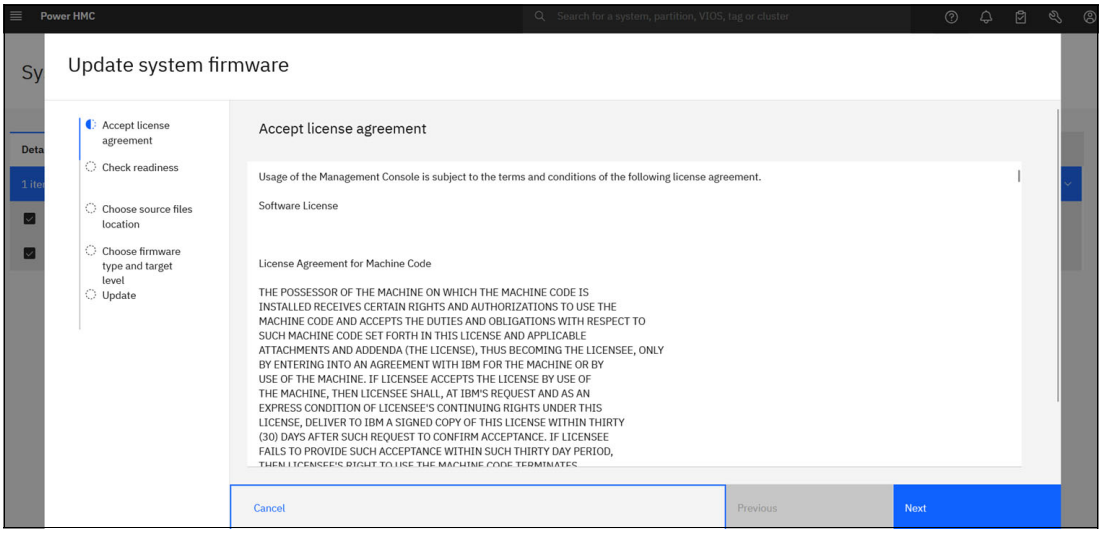

Figure 8-7 The license agreement is displayed for the user to review and accept the terms.

Figure 8-8 on page 165 shows that the system readiness check has completed successfully for the IBM Power11 system firmware update process.

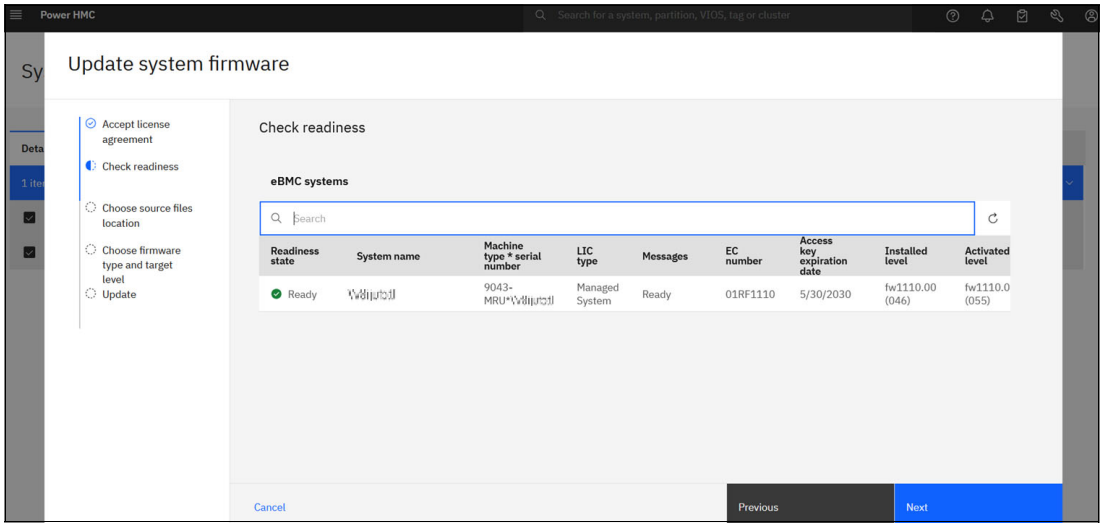

Figure 8-8 The system readiness check completed successfully.

7. Figure 8-9 illustrates the Choose source files location drop-down menu, which is used 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

   Select a location for the source files and click **Next**.

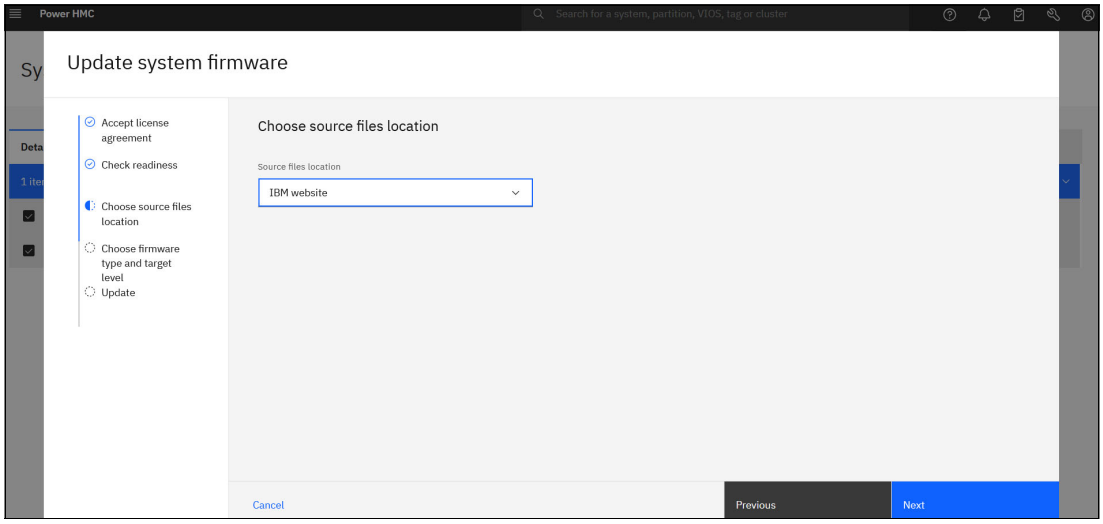

Figure 8-9 The Choose source file location drop-down menu used for updating the system firmware

8. As shown in Figure 8-10 on page 166, click **Search available levels** to locate the new system firmware on the installation media.
9. On the following screen, use the drop-down menu under the Target eBMC Level column to select the system firmware level.
10. From the drop-down menu under the Type column, select **Update** or **Upgrade.**

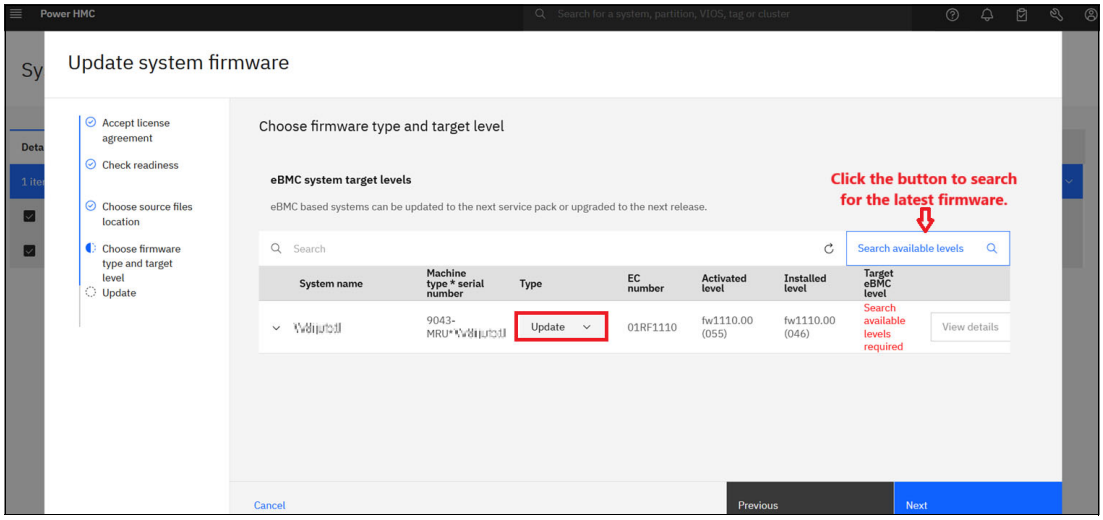

Figure 8-10 System firmware target level selection screen. Use this screen to choose between a firmware update or upgrade.

Importing updates from HMC

The Import updates wizard in the HMC GUI, as shown in Figure 8-11, guides you through the necessary steps to import firmware update or upgrade files. These files can then be reused to update the firmware on other IBM Power11 systems.

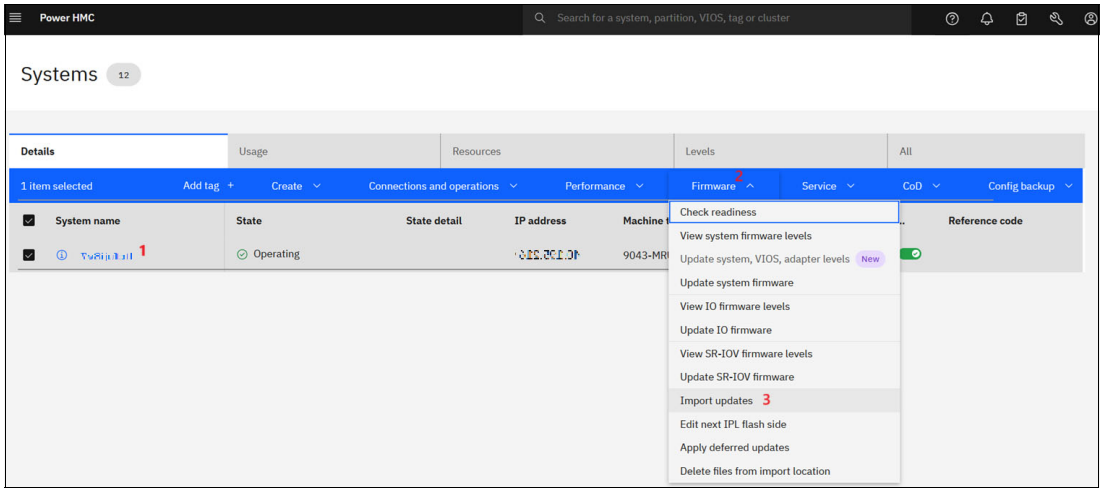

Figure 8-11 The Import updates option in the HMC GUI.

Deleting files from import location

The Delete files from import location wizard in the HMC GUI, as shown in Figure 8-12 on page 167, guides you through the steps to delete imported firmware files. This allows you to manage firmware files that are stored on the HMC.

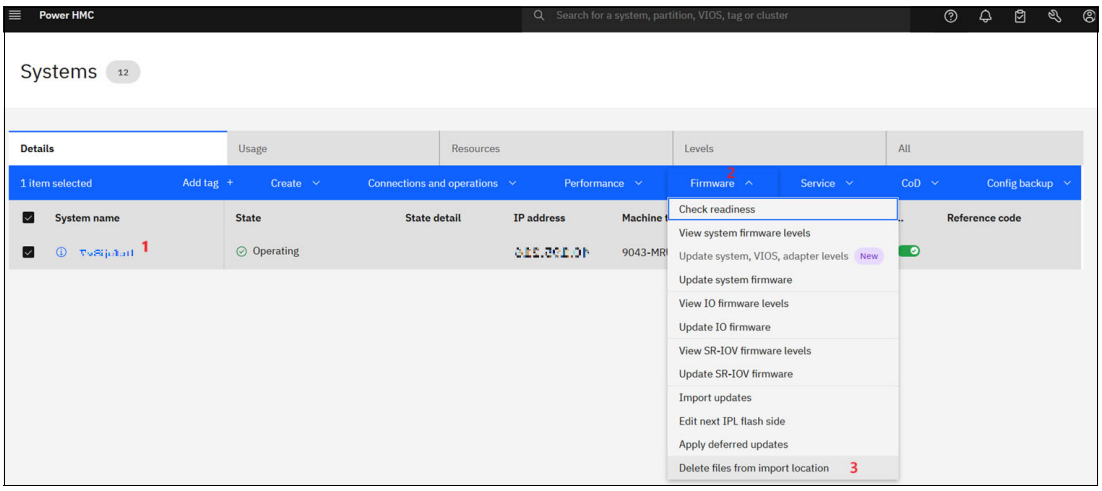

Figure 8-12 The Delete files from import location option in the HMC GUI.

SR-IOV shared mode adapter firmware management

SR-IOV capable adapters running in shared mode use a different firmware versioning mechanism compared to those operating in dedicated mode.

Verity the configuration requirements:

- Ensure that the managed system includes PCIe adapter(s) that support SR-IOV functionality.
- Each SR-IOV-capable adapter must be installed in a PCIe slot that supports SR-IOV. For more information, see Adapter placement for the 9043-MRU.
- Verify that the system is running supported levels for the required software:
  - System firmware
  - HMC or PowerVM NovaLink
  - Operating system (with SR-IOV driver support)
- Configure the adapter(s) in SR-IOV shared mode by using either the HMC or NovaLink interface.

For an SR-IOV adapter in shared mode, the adapter firmware and driver are integrated with the system firmware. Firmware updates and accompanying ReadMe documentation are available on IBM Fix Central.

When an adapter is configured in SR-IOV Shared mode, its firmware and driver are automatically updated to the latest versions 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. Additionally, if system firmware updates are installed concurrently, selective manual updates of SR-IOV adapter firmware can be performed using the Hardware Management Console (HMC).

For more information, see Update SR-IOV Firmware and Updating the SR-IOV adapter firmware.

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 8-13 displays the View SR-IOV firmware levels and Update SR-IOV firmware options that are available in the HMC GUI.

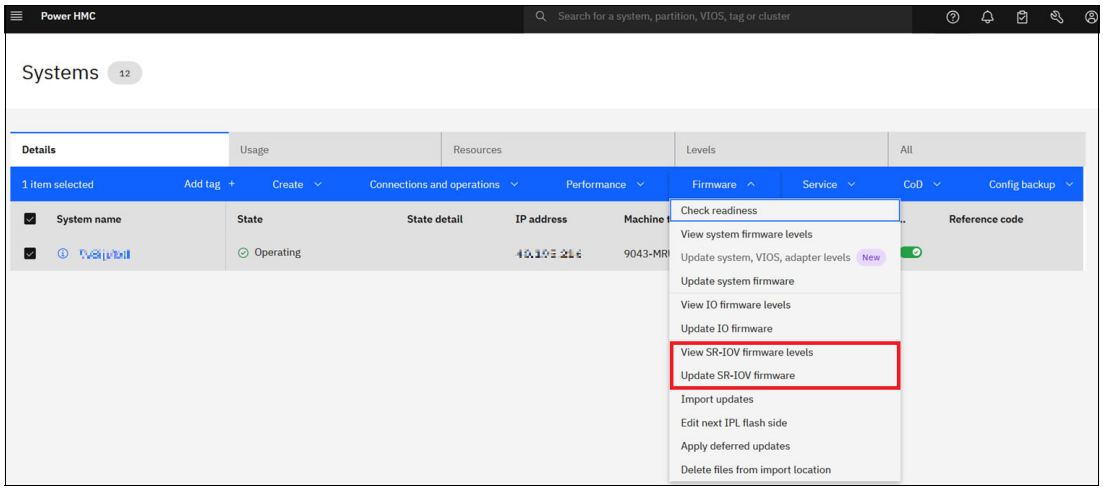

Figure 8-13 View SR-IOV firmware levels and Update SR-IOV firmware options available in the HMC GUI.

Updating system firmware by using the eBMC

You can use the eBMC Advanced System Management Interface (ASMI) to update the system firmware on both managed and unmanaged systems. To perform a firmware update by using the eBMC ASMI, the system must be powered off. The firmware package includes multiple files, but only the firmware image with the .tar file extension is required for the update process. See Figure 8-14.

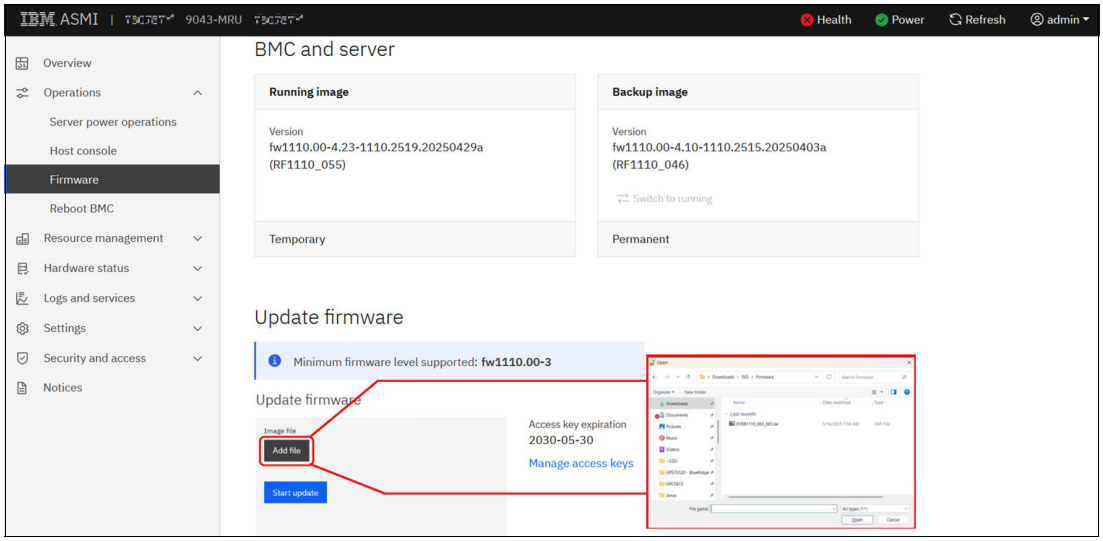

Figure 8-14 Firmware Update by using the eBMC ASMI.

I/O microcode update strategy

Microcode updates are the responsibility of the customer, unless the services are provided under a Microcode Support Services contract.

Device microcode can be updated by using the following methods:

- HMC: The Update I/O firmware feature available in the HMC GUI allows administrators to apply microcode updates to supported I/O devices.

- Diagnostic menus (AIX and VIOS): Update the microcode 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. For more information, see Standalone Diagnostics CD Overview and Download.
- Linux: Refer to the microcode README files available on IBM FixCentral for detailed instructions about updating microcode on Linux systems. These instructions can include the use of additional 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 Program Temporary Fixes (PTFs), which can be downloaded from IBM FixCentral.

Note: IBM i is not supported on E1150 servers. The information provided in this section is included for reference purposes only and to highlight IBM i options available on other IBM Power11 systems that do support IBM i.

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 to update those levels by using a firmware repository. The HMC does not have direct access to log in to the partitions to perform I/O firmware updates. Instead, it relies on the Resource Monitoring and Control (RMC) infrastructure. The RMC infrastructure is also used for features like Dynamic Logical Partitioning (DLPAR) and Service Focal Point to facilitate communication between the HMC and the partitions. On each partition, the **`invscout`** command handles the query and update process. The **`invscout`** command exchanges inventory and update-related files with the HMC over the RMC interface.

Figure 8-15 displays the View I/O firmware levels and Update I/O firmware options available in the HMC GUI.

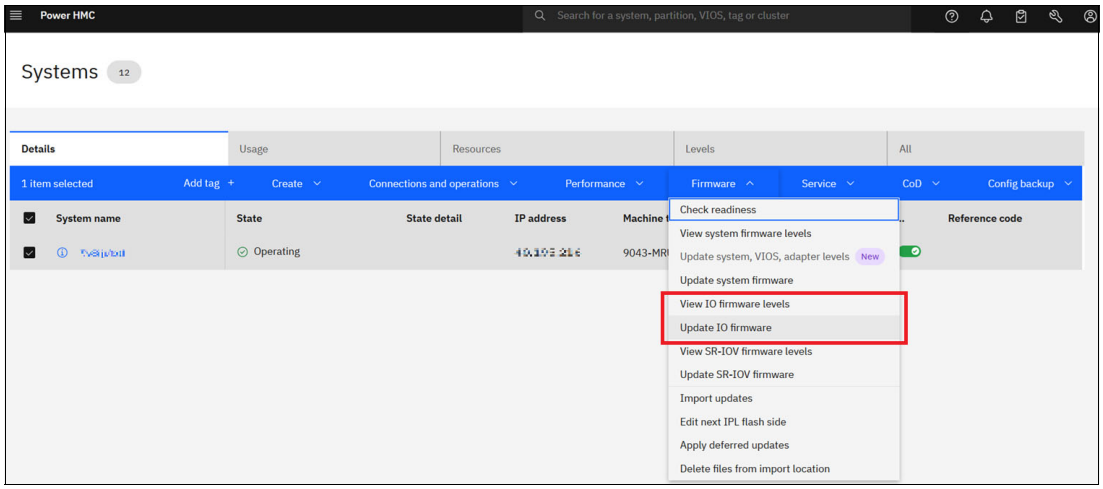

Figure 8-15 View I/O firmware levels and Update I/O firmware options available in the HMC GUI.

The View I/O firmware levels panel in Figure 8-16 on page 170 shows I/O firmware information for both FSP and eBMC-based systems, as available in the HMC GUI.

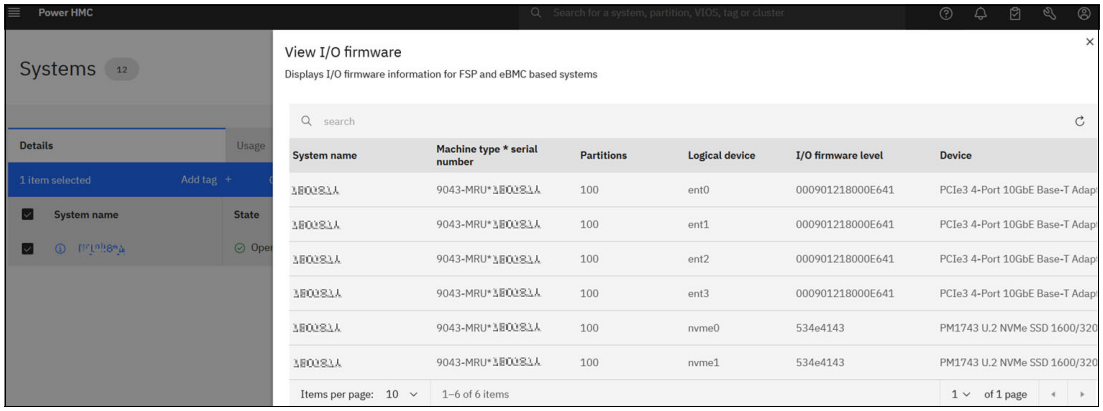

Figure 8-16 View I/O firmware levels panel available in the HMC GUI.

Figure 8-17 illustrates the Choose source files location drop-down menu, which is used 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

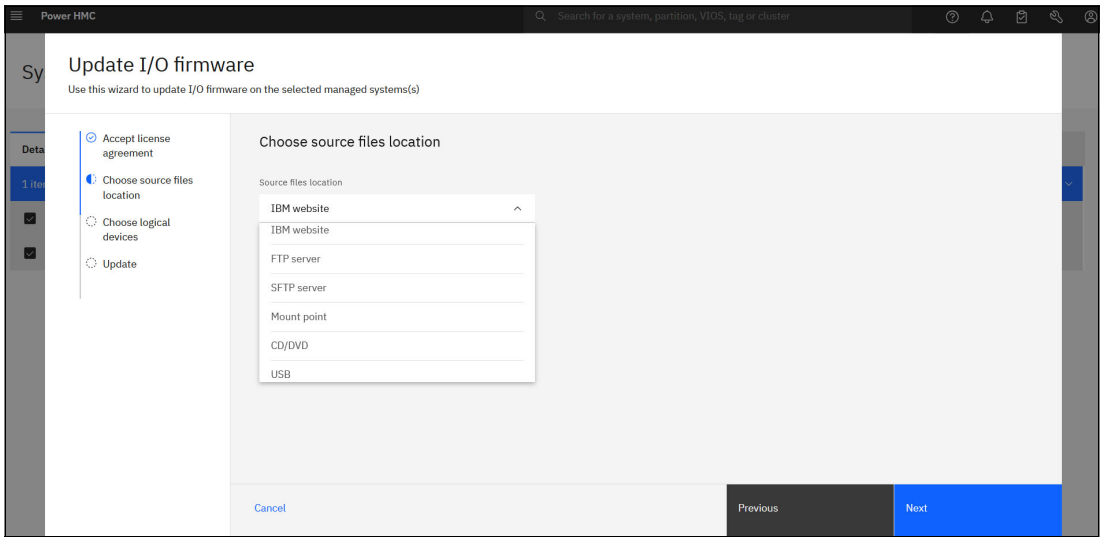

Figure 8-17 Update I/O firmware screen

For additional information on how to view or update the I/O firmware from the HMC, see View or update the I/O firmware from the HMC.

8.7.3 Updating access keys

Since the introduction of the IBM Power8 processor-based servers, IBM uses the concept of an update access key (UAK) for each server.

When system firmware updates are applied to the system, the UAK and its expiration date are checked. System firmware updates include a release date. If the release date for the firmware updates is past the expiration date for the update access key, the updates are not processed.

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 a firmware update is applied, the system checks whether the update's release date is later than the UAK's expiration date. If it is, the firmware update is not processed. When a UAK expires, it must be replaced by using either the HMC or the eBMC ASMI interface. Importantly, the expiration of a UAK does not affect I/O microcode updates these can still be applied even if the UAK has 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 an UAK that expires after three years. Thereafter, the UAK can be extended every six 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.

The validity and expiration date of the current UAK can be verified by using either the HMC or eBMC graphical interfaces or through their command-line interfaces. The expiration date can also be retrieved from the operating system level.

Verifying the expiration date of UAKs Using HMC

The current UAK expiration date is visible on the View current system firmware levels panel of the HMC, as shown in Figure 8-18.

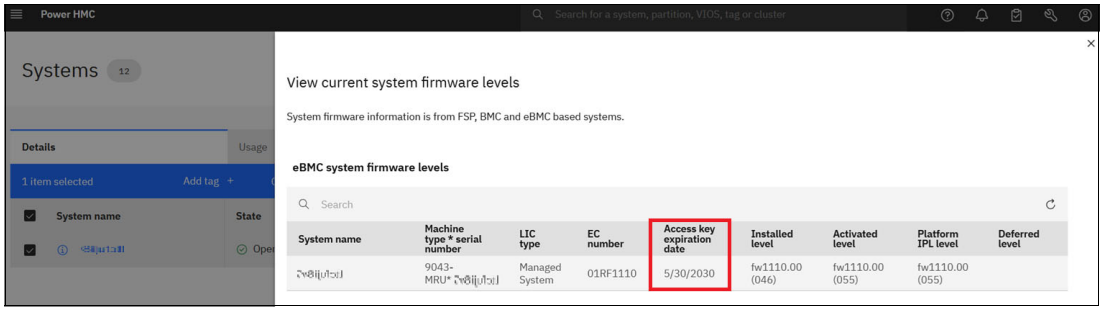

Figure 8-18 UAK expiration date displayed on the HMC View current system firmware levels page.

The HMC or virtual HMC software appliance (vHMC) can be used to configure the automatic firmware update access key. For more information, How to Automatically Update Access Key.

Verifying the expiration date of UAKs by using eBMC ASMI

The current UAK expiration date is visible on the Firmware page and the Overview page of the eBMC ASMI, as shown in Figure 8-19 on page 172 and Figure 8-20 on page 172.

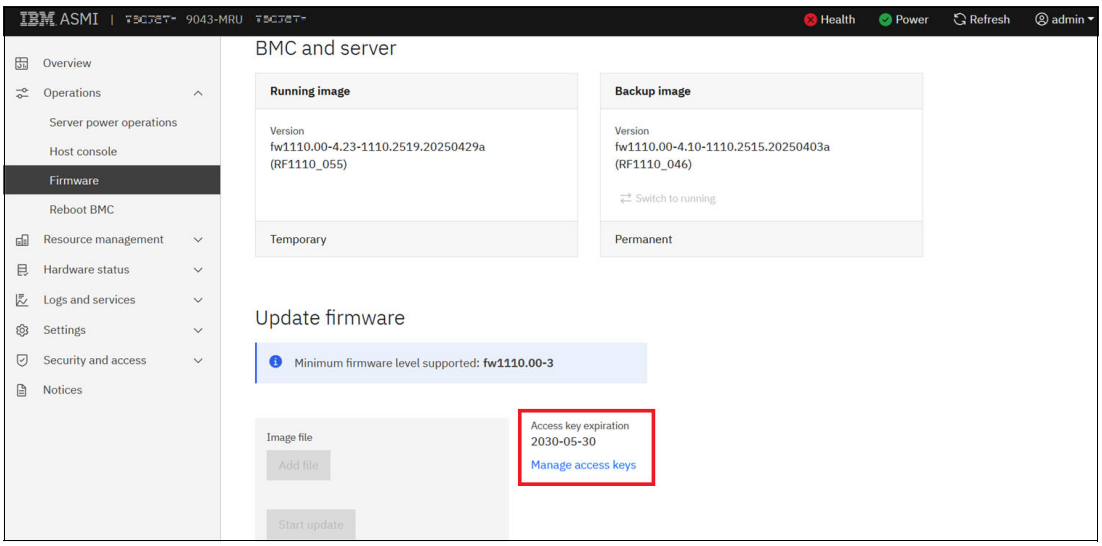

Figure 8-19 UAK expiration date displayed on the eBMC Firmware page.

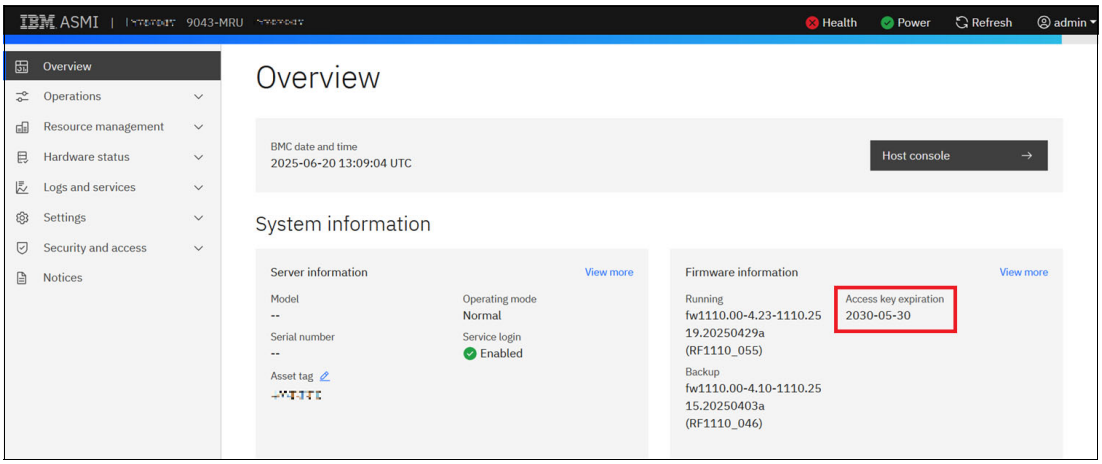

Figure 8-20 UAK expiration date displayed on the eBMC Overview page.

Verifying the expiration date of UAKs by using AIX

There are multiple methods of checking the UAK expiration date within AIX:

- The first option uses the **lscfg** command. Use the following command:

  lscfg -vpl sysplanar0 |grep -p "System Firmware"

  The output is similar to the output that is shown in Example 8-6.

Example 8-6 Output of the lscfg command to check 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
```

- Alternatively, you can **grep** on Access Key as shown in Example 8-7.

Example 8-7 Output of the lscfg command to check UAK expiration date

```
# lscfg -vpl sysplanar0 | grep "Access Key"
        Update Access Key Exp Date..20301014
```

- A third option on AIX 7.3 is the **lparstat** command, which is shown in Example 8-8.

Example 8-8 Output of the lparstat command to check 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 isn't a single Linux command to view a UAK directly. Instead, UAKs are managed through IBM's tools and interfaces, such as ESA, ASMI, and HMC.

9

Virtualization

Virtualization on IBM Power is a cornerstone of its enterprise computing strategy. It offers robust, scalable, and secure environments for running multiple workloads on a single physical server. At the heart of this capability is PowerVM, IBM enterprise-grade virtualization technology. PowerVM enables the creation of logical partitions (LPARs), allowing multiple operating systems such as AIX, IBM i, and Linux to run concurrently on the same hardware. It supports advanced features such as live partition mobility, dynamic resource allocation, and Micro-Partitioning, which help maximize hardware usage and reduce operational costs.

In addition to PowerVM, IBM Power systems 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 well with modern orchestration tools and supports features such as nested virtualization and SR-IOV for high-performance networking. KVM is ideal for organizations that want to use open-source technologies but still benefit from the performance and reliability of IBM Power hardware.

A key component of virtualization efficiency on IBM Power is the use of shared processor pools (SPPs). SPPs allow multiple LPARs to share a pool of physical processors, enabling dynamic allocation of CPU resources based on workload demand. Power11 PowerVM adds support for resource groups to enhance processor sharing capabilities and improve efficiency. Processor pools extend the capabilities of shared processor pools to improve processor usage and enhance isolation and control by limiting the maximum CPU resources available to each pool. When combined with PowerVM or KVM, SPPs help optimize performance, reduce licensing costs, and ensure consistent service levels across virtual environments.

This chapter contains the following topics:

- PowerVM
- KVM support

9.1 PowerVM

The PowerVM platform is the family of technologies, capabilities, and offerings that delivers industry-leading virtualization for enterprises. It is the term that describes Power processor-based server virtualization, which includes IBM Power Hypervisor, logical partitioning, IBM Micro-Partitioning®, Virtual I/O Server (VIOS), Live Partition Mobility (LPM), and more. PowerVM is a combination of hardware enablement and software.

PowerVM Enterprise Edition is included with every IBM Power E1150 server. It is offered as hardware feature code #EPVW and provides the following features:

- Support for up to 20 partitions per core
- Virtual I/O Server (VIOS)
- Resource groups that can be combined with Shared Processor Pools (SPPs)
- Live Partition Mobility (LPM) for seamless workload movement

9.1.1 IBM Power Hypervisor

Power processor-based servers are combined with PowerVM technology and offer the following key capabilities that can help to consolidate and simplify IT environments:

- Improve server usage and share I/O resources to reduce the total cost of ownership (TCO) and better use IT assets
- Improve business responsiveness and operational speed by dynamically reallocating resources to applications as needed to better match changing business needs or handle unexpected changes in demand
- Simplify IT infrastructure management by making workloads independent of hardware resources so that business-driven policies can be used to deliver resources that are based on time, cost, and service-level requirements

Combined with features in the Power E1180, the IBM Power Hypervisor delivers functions that enable other system technologies, including logical partition (LPAR) technology, virtualized processors, IEEE VLAN-compatible virtual switch, virtual SCSI adapters, virtual Fibre Channel adapters, and virtual consoles.

The Power Hypervisor is a basic component of the system's firmware and offers the following functions:

- Provides an abstraction between the physical hardware resources and the LPARs that use them.
- Enforces partition integrity by providing a security layer between LPARs
- Controls the dispatch of virtual processors to physical processors
- Saves and restores all processor state information during a logical processor context switch
- Controls hardware I/O interrupt management facilities for LPARs
- Provides VLAN channels between LPARs that help reduce the need for physical Ethernet adapters for inter-partition communication
- Monitors the flexible service processor (FSP) and performs a reset or reload if it detects the loss of an FSP, notifying the operating system if the problem is not corrected

The Power Hypervisor 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 is required by the Power Hypervisor firmware 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 hardware page table (HPT) that contributes to hypervisor memory usage. The HPT is used by the operating system to translate from effective addresses (EAs) to physical real addresses in the hardware. This translation from effective to real addresses allows multiple operating systems to run simultaneously in their own logical address space. Whenever a virtual processor for a partition is dispatched on a physical processor, the hypervisor indicates to the hardware the location of the partition HPT that can be used when translating addresses.

The amount of memory for the HPT is based on the maximum memory size of the partition and the HPT ratio. For AIX, VIOS, and Linux partitions, the default HPT ratio is 1/128th of the maximum memory size of the partition. AIX, VIOS, and Linux use larger page sizes of 16 KB and 64 KB instead of using 4 KB pages. The use of 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 maximum memory size that is specified 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.

The HPT ratio can also be configured. Use the **hpt_ratio** parameter for the **chsyscfg** HMC command to define the HPT ratio that is used for a partition profile. The valid values are 1:32, 1:64, 1:128, 1:256, or 1:512. Specifying a smaller absolute ratio decreases the overall memory that is assigned to the HPT.

Testing is required when changing the HPT ratio. A smaller HPT might incur more CPU usage because the operating system might need to reload the entries in the HPT more frequently. Most customers choose to use the IBM provided default values for the HPT ratios.

Memory usage for I/O devices

In support of I/O operations, the hypervisor maintains structures that are called *translation control entities* (TCEs), which provide an information path between I/O devices and partitions. The TCEs provide the address of the I/O buffer, indications of read versus write requests, and other I/O-related attributes. Many TCEs are used per I/O device, so multiple requests can be active simultaneously to the same physical device. To provide better affinity, the TCEs are spread across multiple processor chips or drawers to improve performance while accessing the TCEs.

For physical I/O devices, the base amount of space for the TCEs is defined by the hypervisor and is based on the number of I/O devices that are supported. A system that supports high-speed adapters can also be configured to allocate more memory to improve I/O performance. Linux is the only operating system that uses these extra TCEs so that the memory can be freed for use by partitions if the system uses only AIX.

Memory usage for virtualization features

Virtualization requires that the Power Hypervisor allocate more memory for hardware statesave areas and various virtualization technologies. For example, on Power11 processor-based systems, each processor core supports a maximum of eight simultaneous multithreading (SMT) threads, and each thread contains more than 80 different registers.

The Power Hypervisor must set aside save areas for the register contents of the maximum number of configured virtual processors. Because of the greater the number of physical hardware devices, then the greater the number of virtual devices, the greater the amount of virtualization, and the more hypervisor memory is required. For efficient memory usage, the preferred and maximum values for various attributes (processors, memory, and virtual adapters) must be based on business needs and not set higher than actual requirements.

Predicting memory that is used by the Power Hypervisor

You can use the IBM System Planning Tool (SPT) to estimate the amount of hypervisor memory that is required for a specific server configuration. After the SPT run file is downloaded and installed, you can define a configuration by selecting the correct hardware platform and the installed processors and memory, and defining partitions and partition attributes. SPT can estimate the amount of memory that is assigned to the hypervisor, which assists you when you change a configuration or deploy new servers.

The Power Hypervisor provides the following types of virtual I/O adapters:

- Virtual SCSI

  The Power Hypervisor provides a virtual SCSI mechanism for the virtualization of storage devices. The storage virtualization is accomplished by using two paired adapters: a virtual SCSI server adapter and a virtual SCSI customer adapter.

- Virtual Ethernet

  The Power Hypervisor provides a virtual Ethernet switch function that allows partitions fast and secure communication on the same server. The partitions can communicate without a need for physical interconnection or connectivity outside of the server if a Layer 2 bridge to a physical Ethernet adapter is set in one VIOS partition, also known as *Shared Ethernet adapter* (SEA).

- Virtual Fibre Channel

  A virtual Fibre Channel adapter provides customer LPARs with a Fibre Channel connection to a storage area network through the VIOS partition. The VIOS partition provides the connection between the virtual Fibre Channel adapters on the VIOS partition and the physical Fibre Channel adapters on the managed system.

- Virtual (TTY) console

  Each partition must have access to a system console. Tasks, such as operating system installation, network setup, and various problem analysis activities, require a dedicated system console. The Power Hypervisor provides the virtual console by using a virtual TTY or serial adapter and a set of hypervisor calls to operate on them. Virtual TTY does not require the purchase of any other features or software, such as the PowerVM Edition features.

Logical partitions

LPARs and virtualization increase the usage of system resources and add a level of configuration possibilities.

Logical partitioning is the ability to make a server run as though it were two or more independent servers. When you logically partition a server, you divide the resources on the

server into subsets, called *LPARs*. You can install software on an LPAR, and the LPAR runs as an independent logical server with the resources that you allocated to the LPAR.

In some documentation, LPAR is also referred to as a *virtual machine* (VM), which makes it look similar to what other hypervisors offer. However, compared to other hypervisors, LPARs provide a higher level of security and isolation and other features that are described in this chapter.

Processors, memory, and I/O devices can be assigned to LPARs. AIX, IBM i, Linux, and VIOS can run on LPARs. VIOS provides virtual I/O resources to other LPARs with general-purpose operating systems.

LPARs share a few system attributes, such as the system serial number, system model, and processor Feature Codes. All other system attributes can vary from one LPAR to another.

Micro-Partitioning

When you use the 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 on a set of processors that is called a *shared processor pool* (SPP). When virtual processors are used, the operating system manages the fractions of processing power that are assigned to the LPAR.

From an operating system perspective, a virtual processor cannot be distinguished from a physical processor, unless the operating system is enhanced to determine the difference. Physical processors are abstracted into virtual processors that are available to partitions.

On the Power11 processor-based server, a partition can be defined with a processor capacity as small as 0.05processing 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 shared processor partitions are dispatched and time-sliced on the physical processors under the control of the Power Hypervisor. The shared processor partitions are created and managed by the HMC.

The Power E1150 supports up to 256 cores in a single system and 1000 micropartitions, which is the maximum that PowerVM supports.

Note: Although the Power E1150 supports up to 1000 micropartitions, the real limit depends on application workload demands in use on the server.

Processing mode

When you create an LPAR, you can assign entire processors for dedicated use, or you can assign partial processing units from an SPP. This setting defines the processing mode of the LPAR.

Dedicated mode

In dedicated mode, physical processors are assigned as a whole to partitions. The SMT feature in the Power11 processor core allows 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 become processor donors for idle processors that they own, which allows for the donation of spare CPU cycles from dedicated processor partitions to an SPP. The dedicated partition maintains

absolute priority for dedicated CPU cycles. Enabling this feature can help increase system usage without compromising the computing power for critical workloads in a dedicated processor 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 (the maximum number is 20 times the number of processing units that are assigned to the partition). The Power Hypervisor dispatches virtual processors to physical processors according to the partition's processing units entitlement. One processing unit represents one physical processor's processing capacity. All partitions receive a total CPU time equal to their processing unit's entitlement. The logical processors are defined on top of virtual processors. Therefore, even with a virtual processor, the concept of a logical processor exists, and the number of logical processors depends on whether SMT is turned on or off.

9.1.2 Multiple shared processor pools

MSPPs are supported on Power11 processor-based servers. This capability allows a system administrator to create a set of micropartitions with the purpose of controlling the processor capacity that can be used from the physical SPP.

Micropartitions are created and then identified as members of the default processor pool or a user-defined SPP. The Power Hypervisor monitors the virtual processors that exist within the set of micropartitions. Processor capacity is managed according to user-defined attributes.

If the 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 capacity 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 to the set of micropartitions within that SPP.

All Power servers that support the MSPP capability have a minimum of one (the default) SPP and up to a maximum of 64 SPPs.

This capability helps customers reduce TCO significantly when the cost of software or database licenses depends on the number of assigned processor-cores.

Note: SPPs can be used with Resource Groups. Such configuration can improve performance and isolation significantly.

9.1.3 Virtual I/O Server

The Virtual I/O Server (VIOS) is part of PowerVM. It is the specific appliance that allows the sharing of physical resources among LPARs to allow more efficient usage (for example, consolidation). When using device virtualization through the VIOS, the VIOS owns the physical I/O resources, which include SCSI, Fibre Channel, network adapters, or optical devices. The VIOS allows customer partitions to share access to the physical resources, which minimizes and optimizes the number of physical adapters in the system.

The VIOS eliminates the requirement that every partition owns a dedicated network adapter, disk adapter, and disk drive. The VIOS supports OpenSSH for secure remote logins. It also provides a firewall for limiting access by ports, network services, and IP addresses.

Figure 9-1 shows an overview of a VIOS configuration.

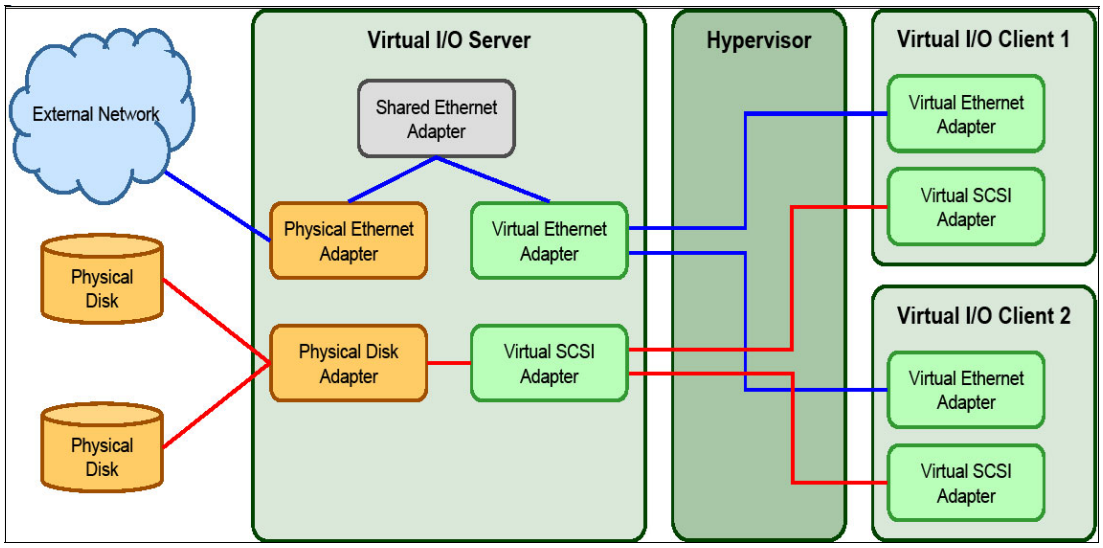

Figure 9-1 Architectural view of the VIOS

It is a best practice to run dual VIO servers per physical server.

Shared Ethernet Adapter

Use a Shared Ethernet adapter (SEA) to connect a physical Ethernet network to a virtual Ethernet network. The SEA provides this access by connecting the Power Hypervisor VLANs to the VLANs on the external switches. Because the SEA processes packets at Layer 2, the original MAC address and VLAN tags of the packet are visible to other systems on the physical network. IEEE 802.1 VLAN tagging is supported.

By using the SEA, several customer partitions 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 virtual Ethernet networks (internal) and one or more (Etherchannel) physical network adapters (external). These virtual Ethernet network adapters are defined by the Power Hypervisor on the VIOS.

Virtual SCSI

Virtual SCSI is used to view a virtualized implementation of the SCSI protocol. Virtual SCSI is based on a customer/server relationship. The VIOS LPAR owns the physical I/O resources and acts as a server or, in SCSI terms, a target device. The client LPARs access the virtual SCSI backing storage devices that are provided by the VIOS as clients.

Use the HMC to configure the virtual I/O adapters, a virtual SCSI server adapter and a virtual SCSI client adapter. 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 allows a client partition to access physical SCSI and SAN-attached devices and LUNs that are mapped to be used by the client partitions. The provisioning of virtual disk resources is provided by the VIOS.

Internet Small Computer Systems Interface

The internet small computer systems interface (iSCSI) disk is supported in the Virtual I/O Server (VIOS) 3.1.0, or later

The iSCSI disk provides block-level access to storage devices by carrying SCSI commands over an IP network. The iSCSI disk is used to facilitate data transfers over the network by using TCP, a reliable transport mechanism that uses either IPV6 or IPV4 protocols. The iSCSI disk is used to manage storage over long distances.

The iSCSI support in VIOS allows iSCSI disks to be exported to client logical partitions as virtual disks (vSCSI disks). This support is available in VIOS version 3.1 and later.

VIOS version 3.1 enables multipath I/O (MPIO) support for the iSCSI initiator. With MPIO support, you can configure and create multiple paths to an iSCSI disk, similar to other protocols. The client logical partition can run either an AIX or Linux operating system.

VIOS version 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, you can create multiple iSCSI software initiator devices on a single AIX operating system instance.

N_Port ID Virtualization

N_Port ID Virtualization (NPIV) is a technology that allows multiple LPARs to access one or more external physical storage devices through the same physical Fibre Channel adapter. This adapter is attached to a VIOS partition that acts only as a pass-through that manages the data transfer through the Power Hypervisor.

Each partition features one or more virtual Fibre Channel adapters, each with their own pair of unique worldwide port names. With this configuration, you can connect each partition to independent physical storage on a SAN. Unlike virtual SCSI, only the client partitions see the disk.

For more information and requirements for NPIV, see *IBM PowerVM Virtualization Managing and Monitoring,* SG24-7590.

9.1.4 Live Partition Mobility

Use Live Partition Mobility (LPM) to move a running LPAR from one system to another without disruption. Inactive partition mobility allows you to move a powered-off LPAR from one system to another one.

LPM provides systems management flexibility and improves system availability by avoiding the following situations:

- Planned outages for hardware upgrade or firmware maintenance.
- Unplanned downtime. With preventive failure management, if a server indicates a potential failure, you can move its LPARs to another server before the failure occurs.

For more information and requirements for LPM, see *IBM PowerVM Live Partition Mobility*, SG24-7460.

HMCV10R1 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 (#EMAM) under the PowerVM umbrella that enables memory expansion on the system.

AME is an innovative technology that supports the AIX operating system. 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 100% or more. This expansion can allow a partition to complete more work or support more users with the same physical amount of memory. Similarly, it can allow 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 beneficial, but the degree of expansion depends on how compressible the memory content is. It also depends on having adequate spare CPU capacity available for compression and decompression.

The Power E1080 includes a hardware accelerator that is designed to boost AME efficiency and use fewer processor core resources. In each AIX partition, you can enable or disable AME. Control parameters set the amount of expansion in each partition to help control the amount of CPU used by the AME function.

You must IPL the specific partition where you enable memory expansion. When enabled, monitoring capabilities are available in standard AIX performance tools, such as lparstat, vmstat, topas, and svmon.

A planning tool is included with AIX, so you can sample workloads and estimate how expandable the partition's memory is and how much CPU resource is 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 and is permanent. It cannot be moved to a different server.

9.1.6 Remote Restart function

Remote Restart provides a high availability option for partitions. If an error causes a server outage, you can restart a partition configured for Remote Restart on a different physical server. After an outage, sometimes starting the server takes longer. In such cases, you can use the Remote Restart function to reprovision the partition more quickly to a different server. You can perform Remote Restart faster than restarting the failed server and then restarting its partitions.

The Remote Restart function relies on technology similar to Live Partition Mobility (LPM), where a partition is configured with storage on a shared SAN that is accessible by the server hosting the partition.

HMC Version 10 Release 1 introduces 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, the Power processor modes affect various virtualization features.

On Power servers, partitions can be configured to run in several modes, including the following modes:

- Power9: This native mode for Power9 processors implements Version 3.0 of the IBM Power ISA.
- Power10: This native mode for Power10 processors implements Version 3.1 of the IBM Power ISA.
- Power11: This native mode for Power11 processors implements Version 3.1 of the IBM Power ISA.

For more information, see Processor compatibility mode definitions.

Figure 9-2 shows the available processor modes on a Power E1180.

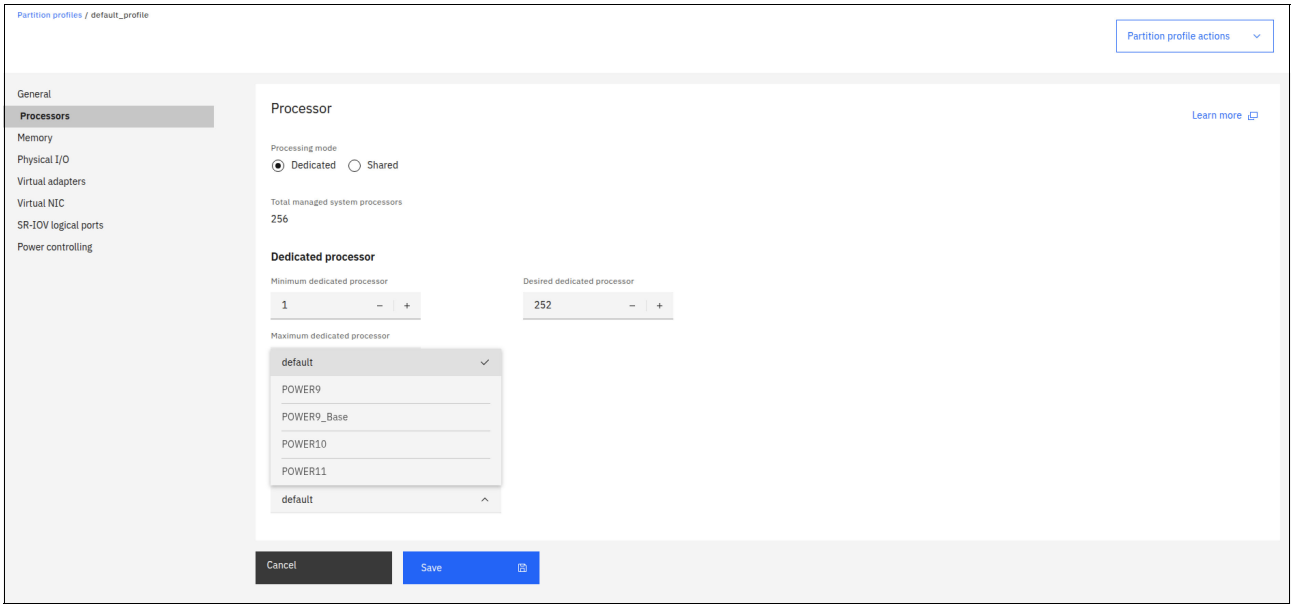

Figure 9-2 Processor modes

Processor compatibility mode is important when LPM migration is planned between different generations of servers. An LPAR that might be migrated to a machine that is managed by a processor from another generation must be activated in a specific compatibility mode.

Single root I/O virtualization

Single root I/O virtualization (SR-IOV) is an extension to the PCIe specification with which multiple operating systems can simultaneously share a PCIe adapter 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 much like logical partitioning enables server consolidation. With an adapter capable of SR-IOV, you can assign virtual slices of a single physical adapter to multiple partitions through logical ports without using a VIOS.

9.1.8 More information about virtualization features

The following IBM Redbooks publications provide more information about the virtualization features:

- *IBM PowerVM Best Practices*, SG24-8062
- *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- *IBM PowerVM Virtualization Managing and Monitoring,* SG24-7590
- *IBM Power Systems SR-IOV: Technical Overview and Introduction*, REDP-5065

9.1.9 Resource groups

Resource groups are a feature introduced with IBM Power11 that is designed to enhance system performance by up to 25% through improved workload optimization and resource affinity. Although shared processor pools (SPPs) already provide compute capacity isolation by capping the maximum resources available to each pool, resource groups take this further by introducing advanced affinity-based optimizations for more efficient workload dispatching.

Early performance evaluations indicate that, when configured effectively, resource groups can deliver shared processor performance that closely matches that of dedicated processors, particularly in large, partitioned environments.

Resource groups properties

With the introduction of Resource Groups in IBM Power11, system administrators gain a powerful new tool for organizing and optimizing compute resources. Designed for flexibility and performance, Resource Groups allow for the grouping of both dedicated and shared processor partitions, enabling more granular control over resource allocation and workload management. This feature enhances system efficiency, supports dynamic reconfiguration, and integrates with existing technologies like SPPs, live partition mobility, and PEP 2.0.

Resource groups include the following core features:

- Performance optimization: Resource groups improve workload affinity, enabling shared processor performance nearly equivalent to dedicated processors.
- Flexible configuration:
  - Can include both dedicated processor partitions and shared processor partitions.
  - Resource group configuration specifies the number of general-purpose cores (AIX/IBM i/Linux/VIOS) and IFL cores (Linux/VIOS).
  - Powered-off partitions can be reassigned between groups.
  - Resources not assigned to a user-defined group are placed in the default resource group.
  - Cores can be dynamically reallocated among groups.
- Affinity and isolation:
  - Each resource group has its own set of shared processor pools.
  - Enhances compute capacity isolation and affinity-based dispatching.
- Monitoring and management:
  - Usage metrics available for resource groups, SPPs, and the overall system.
  - Dynamic Platform Optimizer can run at system or group level.
- Mobility and compatibility:
  - When using Live Partition Mobility, the resource group on the target system can be selected.
  - 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 includes several example use cases, but use cases can also be combined to maximize functionality and tailor performance to specific business needs.

Consolidation across multiple lines of business

Figure 9-3 illustrates how resources are allocated and shared across multiple lines of business on Power10 systems. With the introduction of resource groups in Power11, resource allocation becomes more efficient, enabling better workload isolation and improved performance across the system.

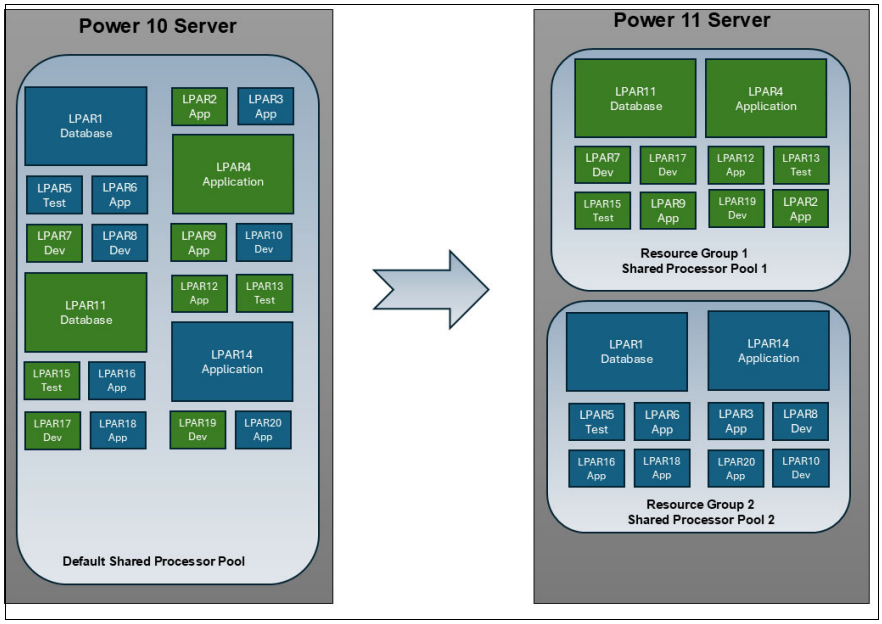

Figure 9-3 Example where resource groups consolidate multiple lines of business

Isolation of production workloads from test and dev workloads

In this example, the development and test LPARs are restricted to using resources only from the default resource group, limiting their access to shared system resources and isolating them from production workloads.

Figure 9-4 on page 187 shows how resource groups are used to provide this isolation of development and test workloads.

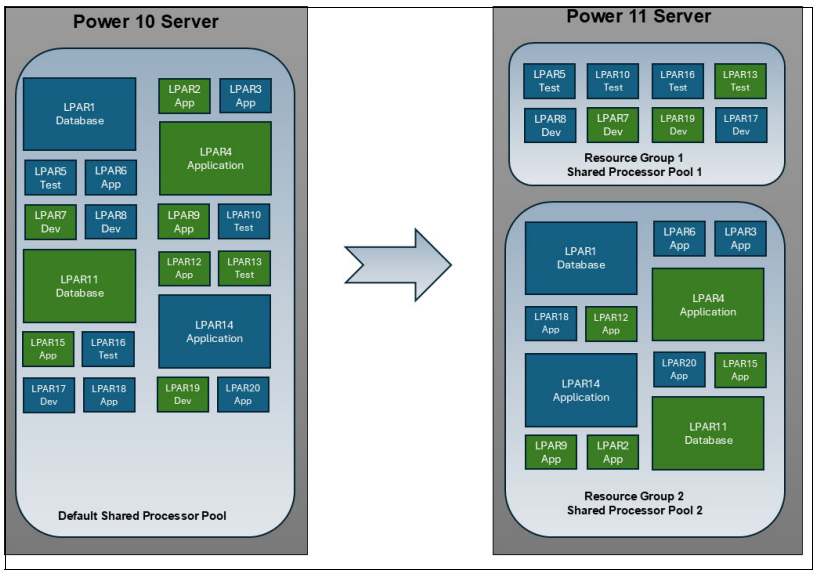

Figure 9-4 Example where resource groups isolate workloads

Improve application performance by grouping workload tiers into resource groups

In this example, performance is enhanced through improved affinity between the database and application server, allowing for more efficient resource utilization and faster communication within the same resource group.Figure 9-5 shows how grouping workloads into different resource groups can improve application affinity and improve performance.

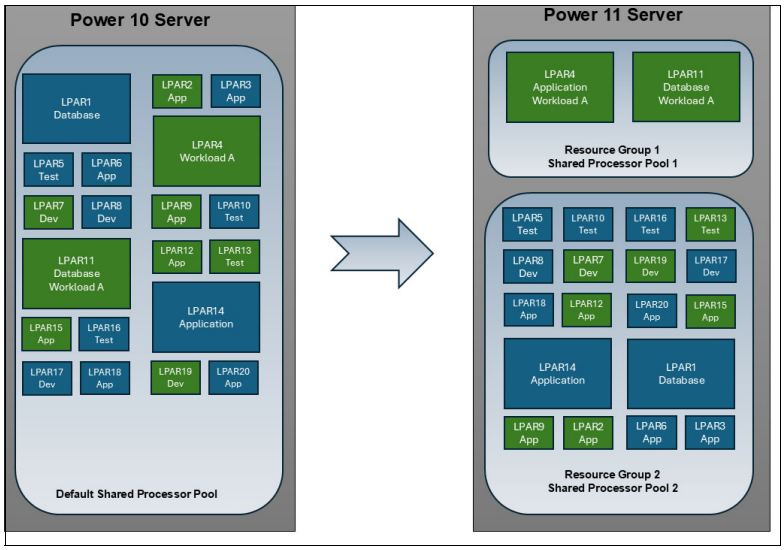

Figure 9-5 Example where resource groups improve performance by grouping workload tiers

System-level isolation in multi-server consolidation scenarios

Resource groups enable effective workload isolation on IBM Power11 systems. This allows organizations to consolidate multiple smaller servers into a single, more powerful Power11 system while logically separating workloads into distinct groups. Each group can be managed independently, ensuring performance, security, and resource control across different business functions or environments. Figure 9-6 on page 188 shows how using resource groups can help maintain system level isolation while doing server consolidation.

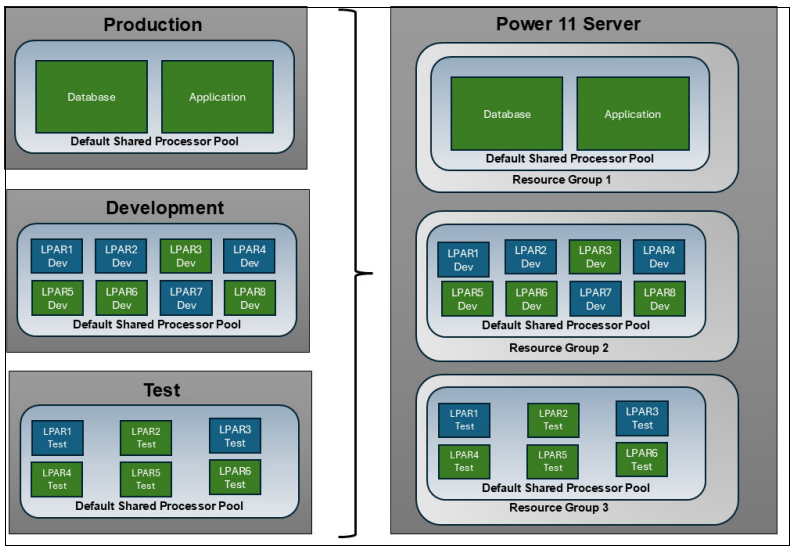

Figure 9-6 Example where resource groups provide an additional layer of workload isolation

Improved performance by mapping Shared Processor Pools into resource groups

Shared Processor Pools (SPPs) and resource groups can be used together on IBM Power11 systems to simultaneously optimize application performance and reduce total cost of ownership (TCO). This combination allows for precise control over resource allocation and workload isolation, ensuring efficient usage of compute capacity while maintaining performance consistency across diverse workloads.

Figure 9-7 shows how SPPs and resource groups enhance your performance and reduces your TCO.

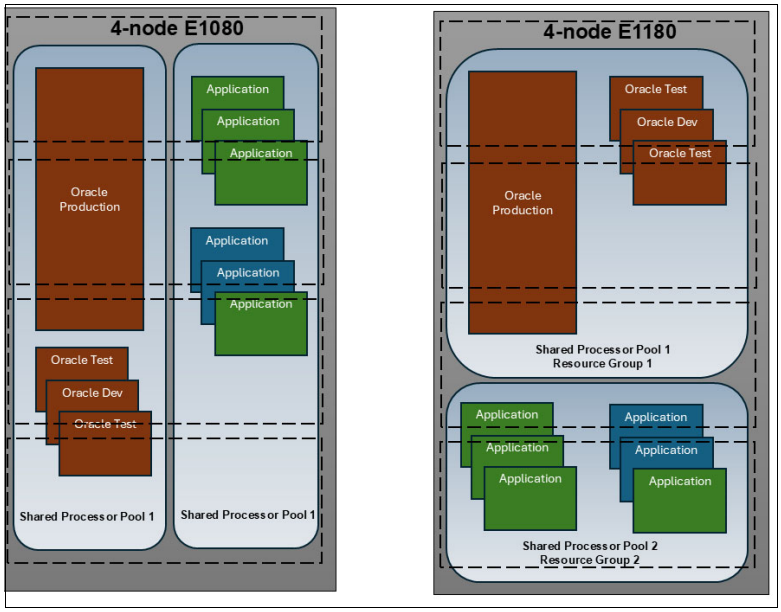

Figure 9-7 Example where resource groups improve performance by mapping Shared Processor Pools into resource groups.

Resource groups advisor

Resource groups advisor (RGA) is a web-based tool designed to assist with the configuration of resource groups on IBM Power. Offered as a free service, RGA analyzes a customer's server setup and provides tailored recommendations for optimal resource group configurations. It helps model and validate configurations to ensure efficient resource allocation and system performance.

For more information, see IBM Resource Groups Advisor for IBM Power.

The welcome screen for RGA is shown in Figure 9-8.

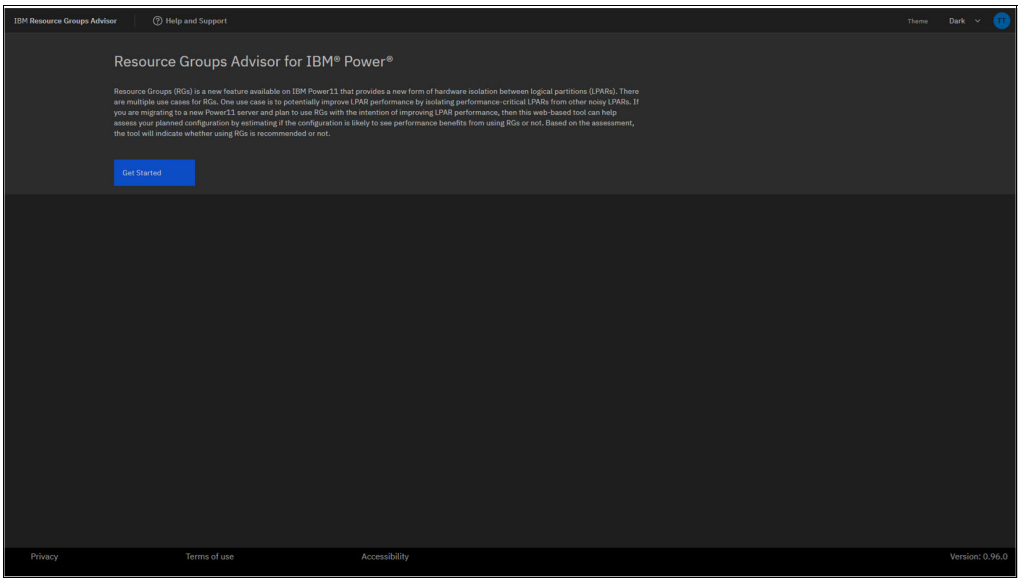

Figure 9-8 Welcome screen

Use the following instructions to run RGA.

1. Select **Get Started** to open a page where you can choose to create a new configuration or upload a previously saved version
2. Select **New Configuration**.
3. In the New Configuration screen (Figure 9-9 on page 190), provide the system configuration parameters.
4. Press **Next.**
5. For resource group configuration (Figure 9-10 on page 190), you can accept the default or you can add additional resource groups to your configuration.

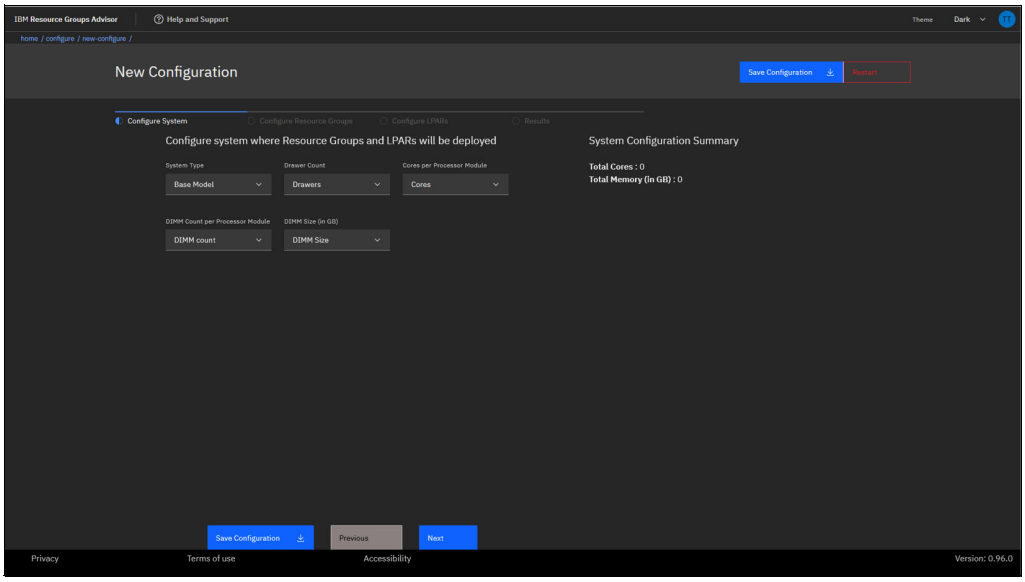

Figure 9-9 New Configuration screen

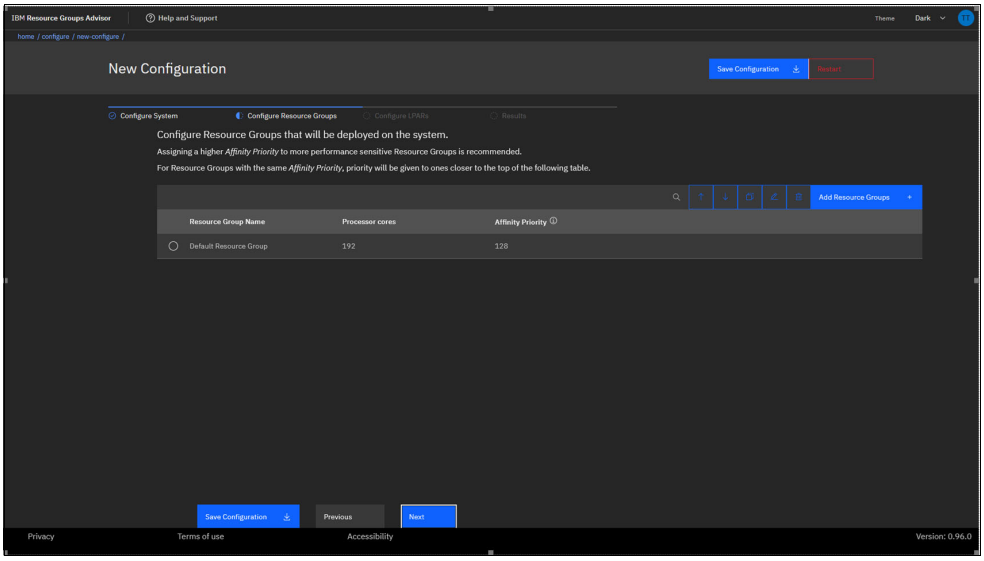

Figure 9-10 Resource group definitions

6. After defining your resource groups, click **Next** and define the LPAR details.
7. Select **Add LPAR** to enter the details about each LPAR, such as whether it is running in Shared or Dedicated mode, the processor and memory allocations, and the assigned resource group for the LPAR. You can also use the IBM System Planning Tool to add LPARs.
8. Figure 9-11 on page 191 shows an LPAR summary screen.

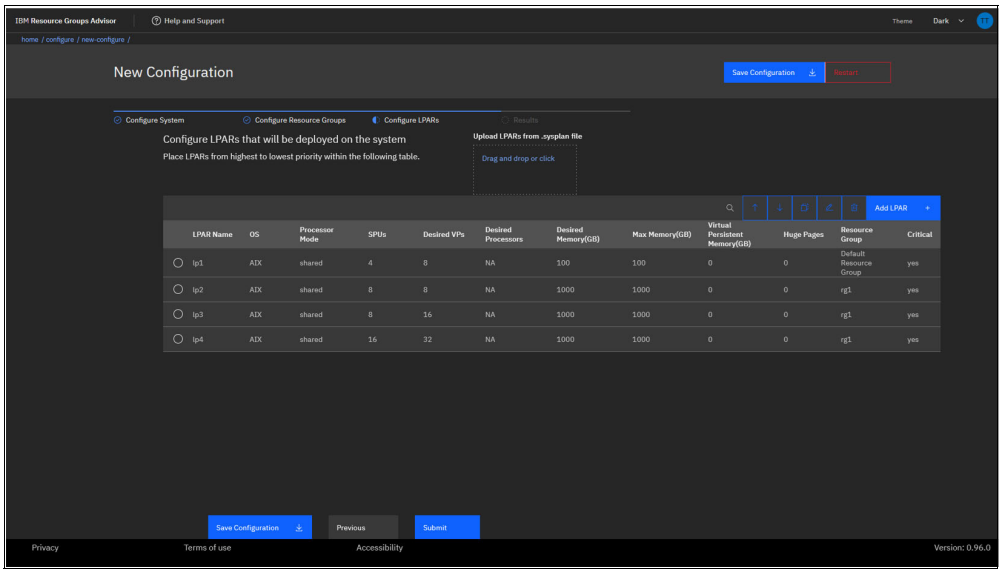

Figure 9-11 LPAR configuration

9. After entering all of the LPAR information, press **Submit** to get the model's output similar to what is shown in Figure 9-12.

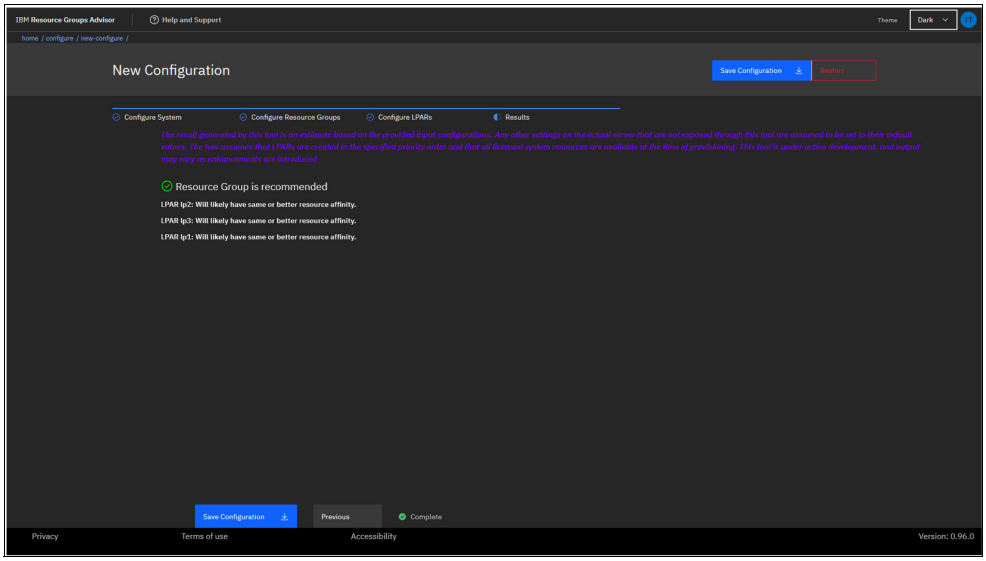

Figure 9-12 Output from the tool

10. You can save the configuration or start a new model.

Note: RGA simulates how PowerVM creates resource groups and allocates resources to LPARs on a system, all based on the provided input configurations. No customer data is pulled from actual servers.

9.2 KVM support

IBM Power11 processor based servers can use Kernel-based Virtual Machine (KVM) within a PowerVM logical partition (LPAR). This allows for the creation and management of lightweight Linux virtual machines (VMs) using standard KVM tools, while still using the existing resources of the PowerVM LPAR. Essentially, KVM becomes an additional virtualization option alongside PowerVM on IBM Power.

KVM on IBM Power11 processor-based servers is not a replacement for PowerVM but is an additional capability that brings the power, speed, and flexibility of the KVM virtualization technology to a PowerVM LPAR.

KVM-enabled LPARs can host PPC64-LE KVM guests, which are essentially Linux VMs. The KVM guests within the LPAR use resources (CPU, memory, I/O) that have been allocated to the LPAR by the Power Hypervisor. This approach offers flexibility in deploying Linux workloads and can be more cost-effective than other virtualization solutions, especially for organizations already invested in the Linux ecosystem. The integration of KVM with the Linux kernel can lead to high performance, especially 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. You can also consolidate different kind of workload in one or more IBM Power11 processor based servers s (see Figure 9-13 on page 193 for an example). In essence, KVM on Power11 provides a way to make use of the strengths of both PowerVM and KVM, offering a powerful and flexible virtualization environment for Linux workloads.

KVM in a PowerVM LPAR uses the industry-standard Linux KVM virtualization stack (Figure 9-14 on page 193) and can integrate within an existing Linux virtualization configuration. KVM can run in an LPAR because of the following features:

- Because of IBM Power architecture and Power11 advanced virtualization capabilities, multiple operating system (OS) instances can share hardware resources while providing isolation. The Radix MMU architecture provides the capability to 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. These changes extend the Power platform architecture to include new hypervisor interfaces.
- The Linux kernel includes the KVM kernel module (KVM), which provides core virtualization infrastructure to run multiple virtual machines in a Linux host LPAR. Upstream kernels and enabled downstream distributions such as Fedora and Ubuntu use the Power architecture extensions to create and manage KVM guests in the PowerLinux LPAR.
- QEMU: User space component that implements virtual machines on the host that use KVM functions.
- LibVirt provides a toolkit for virtual machine management.

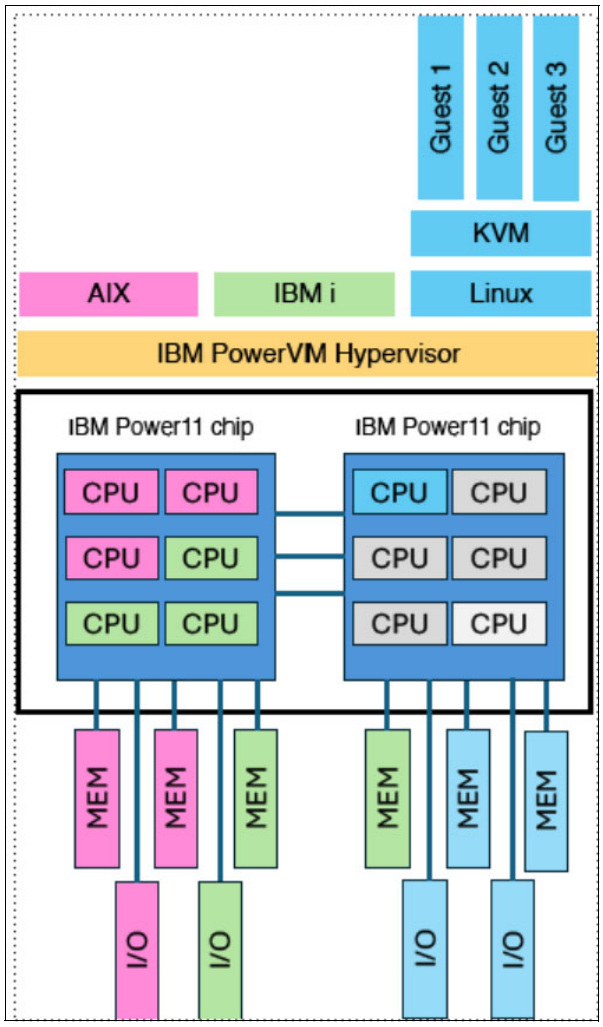

Figure 9-13 IBM Power stack for KVM

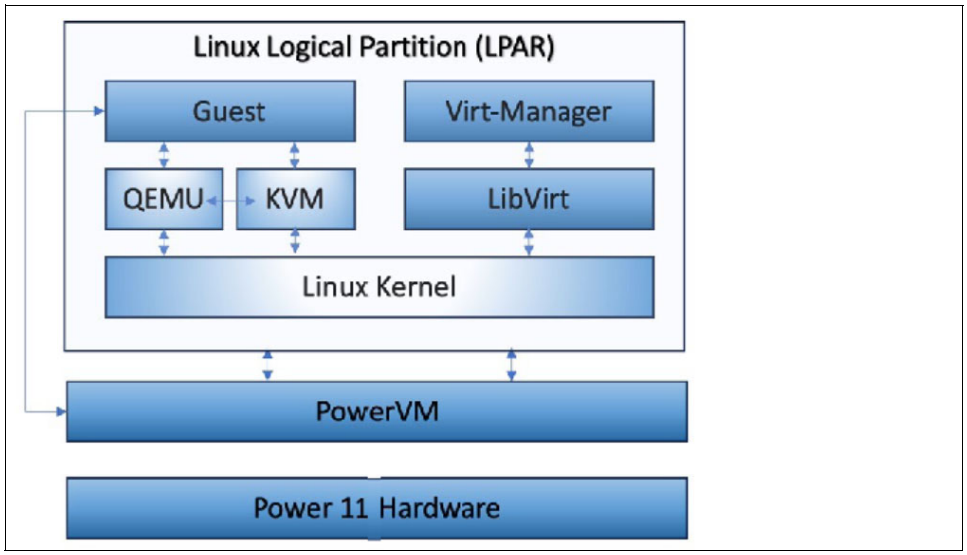

Figure 9-14 Industry Standard Linux Virtualization Stack

KVM support on IBM Power11 processor based server has the following requirements:

- Partition must be a Linux partition that runs in Power10 processor compatibility mode.
- 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 BMC.
- Partition must be running in Radix mode (default MMU mode for Linux LPARs).
- Partition must be assigned with dedicated CPUs with processor sharing set to Never Allow.

The following features are not supported on KVM logical partitions:

- Shared processors
- vPMEM LUNs
- Platform keystore
- Live partition migration (LPM)
- Dynamic platform optimization
- Add or remove memory, processor, and I/O DLPAR

For guests, KVM does not support PCI pass-through of LPAR-attached PCI devices to KVM guests. The plan is to support the feature in future releases.

For more information, see KVM in a PowerVM LPAR.

10

Hybrid cloud solutions

This chapter introduces IBM Power servers as a foundational platform for hybrid cloud environments. These servers support mission-critical workloads with exceptional performance, security, and flexibility.

IBM Power systems, built on the advanced Power10 architecture, deliver significant improvements in compute efficiency and operational cost savings. These systems integrate seamlessly across on-premises, private, and public cloud infrastructures, and support a wide range of enterprise operating systems, including AIX, IBM i, and Linux.

This chapter explains how Power servers enable scalable, cloud-native application development using Kubernetes and Red Hat OpenShift. It also describes key management tools such as IBM PowerVC, IBM Cloud Management Console, and Power Enterprise Pools 2.0, which simplify virtualization, optimize resource allocation, and enhance operational agility.

You can access IBM Power Virtual Server (PowerVS), a cloud-based extension of Power systems, through a flexible, pay-as-you-go model. PowerVS enables workload migration, rapid provisioning, and robust business continuity, allowing organizations to modernize at their own pace.

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 Power Enterprise Pools (PEP) 2.0. Because of this innovation, multiple IBM Power systems can operate as a unified resource pool, enabling shared 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.

10.1.1 Power Enterprise Pools

In this model, each system within the pool is provisioned with a set number of base processor and memory activations. When configured as a Power Enterprise Pool, these base activations and their associated operating system entitlements are aggregated to define the pool's baseline capacity. Simultaneously, all hardware resources across the pool are fully activated and made available for workload deployment allowing usage to exceed the base capacity when needed. IBM Cloud Management Console 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, offering granular resource sharing, improved cost efficiency, and simplified management for enterprise IT environments.

PEP with Shared Utility Capacity on Power E1150 provides enhanced multi-system resource sharing and by-the-minute tracking and consumption of compute resources across a collection of Power E1150 and E1050 systems within a single PEP 2.0. Shared Utility Capacity provides a complete range of flexibility to tailor initial system configurations with the right mix of purchased and pay-for-use consumption of processor, memory, and software. Any Power E1150 and Power E1050 systems can seamlessly interoperate and share compute resources within the same pool. Clients with existing PEP 2.0 of Power E1050 systems can add one or more Power E1150 systems into their pool and migrate to them at the rate and pace of their choosing. Clients with Mobile Capacity on a Power E1150 can upgrade their system to support PEP 2.0 to make use of Shared Utility Capacity resources.

A Power infrastructure that is consolidated onto Power E1150 systems has the potential to greatly simplify system management, so IT teams can focus on optimizing their business results instead of moving resources within their data center. Shared Utility Capacity resources are easily tracked by a virtual machine (VM) and monitored by a CMC, which integrates with local HMCs to manage the pool and track resource use by system and VM, by the minute, across a pool.

Clients need not 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 E1150 systems, Mobile and Shared Utility Capacity capabilities are now offered through the purchase of a PEP Subscription (5765-P2E) for each Power E1150 system being configured for use in a Power Enterprise Pool.

When a PEP 2.0 pool is started, each eligible Power E1150 system's purchased processor activations, memory activations, and supported operating system entitlement resources become Base Capacity resources as part of the Power Enterprise Pool. They are aggregated across a defined pool of systems for consumption monitoring. Metered Capacity is the additional installed processor and memory resource that exceeds each system's Base Capacity. It is activated and made available for immediate use when a pool is started, then monitored by the minute by a CMC.

Metered resource usage is charged only for minutes exceeding the pool's aggregate Base resources, and usage charges are debited in real-time against a client's Capacity Credits (5819-CRD) on account or can 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 Power11 to their existing pool with Power9 and Power10 processor-based systems, are recommended to reach out to their IBM representatives or IBM Expert Labs to discuss transition planning and best practices.

10.1.2 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 ensuring smooth and efficient infrastructure operations.

The IBM Cloud Management Console (CMC) for Power offers a comprehensive view of your IBM Power cloud environment, regardless of the number of systems or data centers involved. It provides several benefits:

- 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 empowers IT teams to manage their infrastructure more effectively and make informed decisions with confidence. The IBM CMC now offers full support for IBM Power11 systems, ensuring compatibility with the latest advancements in Power architecture.

The IBM CMC is also used to monitor and manage PEP 2.0 pools in your enterprise. The CMC Enterprise Pools 2.0 application helps you to monitor base and metered capacity across a Power Enterprise Pool 2.0, with summary and drill-down views of real-time and historical resource consumption by the logical partition.

CMC provides a cloud-based interface to monitor, manage, and optimize IBM Power infrastructure. It supports AIX, IBM i, and Linux VM workloads, and is accessible through a secure, web-based dashboard. The IBM Cloud Management Console is a cloud-based portal that includes the following features:

- View health and performance of IBM Power servers
- Monitor capacity and performance metrics
- Get insights into firmware levels and configurations
- Receive proactive support and alerts

For organizations with limited IT staff or those looking to streamline infrastructure oversight, CMC offers a way to perform essential monitoring and management tasks without requiring complex, locally hosted solutions.

CMC integrates with the Hardware Management Console (HMC) and provides advanced insights into performance, capacity, firmware compliance, patch planning, and system health. With the scalability and built-in virtualization of Power11, CMC helps administrators maintain control over increasingly complex hybrid IT environments, whether on-premises or connected to IBM Cloud or other supported clouds.

CMC is offered through a software-as-a-service (SaaS) model and is licensed on a per-server subscription basis. Power11 systems are registered to IBM CMC through their HMCs, which securely transmit telemetry and configuration data to the cloud portal. IBM provides both a no-cost base offering and optional chargeable add-ons that include more advanced features, such as predictive analytics, patch automation, and longer 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 services. To meet these demands, they must innovate by developing new applications, by modernizing existing ones, and by ensuring seamless scalability.

Organizations achieve this 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 empowers both developers and IT operations teams with the flexibility and speed needed to build, deploy, and manage applications across diverse environments – on-premises, in the cloud, or across multiple infrastructures. This platform accelerates digital transformation by simplifying scalability and enhancing security.

10.2.1 Red Hat OpenShift on Power

Red Hat OpenShift Container Platform is fully supported on IBM Power servers, providing a robust and scalable environment for running cloud-native applications and modernizing traditional workloads on the IBM Power architecture (ppc64le). This integration enables enterprise-grade container orchestration for building, deploying, and managing containerized applications with Red Hat OpenShift.

Red Hat OpenShift on IBM Power offers a flexible and efficient platform for hybrid IT environments. You can run cloud-native applications alongside traditional VM-based workloads on the same infrastructure, allowing colocation of containerized workloads with existing applications that are running on AIX, IBM i, or Linux. This configuration helps ensure low-latency communication between applications and data.

IBM Power servers provide a resilient and secure foundation for your cloud implementation using Red Hat OpenShift. This chapter describes the benefits of using a multi-architecture cluster and provides implementation guidelines to help you deploy a multi-architecture cluster using IBM Power control plane nodes and x86-based or AMD-based worker nodes.

Red Hat OpenShift is a leading enterprise Kubernetes platform that offers a robust foundation for developing, deploying, and scaling cloud-native applications. It extends Kubernetes with additional features and tools that enhance productivity and security, making it a strong choice for businesses adopting container technology at scale.

Red Hat OpenShift provides a unified platform to build, modernize, and deploy applications. You can work more efficiently with a complete set of services for bringing applications to market on your preferred infrastructure. Red Hat OpenShift delivers a consistent experience across public cloud, on-premises, hybrid cloud, and edge environments.

Red Hat OpenShift offers a unified, flexible platform to meet a variety of business needs, from an enterprise-ready Kubernetes orchestrator to a comprehensive cloud-native application development solution. You can choose to self-manage the platform or use it as a fully managed cloud service.

Red Hat OpenShift clusters

A Red Hat OpenShift Container Platform cluster consists of multiple nodes, which can run on either physical or virtual machines. The following node requirements apply:

- Three control plane nodes to manage and maintain the cluster
- Two worker nodes to run containerized workloads
- A temporary bootstrap node used during installation to host configuration and installation files

In OpenShift Container Platform 4.18, the bootstrap and control plane nodes must run Red Hat Enterprise Linux CoreOS (RHCOS), a minimal, immutable container host OS 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 OpenShift Container Platform 4.16 and will be removed in a future release.

Multiple architecture compute clusters

Enterprises that are expanding their operations often deploy applications in heterogeneous environments by using different types of hardware. A prime example includes data centers using x86_64 servers and Power servers, but edge locations might feature ARM-based devices because of their power efficiency. Although IBM Power is an ideal choice for running Red Hat OpenShift clusters, it is possible that some applications within the environment aren't compatible with IBM Power architecture nodes. You can retain the benefits of IBM Power and retain simpler cluster management and efficient resource usage by incorporating both IBM Power and x86_64 architecture nodes in the same cluster.

The release of Red Hat OpenShift 4.14 brought the OpenShift Container Platform multiple-architecture compute feature to IBM Power. Multiple-architecture compute provides a single heterogeneous cluster, enabling fit-for-purpose computing so clients can align tasks and applications to CPU strengths and software availability rather than one architecture. This support was expanded in Red Hat OpenShift 4.15, which enabled a Red Hat OpenShift cluster to support an IBM Power control plane and add x86 architecture worker nodes.

By using multiple-architecture compute for OpenShift Container Platform, you 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 for composite solutions that span multiple architectures.

Multi-architecture capability 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 IBM Power in corporate settings.
- Reduced complexity through standardizing application deployment across architectures
- Streamlines operations and eliminates the need to maintain separate stacks for different hardware.

- Cost efficiency

  Differing hardware architectures provide varying cost-to-performance ratios, which can help minimize overall infrastructure expenses.

- Multiple-architecture support facilitates optimal resource usage

  Organizations can select the most cost-effective architecture for each specific workload. For instance, ARM servers could be more affordable for lightweight services, whereas x86_64 or IBM Power servers might excel at handling heavy computational tasks.

- Energy efficiency

  Specific architectures, like ARM, are known for their low power consumption, which can substantially decrease energy costs, notably 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 provides benefits for companies:

- Scale elastically across platforms: Multiple-architecture compute enables businesses to dynamically allocate resources among different types of hardware to handle surges in demand without being restricted to a single architecture.
- Escape vendor lock-in: With multiple-architecture compute, corporations are not tied to a single supplier or type of hardware, thus avoiding vendor lock-in and empowering more bargaining power during procurement decisions.
- Optimize performance: Each architecture boasts unique strengths and weaknesses contingent upon the application or workload. Multiple-architecture compute maximizes performance by aligning application requirements with the architectural advantages.
- Create tailored solutions: Select applications might gain from the high I/O throughput of IBM Power, but others might perform optimally on the high-throughput, multi-core processors of x86_64 architectures.
- Meet specialized computing needs: Certain tasks might need specialized hardware, such as GPUs for machine learning work flows, 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 where all of the management and worker functions are installed on a single node. This approach is especially suitable for development, testing, or edge computing environments.

However, it is important to recognize the inherent limitations of a single-node configuration, which includes the absence of built-in high availability. If the node fails, the entire cluster becomes unavailable, potentially resulting in downtime and data loss. However, the robust high availability features of IBM Power infrastructure, when properly designed and built with redundant power, networking, and storage, can significantly mitigate these risks.

10.2.2 Red Hat OpenShift AI on IBM Power servers

Red Hat OpenShift AI brings enterprise-grade capabilities for artificial intelligence and machine learning (ML) 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 ML workloads across hybrid and multi-cloud infrastructures. By integrating AI workflows into the same infrastructure 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 (RHOAI) enables organizations to build, deploy, and manage AI-enabled applications at scale across hybrid cloud environments.

RHOAI Self-Managed is available by installing the Red Hat OpenShift AI Operator and configuring it to manage the stand-alone components of the platform. RHOAI Self-Managed is supported on Red Hat OpenShift Container Platform across multiple architectures, including ppc64le (IBM Power), s390x (IBM Z), and x86_64. This includes 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
- Google Cloud Platform
- Microsoft Azure
- VMware vSphere
- Oracle Cloud

IBM Power Servers, renowned for their high performance, scalability, and reliability, are well suited for compute-intensive AI workloads. The integration of RHOAI with IBM Power Servers enables enterprises to take advantage of both platforms' capabilities for robust, production-grade AI deployments.

As of RHOAI version 2.20 (Self-Managed), the IBM Power architecture (ppc64le) is available as a Technology Preview. Currently, model deployments on IBM Power are supported only in standard mode. RHOAI provides an integrated platform for developing, training, serving, and monitoring AI and ML models. I Support for vLLM is also available on the IBM Power architecture as a Technology Preview, with vLLM runtime templates accessible for experimentation and development.

When deployed on IBM Power, Red Hat OpenShift AI benefits from the platform’s high-performance architecture, which is optimized for data-intensive and compute-heavy workloads. The latest IBM Power processors feature built-in AI inferencing capabilities, delivering faster insights and reduced latency for real-time applications. This will be further enhanced with the planned availability and support of the IBM Spyre Accelerator in IBM Power servers. The Spyre adapter is a PCIe-attached AI card that is engineered 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 IBM Power 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 use AI with confidence, ensuring performance, efficiency, and flexibility across diverse workloads and environments.

10.2.3 IBM Cloud Paks

IBM Cloud Paks are a suite of modular, containerized software solutions 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 Cloud Pak 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 Paks take a bundled approach so that you can accelerate your modernization journey by packaging everything you need to get started.

There are three main benefits of IBM Cloud Paks:

- They are comprehensive and can be easy to use.
- They are supported by IBM.
- They run anywhere Red Hat OpenShift runs.

IBM Cloud Paks on IBM Power

Optimized for deployment on IBM Power, IBM Cloud Paks 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.

IBM Power enhances the value of Cloud Paks by offering superior performance, scalability, and cost-efficiency for modern workloads. Combined with Red Hat's open-source ecosystem, businesses can both modernize legacy applications and build new cloud-native solutions on a unified, enterprise-ready platform.

There are several IBM Cloud Paks that are currently supported on IBM Power, each focused on a specific domain:

- Cloud Pak for Applications
- Cloud Pak for Data
- Cloud Pak for Integration
- Cloud Pak for AIops

IBM Cloud Pak for Applications

IBM Cloud Pak for Applications is an enterprise-ready, containerized software solution designed to modernize existing applications and develop new cloud-native apps. Built on IBM WebSphere® offerings and Red Hat OpenShift Container Platform, it provides a comprehensive set of tools to help organizations transition between public, private, and hybrid clouds.

IBM Cloud Pak for Applications includes IBM Cloud Transformation Advisor, an AI-powered tool that assists in refactoring and rearchitecting legacy applications. The solution includes automated vulnerability assessment and identification, helping ensure continuous security compliance across all deployment environments. It also automates audit reporting, simplifying compliance management. Developers can use their preferred IDEs to build and deploy applications, with support for modern runtimes and DevOps workflows. This integration streamlines the development process and enhances productivity.

IBM Cloud Pak for Data

You can use IBM Cloud Pak for Data to unify and simplify data collection, organization, and analysis. It is ideal for AI and analytics workloads. IBM Cloud Pak for Data is a unified, preintegrated data and AI platform designed to help 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, Amazon Web Services (AWS), and Microsoft Azure.

The platform allows secure access to data at its source, eliminating the need for data migration and reducing 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 including AI throughout the business with trust and transparency, the platform supports the end-to-end AI workflow, ensuring 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. Also, its scalable Kubernetes environment allows 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 for AI-driven digital transformation.

IBM Cloud Pak for Integration

IBM Cloud Pak for Integration is a comprehensive, AI-powered hybrid integration platform that is designed to connect 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 integrations. Running on Red Hat OpenShift, IBM Cloud Pak for Integration supports both cloud and on-premises deployments, ensuring 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 the integration process, reduces manual workflows, and enhances responsiveness to real-time events. This makes it an ideal solution for organizations looking to modernize their 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 designed to automate work and accelerate business growth. Built for any hybrid cloud, it helps simplify complex workflows, facilitates records management, and enhances overall productivity. The platform uses AI to identify gaps and build low-code and no-code automations, making it easier to streamline 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 automating case and process workflows, converting unstructured content into valuable data, and using software robots to complete tasks based on AI insights. This comprehensive automation solution helps organizations 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 designed to enhance IT operations (ITOps) by using artificial intelligence and machine learning. It integrates 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 helps IT teams 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, which are all aimed at reducing mean time to resolution (MTTR) and improving overall operational efficiency. With its collaborative tools and real-time insights, IBM Cloud Pak for AIOps help IT teams innovate faster, reduce operational costs, and ensure the reliability of mission-critical workloads.

For more information, see IBM Cloud Paks documentation.

10.2.4 Other cloud enablement solutions

From a Red Hat software perspective, there is also a comprehensive set of software solutions to accelerate your modernization efforts, including Red Hat Runtimes, Red Hat 3scale API Management, Red Hat Fuse, and Red Hat AMQ.

10.3 IBM Power Virtual Server

IBM Power Virtual Server (PowerVS) is a cloud-based infrastructure-as-a-service (IaaS) offering that enables businesses to run IBM Power workloads in a flexible, scalable, and secure virtual environment. Because it is built on the same architecture as IBM Power used on-premises, IBM PowerVS supports seamless hybrid cloud integration. This makes it suitable for enterprises that want to modernize their IT infrastructure without abandoning existing investments.

With support for AIX, IBM i, and Linux operating systems, IBM PowerVS provides a versatile platform for mission-critical applications. IBM PowerVS offers a reliable solution for moving AIX, IBM i, and Linux on Power workloads to the cloud without requiring a time-consuming or risky migration to a different platform.

One of the key advantages of IBM PowerVS is its ability to deliver high performance and reliability for enterprise workloads. IBM PowerVS uses IBM Power processors, which are optimized for data-intensive tasks and high-throughput computing. The platform includes features such as dynamic resource allocation, automated scaling, and integrated backup and disaster recovery options. These capabilities help organizations maintain business continuity and meet demanding service-level agreements (SLAs).

IBM PowerVS also integrates with a wide range of IBM Cloud services and third-party tools, so you can build and manage hybrid cloud environments efficiently. It supports DevOps practices, containerization with Red Hat OpenShift, and AI and ML 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, IBM PowerVS offers a cost-effective and agile solution for businesses aiming to innovate and grow in a digital-first world.

You can run IBM i, AIX, and Linux workloads in a cloud environment using IBM Power Virtual Server. This platform provides fast, self-service provisioning and flexible compute, memory, and storage resources. IBM PowerVS is designed to help organizations modernize their infrastructure, migrate to hybrid cloud models, and optimize data center resources.

Key features and benefits of PowerVS include:

- Flexible and scalable: Users can adjust compute, memory, and storage resources on demand, scaling up or down as needed.
- Pay-as-you-use: Billing is based on consumption, allowing organizations to manage cost effectively.

- Hybrid cloud integration: PowerVS facilitates integration of AIX and IBM i workloads with cloud-native applications and services.
- Performance and security: PowerVS uses the performance and security features of IBM Power while offering the agility of cloud computing.
- Managed infrastructure: PowerVS offers managed infrastructure, where IBM handles the hardware maintenance, allowing organizations to focus on their workloads from a software and application standpoint.
- Disaster recovery: You can use PowerVS to implement disaster recovery solutions for critical applications.
- Simplified cloud adoption: PowerVS helps streamline the cloud adoption process for IBM i, AIX, and other Power systems workloads.

10.3.1 IBM PowerVS options

PowerVS is a single offering that can be delivered in two variations: off-premises, where the infrastructure components are located in IBM data centers, and on-premises, where the infrastructure components are located in the client's data center. The on-premises variation is known as IBM PowerVS Private Cloud. Both variations provide a cloud-based consumption model, where you pay for resources as they are consumed, and are managed with the same management interfaces.

Figure 10-1 shows the two different implementations of IBM PowerVS.

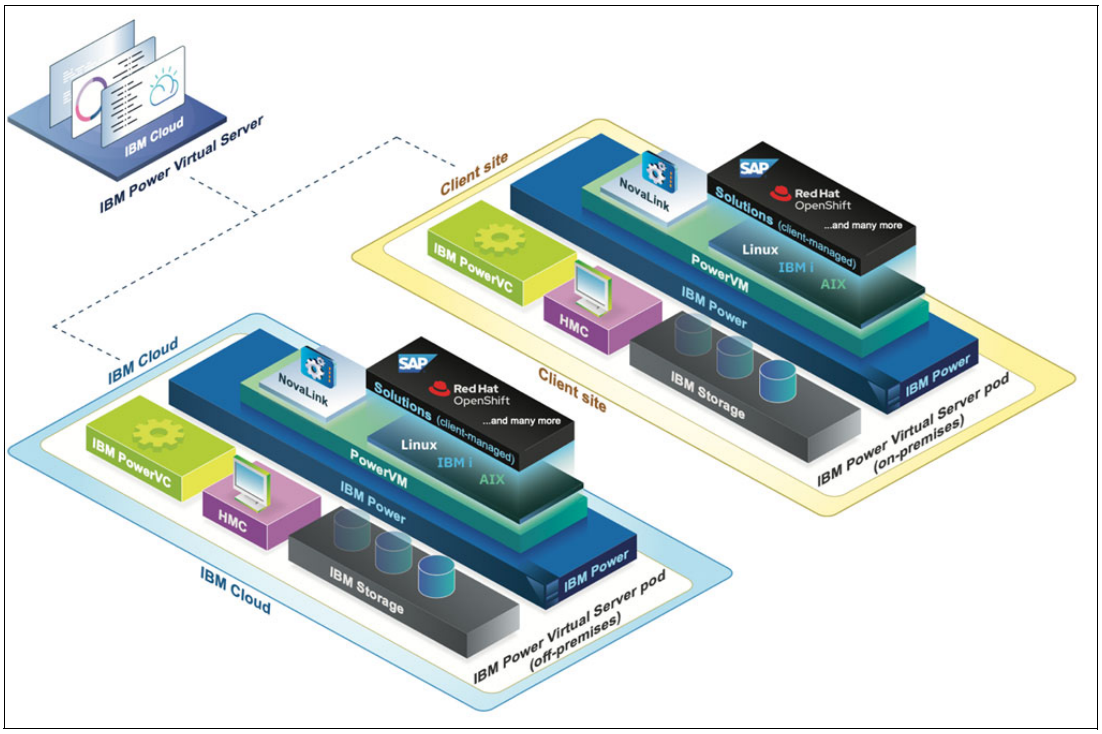

Figure 10-1 IBM PowerVS options

PowerVS Private Cloud is discussed in section 10.3.3, "PowerVS Private Cloud" on page 207.

Use cases for PowerVS

IBM Power Virtual Server (PowerVS) supports a wide range of use cases, particularly for enterprises that rely on IBM Power for mission-critical workloads. The following list provides some of the most common and impactful scenarios:

- Hybrid cloud modernization

  Organizations with on-premises IBM Power infrastructure can extend their environment to the cloud by using PowerVS. This enables a hybrid cloud model in which workloads can be moved or replicated to the cloud for scalability, disaster recovery, or testing, without requiring you to replatform or rewrite applications. This approach is especially useful for businesses running AIX or IBM i environments that want to modernize gradually.

- Disaster recovery and high availability

  PowerVS provides a reliable platform for disaster recovery (DR). You can replicate on-premises workloads to the cloud. After a failure, workloads can be quickly restored in the cloud, minimizing downtime. PowerVS also supports high availability configurations to ensure business continuity for critical applications.

- Development and testing environments

  You can use PowerVS to quickly create isolated environments for application development, testing, and quality assurance. This 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 ERP workloads

  IBM PowerVS is certified to run SAP HANA and other SAP ERP applications. Enterprises can migrate or extend their SAP environments to the cloud for improved scalability, performance, and cost-efficiency, especially when paired with IBM's high-performance Power10 processors.

- AI and machine learning

  PowerVS supports AI and machine learning workloads that require high compute power and memory bandwidth. It integrates with IBM Watson® and Red Hat OpenShift, allowing data scientists to build, train, and deploy models using familiar tools while using the performance of IBM Power.

- Database hosting

  Many organizations use PowerVS to host enterprise-grade databases such as IBM Db2, Oracle, or PostgreSQL. The platform offers high I/O throughput and reliability, making it ideal for transaction-heavy applications in industries, such as banking, retail, and logistics.

- Legacy application hosting

  For businesses with legacy applications running on AIX or IBM i, PowerVS provides a cloud-based alternative to aging on-premises hardware. This allows you to extend the life of critical applications while gradually modernizing your infrastructure.

- Compliance and data sovereignty

  Organizations in regulated industries (for example, healthcare, finance, and government) use PowerVS to meet strict compliance and data residency requirements. IBM offers regional data centers and compliance certifications to help meet these needs.

10.3.2 PowerVS in the Cloud

In this variation, IBM PowerVS resources reside in IBM data centers with dedicated networking and storage area network (SAN)-attached storage. You can choose from 21 IBM data centers and select the one that is closest to your users. IBM Power clients who rely on

private cloud infrastructure can now quickly and economically extend their Power IT resources to the cloud.

In these IBM data centers, PowerVS resources are separated from the rest of IBM Cloud servers by using separate networks and direct-attached storage. In this offering, you have the flexibility to choose from different IBM Power server models, an IBM-supplied operating system image, and multiple storage tiers based on your specific workload requirements. You can also choose your own image by using Bring Your Own License (BYOL) to bring preconfigured images from your enterprise infrastructure to Power Virtual Server. This offering also supports features such as shared processor pools, public connectivity for virtual machines (VMs), placement groups, and Global Replication Service (GRS), which you can use based on your requirements.

For more information about IBM PowerVS, see Getting Started with IBM Power Virtual Server.

10.3.3 PowerVS Private Cloud

Some clients want a cloud experience on IBM Power servers but must maintain control of their data and applications, whether those requirements are regulatory, security-based, or performance-based. For those clients, IBM designed an offering that brings the PowerVS solution to a data center of your choice.

PowerVS Private Cloud extends all the benefits of IBM PowerVS into your data center or a partner's data center. The enhanced capabilities of IBM PowerVS Private Cloud provide managed infrastructure as a service at client locations, with metered consumption and no upfront costs to support Hybrid by Design delivery of services.

The IBM PowerVS Private Cloud offering includes 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. In some cases, enterprises have workloads that are sensitive or have ultra-low latency requirements that are better served on site and in close proximity to other on-site workloads.

- Maintain customer data in region and specific geographies in the location of their choice

  Country sovereignty regulations require some data and workloads to stay in country. 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.

- Provide a hybrid cloud experience

  You can create a unified hybrid cloud landscape by integrating PowerVS running both at an IBM site and at your site, with the ability to manage all virtual machines (VMs) and infrastructure through a unified user interface. You gain flexibility by using an as-a-service model with intentional workload placement on and off premises.

- Deliver a predictable charging model with committed monthly spend combined with flexible consumption and metered usage-based pricing

  Both PowerVS offerings, off-premises at an IBM location and on-premises at your site, include compute, memory, storage, and operating system licenses that are fully metered by the hour. This model allows you to pay for what you use each month with no upfront payment.

- Streamline IT operations

  Whether in the cloud or at your site, IBM manages the infrastructure, freeing you to focus on business outcomes instead of infrastructure management. IBM owns, delivers, and sets up PowerVS in the data center of your choice and provides a fully managed solution, including monitoring, security, firmware updates, and infrastructure management.

- Provide enhanced security and control of data

  IBM PowerVS is designed to provide comprehensive security for IBM Power infrastructure by integrating with IBM Cloud tooling to manage security. This approach eliminates the need to manage Power infrastructure security and maintains sensitive data and workloads on-premises.

The physical infrastructure is delivered as a point of delivery (pod) that is deployed in your data center. A pod is the physical component that resides within the client data center and contains the compute, storage, and network components. A pod includes one or more racks where each component 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 (SREs) 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 your IBM Cloud account. This architecture enables you to scale your private cloud infrastructure horizontally by adding more pods to meet your workload requirements.

For more information about PowerVS Private Cloud, see the IBM Redpaper *Introduction to IBM Power Virtual Server Private Cloud*, REDP-5745.

10.3.4 Introducing IBM Power11 in the Cloud

Power11 is available in IBM PowerVS and IBM PowerVS Private Cloud. With the launch of IBM Power11, you and independent software vendors (ISVs) gain immediate access to the latest Power hardware in the cloud, which enables faster innovation and greater flexibility. Whether you are modernizing applications or enhancing business continuity, Power11 in PowerVS delivers a seamless hybrid cloud experience and includes several key benefits:

- Rapid deployment: Launch Power11 virtual servers in the cloud in under 10 minutes
- Accelerated modernization: Enable agile development and testing on the latest Power infrastructure
- Enhanced Business Resiliency: Strengthen continuity strategies with cloud-based Power11 resources
- Data sovereignty and compliance: Meet regulatory and sensitive data requirements with on-premises options
- Consistent hybrid cloud experience: Gain unified operations across on-premises and cloud environments.

Figure 10-2 on page 209 shows the initial implementation of Power11 processors in PowerVS on the day of announcement.

Power11 technology in PowerVS allows you to immediately take advantage of the new functionality without a requirement for initial capital investment. You can safely implement workloads that take advantage of the trusted security capabilities of Power11:

- Quantum-safe infrastructure compliance with optional workload cryptographic inventory discovery

- A unified cyber resiliency solution responding to evolving cyberthreats and regulatory standards with optional 3rd party services

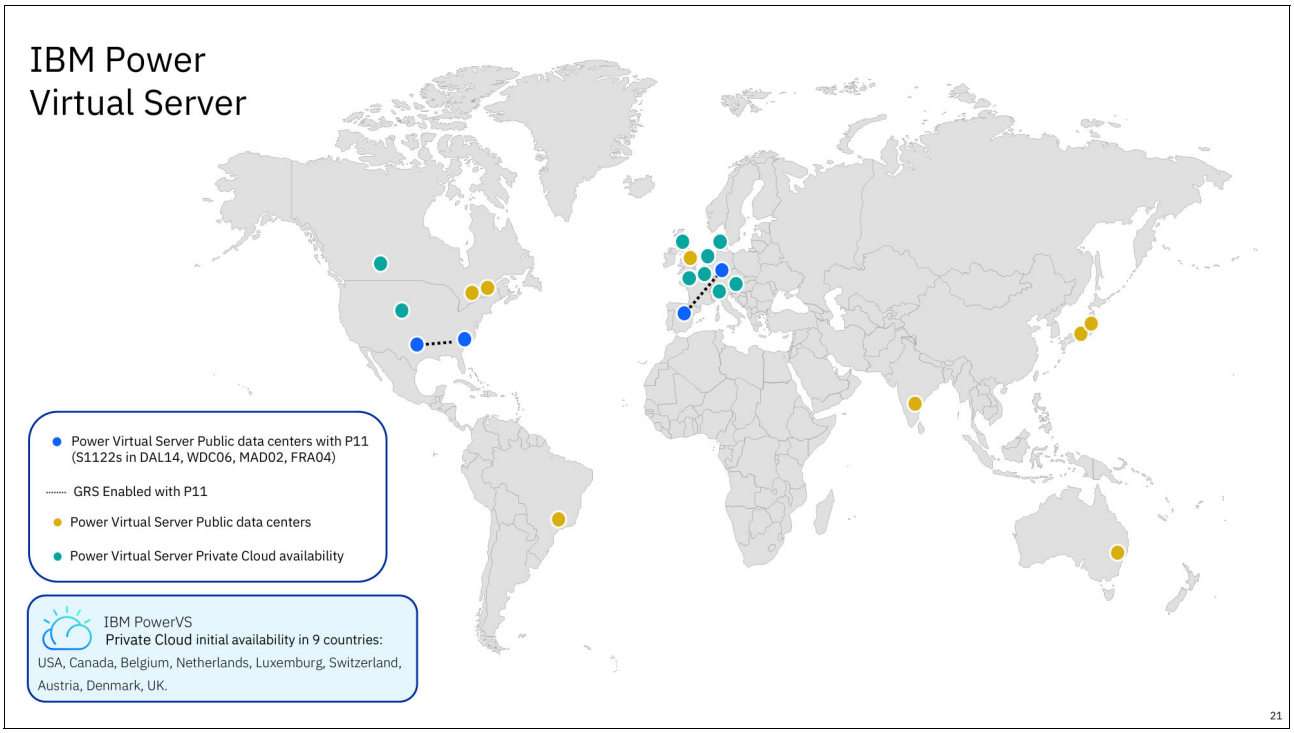

Figure 10-2 Availability of Power11 in Power Virtual Server

In addition, you can immediately take advantage of the improved performance of Power11 as shown in Figure 10-3.

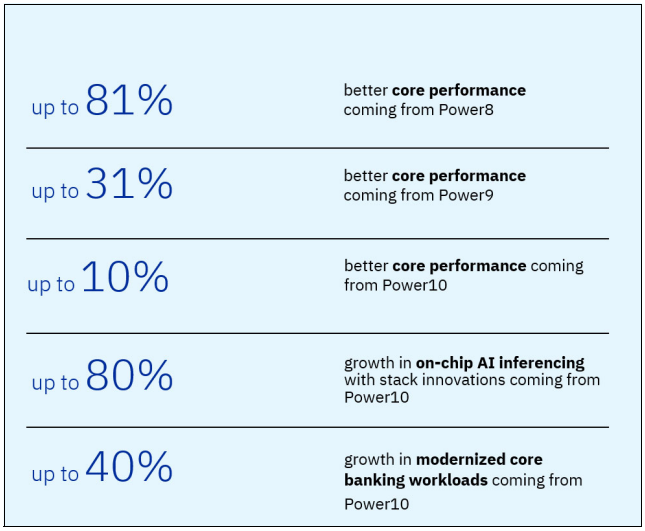

Figure 10-3 Potential performance benefits from Power in PowerVS

Abbreviations and acronyms

42U	42 EIA units
4U	4 EIA
ACF	Access Control File
AER	Autonomous error resolution
AES	Advanced Encryption Standard
AI	Artificial intelligence
AME	Active Memory Expansion
AML	Anti-money laundering
AMM	Active memory mirroring
AOC	Active optical cable
API	Application programming interface
ARI	analog rack interface
ASMI	Advanced System Management Interface
AWS	Amazon Web Services
BIKE	Bit Flipping Key Encapsulation
BMC	Baseboard management controller
BSC	Blind swap cassette
BYOL	Bring Your Own License
CCA	Common Cryptographic Architecture
CCIN	Customer Card Identification Number
CD	Continuous deployment
CLI	Command-line interface
CMC	Cloud Management Console
CRC	Cyclic redundancy check
CRM	Customer relationship management
CRU	Customer-replaceable unit
CTR	Counter
CVEs	Critical vulnerability exploits
CoD	Capacity on Demand
DCM	Dual Chip Module
DDIMM	Differential DIMM
DLPAR	Dynamic Logical Partitioning
DORA	Digital Operational Resilience Act
DPO	Dynamic Platform Optimizer
DR	Disaster recovery
DRO	Dynamic Resource Optimization
DWPD	Drive Writes Per Day
EAs	Effective addresses
ECC	Error detection code
EEH	Enhanced error handling
EEH	Extended error handling
EHR	Electronic health record
EMEA	Europe, Middle East, and Asia
EMR	Electronic medical record
EOL	End-of-life
EOS	End of Service
EOS	End-of-support
ERP	Enterprise resource planning
ESA	Electronic Service Agent
ESM	Enclosure Services Manager
ESM	Expansion Service Manager
ESR0	E1150 and the NED24
ESS	Entitled Systems Support
EUS	Extended update support
EoFS	End of Fix Support
FC	Feature Code
FFDC	First Failure Data Capture
FHE	Fully Homomorphic Encryption
FIR	Fault Isolation Register
FLRT	Fix Level Recommendation Tool
FRU	Field Replaceable Unit
FSP	Flexible Service Processor
GLVM	Geographic Logical Volume Manager
GPLv2	GNU General Public License version 2
GRS	GLOBAL replication service
GUI	graphical user interface
GenAI	GPUs in certain generative AI
HCL	HashiCorp Configuration Language
HIS	health care information system
HMC	Hardware management console
HNDL	Harvest Now, Decrypt Later
HPC	High-performance Computing
HPTs	Hardware page tables
HQC	Hamming Quasi-Cyclic
HSMs	Hardware security modules
IASP	Independent auxiliary storage pool

IBM	International Business Machines Corporation
IPL	initial program load
IPMI	Intelligent Platform Management Interface
ISA	Instruction set architecture
ISC	Integrated stack capacitor
ISV	Independent software vendor
ITOps	IT operations
IaC	Infrastructure as Code
IaaS	infrastructure as a service
KVM	Kernel-based Virtual Machine
KVM	Keyboard/Video/Mouse
LA	Latin America
LAN	local area network
LDAP	Lightweight directory access protocol
LLMs	Large language models
LPAR	logical partition
LPM	Live Partition Mobility
LWE	Learning With Errors
MCU	Memory Controller Unit
MES	Miscellaneous Equipment Specification
ML	Machine learning
MMA	Matrix Math Assist
MMA	Matrix-Math Accelerator
MPIO	Multipath I/O
MSP	Mover Server Partition
MSS	Merkle Signature Scheme
MTTR	Mean Time to Resolution
NPIV	N_Port ID Virtualization
NTP	Network Time Protocol
OCC	On-Chip Controller
OCP	OpenShift Container Platform
ODB	Operational database
OLM	Operator Lifecycle Management
OMI	Open Memory Interface
OS	Operating system
PCIe	Peripheral Component Interconnect Express
PDUs	Power distribution units
PECs	PCI Express controllers
PEP	Power Enterprise Pools
PFA	Predictive failure analysis
PHBs	PCI host bridges
PHYP	Power Hypervisor
PQC	Post-Quantum Cryptography
PSU	Power supply unit
PTFs	Program Temporary Fixes
PoC	Proof of concept
QSE	Quantum-safe encryption
RAG	Retrieval-augmented generation
RAS	Reliability, Availability, and Serviceability
RBAC	Role-based access control
RDHX	Rear Door Heat Exchanger
REST	Representational State Transfer
RGA	Resource Groups Advisor
RHCOS	Red Hat Enterprise Linux CoreOS
RHEL	Red Hat Enterprise Linux
RHOAI	Red Hat OpenShift AI
RHOCP	Red Hat OpenShift Container Platform
RMC	Resource Monitoring and Control
ROP	Return-oriented programming
SAN	Storage area network
SCM	Single Chip Module
SE2	Standard Edition 2
SEA	Shared Ethernet Adapter
SEC	Securities and Exchange Commission
SEL	System event log
SFPs	Supports the replacement of transceivers
SLA	Service-level agreement
SLES	SUSE Linux Enterprise Server
SMT	Simultaneous multithreading
SP	Service Pack
SPP	Shared processor pool
SPT	System Planning Tool
SPs	Service Packs
SREs	Site reliability engineering
SR-IOV	Single Root I/O Virtualization
SSH	Secure Shell
SSIC	BM System Storage Interoperation Center
SSR	Support Service Representative
SUE	Special Uncorrectable Error

SUMA	Service Update Management Assistant
SWMA	Software Maintenance Agreement
SaaS	Software-as-a-service
TAM	The Technical Account Manager
TCEs	Translation control entities
TCO	Total Cost of Ownership
TDP	Thermal design point
TE	Trusted Execution
TL	Technology Levels
TME	Transparent Memory Encryption
TPM	Trusted platform module
UAK	update access key
UDNs	User Defined Networks
UOV	Unbalanced Oil and Vinegar
VIOS	Virtual I/O Server
VM	Virtual machine
VMI	Virtualization Management Interface
VMRM	VM Recovery Manager
VMs	Virtual machines
WOF	Workload Optimized Frequency
ZPD	Zero Planned Downtime
eBMC	Enterprise baseboard management controller
eSCM	Entry Single Core Module
iPDUs	intelligent PDUs
iSCSI	Internet Small Computer Systems Interface
vHMC	Virtual HMC appliance

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Modernization Techniques for IBM Power*, SG24-8582
- *IBM i 7.6 features and function*, SG24-8588
- *IBM Power Security Catalog*, SG24-8568
- *Creating Red Hat OpenShift Multiple Architecture Clusters with IBM Power*, SG24-8565
- *Introduction to IBM Power Virtual Server Private Cloud*, REDP-5745
- *Using Ansible for Automation in IBM Power Environments*, SG24-8551
- *IBM PowerHA SystemMirror and IBM VM Recovery Manager Solutions Updates*, REDP-5694

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM Power

  https://www.ibm.com/products/power
- IBM Documents Power11

  https://www.ibm.com/docs/en/power11
- IBM Power11 Raises the Bar for Enterprise IT

  https://newsroom.ibm.com/2025-07-08-ibm-power11-raises-the-bar-for-enterprise-it
- Maximize IBM Power11 with IBM Technology Lifecycle Services

  https://www.ibm.com/new/announcements/maximize-ibm-power-11-with-ibm-technology-lifecycle-services
- Linux on Power

  https://www.ibm.com/products/power/linux
- AIX

  https://www.ibm.com/products/aix

- IBM i

  https://www.ibm.com/products/ibm-i
- Red Hat OpenShift AI: Supported Configurations

  https://access.redhat.com/articles/rhoai-supported-configs
- 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 Self-Managed 2.20

  https://docs.redhat.com/en/documentation/red_hat_openshift_ai_self-managed/2.20
- Red Hat OpenShift AI Service Definition

  https://access.redhat.com/articles/7102731
- OpenShift Operator Life Cycles

  https://access.redhat.com/support/policy/updates/openshift_operators
- RHEL Versions Utilized by RHEL CoreOS and OCP

  https://access.redhat.com/articles/6907891
- 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

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8589-00

ISBN 0738462268

Printed in U.S.A.